

JONATHAN A. PARKER

Princeton University

The Consumption Risk of the Stock Market

OVER THE PAST century in the United States, the average annual return on the stock market has exceeded that on short-term government bonds by 6 percentage points. The natural economic explanation for the premium on equity is the greater risks associated with investing in the stock market. However, the large premium that we observe cannot be explained by the canonical, consumption-based asset pricing model. Risk is best measured as the extent to which a return alters marginal utility. Since marginal utility is closely related to consumption, and consumption moves little with returns, the measured risk of the stock market is small.¹

One common informal interpretation of this equity premium puzzle is that stocks are a good deal. In this view, the model is taken as a reasonably accurate description of optimal behavior and a poor description of actual behavior. This normative view of the model and the data implies that households should increase their holdings of equity and even borrow to invest in the stock market.² Such thinking has also entered important areas

I thank Pierre-Olivier Gourinchas, Christian Julliard, David Laibson, Monika Piazzesi, Bruce Preston, Christopher Sims, Nicholas Souleles, Annette Vissing-Jørgensen, Mark Watson, and Michael Woodford for helpful discussions. I am also grateful to the discussants and participants in the Brookings Panel for their comments. Motohiro Yogo and Christian Julliard provided excellent research assistance in the summers of 2000 and 2001, respectively. Financial support was provided by National Science Foundation grant SES-0096076 and by a National Bureau of Economic Research Aging and Health Economics Fellowship through National Institute on Aging grant number T32 AG00186.

1. See Grossman and Shiller (1981), Shiller (1982), and Mehra and Prescott (1985).

2. This view is endorsed in such widely cited books as Glassman and Hassett (1999).

of public policy, most notably in proposals to allow funds from the Social Security system—whether the current \$1 trillion surplus in the trust fund or the entire \$10 trillion in implicit liabilities—to be invested in the stock market rather than entirely in government bonds as is currently done.

The positive view of the equity premium puzzle is that we simply do not understand asset prices. Since the puzzle was discovered, economists' efforts to find a model that rationalizes the premium have yielded little success. That is, there is as yet no model of a household investment problem with reasonable levels of risk aversion that explains the variation in returns over time, and the difference in returns between stocks and bonds in particular. This leaves economists largely unable to model investment behavior and largely unable to provide policymakers with guidance for the diversification of the Social Security system.

This paper proposes an understanding of the risk-return trade-off between stocks and bonds that departs from the canonical model in two ways. First, I ignore many issues in asset pricing and focus solely on the ultimate risk to consumption of a given portfolio choice. That is, rather than measure the risk to consumption as the contemporaneous response of consumption to returns on the stock market, this paper measures the risk as the medium-term impact of stock market returns on consumption.³ Second, in addition to studying the medium-term risk of equity as measured by aggregate consumption data, I follow Gregory Mankiw and Stephen Zeldes and ask whether the risk of equity justifies its return for the subset of households that hold equity.⁴ The main finding of the paper is that the medium-term risk of equity is much greater than the contemporaneous risk, both for the representative household and for the representative stockholder. For households that hold equity, the medium-term risk is largely sufficient to justify the high relative return of equity.

Measuring the risk of equity as the medium-term impact of a return on consumption has several appealing features. First, this approach maintains the assumption that the primary determinant of utility is the level of consumption. This assumption is intuitive and has proved useful and successful in many other branches of economics. Second, this approach is consistent with the theory of portfolio choice in that the medium-term

3. This approach is also pursued in contemporaneous work by Gabaix and Laibson (forthcoming).

4. Mankiw and Zeldes (1991).

risk and the contemporaneous risk should be approximately the same according to the canonical model.

Most important, the medium-term risk is a better measure of the true risk of the stock market under a wide class of extant models used in the study of household consumption and saving. If consumption responds with a lag to changes in wealth, then the contemporaneous covariance of consumption and wealth understates the risk of equity, and the medium-term risk provides the correct measure. This slow adjustment is a well-documented feature of consumption data: consumption displays excess smoothness in response to wealth shocks, a result that predates the equity premium puzzle.⁵ Existing explanations for the slow adjustment of consumption include direct costs of adjusting consumption, nonseparability of the marginal utility of consumption from factors such as hours worked that themselves adjust only slowly, constraints on borrowing or changes in risk that hinder consumption smoothing, and constraints on information flow or calculation such that household behavior is “near-rational.” The common feature of these models is that consumption responds slowly to an unanticipated change in stock market wealth, so that only after some time is the full impact observed in the movement of consumption.

The medium-term risk provides a robust measure of the risk of equity in that it allows us to remain to some extent agnostic about the particular optimization problem faced by households. This robustness feature is valuable because the correct model of household saving and portfolio choice has to date escaped discovery. But robustness comes at the price of not understanding the time variation in the process for stock returns, which has been the focus of much recent research on the equity premium. Instead, this paper addresses the question of whether stocks, given their riskiness, provide reasonable or exceptional returns within the context of a wide class of models. This approach leads to three main findings.

First, the contemporaneous covariance of consumption and stock market returns is misleading. The medium-term risk of equity is close to an order of magnitude greater than that implied by the contemporaneous covariance. This result is related to the fact that the consumption-based

5. The slow adjustment of consumption has a long history starting with Flavin (1981) and Hall and Mishkin (1982). See the surveys in Deaton (1992) and Browning and Lusardi (1996). Three recent papers that explore the consumption response to the stock market more generally are Parker (1999b), Ludvigson and Steindel (1999), and Dynan and Maki (2001).

capital asset pricing model performs better at long horizons.⁶ This finding is consistent with the slow adjustment of consumption to a change in wealth, and inconsistent with the textbook model in which consumption adjusts instantaneously to a return.

Second, even when measured after consumption has adjusted to a return, the consumption risk of equity remains insufficient to justify the return. Risk aversion estimated from the medium-term risk, despite being much lower than risk aversion estimated from the contemporaneous risk, is still implausibly high and so rejects the set of assumptions embodied in this measure. For the typical household, consumption rises by less than 1 percent over a horizon of two years following a 10 percent innovation in stock market prices. Estimated levels of medium-term risk imply that the coefficient of relative risk aversion for the representative household would have to be around 40 to rationalize the equity premium. This is true whether consumption is measured as flow consumption or as total consumption expenditure.

The third finding, and the main result of this paper, is that the medium-term consumption risk of equity and the return of equity are consistent with reasonable levels of risk aversion for those households that hold stock directly. The marginal investment decision that determines the risk-return trade-off in equilibrium may not be faced by households that do not hold equity. Mankiw and Zeldes and subsequent papers investigate the contemporaneous risk of equity for stockholders and conclude that the consumption of stockholders covaries more with returns than does the consumption of the typical household.⁷ Using data from the Consumer Expenditure Survey (CEX) of the Bureau of Labor Statistics, I show that the covariance of asset returns and the consumption growth of stockholders over periods from one and a half to two years is close to two orders of magnitude greater than the contemporaneous covariance in aggregate data on consumption. Although these results are subject to considerable statistical uncertainty, a reasonable decomposition is that half of this increase comes from measuring medium-term risk rather than contemporaneous risk, and half from measuring the risk for stockholders using the CEX rather than the risk for the representative agent using the National Income

6. See Brainard, Nelson, and Shapiro (1991), Cochrane and Hansen (1992), Daniel and Marshall (1997), and Piazzesi (forthcoming).

7. See Vissing-Jørgensen (1998), Attanasio, Banks, and Tanner (1998), and Brav, Constantinides, and Geczy (1999).

and Product Accounts (NIPA). With respect to this second step, there are several differences between the CEX consumption data on stockholders and the NIPA consumption data, but the data do not allow a clear decomposition of the increase in measured risk among these differences.⁸ It is clear, however, that although stockholders do face more equity risk, the method of aggregation and the population covered in the NIPA data seem to be important reasons for the lower measured contemporaneous risk of equity in the NIPA data.

Are the levels of risk aversion implied by average returns and these levels of risk plausible? In the first set of analyses using the CEX data, the risk aversion coefficient for stockholders is estimated to be around 10 to 20. But the period covered by the CEX is one of unusually high returns in many years. When adjustment is made for the fact that aggregate measures show an unusually low covariance between returns and consumption in this period, the return on equity and the medium-term risk of equity are more plausible. For the preferred measure of consumption, point estimates of the risk aversion of the typical stockholder lie between 4 and 8.⁹

Since the consumption risk of equity for the typical household is significantly less than that faced by households holding equity, stocks are in some ways a good deal for many households. That is, for the typical household not now invested in stocks, the expected return on an investment in equity at the margin outweighs the risks to consumption. The final section of the paper returns to the implications of these findings and the questions raised by limited participation in the stock market.

Measuring the Riskiness of Stocks

Are households allocating their wealth optimally between stocks and bonds? Or are stocks an undiscovered bargain for most households? The canonical model for addressing these questions assumes that all house-

8. In part this is because the order in which one takes the steps from one series to the other affects the relative importance of each difference, and in part it is because, consistent with substantial statistical uncertainty, the importance of different sources is not robust to small changes in data construction.

9. This finding is closely related to the finding that the risk of equity for the consumption of luxury goods is sufficient to rationalize the average premium on equity. See Ait-Sahalia, Parker, and Yogo (2001).

holds seek to maximize the expected present discounted value of utility flows from consumption. Because consumption has diminishing marginal benefits in any period, households want similar levels of consumption over time and over future events. When households are optimally allocating their wealth to consumption and different forms of saving, an extra dollar invested in stocks instead of bonds increases the future consumption level that the household expects on average, and this increase is exactly offset by the increased risk to future consumption that the extra dollar invested in stock brings. The puzzle is that if household utility increases with consumption, and only with consumption in the present, this optimization condition is far from being met in observed data on consumption and returns. At reasonable levels of risk aversion, the risk of equity is very small relative to its return.

To put the question more formally, assume that the representative household has a utility function of the constant-relative-risk-aversion form with a coefficient of relative risk aversion γ and the opportunity to allocate savings between a risk-free asset and a portfolio of equities that earns the return on the stock market. The stated optimality condition is

$$(1) \quad E_t[C_{t+1}^{-\gamma}(1 + z_{t,t+1})] - E_t[C_{t+1}^{-\gamma}](1 + r_{t,t+1}^f) = 0,$$

where C is aggregate consumption per capita, $r_{t,t+1}^f$ is the risk-free real interest rate in the economy between periods t and $t + 1$, and $z_{t,t+1}$ is the return on stocks between t and $t + 1$. A dollar invested in equity increases utility in period $t + 1$ by the gross payoff of the asset in dollars, $1 + z_{t,t+1}$ or $1 + r_{t,t+1}^f$, times the marginal utility of a dollar, $C_{t+1}^{-\gamma}$. When the investment decision is optimal, the expected increase in utility from investing one dollar more in stocks is exactly offset by the expected decline in utility from investing one dollar less in the risk-free asset.

Following Lars Peter Hansen and Kenneth Singleton,¹⁰ assume as an approximation that returns and consumption growth are jointly distributed log-normally conditional on information available in period t , so that equation 1 can be rewritten as

$$(2) \quad E_t[r_{t,t+1}] + \frac{1}{2} \text{var}_t(r_{t,t+1}) - \gamma \text{cov}_t(\Delta \ln C_{t+1}, r_{t,t+1}) = 0,$$

where $r_{t,t+1} = \ln(1 + z_{t,t+1} - r_{t,t+1}^f)$, the logarithm of the gross excess return of stocks over the risk-free rate. Taking the unconditional expectation of

10. Hansen and Singleton (1983).

equation 2 and reorganizing yields an equation that can be used to estimate the relative risk aversion of the representative agent:

$$(3) \quad \gamma = \frac{E[r_{t,t+1}] + \frac{1}{2} \text{var}_c(r_{t,t+1})}{\text{cov}_c(\Delta \ln C_{t+1}, r_{t,t+1})},$$

where the subscript c denotes the average conditional second moment: $\text{var}_c(r_{t,t+1}) \equiv E[\text{var}(r_{t,t+1})]$ and $\text{cov}_c(\Delta \ln C_{t+1}, r_{t,t+1}) \equiv E[\text{cov}_t(\Delta \ln C_{t+1}, r_{t,t+1})]$. Estimation of equation 3 requires calculation of conditional moments, and therefore a choice of conditioning information. An alternative is to assume that the joint unconditional distribution of consumption and returns is log normal, so that taking the unconditional expectation of equation 1 yields

$$(4) \quad \gamma = \frac{E[r_{t,t+1}] + \frac{1}{2} \text{var}(r_{t,t+1})}{\text{cov}(\Delta \ln C_{t+1}, r_{t,t+1})}.$$

The average equity premium in U.S. quarterly data from 1959 to 2000 is 0.0529 (5.29 percentage points) with an unconditional standard deviation of 0.1638.¹¹ These and all such numbers throughout the paper are reported at annual rates. The unconditional covariance of excess returns and the growth rate of real flow consumption over the same period is 0.00017. These levels of risk and return imply, from equation 4, that the risk aversion coefficient for the representative agent is an implausible 379. To put this number in perspective, a household this risk-averse would be willing to give up more than 24 percent of its consumption to avoid an even-odds gamble in which it would either win or lose 25 percent of its consumption. Even with a risk aversion coefficient as low as 10, a household would choose a 19 percent sure decline in consumption over this gamble. Levels of risk aversion around 4 are widely considered plausible, and levels above 10 highly implausible.¹²

11. The data and calculations in this paragraph are described in more detail below. The result is standard and is representative of any reasonable use of the data.

12. See, for example, the discussion in Mehra and Prescott (1985) and the thought experiment in Mankiw and Zeldes (1991, p. 105). Further, large levels of risk aversion pose other problems for the model, most notably in matching the observed risk-free real interest rate. See Weil (1989).

We can restate the puzzle by assuming a reasonable level of risk aversion and asking what the observed covariance of consumption and returns implies for the equity premium. If the typical household's risk aversion coefficient were 4, the household would be indifferent at the margin between stocks and bonds only if $E[r_{t,t+1}] + \frac{1}{2} \text{var}(r_{t,t+1}) = 0.00069$, or less than one-tenth of 1 percent. Because stocks return far more, from the perspective of this basic model, stocks appear to be an amazingly good deal.

Moving beyond the canonical model, first note that the risk of stocks, the premium on equity, and the intertemporal allocation of consumption are all jointly evaluated by equations 4 and 3. This is at odds with the literature on the intertemporal allocation of consumption, which is in wide agreement that the simple, textbook model of a representative consumer is false.¹³ In particular, the empirical literature studying consumption and saving behavior suggests that the following assumptions are at least questionable and at worst quite misleading: first, that utility is additively separable from factors that adjust slowly and covary with returns, such as leisure; second, that uninsurable idiosyncratic risk and borrowing constraints are not important; third, that consumption can be instantaneously adjusted or, if there are adjustment costs on some items, such as durable goods, the utilities derived from these categories are additively separable from the utility of other consumption; fourth, that aggregate consumption data accurately measure movements in flow consumption; and fifth, that households perfectly optimize without informational or calculation constraints.¹⁴ These findings have not escaped the notice of the literature on asset pricing. But ever since the discovery of the equity premium puzzle, macroeconomists have struggled to understand this combination of high stock returns and low stock risk, to little avail. There is currently no empirically reasonable, structural model of household behavior that matches the facts.¹⁵

13. This literature is not in agreement about the explanation for this failure, but this literature has the problem of too many models fitting the time-series data on consumption and risk-free returns rather than none.

14. On the first assumption see Attanasio and Weber (1995) and Basu and Kimball (2000); on the second see Zeldes (1989), Caballero (1990), Carroll (1997), and Gourinchas and Parker (forthcoming); on the third see Grossman and Laroque (1990), Caballero (1993), Ogaki and Reinhart (1998), Attanasio and Weber (1995), and Baxter and Jermann (1999); on the fourth see Attanasio and Weber (1993) and Wilcox (1992); and on the fifth see Caballero (1995), Parker (1999a), Souleles (forthcoming), and Sims (2001).

15. I do not discuss these models, because good surveys already exist—see Campbell (1999) and Kocherlakota (1996)—and because discussing any one model opens a Pandora's

It is worth noting that finance theory also does not provide an understanding of aggregate stock returns and risks. Modern finance theory prices assets from the assumption of no arbitrage, which means that risky assets can be priced only in reference to other risky or nonrisky assets. The return on the aggregate portfolio of stocks can only be deemed reasonable or unreasonable relative to the prices of a set of assets that span the returns on this portfolio, that is, relative to the prices of the assets in this portfolio.¹⁶

This paper evaluates the risk-return trade-off between stocks and bonds by focusing on the medium-term risk to consumption rather than searching for the correct stochastic discount factor to be used to price assets. Rather than measure the risk to consumption from the contemporaneous co-movement of consumption and returns, I measure the risk to consumption by the response of consumption to a return over a longer horizon, as given by

$$(5) \quad \text{cov}_c \left[\ln \left(\frac{C_{t+1+S}}{C_t} \right), r_{t,t+1} \right] \text{ or } \text{cov} \left[\ln \left(\frac{C_{t+1+S}}{C_t} \right), r_{t,t+1} \right].$$

Consumption risk is measured by the covariance of the excess return of stocks at $t + 1$ and the change in consumption from t to $t + 1 + S$, where S is the horizon over which the consumption response is studied. To judge whether the medium-term risk is sufficient to rationalize the equity premium, I calculate the level of risk aversion of the representative agent implied by each measure. Paralleling much of the literature on the equity

box. That said, perhaps the leading candidate for a successful model is that based on habit formation (Constantinides, 1990). Many versions of the habit formation model now exist, but all have at least one of three major shortcomings. First, many models (for example, Campbell and Cochrane, 1999) still require extremely high risk aversion coefficients to match the data. Second, models constructed to fit some aspects of asset data are rejected when tested in other contexts (see Dynan, 2000, and Otrok, Ravikumar, and Whiteman, 2001). Third, these models suggest that the level of consumption in the medium or the long term is largely irrelevant for welfare, a difficult assumption to square with results and successful models in most other areas of economics.

16. Factor models that omit asset prices can price assets without reference to other assets. Although many observed factors that predict time variation in returns, like those of Lettau and Ludvigson (2001), Lamont (1998), and Campbell (1987), are based on asset prices, factor models lack the structure necessary to judge the riskiness of stocks or the sensibility of implied preferences. Factor models provide clues about the structure of the correct model but cannot address the question at hand.

premium puzzle, I calculate risk aversion from the unconditional moments as

$$(6) \quad \gamma_s = \frac{E[r_{t,t+1}] + \frac{1}{2} \text{var}(r_{t,t+1})}{\text{cov}\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right), r_{t,t+1}\right]},$$

where γ_s denotes risk aversion based on consumption growth over a horizon of S periods. I also present estimates from the same equation with the conditional covariance in the denominator. When $S = 0$, one recovers the usual estimate, as given by equation 4 for the unconditional case.

Why evaluate risk using the medium-term risk, as in equations 5 and 6? If households choose their portfolio at time t , and the impact of this choice and the realized return on stocks take time to appear in the consumption data, then this measure provides a better measure of the risk of stocks than does equation 3. To be more concrete, I show that the medium-term risk of equity would be approximately correct if the textbook model were true, or if the marginal utility of consumption is shifted by a stationary variable that covaries with returns. The medium-term risk is also a better measure when constraints on information flow slow consumption movements. Finally, aggregate consumption is constructed using some lagged data. The medium-term risk correctly measures consumption risk in the presence of this measurement error.

First consider the textbook model. Households seek to smooth consumption over time, which is captured by the following consumption Euler equation for the risk-free rate between $t + 1$ and $t + 1 + S$:

$$(7) \quad E_{t+1} \left[\frac{\beta^S (1 + r_{t+1,t+1+S}^f) C_{t+1+S}^{-\gamma}}{C_{t+1}^{-\gamma}} \right] = 1,$$

which implies

$$(8) \quad C_{t+1}^{-\gamma} = \beta^S (1 + r_{t+1,t+1+S}^f) C_{t+1+S}^{-\gamma} - \varepsilon_{t+2,t+1+S},$$

where β is the factor by which households discount the future, and $\varepsilon_{t+2,t+1+S}$ is the expectation error between actual marginal utility in $t + 1 + S$ and its expected value in $t + 1$. Substituting into equation 1 and noting that the expectation error has a mean of zero and is uncorrelated with information known at time $t + 1$ yields

$$(9) \quad E_t[(1 + r_{t+1,t+1+S}^f) C_{t+1+S}^{-\gamma} Z_{t,t+1}] - E_t[(1 + r_{t+1,t+1+S}^f) C_{t+1+S}^{-\gamma}](1 + r_{t,t+1}^f) = 0.$$

Under the assumption that consumption and returns are unconditionally distributed joint normally, risk aversion is equal to¹⁷

$$(10) \quad \gamma_s = \frac{E[r_{t,t+1}] + \frac{1}{2} \text{var}(r_{t,t+1}) + \text{cov}[\ln(1 + r_{t+1,t+1+S}^f), r_{t,t+1}]}{\text{cov}\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right), r_{t,t+1}\right]}.$$

The extent to which risk aversion calculated directly from the medium-term risk as in equation 6 differs from that calculated from the textbook model depends on the extent to which an innovation to returns leads to a change in future risk-free rates. Intuitively, according to the textbook model, if an innovation to returns were to lead to a significant revision in planned intertemporal substitution in consumption over the next S periods, then looking S periods out could be quite misleading. If the only reason for consumption to diverge from planned consumption between periods $t + 1$ and $t + 1 + S$ were future innovations, the medium-term risk would exactly measure consumption risk.

The size of this additional term is negligible. An upper bound on this covariance is given by the average unconditional covariance of the cumulative return from one-period risk-free rates. This upper bound is two orders of magnitude less than the equity premium at an annual rate, ranging from -0.0002 at a horizon of one quarter to -0.0008 at a horizon of ten quarters. Using equation 10 instead of equation 6 throughout the analysis would lower estimates of risk aversion by 1 percent. In sum, if the textbook model were true, estimates of risk aversion from the medium-term risk and from the contemporaneous risk would be very close.

Turning next to more general models in which the medium-term risk of equity provides an accurate characterization of the risk of equity, suppose that the marginal utility of the representative agent is shifted by a variable, Ψ_t , so that marginal utility is given by $C_{t+1}^{-\gamma} \Psi_{t+1}$. Suppose further that this variable has a stationary distribution, so that as S gets large, the

17. This derivation and those in the rest of this section are done for the case of unconditional moments only. However, derivations of risk aversion in terms of conditional moments are completely analogous.

distribution of Ψ_{t+1+S} from the perspective of period t is its unconditional distribution, $F(x)$. That is, for large S , $\Psi_{t+1+S}|t \sim F(x)$. The factor Ψ_{t+1} captures many of the extant models of slow adjustment of consumption listed above. This factor can represent transitory movement, following a market return, in the share of hours devoted to market work, in the relative productivity of home production, in the stock of durable goods relative to flow consumption, in individual consumption risk, in the cross-sectional distribution of marginal utilities, and so forth. A model in which marginal utility adjusts at the time of the return, but in which a stationary confounding variable implies that the contemporaneous change in consumption does not accurately measure this change, has the property that, for large enough S , the medium-term risk measures the consumption risk of equity.

Using the same derivation as for the textbook model, it follows that risk aversion is given by

$$(11) \quad \gamma_s = \frac{E[r_{t,t+1}] + \frac{1}{2} \text{var}(r_{t,t+1}) + \text{cov}[\ln(1 + r_{t+1,t+1+S}^f), r_{t,t+1}] + \text{cov}[\ln \Psi_{t+1+S}, r_{t,t+1}]}{\text{cov}\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right), r_{t,t+1}\right]}$$

$$(12) \quad = \frac{E[r_{t,t+1}] + \frac{1}{2} \text{var}(r_{t,t+1}) + \text{cov}[\ln(1 + r_{t+1,t+1+S}^f), r_{t,t+1}]}{\text{cov}\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right), r_{t,t+1}\right]}$$

since the realization of Ψ_{t+1+S} is independent of information available at time $t + 1$. Thus the medium-term risk is as valid a measure of consumption risk in the presence of stationary utility shifters as it is for the textbook model.

The medium-term risk is also useful for measuring the risk of equity for some models in which households face constraints on information, calculation, or adjustment of consumption so that consumption and marginal utility move only slowly to the new optimal level following a shock. As a specific example, Xavier Gabaix and David Laibson add to the canonical model of Robert Merton the assumption that households face costs of monitoring their portfolio balances, and therefore check and learn their

wealth only infrequently, once every D periods.¹⁸ Each time a household learns its wealth, it adjusts its consumption in response to all the market returns during the interval since it last did so. Assuming that a constant measure of households are learning their balances and adjusting their consumption at every instant, aggregate consumption adjusts smoothly and slowly over D periods to reflect a given return. To evaluate the risk of equity then requires studying the medium-term risk.

A final reason to evaluate the risk of equity with the medium-term risk is that aggregate consumption data may measure consumption responses with delay, even if the true consumption response is instantaneous. As demonstrated by David Wilcox,¹⁹ serially correlated measurement error is introduced into aggregate consumption data by sampling error, by the imputation procedures used, and by definitional difficulties involved in constructing measures of aggregate consumption from monthly survey data on retail sales.

Measuring Aggregate Consumption and Returns

This section describes the important issues that arise in the use of consumption and returns data to measure the medium-term risk of equity. The appendix contains a complete description of the data used.

Consumption is measured in two different ways. First, flow consumption is defined as real consumption expenditure per capita on nondurable goods and services less expenditure on education services, medical care services, housing services, personal business services, and footwear. This measure of consumption corresponds closely to the theoretical concept but requires ignoring the remainder of goods and services that households purchase. Theoretically, this can be justified if these other goods enter utility in an additively separable manner, but this is unlikely to be true. In fact, flow consumption overrepresents goods that are necessities, such as food, and underrepresents those that are luxury goods, such as household appliances, medical care, jewelry, and electronics. This is a concern in measuring the impact of the stock market on consumption because the response of expenditure should fall more heavily on luxury

18. Gabaix and Laibson (forthcoming). See also Lynch (1996) and Marshall and Parekh (1999).

19. Wilcox (1992).

goods than on necessities. The typical household reduces its expenditure on luxury goods by more than on necessities when its wealth declines. This problem is severe for studying the consumption response to equity because stockholding is heavily concentrated among wealthier households, for whom a larger share of expenditure, and of expenditure variability, is luxury consumption.²⁰

To partly deal with this concern, the analysis is also conducted defining total consumption expenditure as total personal consumption expenditure less expenditure on education, medical care, and personal business expenses. This approach is atypical because utility comes from the service flow from durable goods, not from expenditure on durable goods. However, the medium-term risk provides the correct measure of consumption risk even when expenditure on durable goods is included. Suppose that utility comes also from the service flow from the stock of a durable good, K_t . The stock of the durable good is related to expenditure as

$$(13) \quad K_{t+1} = (1 - \delta) K_t + L_{t+1},$$

where δ is the rate at which the durable good depreciates. If there are no costs to households of adjusting the stock of durable goods that they hold, then expenditure (L_t) will be volatile as households increase or decrease expenditure to adjust their stock, whereas consumption will remain relatively stable. After the adjustment, the new level of expenditure will be proportional to the new level of the stock of durable goods. Using the canonical measure would underestimate risk aversion, since the contemporaneous covariance of expenditure growth and excess returns is large, whereas the actual covariance of consumption growth and returns is significantly lower.

If the growth rate of consumption is stationary, equation 13 implies that the stock of the good, and thus its service flow, is cointegrated with expenditure. To the extent that a large positive return leads to an upward revision in the stock of a durable good, this will still be apparent a few years later in higher expenditure, which is proportional to the higher consumption flow from the services of this durable good. This suggests looking at the medium-term response of expenditure to an excess return, which is exactly what this paper does. A possible complication is that households seem to face costs associated with adjusting their stocks of durable goods, and

20. See Ait-Sahalia, Parker, and Yogo (2001).

this delays the adjustment. However, as long as one looks at the response of expenditure after enough time has elapsed, the adjustment will be complete, and the change in total consumption will measure the change in utility from both the service flow from the stock of durable goods and the purchases of nondurable goods and services.

Both consumption series are deflated by chain-weighted price indexes constructed from published series. Data cover the first quarter of 1959 to the fourth quarter of 2000. Excess returns are calculated over the same period as the difference between the Fama risk-free rate and the return on the New York Stock Exchange and the American Stock Exchange composite indexes. The excess return dated $t + 1$ is the excess return during the period $t + 1$.

The Risk of Stocks for the Representative Household

This section uses aggregate consumption per capita to measure the medium-term risk of equity. Paralleling the previous literature on the equity premium puzzle, I first estimate unconditional covariances of excess returns and consumption growth. Second, I analyze conditional risk both by estimating the response of consumption to innovations in returns in a vector autoregression and by calculating covariances using innovations to excess returns. The measures of risk aversion implied by conditional and unconditional covariances are used to determine whether the consumption risk of equity is sufficient to rationalize its average excess return.

The unconditional covariance of consumption growth and excess returns at horizon S is estimated as

$$(14) \quad \frac{1}{T-1-S} \sum_{t=1}^{T-1-S} \ln\left(\frac{C_{t+1+S}}{C_t}\right) \left(r_{t,t+1} - \frac{1}{T-1-S} \sum_{s=1}^{T-1-S} r_{s,s+1} \right),$$

for which the expected value is the covariance of interest. Standard errors are calculated as Newey-West standard errors, with the number of lags equal to the horizon plus one. The risk aversion implied by each covariance is calculated by replacing the theoretical moments in equation 6 with their empirical counterparts, and the associated statistical uncertainty is calculated using the delta method. The numerator of equation 6 is

estimated as 6.63 percent. Since the focus of this paper is on the consumption risk of equity, this number is taken as given when making inferences. Standard errors reported for risk aversion reflect only the uncertainty about the covariance of consumption and excess returns and not uncertainty about the mean excess return or its own variance. Because risk aversion is a nonlinear transformation of the covariance, the standard errors on risk aversion estimates are useful for inferring the statistical distance from a reasonable (low) level of risk aversion, and not useful for estimating the statistical difference from an extremely high level of risk aversion. The latter inference can be made by examining whether the covariance is statistically different from zero.

The first row of table 1 reports the unconditional covariances for each consumption measure and horizon and the risk aversion coefficients implied by these covariances; standard errors are reported below each estimate. As discussed in the previous section, the contemporaneous covariances are extremely low. Since the variance of returns is 0.027, the covariance implies that flow consumption growth is 0.06 percent above average when returns are 10 percent above normal. As the table also shows, such small contemporaneous covariances imply implausibly high levels of risk aversion. Interestingly, total consumption expenditure has a lower contemporaneous covariance with excess returns than does flow consumption, consistent with an important role for adjustment costs in the dynamics of expenditure on durable goods.

The remaining rows of table 1 show that the covariance of returns and consumption growth rises significantly with the horizon over which consumption growth is measured, although this rise is not monotonic. The covariance of the change in flow consumption over three years and excess returns is an order of magnitude larger than the contemporaneous covariance. For total expenditure the increase in covariance is more striking, but the majority of this increase occurs in moving from the contemporaneous covariance to a horizon of one quarter. After this, the covariance doubles as the horizon is increased. The impact of horizon on the risk aversion coefficient is the inverse of its impact on the covariance. As risk is measured over longer horizons, the implied levels of risk aversion decline by an order of magnitude or more for each series.

One reason this decline might be large is that the economy does not proceed in three-month-long units of time that neatly align with the quarters in which the data are measured. If the textbook model were correct and if

Table 1. Consumption Risk of Equity and Implied Risk Aversion, 1959–2000^a

<i>Horizon (quarters)</i>	<i>Flow consumption</i>		<i>Total consumption expenditure</i>	
	<i>Unconditional covariance^b</i>	<i>Implied risk aversion^c</i>	<i>Unconditional covariance^b</i>	<i>Implied risk aversion^c</i>
0	0.00017 (0.00016)	379.3 (342.0)	0.00001 (0.00021)	12,067.0 (453,259.3)
1	0.00080 (0.00031)	83.0 (32.5)	0.00088 (0.00037)	75.4 (31.7)
2	0.00104 (0.00043)	63.7 (26.1)	0.00151 (0.00053)	44.0 (15.5)
3	0.00112 (0.00051)	59.2 (27.0)	0.00161 (0.00064)	41.2 (16.3)
4	0.00155 (0.00064)	42.7 (17.6)	0.00207 (0.00075)	32.0 (11.6)
5	0.00170 (0.00070)	39.1 (16.2)	0.00210 (0.00082)	31.5 (12.3)
6	0.00186 (0.00074)	35.6 (14.0)	0.00212 (0.00086)	31.2 (12.7)
7	0.00198 (0.00084)	33.4 (14.1)	0.00215 (0.00099)	30.8 (14.2)
8	0.00163 (0.00096)	40.6 (23.9)	0.00188 (0.00114)	35.3 (21.5)
9	0.00129 (0.00103)	51.5 (41.4)	0.00159 (0.00124)	41.6 (32.2)
10	0.00167 (0.00125)	39.7 (29.7)	0.00194 (0.00147)	34.1 (25.7)
11	0.00175 (0.00145)	37.8 (31.4)	0.00171 (0.00170)	38.8 (38.6)

Source: Author's calculations. See appendix for data sources.

a. Standard errors are reported in parentheses; for covariance estimates they are calculated using the Newey-West procedure, and for risk aversion estimates they do not reflect uncertainty in the numerator and are calculated by the delta method.

b. Covariance of the logarithm of excess returns of stocks over the risk-free interest rate during time $t + 1$ and the logarithm of consumption growth from time t to $t + 1 + S$, where S is the horizon.

c. Calculated as the sum of the mean log excess return and one-half its variance, divided by their covariance.

consumption portfolio decisions were continuously reoptimized, the quarterly frequency of the data would imply that the contemporaneous covariance for flow consumption would be understated by a factor of two, given the definition of returns.²¹ But this adjustment would not alter the main conclusion that the medium-term risk is much greater than the contemporaneous risk for flow consumption. Given adjustment costs, it is unclear what the correct adjustment, if any, is for total expenditure.

Another possible concern with these results is that the declining estimates of risk aversion could be due to the fact that estimates using longer

21. I thank Gregory Mankiw for raising this point. See Grossman, Melino, and Shiller (1987).

horizons must omit more recent stock market data, for which consumption data $t + 1 + S$ periods later are not yet available. Thus estimates with longer horizons do not use data on some of the spectacular returns on equity in recent periods. This is in fact driving none of the results. Holding the sample of returns constant across all horizons—omitting the most recent eleven return observations in all calculations—makes very little difference. The conclusions are also the same if the estimates of the mean consumption growth and mean returns used in calculating the covariance of interest are held constant at their values calculated for the longest possible sample.

A less skeptical reaction to these results is that reasonable levels of risk aversion are not far from a 95 percent confidence interval surrounding these point estimates. One way to assess this argument is to estimate risk aversion from conditional moments, which is both a different cut of the data and potentially more precise. Empirically, the stochastic process of equity returns is such that after a series of particularly high returns, returns are on average lower, and vice versa. That is, at horizons of a few years, there is a negative correlation in returns. Although the predictable component is small relative to the uncertainty in excess returns, eliminating the predictability of returns from the covariance of interest may sharpen inference. However, as noted in the previous section, estimation using conditional moments requires specifying conditioning information. This is approached in two ways.

First, I calculate the covariance and risk aversion using the impulse responses to returns in a vector autoregression (VAR). This approach has the advantage of providing a clear picture of the consumption movement that follows an innovation in excess returns and the advantage (or weakness) of imposing a smooth response to the innovation. I estimate a three-variable VAR in excess returns $r_{t-1,t}$, the logarithm of consumption ($\ln C_t$), and the dividend-to-price ratio d_t/P_{t-1} , each with four lags. The dividend-to-price ratio is included because it is a good predictor of future returns. I take a Bayesian approach, asking the VAR to fit the unit root in consumption and treating parameters as random. Impulse responses and confidence intervals are constructed from the estimated parameters by Monte Carlo methods rather than inversion. The appendix contains additional details of the method.

Excess returns are ordered first so that the impulse responses measure the impact of an innovation to returns on consumption. It is important to note that, as for the covariance, the innovation to returns is not structural

but is an amalgam of structural shocks to the economy such as news about current labor income or future rates of return. Thus the estimated impulse response does not measure the marginal propensity to consume out of stock wealth. It does measure the medium-term risk.

Figures 1 and 2 display the impulse responses of each variable to an innovation in excess returns. Figure 1 displays the responses of all three variables for the system that includes flow consumption, and figure 2 does the same for the system that includes total expenditure. Both figures display point estimates of the impulse response functions as well as two-standard-error bands. Both systems show less than a 1 percent change in consumption in response to a 1-standard-deviation innovation in excess returns (which in each VAR is roughly 7.8 percent at a quarterly rate). Most of this movement occurs in the first few quarters following the innovation. In fact, the slope of the impulse response is not statistically different from zero at five quarters and beyond.²² Consumption rises by less than 0.1 percent for a 1 percent innovation in excess returns at this horizon.

The impulse response of consumption to an innovation in returns measures the medium-term risk to consumption. The covariance of interest (at a quarterly rate) is

$$(15) \quad \widehat{\text{cov}}_c \left[\ln \left(\frac{C_{t+5}}{C_t} \right), r_{t,t+1} \right] = \hat{\sigma}_c(r_{t,t+1}) \cdot IRF_5(\ln C),$$


where $IRF_5(\ln C)$ denotes the impulse response of the log of consumption at horizon 5 to an innovation to returns, and $\hat{\sigma}_c(r_{t,t+1})$ is the standard deviation of the innovation to returns in the VAR. To gauge whether the medium-term risk is enough to justify the average return on equity, I calculate the implied risk aversion of the representative household by substituting this covariance into equation 6.²³

Table 2 shows the point estimates of the medium-term risk and the implied coefficients of relative risk aversion estimated from each mea-

22. This statement is based upon confidence intervals in similar VARs in which consumption growth is included in place of the level of consumption.

23. One could also adjust the numerator so that the variance of returns is reduced by the extent to which they are one-step-ahead predictable. This is not done for three reasons: first, this adjustment in practice is small; second, the bias from not making this adjustment is toward making risk aversion larger; and finally, this makes clear that differences in risk aversion across horizon and method are due to differences in the estimated consumption risk of equity.


Figure 1. Impulse Responses to a One-Standard-Deviation Shock to Excess Returns, Using Flow Consumption^a



Source: Author's calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.

a. Shown with 95 percent confidence bands calculated by Monte Carlo methods from the estimated asymptotic distribution of the parameter vector.

Figure 2. Impulse Responses to a One-Standard-Deviation Shock to Excess Returns, Using Total Consumption Expenditure^a



Source: Author's calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
a. Shown with 95 percent confidence bands calculated by Monte Carlo methods from the estimated asymptotic distribution of the parameter vector.

Table 2. Consumption Risk of Equity and Implied Risk Aversion Estimated from a Vector Autoregression, 1959–2000^a

<i>Horizon (quarters)</i>	<i>Flow consumption</i>		<i>Total consumption expenditure</i>	
	<i>Conditional covariance^b</i>	<i>Implied risk aversion^c</i>	<i>Conditional covariance^b</i>	<i>Implied risk aversion^c</i>
0	0.00021 (0.00014)	323.1 (224.1)	0.00009 (0.00017)	757.9 (1,498.5)
1	0.00084 (0.00022)	78.9 (20.8)	0.00101 (0.00026)	65.8 (16.6)
2	0.00117 (0.00030)	56.6 (14.3)	0.00175 (0.00034)	37.9 (7.4)
3	0.00128 (0.00038)	51.6 (15.2)	0.00189 (0.00045)	35.1 (8.3)
4	0.00171 (0.00046)	38.9 (10.5)	0.00240 (0.00055)	27.6 (6.3)
5	0.00186 (0.00052)	35.7 (9.9)	0.00258 (0.00062)	25.7 (6.2)
6	0.00195 (0.00057)	34.0 (9.9)	0.00270 (0.00069)	24.6 (6.3)
7	0.00210 (0.00062)	31.5 (9.3)	0.00287 (0.00076)	23.1 (6.1)
8	0.00220 (0.00068)	30.1 (9.2)	0.00299 (0.00083)	22.2 (6.2)
9	0.00228 (0.00073)	29.1 (9.3)	0.00309 (0.00090)	21.5 (6.2)
10	0.00237 (0.00078)	27.9 (9.2)	0.00318 (0.00097)	20.9 (6.4)
11	0.00244 (0.00084)	27.1 (9.3)	0.00324 (0.00104)	20.5 (6.6)

Source: Author's calculations. See appendix for data sources.

a. Based on impulse responses from a VAR (see text and appendix for details). Standard errors are reported in parentheses; for covariance estimates they are bootstrapped from the estimated VAR coefficients, and for risk aversion estimates they do not reflect uncertainty in the numerator and are calculated by the delta method.

b. Estimated as the impulse response of log consumption at horizon S times the standard deviation of returns at an annual rate.

c. Calculated as the sum of the mean log excess return and one-half its variance, divided by their covariance.

sure of consumption. Just as for the unconditional covariance, measured consumption risk rises and estimated risk aversion declines as the horizon over which consumption risk is measured increases. At a horizon of a year, the risk aversion coefficient necessary to rationalize the equity premium is between 30 and 40. Estimates for horizons beyond three years (not shown) fall to between 20 and 25, although statistical uncertainty is substantial at this horizon.

Even these implausibly high estimates of risk aversion are on the low side of those implied by other estimates of the consumption response to

equity innovations in the consumption literature.²⁴ By estimating and inverting a VAR, a degree of smoothness is imposed that is not present in the estimates based on the medium-term risk in table 1. Therefore I also perform a similar exercise in which no smoothness is imposed on the covariances. First, innovations to returns are calculated from rolling regressions on returns, using the same predicting variables as in the VAR. (The appendix provides details.) By using rolling regressions, for the majority of the sample, the information set used to predict returns includes only data known to the agents at the time; the same cannot be said of the VAR, whose coefficients are estimated from the entire sample of data. Second, and to be clear, the conditional covariance is calculated as

$$(16) \quad \frac{1}{T-1-S} \sum_{t=1}^{T-1-S} \ln \left(\frac{C_{t+1+S}}{C_t} \right) \left(r_{t,t+1} - \hat{E}_t[r_{t,t+1}] \right),$$

where $\hat{E}_t[r_{t,t+1}]$ is the prediction of log excess returns from the rolling regression.

Table 3 presents the conditional covariances and implied estimates of risk aversion. These estimates are broadly comparable to those in tables 1 and 2: the consumption risk of equity rises by close to an order of magnitude or more as one considers risk over longer horizons. However, the conditional covariances do not rise as high as the earlier estimates, and at longer horizons they fall back significantly, although again sampling uncertainty grows substantially with the horizon.

Figure 3 summarizes the coefficients of relative risk aversion implied by each estimation method for horizons from zero to nineteen quarters (five years). Although measures of risk aversion estimated from the unconditional covariances, from the VAR, and from the conditional covariances are all in broad agreement at horizons up to two years, they diverge thereafter. At longer horizons, the smoothing imposed by a VAR alters estimated risk aversion by an order of magnitude or more. Tables 1 and 3 show why. Even at a horizon of eleven quarters, the estimated covariances are only one standard error away from zero. There is simply not enough information in the data to select among point estimates at horizons

24. Ludvigson and Steindel (1999) estimate a more structural VAR and find a long-run elasticity of consumption to equity returns of only 5 percent. Parker (1999b) also studies the consumption response to the stock market using survey data from the Panel Study of Income Dynamics and finds elasticities of less than 5 percent.

Table 3. Conditional Consumption Risk of Equity and Implied Risk Aversion, 1959–2000^a

<i>Horizon (quarters)</i>	<i>Flow consumption</i>		<i>Total consumption expenditure</i>	
	<i>Conditional covariance^b</i>	<i>Implied risk aversion^c</i>	<i>Conditional covariance^b</i>	<i>Implied risk aversion^c</i>
0	0.00020 (0.00015)	329.4 (238.5)	0.00005 (0.00020)	1,249.0 (4,715.7)
1	0.00089 (0.00030)	74.5 (25.4)	0.00090 (0.00037)	73.7 (30.5)
2	0.00114 (0.00040)	58.3 (20.5)	0.00155 (0.00050)	42.7 (13.7)
3	0.00113 (0.00051)	58.5 (26.5)	0.00149 (0.00063)	44.5 (18.7)
4	0.00147 (0.00066)	45.0 (20.0)	0.00184 (0.00075)	36.1 (14.8)
5	0.00153 (0.00079)	43.3 (22.3)	0.00175 (0.00092)	38.0 (20.0)
6	0.00166 (0.00097)	39.9 (23.3)	0.00171 (0.00113)	38.8 (25.5)
7	0.00173 (0.00110)	38.3 (24.4)	0.00167 (0.00131)	39.7 (31.1)
8	0.00136 (0.00130)	48.6 (46.4)	0.00135 (0.00153)	49.0 (55.6)
9	0.00077 (0.00143)	85.6 (157.5)	0.00081 (0.00164)	81.9 (166.2)
10	0.00127 (0.00171)	52.4 (70.7)	0.00135 (0.00196)	48.9 (71.0)
11	0.00139 (0.00196)	47.5 (66.9)	0.00121 (0.00227)	54.8 (102.9)

Source: Author's calculations. See appendix for data sources.

a. Standard errors are reported in parentheses; for covariance estimates they are calculated using the Newey-West procedure, and for risk aversion estimates they do not reflect uncertainty in the numerator and are calculated by the delta method.


b. Covariance of the logarithm of excess returns of stocks over the risk-free interest rate during time $t + 1$ and the logarithm of consumption growth from time t to $t + 1 + S$, where S is the horizon.

c. Calculated as the sum of the mean log excess return and one-half its variance, divided by their covariance.

greater than two or three years. The discussion and analysis therefore focus on medium-term risk from one to three years, where there is some degree of agreement across methods and lower, although still quite large, statistical uncertainty.

To sum up, the results so far have three main implications. First, the declining pattern of risk aversion with the horizon over which consumption risk is measured is consistent with the slow adjustment of consumption to changes in wealth, and thus with the literature on consumption behavior discussed in earlier sections. This pattern is inconsistent with the textbook model. The contemporaneous covariance is misleading, and

Figure 3. Risk Aversion Implied by Estimates of Medium-Term Consumption Risk



Source: Author's calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
a. Calculated from the covariance implied by the impulse response of the VAR used in figures 1 and 2.

portfolio recommendations, economic models, or projections that calibrate the risk of equity based on the contemporaneous covariance underestimate the risk of equity.

Second, the equity premium measured by the medium-term risk to consumption is less of a puzzle than one would be led to believe based on the contemporaneous covariance. The covariance between consumption and excess returns rises by a factor of about ten for flow consumption and a

much larger amount for total expenditure, although the latter occurs mostly in extending the horizon for consumption by one quarter from the contemporaneous covariance.

Third, even the lowest point estimate of risk aversion is too high to be plausible. The medium-term risk to consumption is insufficient to explain the mean return on equity over the risk-free rate at reasonable levels of risk aversion. Although the equity premium is now less of a puzzle, it remains a puzzle.

A potentially important feature of the stock market is that a large (but declining) fraction of U.S. households do not hold any equity at all. Thus a leading explanation for the failure of the textbook model in aggregate data on consumption is that many households face costs of entering the stock market and so do not hold equity. If some households are excluded from the equity market, stock market risk will be concentrated among those households that participate. As a result, these households face more risk from equities, and the equity premium in equilibrium is larger, while the covariance of aggregate consumption and returns can remain low.²⁵

Thus, following Mankiw and Zeldes, the remainder of the paper addresses whether the medium-term risk of equity for the subpopulation of households that hold equity is sufficient to rationalize the premium on equity. For these households the risk of equity for wealth and consumption may be more substantial. However, to evaluate this hypothesis requires household survey data on consumption and stockholding.

The Consumer Expenditure Survey

This section describes how the CEX is used to construct a series measuring the consumption growth of households that hold equity. I then deal with two issues in estimating the medium-term risk of equity, and after that I present the main results of the estimation.

Interview data from family files of the CEX are used to construct a series on consumption growth for households holding equity. The CEX contains the best household-level data on consumption over time in the

25. See Mankiw (1986). An alternative way to cast the puzzle is to say that the return on equity is too high relative to the covariance of aggregate wealth and equity returns, for reasonable levels of risk aversion. See Davis, Nalewaik, and Willen (2000) and the discussion in sections 4.3 and 4.4 of Campbell (1999).

United States. The Bureau of Labor Statistics constructs the CEX data from a series of interviews based on a stratified random sample of the U.S. population. Each household is interviewed five times, once every three months, and new households replace, on a monthly basis, those that leave the sample. In a household's first interview, the CEX procedures are explained to the household members, and they are asked to keep track of their expenditure for future interviews. Each subsequent interview collects detailed information on the past three months' consumption expenditure. In each household's second and fifth interviews, demographic and income data are collected, including information about income and earnings during the previous twelve months. This information is updated if it changes during the course of the survey year. In a household's final interview, a set of questions about assets is asked, including the current "estimated market value of all stocks, bonds [private bonds only], mutual funds and other such securities" owned by the household and the amount by which these holdings have changed over the previous twelve months.

The appendix provides a more complete description of how the series is constructed, but three features deserve note here. First, I construct both flow consumption and total expenditure per effective person, but these concepts do not exactly match the same concepts in the aggregate data. Along most dimensions, the CEX measure is a closer match to the theory. Second, some cleaning of the data is undertaken; primarily, observations with extremely low levels of consumption are dropped, and the largest and the smallest growth rates of consumption in each period are dropped, as are all observations in several periods in which survey changes were made that result in consumption growth not being correctly measured. Third, households are categorized as holding equity only if they are holding a positive amount of the types of securities listed above before their first interview. Households with missing or miscoded data are not considered stockholders. This definition of households that own stock excludes many that hold equity indirectly; the CEX does not ask for information on stockholding in pension funds, although some households may include investments in Individual Retirement Accounts, and some may even include pension funds in their response to the survey question.²⁶

26. This measure is not conservative in that it is possible that a household included as a stockholder might in fact hold a corporate bond mutual fund, or corporate bonds directly, and not hold any stock. However, it seems reasonable that households who hold such assets also have easy access to stock funds. In this case the Euler equation for consumption and

In short, for every pair of consecutive observations on consumption for a household, I construct the change in log consumption per effective person, for both flow consumption and total expenditure. These household-level observations are averaged to create a measure of the growth in consumption per effective person. Averaging only over those households that report holding equity before their first observation on consumption gives the consumption growth of households holding equity. The final series has monthly observations on consumption growth over three-month-to-three-month periods from the period October to December 1979 to the period December 1997 to February 1998, with eleven dates missing because of survey changes, for a total of 205 observations on consumption growth.²⁷ Finally, the overlapping nature of the data makes the correlation of covariances quite complex. All inferences are made using Newey-West standard errors allowing for correlations up to lags of $3(S + 1)$ months.

Estimation

There are two important differences between the calculation of risk from aggregate consumption data and the calculation of risk for stockholders in the CEX. First, the CEX consumption series has significantly more measurement error than the aggregate series. Second, the CEX consumption series covers a shorter time period and a period that includes unusually high market returns.

The first practical concern about estimation using the CEX is that there is significant mismeasurement of consumption in the CEX, even after averaging over a large number of households.²⁸ I assume that the mea-

stock return should hold for these households, and they belong in the sample of stockholders.

27. Using all possible three-month-to-three-month changes uses more of the information available in the survey than do papers that collapse the data to a quarterly frequency. This approach increases the amount of data and so aids inference in the presence of measurement error.

28. Measurement error also motivates the linear approach used throughout the paper. First, I do not estimate nonlinear Euler equations, because Vissing-Jørgensen (1998) and Brav, Constantinides, and Geczy (1999) find problems with measurement error for estimating nonlinear Euler equations. Second, I do not use Hansen and Jagannathan bounds because they require that one estimate the variance of consumption growth rather than its covariance with returns. Given both mismeasurement and clustering in the survey design, this is an uncertain exercise.

surement error in consumption is classical; that is, it is additive when consumption is expressed in logarithms and has a mean of zero conditional on true consumption and returns at all leads and lags.²⁹ Let time be measured in months, let t denote the last month of an observation, and let Δ_3 denote the difference operator across a three-month period. Observed consumption growth is a monthly series of three-month averages:

$$\begin{aligned} (17) \quad \Delta_3 \ln C_t &= \Delta_3 \ln C_t^* + \Delta_3 \eta_t \\ (18) &= \ln C_t^* - \ln C_{t-3}^* + \eta_t - \eta_{t-3}, \end{aligned}$$

where C^* is true consumption, C is the observed measure, and η is the measurement error. Because the measurement error is uncorrelated with returns, it does not bias estimation of the risk of the stock market. The theoretical covariance of observed consumption growth and returns is

$$\begin{aligned} \text{cov} \left[\ln \left(\frac{C_{t+3+3S}}{C_t} \right), r_{t,t+3} \right] &= \text{cov} \left[\ln \left(\frac{C_{t+3+3S}^*}{C_t^*} \right) + \ln \left(\frac{\eta_{t+3+3S}}{\eta_t} \right), r_{t,t+3} \right] \\ (19) \quad &= \text{cov} \left[\ln \left(\frac{C_{t+3+3S}^*}{C_t^*} \right), r_{t,t+3} \right] + \text{cov} \left[\ln \left(\frac{\eta_{t+3+3S}}{\eta_t} \right), r_{t,t+3} \right] \\ &= \text{cov} \left[\ln \left(\frac{C_{t+3+3S}^*}{C_t^*} \right), r_{t,t+3} \right], \end{aligned}$$

where S is still measured in quarters. Thus calculation of the covariance using mismeasured data is unbiased.³⁰

There is also practical evidence that bias in this context is small. First note that, ignoring means, the covariance is the coefficient in an ordinary least squares regression of returns on consumption growth times the standard deviation of returns. Given this interpretation, the estimation is akin

29. Because of the overlapping nature of the data, sampling generates correlation across measurement error in different periods. However, this mismeasurement should still be uncorrelated with returns. Further, since any given household is in the survey for at most three consecutive consumption growth rate observations, correlation across observations is short-lived, and inference with the Newey-West standard errors should provide accurate measures of statistical uncertainty.

30. This result also follows directly from the result that uncertainty in a stochastic discount factor that is uncorrelated with returns does not affect the price of the asset. This result, in turn, is the reverse of the well-known result that uncertainty in returns that is uncorrelated with consumption does not affect the price of the asset.

to the estimation of a consumption Euler equation, with actual rather than expected returns on the right-hand side. The CEX data in grouped form have been used in several contexts to estimate linear consumption Euler equations, and the consensus of this literature is that there is little finite-sample bias.³¹

The second practical concern is that the CEX data cover a significantly shorter time period than the aggregate data and, in particular, that this period is one of unusually high returns. The average excess return during this period is 3½ percentage points higher at an annual rate than during the entire sample. The resulting movements of household wealth are also striking. During the period for which the CEX data are available, the ratio of equity wealth to disposable income rises from 2.07 to 3.72.³² From 1952 to 1980 the ratio is relatively stable: it does not rise above 2.79 and has a standard deviation of 0.23.

Given these extraordinary returns, it is possible that the covariance in this period is not representative of the covariance over the entire sample. That is, inference that assumes that this high level of average returns is expected (as does inference based on equation 14, for example) will likely underestimate the risk of equity.

To evaluate the importance of this issue, I redo the analysis of the medium-term risk of equity in aggregate data for the period covered by the CEX data: the last quarter of 1979 to the first quarter of 1998. Tables 4 and 5 present, for flow consumption and total consumption expenditure, respectively, the unconditional and conditional covariances and the coefficients of relative risk aversion that these levels imply. The format of these tables is similar to that of tables 1 to 3, although the horizon does not extend as far. Since the sample size is significantly reduced for the CEX period, only shorter horizons are considered. Statistical uncertainty is already quite high at a horizon of seven quarters.

The booming stock market of the 1980s and 1990s leads to lower estimates of the covariance of consumption growth and to higher estimates of risk aversion in this subsample. The contemporaneous covariances are

31. See in particular Vissing-Jørgensen (forthcoming). My own experience is that grouped linear estimators perform quite similarly as one reduces the sample size until the average group size falls below seventy-five. Of course, this may differ from application to application.

32. Data for these calculations are quarterly Flow of Funds data from the Federal Reserve.

Table 4. Consumption Risk of Equity and Implied Risk Aversion, Using NIPA Flow Consumption Data, 1979–98^a

Horizon (quarters)	Unconditional estimates			Conditional estimates		
	Covariance ^b	Implied level of risk aversion ^c	Ratio of implied risk aversion to estimate for 1959–2000 ^d	Covariance ^b	Implied level of risk aversion ^c	Ratio of implied risk aversion to estimate for 1959–2000 ^d
0	-0.00002 (0.00023)	n.a. n.a.	n.a.	0.00009 (0.00021)	702.5 (1,550.4)	2.1
1	0.00031 (0.00047)	210.9 (318.0)	2.5	0.00056 (0.00044)	118.0 (92.2)	1.6
2	0.00058 (0.00065)	114.5 (128.6)	1.8	0.00096 (0.00056)	69.0 (40.4)	1.2
3	0.00045 (0.00083)	147.3 (272.4)	2.5	0.00071 (0.00074)	93.6 (98.2)	1.6
4	0.00088 (0.00100)	75.7 (86.8)	1.8	0.00098 (0.00088)	67.4 (60.1)	1.5
5	0.00091 (0.00108)	72.8 (86.5)	1.9	0.00101 (0.00106)	65.6 (69.0)	1.5
6	0.00073 (0.00106)	91.2 (132.7)	2.6	0.00058 (0.00117)	113.6 (227.6)	2.8
7	0.00077 (0.00116)	85.9 (128.6)	2.6	0.00059 (0.00135)	113.1 (260.8)	3.0

Source: Author's calculations. See appendix for data sources.

a. Standard errors are reported in parentheses; for covariance estimates they are calculated using the Newey-West procedure, and for risk aversion estimates they do not reflect uncertainty in the numerator and are calculated by the delta method.

b. Covariance of the logarithm of excess returns of stocks over the risk-free interest rate during time $t + 1$ and the logarithm of consumption growth from time t to $t + 1 + S$, where S is the horizon. Calculated as the sum of the mean log excess return and one-half its variance (both estimated over 1959–2000), divided by their covariance.

c. Ratio of implied risk aversion shown here to that of the corresponding horizon from table 1 (unconditional covariance) or table 3 (conditional covariance).

d. Ratio of implied risk aversion shown here to that of the corresponding horizon from table 1 (unconditional covariance) or table 3 (conditional covariance).

Table 5. Consumption Risk of Equity and Implied Risk Aversion, Using NIPA Total Consumption Expenditure Data, 1979–98^a

<i>Horizon (quarters)</i>	<i>Unconditional estimates</i>			<i>Conditional estimates</i>		
	<i>Covariance</i>	<i>Implied level of risk aversion</i>	<i>Ratio of implied risk aversion to estimate for 1959–2000</i>	<i>Covariance</i>	<i>Implied level of risk aversion</i>	<i>Ratio of implied risk aversion to estimate for 1959–2000</i>
0	-0.00031 (0.00030)	n.a. n.a.	n.a.	-0.00020 (0.00028)	n.a. n.a.	n.a.
1	0.00045 (0.00053)	148.3 (175.3)	2.0	0.00058 (0.00048)	114.2 (93.6)	1.5
2	0.00110 (0.00076)	60.1 (41.6)	1.4	0.00138 (0.00061)	47.9 (21.3)	1.1
3	0.00107 (0.00096)	62.0 (55.8)	1.5	0.00113 (0.00078)	58.6 (40.3)	1.3
4	0.00120 (0.00111)	55.5 (51.6)	1.7	0.00108 (0.00084)	61.6 (48.0)	1.7
5	0.00117 (0.00117)	56.5 (56.2)	1.8	0.00102 (0.00104)	65.0 (66.6)	1.7
6	0.00085 (0.00120)	78.4 (110.9)	2.5	0.00038 (0.00122)	174.2 (559.2)	4.5
7	0.00086 (0.00131)	77.5 (118.2)	2.5	0.00036 (0.00152)	184.5 (779.4)	4.7

Source: Author's calculations. See appendix for data sources.

a. See table 4 for details.

negative for all cases except the conditional covariance of flow consumption, which is half its value for the entire sample. Turning to the medium-term risk, the covariances rise with the horizon, as was the case in the entire sample, and although the increase is large, the estimates of medium-term risk are significantly smaller than those that use the entire sample. At even longer horizons, the point estimates of the covariances capturing the medium-term risk in the 1980s and 1990s are negative and thus inconsistent with the textbook model. Note that, for reasons of comparability, as discussed above, the numerator of the estimating equation for risk aversion (equation 6) is calculated from the entire sample. Thus risk aversion is higher in this period not directly because of higher average returns during the period but because of lower observed medium-term risk.

Tables 4 and 5 also report, for each consumption measure, the ratio of the risk aversion coefficient for the medium-term risk of equity in the 1980s and 1990s to that in the entire sample (1959–2000, from tables 1 and 3). For future use in adjusting estimates from the CEX data, let these ratios be denoted by Φ_s . Risk aversion coefficients estimated from the smaller sample are typically two times those estimated in the entire sample but range between 1.1 and 4.7.³³

In sum, the period over which household-level data are available on the consumption of stockholders is also a period in which equity markets performed unusually well and in which the covariance of aggregate consumption growth and returns was unusually low. Any comparison of the covariance of stockholder consumption from the CEX with the model must take account of this fact. Not doing so risks rejecting the reasonableness of the model because of the fact that household-level data are available only in an unusual period.

The Risk of Stocks for Stockholders

Is the medium-term risk of equity for stockholders sufficient to rationalize its high relative return? Tables 6 and 7 report covariance estimates using the consumption growth of stockholders and measuring consumption as flow consumption or total consumption expenditure, respectively.

33. This factor also measures any differences in the estimates of the covariance due to different small-sample properties of the covariance measure between the longer and the shorter samples.

Table 6. Consumption Risk of Equity for Stockholders and Implied Risk Aversion, Using CEX Flow Consumption Data^a

<i>Horizon (quarters)</i>	<i>Unconditional estimates</i>			<i>Conditional estimates</i>		
	<i>Covariance, 1979–98^b</i>	<i>Adjusted covariance, 1959–2000^c</i>	<i>Adjusted risk aversion, 1959–2000^c</i>	<i>Covariance, 1979–98^b</i>	<i>Adjusted covariance, 1959–2000^c</i>	<i>Adjusted risk aversion, 1959–2000^c</i>
0	0.00072 (0.00064)	n.a.	n.a.	0.00097 (0.00068)	0.00208 (0.00145)	31.9 (22.3)
1	0.00194 (0.00088)	0.00493 (0.00224)	13.4 (6.1)	0.00244 (0.00082)	0.00387 (0.00130)	17.1 (5.8)
2	0.00251 (0.00115)	0.00451 (0.00207)	14.7 (6.8)	0.00357 (0.00106)	0.00422 (0.00125)	15.7 (4.7)
3	0.00133 (0.00127)	0.00331 (0.00316)	20.0 (19.0)	0.00246 (0.00118)	0.00394 (0.00189)	16.8 (8.1)
4	0.00229 (0.00114)	0.00406 (0.00202)	16.3 (8.1)	0.00373 (0.00123)	0.00559 (0.00184)	11.9 (3.9)
5	0.00210 (0.00126)	0.00391 (0.00235)	16.9 (10.1)	0.00330 (0.00146)	0.00500 (0.00221)	13.2 (5.8)
6	0.00223 (0.00145)	0.00572 (0.00372)	11.6 (7.6)	0.00346 (0.00166)	0.00984 (0.00472)	6.7 (3.2)
7	0.00261 (0.00141)	0.00670 (0.00362)	9.9 (5.3)	0.00412 (0.00163)	0.01216 (0.00481)	5.4 (2.2)

Source: Author's calculations. See appendix for details.

a. Sample periods are 1979:4 through 1998:1 for unadjusted covariance estimates and 1959:1 through 2000:4 for adjusted estimates. Standard errors are reported in parentheses; for covariance estimates they are calculated using the Newey-West procedure, and for risk aversion estimates they do not reflect uncertainty in the numerator and are calculated by the delta method.

b. Covariance of the logarithm of excess returns of stocks over the risk-free interest rate during time $t + 1$ and the logarithm of consumption growth from time t to $t + 1 + S$, where S is the horizon.

c. Scaled by the ratio of aggregate covariance in the two periods.

Table 7. Consumption Risk of Equity for Stockholders and Implied Risk Aversion, Using CEX Total Consumption Expenditure Data^a

<i>Horizon (quarters)</i>	<i>Unconditional estimates</i>			<i>Conditional estimates</i>		
	<i>Covariance, 1979–98</i>	<i>Adjusted covariance, 1959–2000</i>	<i>Adjusted risk aversion, 1959–2000</i>	<i>Covariance, 1979–98</i>	<i>Adjusted covariance, 1959–2000</i>	<i>Adjusted risk aversion, 1959–2000</i>
0	0.00248 (0.00149)	n.a.	n.a.	0.00221 (0.00151)	n.a.	n.a.
1	0.00335 (0.00151)	0.00659 (0.00297)	10.1 (4.5)	0.00316 (0.00150)	0.00490 (0.00232)	13.5 (6.4)
2	0.00377 (0.00197)	0.00515 (0.00269)	12.9 (6.7)	0.00443 (0.00197)	0.00498 (0.00221)	13.3 (5.9)
3	0.00290 (0.00254)	0.00437 (0.00383)	15.2 (13.3)	0.00330 (0.00264)	0.00435 (0.00348)	15.3 (12.2)
4	0.00451 (0.00226)	0.00782 (0.00392)	8.5 (4.2)	0.00516 (0.00256)	0.00880 (0.00437)	7.5 (3.7)
5	0.00554 (0.00266)	0.00993 (0.00477)	6.7 (3.2)	0.00591 (0.00295)	0.01013 (0.00505)	6.5 (3.3)
6	0.00391 (0.00323)	0.00981 (0.00810)	6.8 (5.6)	0.00438 (0.00355)	0.01968 (0.01595)	3.4 (2.7)
7	0.00316 (0.00285)	0.00794 (0.00716)	8.3 (7.5)	0.00435 (0.00317)	0.02024 (0.01475)	3.3 (2.4)

Source: Author's calculations. See appendix for data sources.

a. See table 6 for details.

The first point to note is that the unadjusted contemporaneous covariances for stockholders are at least five times larger than those estimated using aggregate consumption data over the entire sample (first row of tables 1, 2, and 3), and larger still than those estimated using NIPA data and the shorter sample (first row of tables 4 and 5), which are negative in three out of four cases. These raw contemporaneous covariances imply risk aversion coefficients of between 70 and 90 for flow consumption, and between 27 and 30 for total expenditure. Although both of these estimates are still implausibly large, they are much smaller than in the NIPA data on all households.

Second, the unadjusted covariances rise significantly as one extends the horizon over which consumption risk is measured. The medium-term consumption risk of equity is three to four times greater than the contemporaneous risk for flow consumption, and over twice as large for total expenditure. Thus the main finding for the aggregate data holds here as well: the medium-term risk of equity is significantly larger than the contemporaneous risk. Although an increase in risk of two to four times is large, it is interesting to note that the increase implied by these point estimates is not as large as the increase with horizon observed in the aggregate estimates.

Vissing-Jørgensen and Brav, Constantinides, and Geczy also document that the contemporaneous consumption risk of equity for stockholders is greater than that for all CEX households.³⁴ But more important, the consumption risk of stockholders is also larger at horizons of one to two years. The raw estimates of medium-term risk imply levels of risk aversion as low as the teens. Although these estimates are close to reasonable levels of risk aversion and thus close to rationalizing the equity premium (especially if one measures distance in terms of statistical uncertainty), these covariances are not adjusted for the unusual returns of the 1980s and 1990s.³⁵

34. Vissing-Jørgensen (1998); Brav, Constantinides, and Geczy (1999). Vissing-Jørgensen (forthcoming) does not include estimates of risk aversion from this method since the contemporaneous covariances are not large enough to ensure reasonably tight confidence intervals around the risk aversion estimates. Standard errors using the delta method (used here) are sensitive to the value of the covariance.

35. Dynan and Maki (2001) use the CEX to study the consumption response of stockholders to the stock market, conditional on household-level income movements and the mean return in each year. Although this formulation is not theoretically correct for the current exercise, they find levels of medium-term risk much greater than the contemporaneous

To account for the unusual characteristics of the period during which household-level data are available, I adjust the medium-term risk estimated in the CEX by the amount that the medium-term risk of equity rises in aggregate data when one moves from the subsample covered by the CEX to the full sample available. That is, I estimate the risk of returns for consumption and the risk aversion of the typical stockholder as

$$(20) \quad \text{cov}_{S,SH} = \Phi_S \text{cov}_{S,SH,CEX}$$

$$(21) \quad \gamma_{S,SH} = \frac{E[r_{t,t+1}] + \frac{1}{2} \text{var}(r_{t,t+1})}{\text{cov}_{S,SH}},$$

where $\text{cov}_{S,SH,CEX}$ is the covariance of consumption growth over horizon S estimated for stockholders over the short sample covered by the CEX data, and Φ_S is the sample ratio of the medium-term risk estimated using the full sample of aggregate data available to the corresponding estimate using only the period covered by the CEX, as described above and presented in tables 4 and 5.

In tables 6 and 7 the two columns to the right of each raw covariance present these adjusted covariances and the implied levels of risk aversion. The medium-term risk to consumption estimated from the consumption of stockholders and adjusted for the period covered by the CEX is quite reasonable. When marginal utility is measured as flow consumption, which omits many luxuries, the covariances show sufficient risk to consumption that the risk aversion coefficient of the representative stockholder need only be in the neighborhood of 5 to 10. When marginal utility is measured by the medium-term movement in total expenditure, the risk aversion of the representative stockholder need only be in the range of 3 to 8 to rationalize the equity premium.

How important is looking at the medium-term response rather than the contemporaneous response? This question is harder to answer, since the aggregate data in the period covered by the CEX imply a negative contemporaneous covariance. However, if one adjusts the contemporaneous covariance by the same factor used at a horizon of one quarter, then the

risk, their point estimates would imply even higher covariances and lower (more plausible) estimates of risk aversion than found here.

implied risk to consumption rises by 3.6 times and 5.9 times for unconditional and conditional flow consumption, respectively, as one moves from contemporaneous risk to medium-term risk at a horizon of seven quarters. Although not quite as large as the order-of-magnitude increase in the aggregate data, this increase is still large.

In sum, the medium-term risk of equity for stockholders in the CEX is larger than the medium-term risk for the consumption of all households in the NIPA. Adjusted for the low covariance found in this period in the aggregate data, the medium-term consumption risk of equity for stockholders is consistent with the high average rate of return to stocks and reasonable levels of risk aversion for stockholders. According to tables 6 and 7, the combination of limited participation and slow adjustment leaves almost no equity premium puzzle. Before concluding, I explore this result further in the next section.

The Risk of Stocks for Rich Households, Older Households, and All CEX Households

This section addresses two questions. First, the CEX data on stockholders differ from aggregate consumption data in a number of ways. To what extent does each of these differences drive the main result? Second, the medium-term risk of stockholders presumably exceeds the risk for all households because the wealth of stockholders is more highly correlated with returns than that of the typical household. Is the consumption risk of those with little labor income (older households) or with more asset wealth (rich households) larger still? Ultimately, statistical uncertainty hampers the ability of the data to answer these questions. The results presented here are therefore informative but not conclusive.

The CEX consumption series on stockholders differs from the NIPA consumption data in several respects. The CEX series is aggregated in a manner consistent with the theory, whereas the NIPA consumption series is not. NIPA consumption growth is the first difference of the logarithm of the average level of consumption across the population. This not only confounds movements in the distribution of consumption with movements in the typical household's marginal utility, but also includes households that die, immigrate, emigrate, or are "born" into the sample between

period t and period $t + 1$. In the CEX the more theoretically appropriate approach is used of first converting the household data to logarithms, then differencing for households present in both t and $t + 1$, then averaging.³⁶ Correct aggregation can make a significant difference in estimating the response of expected consumption growth to time variation in risk-free real interest rates,³⁷ suggesting that correct aggregation may matter for the present set of results.

Other differences also exist. Consumption in the NIPA includes the spending of nonprofit organizations, whose objectives probably are not captured by the same model we apply to households. As noted, the definitions of consumption are not identical, and the CEX data tend to measure durable purchases more accurately than small nondurable purchases for reasons of recall. Further, the CEX sample covers a different population of households. As discussed in more detail in the appendix, rural households, military households, and students living in dormitories are excluded. The CEX is also a nonrandom sample.

To evaluate the importance of these differences between the CEX and the NIPA data, I construct a series from the CEX that mimics the NIPA series. I build average consumption per person by period using the CEX data and sample weights and then calculate the medium-term risk of equity using this measure of consumption instead of the NIPA data. Tables 8 and 9, which are similar in format to tables 6 and 7, show that the consumption risk of equity as measured in the CEX data aggregated to mimic the NIPA is lower than that measured in correctly aggregated data on stockholders. However, it is also the case that the estimated consumption risk of equity is significantly larger in this CEX consumption series covering all households than it is in the NIPA consumption data in the comparable period (tables 4 and 5). Raw covariances estimated in the NIPA data range up to roughly 0.0012, whereas in this aggregation of the CEX the estimated covariances are as large as 0.0029. Like the NIPA data, the CEX data show a significantly larger risk of consumption at horizons of one or two quarters, but unlike the NIPA data, as the horizon is increased further, the estimated risk declines. At horizons beyond four quarters, the CEX aggregate has a lower estimated risk of consumption.

36. Calculations per effective householder are also done at the household level rather than simply making the data per person.

37. See Attanasio and Weber (1993).

Table 8. Consumption Risk of Equity and Implied Risk Aversion, Using Pseudo-NIPA Flow Consumption Data^a

Horizon (quarters)	Unconditional estimates			Conditional estimates		
	Covariance, 1979–98 ^b	Adjusted covariance, 1959–2000 ^c	Adjusted risk aversion, 1959–2000 ^c	Covariance, 1979–98 ^b	Adjusted covariance, 1959–2000 ^c	Adjusted risk aversion, 1959–2000 ^c
0	-0.00035 (0.00068)	n.a.	n.a.	-0.00035 (0.00068)	-0.00074 (0.00145)	n.a.
1	0.00182 (0.00078)	0.00463 (0.00198)	14.3 (6.1)	0.00186 (0.00085)	0.00295 (0.00135)	n.a.
2	0.00280 (0.00106)	0.00503 (0.00190)	13.2 (5.0)	0.00259 (0.00106)	0.00306 (0.00125)	22.5 (10.3)
3	0.00159 (0.00102)	0.00396 (0.00254)	16.8 (10.8)	0.00134 (0.00084)	0.00214 (0.00135)	21.6 (8.8)
4	0.00184 (0.00129)	0.00326 (0.00228)	20.4 (14.3)	0.00154 (0.00109)	0.00231 (0.00163)	30.9 (19.4)
5	0.00084 (0.00118)	0.00157 (0.00220)	42.3 (59.1)	0.00073 (0.00098)	0.00111 (0.00149)	28.8 (20.4)
6	0.00052 (0.00111)	0.00132 (0.00285)	50.1 (107.5)	0.00027 (0.00099)	0.00076 (0.00281)	59.5 (79.7)
7	0.00064 (0.00108)	0.00164 (0.00277)	40.5 (68.5)	0.00035 (0.00092)	0.00103 (0.00271)	86.9 (320.8)
						64.2 (169.0)

Source: Author's calculations. See appendix for data sources.

a. CEX data are used to construct a series consistent with the NIPA consumption measures (see text and appendix for details). Sample periods are: 1979-4 through 1998:1 for unadjusted covariance estimates and 1959:1 through 2000:4 for adjusted estimates. Standard errors are reported in parentheses; for covariance estimates they are calculated using the Newey-West procedure, and for risk aversion estimates they do not reflect uncertainty in the numerator and are calculated by the delta method.

b. Covariance of the logarithm of excess returns of stocks over the risk-free interest rate during time $t + 1$ and the logarithm of consumption growth from time t to $t + 1 + S$, where S is the horizon.

c. Scaled by the ratio of aggregate covariance in the two periods.

Table 9. Consumption Risk of Equity and Implied Risk Aversion, Using Pseudo-NIPA Total Consumption Expenditure Data^a

<i>Horizon (quarters)</i>	<i>Unconditional estimates</i>			<i>Conditional estimates</i>		
	<i>Covariance, 1979–98</i>	<i>Adjusted covariance, 1959–2000</i>	<i>Adjusted risk aversion, 1959–2000</i>	<i>Covariance, 1979–98</i>	<i>Adjusted covariance, 1959–2000</i>	<i>Adjusted risk aversion, 1959–2000</i>
0	0.00047 (0.00071)	n.a.	n.a.	0.00067 (0.00072)	n.a.	n.a.
1	0.00209 (0.00106)	0.00411 (0.00208)	16.2 (8.2)	0.00211 (0.00110)	0.00327 (0.00170)	n.a. 20.3 (10.5)
2	0.00290 (0.00143)	0.00397 (0.00196)	16.7 (8.3)	0.00274 (0.00123)	0.00308 (0.00138)	21.5 (9.6)
3	0.00242 (0.00147)	0.00365 (0.00221)	18.1 (11.0)	0.00221 (0.00116)	0.00291 (0.00153)	22.7 (12.0)
4	0.00207 (0.00168)	0.00359 (0.00291)	18.4 (14.9)	0.00152 (0.00126)	0.00259 (0.00215)	25.5 (21.2)
5	0.00148 (0.00170)	0.00265 (0.00305)	25.1 (28.8)	0.00112 (0.00122)	0.00192 (0.00209)	34.5 (37.4)
6	0.00091 (0.00167)	0.00228 (0.00419)	29.1 (53.6)	0.00027 (0.00123)	0.00122 (0.00553)	54.2 (245.6)
7	0.00051 (0.00170)	0.00128 (0.00427)	51.9 (173.3)	-0.00036 (0.00132)	n.a. n.a.	n.a. n.a.


Source: Author's calculations. See appendix for data sources.

a. See table 8 for details.

On balance, the CEX pseudo-NIPA data suggest that the medium-term risk rises over a couple of quarters only, and that risk aversion coefficients corresponding to these higher estimated risk levels are around 30 to 50 when the data are not adjusted for the sample, and 15 to 30 when adjusted for the unusual returns of the 1980s and 1990s. These estimates of medium-term risk aversion are about half those in the NIPA data and about twice those for correctly aggregated data on stockholders. These are then the differences due to the definitions and measurement of consumption and the sample differences. Of course, this number has a fair amount of statistical uncertainty in addition to being a rather coarse characterization of many numbers.

How much of this remaining difference is due to restricting the sample to stockholders, and how much due to correct aggregation? Figure 4 shows the ratio of the medium-term risk of equity for stockholders to the medium-term risk of equity for all households in the CEX, aggregated in the same manner as for stockholders. The contemporaneous risk for stockholders is typically positive, whereas for all households, correctly aggregated, it is typically negative. For horizons of one, two, and three quarters, the covariance of stockholder consumption growth with returns is not much larger than that for all households. However, for four quarters and beyond, and with the caveat that statistical uncertainty is high, this ratio is more clearly larger. The medium-term risk to the consumption of stockholders is on average a third larger than that for all households in the CEX. Thus the restriction to stockholders (for the already restricted sample of the CEX) increases the medium-term risk to consumption by about a third, whereas correct aggregation increases the medium-term risk by roughly two-thirds. The restriction to stockholders makes a greater difference for the contemporaneous risk of equity than it does for the medium-term risk.

A second set of questions is whether different subpopulations might have an even higher medium-term risk of equity or give further clues as to the correct theoretical model to explain limited participation in the stock market. One possible reason for limited participation is that some households have little net worth beyond future labor income. If such households face restrictions on borrowing to invest in the stock market, then small costs of entering the market may keep these households from investing in stocks, even though the marginal investment in stock has high returns and little consumption risk. Similarly, some households might hold stocks but be prevented from increasing those holdings because they find

Figure 4. Ratio of Consumption Risk for Stockholders to That for All Households

Source: Author's calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.

themselves up against the constraint on borrowing. To shed light on this explanation, I examine two different populations: older households, for whom future labor income is relatively small, and richer stockholding households, who, if constrained, are at least holding more substantial equity wealth.

Table 10 reports the contemporaneous and the medium-term risk for households in which the average age of the head of the households is sixty-five or greater (top panel) and households holding more than \$25,000 (in 1982–84 dollars) in securities (bottom panel). Point estimates suggest

Table 10. Consumption Risk of Equity and Implied Risk Aversion, for Older and Wealthy Households in the CEX^a

Horizon (quarters)	Total consumption expenditure					
	Flow consumption			Total consumption expenditure		
	Unconditional estimates		Conditional estimates	Unconditional estimates		Conditional estimates
	Covariance, 1979–98 ^b	Adjusted risk aversion, 1959–2000 ^c	Covariance, 1979–98 ^b	Adjusted risk aversion, 1959–2000 ^c	Covariance, 1979–98 ^b	Adjusted risk aversion, 1959–2000 ^c
0	-0.00057 (0.00077)	n.a.	-0.00055 (0.00074)	n.a.	0.00171 (0.00093)	0.00131 (0.00091)
5	0.00104 (0.00105)	34.3 (34.7)	0.00114 (0.00119)	38.3 (39.8)	0.00077 (0.00152)	0.00087 (0.00163)
6	0.00110 (0.00078)	23.4 (16.5)	0.00123 (0.00085)	19.0 (13.1)	0.00140 (0.00153)	0.00168 (0.00152)
<i>Households headed by persons aged sixty-five and older</i>						
0	0.00055 (0.00149)	n.a.	0.00017 (0.00152)	180.1 (1586.4)	0.00265 (0.00262)	0.00211 (0.00275)
5	0.00242 (0.00288)	14.7 (17.5)	0.00271 (0.00296)	16.1 (17.6)	0.00741 (0.00456)	0.00567 (0.00502)
6	0.00180 (0.00305)	14.4 (24.4)	0.00210 (0.00319)	11.1 (16.9)	0.00436 (0.00493)	0.00236 (0.00618)
<i>Households holding more than \$25,000 in equities^d</i>						
0	0.00055 (0.00149)	n.a.	0.00017 (0.00152)	180.1 (1586.4)	0.00265 (0.00262)	0.00211 (0.00275)
5	0.00242 (0.00288)	14.7 (17.5)	0.00271 (0.00296)	16.1 (17.6)	0.00741 (0.00456)	0.00567 (0.00502)
6	0.00180 (0.00305)	14.4 (24.4)	0.00210 (0.00319)	11.1 (16.9)	0.00436 (0.00493)	0.00236 (0.00618)

Source: Author's calculations. See appendix for data sources.

a. Sample periods are 1979:4 through 1998:1 for unadjusted covariance estimates and 1959:1 through 2000:4 for adjusted estimates. Standard errors are reported in parentheses; for covariance estimates they are calculated using the Newey-West procedure, and for risk aversion estimates they do not reflect uncertainty in the numerator and are calculated by the delta method.

b. Covariance of the logarithm of excess returns of stocks over the risk-free interest rate during time $t+1$ and the logarithm of consumption growth from time t to $t+1+S$, where S is the horizon.

c. Scaled by the ratio of aggregate covariance in the two periods.

d. In 1982-\$4 dollars.

that older households do not on average bear a higher consumption risk of equity than does the typical household. Although quite uncertain, these levels of medium-term risk are slightly smaller than those of the typical CEX household. This is consistent with an alternative view of the world in which elderly households hold large amounts of Social Security wealth, which is relatively safe and whose returns are uncorrelated with the market. Although not shown in the table, there is also no evidence that older households that hold stock bear any different degree of consumption risk than does the typical stockholding household. In sum, there is no evidence that households in which the head is sixty-five or older bear more consumption risk of equity than do similar households of all ages.

The bottom panel of table 10 finds no evidence that wealthy households bear more medium-term consumption risk than the typical household holding stock. Thirty-six percent of stockholding households are classified as rich, so that this series has more measurement error than the series for all stockholders. The point estimates vary more with changes in the horizon but show levels of medium-term risk very similar to those found for all stockholders in tables 6 and 7. In sum, there is little evidence of differences among the population of stockholders, but also little evidence against such differences.³⁸

Discussion

This paper judges the reasonableness of equity's risk and return in a way that is valid under a variety of models of consumer behavior. The average risk and return of equity relative to a risk-free investment are close to being consistent with optimal investment behavior for those households that hold equity, given slow adjustment of consumption, reasonable levels of risk aversion, and time-separable expected utility. This main finding raises two questions.

First, what keeps most households out of the stock market? One explanation is that there are significant fixed costs associated with learning how to invest, as well as flow costs that this paper ignores, such as taxes,

38. It is also worth noting that, partly as a result of statistical uncertainty, and unlike the results up to this point, the magnitudes of the comparisons in this section depend to some extent upon relatively minor (although correct) decisions concerning the construction of the data.

commissions, mutual fund management fees, and the time costs of monitoring one's portfolio. However, even given the higher risk of equity found using the medium-term measure, these costs would seem to be small relative to the high returns on equity. Additional market incompleteness, however, may make small costs sufficient. In particular, some households may not invest in equity because they have little liquid wealth and would pay high interest rates if they were to borrow against future income to invest in equity. Complementing this explanation, markets provide incomplete insurance against changes in labor income, and idiosyncratic risk can significantly increase the effective risk of equity at the household level.³⁹ It would also be useful to understand the relationship between limited participation and the distribution of consumption, and in particular whether the distribution of consumption is becoming bipolar, with stockholders increasingly more wealthy than nonstockholders.

Second, how do we incorporate the models of the slow adjustment of consumption from the literature on consumption smoothing into asset pricing models more generally? This paper provides several clues as to what the correct structural model of asset pricing might look like. Many recent models focus on significantly different models of consumer behavior or utility functions. This paper shows that limited participation and power utility rationalize the average differences between stock returns and the risk-free rate in the medium term, but that consumption smoothing does not capture the high-frequency relationship. These results imply that one should focus on modifications of behavior or utility functions that are consistent with power utility and intertemporal optimization in the medium term. It would be useful to further test this view, for example by studying whether the medium-term risk to the consumption of stockholders also explains returns across groups of stocks or between long- and short-term bonds.⁴⁰

39. See, for example, Constantinides and Duffie (1996), Constantinides, Donaldson, and Mehra (forthcoming), and Heaton and Lucas (1996). Parker (1998) shows that consumption and liquid wealth are closely related at low levels of liquid wealth, implying that an investment in equity could significantly raise consumption volatility for households with low liquid wealth.

40. The risk-return trade-off for long-term bonds lies between those of short-term bonds and the stock market, and so is also a puzzle. It seems reasonable to assume that households that hold corporate bonds have easy access to corporate equity, and therefore that limited participation has no role in explaining the risk-return trade-off between these two assets. For this model to work, therefore, it must be that the increase in measured risk that

Turning to the implications of this paper, since the average return on the stock market is consistent with the risk of the stock market for households holding stock, stocks appear to be a good deal for the typical non-stockholder in that the average return on equity outweighs the risks to consumption for a marginal investment. For the typical household not in the stock market, access to equity offers potentially large gains. Thus significant welfare gains could be realized if the government were able to provide access to diversified equity funds in a manner that avoids the costs or concerns that now keep these households out of the market (and does not entail large costs to the government), such as attempted in some proposals to diversify household Social Security wealth.⁴¹ Although the costs and benefits of such a policy depend on far more than the issue considered here, this paper does clarify the current risk of equity that should be used in calculating the potential welfare gains of such policies; it also makes clear that the relative risk and return of equity are reasonable for those households already participating in the stock market.

The findings of this paper support the view that the stock market boom of the 1980s and the 1990s was driven in part by increased stock market participation in response to declining costs of stock ownership.⁴² The cost of investing in the stock market has declined significantly thanks to legal changes and technological improvements in communications and information processing, and the participation rate has risen dramatically. The increase in stock ownership shares the risk of dividends across more households, which, in the steady state, decreases the correlation between consumption and equity returns for those households holding equity. This, in turn, reduces the premium required on equity and so increases its price, leading to a stock market boom. Whether the quantitative effect of this change in participation is large or small, accounting for limited participation in the stock market is important for understanding the true price of equity.

one observes in moving from a contemporaneous measure to a medium-term measure is greater for bonds than for equity, which, if true, is another interesting clue for modeling consumption behavior.

41. Another important caveat is that the canonical models do not always predict that opening a previously closed market gives welfare gains at all. For discussion of the issues, see the analyses of Allen and Gale (1994), Willen (2001), and Abel (2001).

42. See, for example, the discussions in Vissing-Jørgensen (1998) and Heaton and Lucas (2000).

APPENDIX

*Data Sources and Methods***Aggregate Data**

Data on personal consumption expenditure are extracted from the U.S. Central data available through Data Resources International (DRI) and from data available on the World Wide Web site of the Bureau of Economic Analysis. The data are spliced to form consistent series from the first quarter of 1959 to the first quarter of 2001. All real data are chain weighted.⁴³ Real flow consumption is constructed using series “by major type of product,” whereas real consumption expenditure is constructed using series “by major type of expenditure.” Chain-weighted real data on consumption cannot be added and subtracted to generate real measures of different combinations of consumption goods. Biases arise from differences in rates of price change and changing shares of consumption. In practice, however, the exact series can be closely approximated by combining series as one would when constructing a chain-weighted series from the detailed data. Thus I construct consumption series and price indexes as follows.

Let

C_t^{All} be total real consumption

C_t^S be total real consumption of services

C_t^{ND} be total real consumption of nondurable goods

C_t^F be real consumption of footwear (a subcategory of nondurable goods)

C_t^H be total real consumption of housing services

C_t^M be total real consumption of medical care services

C_t^E be total real consumption of education services

C_t^B be total real consumption of personal business services

\tilde{C}_t^M be total real consumption of medical care (category by type of expenditure)

43. Some early real data are available only in chained 1992 dollars. The early portion of the series in 1992 dollars is converted to 1996 chain-weighted dollars by rescaling to make the series averages match in the first four common quarters, which are the four quarters of 1967.

\tilde{C}_t^E be total real consumption of education (category by type of expenditure)

\tilde{C}_t^B be total real consumption of personal business expenses (category by type of expenditure).

Let X_t^i be the nominal consumption of category i in period t , and let P_t^i be the price deflator for series i derived by dividing the nominal quantity by the chain-weighted real quantity. Thus by definition, for all i ,

$$(A1) \quad C_t^i = \frac{X_t^i}{P_t^i}.$$

Let C_t^T be our series real total consumption expenditure (total consumption less medical care, education, and personal business services) and C_t be our series real flow consumption (nondurable and services consumption less footwear, housing, medical care, education, and personal business services). Since these data are chain weighted,

$$(A2) \quad \begin{aligned} C_t &\neq C_t^{ND} + C_t^S - C_t^F - C_t^H - C_t^M - C_t^E - C_t^B \\ C_t^T &\neq C_t^{All} - C_t^M - C_t^E - C_t^B. \end{aligned}$$

For $t = 1996$, that is, any date in 1996, we construct C_t and C_t^T by addition. For $t > 1996$, I generate the series for real flow consumption by iterating forward through time on

$$(A3) \quad C_t = C_{t-1} \sqrt[4]{\frac{\frac{X_t^{ND} + X_t^S - \sum_{i=F,H,M,E,B} X_t^i}{P_t^{ND} C_{t-1}^{ND} + P_t^S C_{t-1}^S - \sum_{i=F,H,M,E,B} P_t^i C_{t-1}^i}}{P_{t-1}^{ND} C_t^{ND} + P_{t-1}^S C_t^S - \sum_{i=F,H,M,E,B} P_{t-1}^i C_t^i}}{\frac{X_{t-1}^{ND} + X_{t-1}^S - \sum_{i=F,H,M,E,B} X_{t-1}^i}{P_{t-1}^{ND} C_{t-1}^{ND} + P_{t-1}^S C_{t-1}^S - \sum_{i=F,H,M,E,B} P_{t-1}^i C_{t-1}^i}}},$$

and for $t < 1996$, I generate this series by iterating backward through time on

$$(A4) \quad C_{t-1} = C_t \sqrt[4]{\frac{\frac{X_{t-1}^{ND} + X_{t-1}^S - \sum_{i=F,H,M,E,B} X_{t-1}^i}{P_{t-1}^{ND} C_t^{ND} + P_{t-1}^S C_t^S - \sum_{i=F,H,M,E,B} P_{t-1}^i C_t^i}}{P_t^{ND} C_{t-1}^{ND} + P_t^S C_{t-1}^S - \sum_{i=F,H,M,E,B} P_t^i C_{t-1}^i}}{\frac{X_t^{ND} + X_t^S - \sum_{i=F,H,M,E,B} X_t^i}{P_t^{ND} C_{t-1}^{ND} + P_t^S C_{t-1}^S - \sum_{i=F,H,M,E,B} P_t^i C_{t-1}^i}}}.$$

The construction of real consumption expenditure follows the same logic. For $t > 1996$, I generate the series for real consumption expenditure by iterating forward through time on

$$(A5) \quad C_t^T = C_{t-1}^T \sqrt{\frac{X_t^{All} - \sum_{i=M,E,B} \tilde{X}_t^i}{P_t^{All} C_{t-1}^{All} - \sum_{i=M,E,B} \tilde{P}_t^i \tilde{C}_{t-1}^i} \times \frac{P_{t-1}^{All} C_t^{All} - \sum_{i=M,E,B} \tilde{P}_{t-1}^i \tilde{C}_t^i}{X_{t-1}^{All} - \sum_{i=M,E,B} \tilde{X}_{t-1}^i}},$$

and for $t < 1996$, I generate this series by iterating backward through time on

$$(A6) \quad C_{t-1}^T = C_t^T \sqrt{\frac{X_{t-1}^{All} - \sum_{i=M,E,B} \tilde{X}_{t-1}^i}{P_{t-1}^{All} C_t^{All} - \sum_{i=M,E,B} \tilde{P}_{t-1}^i \tilde{C}_t^i} \times \frac{P_t^{All} C_{t-1}^{All} - \sum_{i=M,E,B} \tilde{P}_t^i \tilde{C}_{t-1}^i}{X_t^{All} - \sum_{i=M,E,B} \tilde{X}_t^i}}.$$

I divide the resulting series by the U.S. population to arrive at consumption per capita. Population data are three-month averages of monthly data from the Census data from DRI. The series drops by very close to 2 million in the first month of 1983 and rises by very close to 2 million after the end of 1985. I adjust the series upward by 2 million over the period 1983:1 to 1985:4 to avoid the otherwise implied large swings in consumption per capita. The resulting population growth over the break periods is similar to the growth rates in nearby months.

Returns on stocks and the risk-free rate are extracted from the Center for Research on Securities Prices (CRSP) data. The stock market index used is the quarterly CRSP index for the New York Stock Exchange and the American Stock Exchange. The risk-free rate is constructed using three-month averages of the monthly Fama risk-free rate (available in the CRSP data). The Fama risk-free rate is converted from continuous time to a quarterly rate. The return subscripted by $t + 1$ represents the return during period $t + 1$ so that

$$(A7) \quad r_{t,t+1} = \ln\left(1 + \frac{P_{t+1} + d_{t+1}}{P_t} - i_{t,t+1}^f\right)$$

where P_t is the price of the basket of equities at the end of period t , d_t are dividends paid during period t , and $i_{t,t+1}^f$ is the gross nominal risk-free rate. The dividend-price ratio corresponding to period t is d_t/P_{t-1} .

Conditional Estimation and the Vector Autoregression

The VAR confidence intervals are constructed by Monte Carlo integration under the assumption that the innovations are Gaussian. The procedure for this paper follows that described in RATS version 5, section 13.4. The VAR includes a complete set of quarter dummies in addition to the predetermined variables. The impulse responses for the entire system are given in figures A1 and A2.


To construct predicted returns, a regression consisting of the first row of the VAR system is run for the period from 1959 to 1974, and the fitted values from this regression are expected returns. From 1975 on, a rolling regression is used in which only information up to $t - 1$ is used to predict returns at t . For simplicity and consistency, this is done for a sample of data with timing structured as in the CEX so that the regression uses three-month lags and predicts three-month returns but does this at a monthly frequency. Thus the same predicted returns series is used for the aggregate and the CEX analyses. Month dummies are included rather than quarter dummies. The aggregate results are almost identical if only quarterly data are used.

The Consumer Expenditure Survey

Survey data are available only from 1980, so that the earliest consumption observation covers October to December 1979. In 1998 and beyond, the survey significantly altered the categories into which consumption expenditures are grouped, in order to match the extensive restructuring of the consumer price index (CPI) at that time. The 1997 files include data on household expenditure for all three-month interview periods starting in 1997, so that the data used cover a period up to and including February 1998. I use both the raw data files and SAS files available from Lorna Greening at <ftp://elsa.berkeley.edu/pub/ices/>.

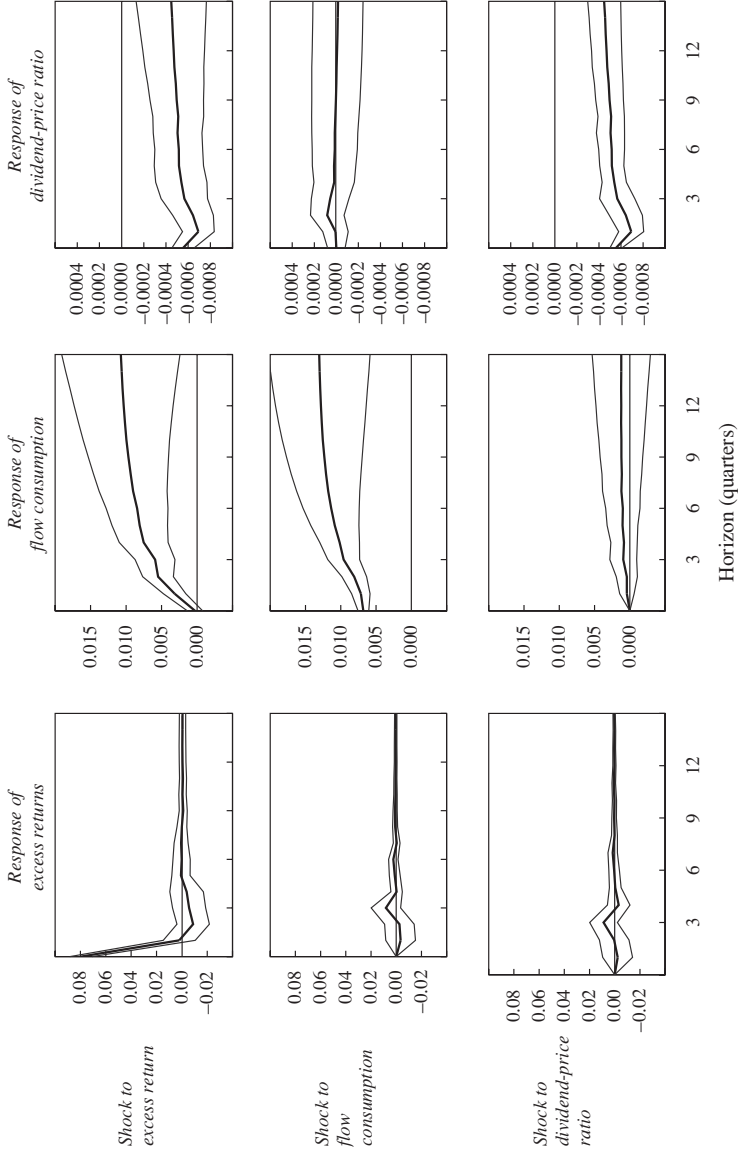
Flow consumption is defined as purchases of food, alcoholic beverages, apparel and apparel services, gasoline and motor oil used in transportation, public transportation, entertainment, personal care, and reading. As in the aggregate data, this definition omits expenditure on health care, housing, education, and financial services. Unlike in the aggregate data, footwear is included. I omit tobacco, because it is addictive; household operations,

Figure A1. Impulse Responses from Quarterly VAR Estimations, Using Flow Consumption^a



Source: Author's calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
 a. All variables include four lagged values. Shown with 95 percent confidence bands calculated by Monte Carlo methods from the estimated asymptotic distribution of the parameter vector.

Figure A2. Impulse Responses from Quarterly VAR Estimations, Using Total Consumption Expenditure^a



Source: Author's calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
 a. All variables include four lagged values. Shown with 95 percent confidence bands calculated by Monte Carlo methods from the estimated asymptotic distribution of the parameter vector.

because it includes repairs of furniture, appliances, and computers, as well as day care expenses, including tuition; and utilities, because, apart from telephone service, these are to a large extent determined by one's housing choice. Total expenditure omits spending on health care, charitable contributions, and education. Both series further differ from NIPA data in excluding spending by nonprofit organizations.

Flow consumption is converted to real terms using the CPI for each category of consumption for the census region in which the household resides. The CPI categories match the CEX categories and provide a reasonable approximation to the optimal chain-weighted index, although each subcategory of the CPI is likely to overstate inflation. Total expenditure is deflated by the CPI for all items less medical care. Similar results obtain if one instead uses the NIPA chain-weighted deflator for total consumption. Consumption per effective householder is calculated by dividing by the number of heads (one or two) plus 0.4 times the number of children.

Because of some implausibly low reports of consumption, I drop the bottom 1 percent of households in the distribution of real flow consumption per effective householder at each date. Rural households are dropped, as are households living in student housing and observations in which family size changes are greater than 3. Finally, the top and bottom 5 percent of the distribution of growth rates of flow consumption are dropped in each period. This trimming occurs at log growth rates of between 50 and 60 percent. Although trimming should not alter inferences if markets are complete among stockholders, trimming could understate the risk of equity, but in fact it does not. Experimenting with no trimming and with trimming 1 percent instead of 5 percent tails gave similar estimates for stockholders, although with slightly greater standard errors and smaller differences between stockholders and all CEX households correctly aggregated.

Consumption series are constructed as average consumption growth for all households and for only those households that report positive holdings of stocks, private bonds, or mutual funds immediately before the first observation on consumption. The data are seasonally adjusted by regressing them on dummies for each month, and the residuals are used in the analysis. Households that have pension wealth but do not directly hold equity are likely to report that they hold no stocks. Thus the series on the consumption of stockholders consists primarily of households that hold stock directly. Households with top-coded amounts in the final interview are deemed stockholders. Errors in processing the survey assigned bal-

ances of \$1 to a significant number of households that actually have balances well in excess of this amount. To be conservative in labeling households as stockholders, these balances are treated as if the household indeed had wealth in these types of assets of \$1. There are an average of 115 valid consumption growth rates per period for stockholders, and 908 average consumption growth rates per period for all households (both after trimming the tails of the distribution).

Because of survey changes, there are eleven missing observations on consumption growth during the sample, so that ultimately valid data on consumption growth are available for 205 partly overlapping three-month periods. Because of decennial survey changes, one cannot construct consumption growth for any households across the last three months of 1985 to the first three months of 1986, and similarly for 1995 to 1996. Further, because of changes in the survey methodology, the three observations on consumption growth ending in December 1987 and January and February 1988 exhibit large changes in the mean and variance of consumption growth and are dropped. Similar survey changes lead to dropping the same three months across 1995 to 1996 and across 1981 to 1982. Thus, from a possible 216 observations on consumption growth, 205 valid observations are available. Because the missing data occur in the middle of the sample, it could significantly reduce the power of our tests. To mitigate this effect, consumption growth rates are imputed to these dates that have the same covariance structure with lagged returns as the remainder of the data, and the standard errors are adjusted to account for the fact that some data are imputed.

Comments and Discussion

N. Gregory Mankiw: About twenty years ago, the consumption-based capital asset pricing model took center stage in discussions of asset pricing and economic fluctuations. Since then many economists have puzzled about the economy's most important risk premium, the spread between the equity return and the risk-free return. According to the model, the right measure of risk is consumption risk, but the measured consumption risk associated with the stock market seems too small to explain an equity premium of 6 percentage points, unless consumers are extraordinarily risk averse.

Fortunately, we have made progress toward explaining the equity premium, and this paper by Jonathan Parker makes a significant contribution to that effort. The resolution to the puzzle rests, at least in part, on two facts that several studies have now documented.

The first fact is that the risk of equities measured using consumption by stockholders alone is greater than the risk measured using aggregate consumption. Stephen Zeldes and I first documented this using data on food consumption from the Michigan Panel Study of Income Dynamics.¹ This finding has been confirmed in other, arguably better data sets by Annette Vissing-Jørgensen and by Parker in this paper.²

Of course, this fact does not explain *why* so many people do not hold stock, and this can be viewed as a puzzle in its own right. Part of the answer is that many people live hand to mouth and own hardly any assets

1. Mankiw and Zeldes (1991).

2. Vissing-Jørgensen (1998).

at all. But that is not the whole story. Zeldes and I documented that many people with sizable liquid assets hold no stock. The most plausible answer for these people is information costs: many people just do not understand stocks and bonds and do not feel comfortable buying them. Imagine how the typical economist would feel if advised to invest his or her retirement funds in rare stamps or seventeenth-century artwork. Most of us probably know nothing about these markets, but we know to stay away from things we do not understand.

The second fact that helps explain the equity premium is that consumption risk is higher when measured using medium-term changes in consumption than when using only the contemporaneous co-movement of consumption with stock market returns. This fact is documented both here and in a parallel paper by Xavier Gabaix and David Laibson.³ Parker estimates (in the first column of his table 1) that lengthening the time horizon raises the measured consumption risk of equities by a factor of ten.

With the benefit of hindsight, it seems that both of these facts should have been obvious. Regarding the fact of limited stock market participation, it is hardly a shock that many people do not hold equities and that those who do hold them face more equity risk than those who do not. Regarding the time horizon, the stock market is widely viewed as a leading indicator of economic activity, two-thirds of which is consumer spending. Thus, it is no surprise that increasing the time horizon raises the measured covariance.

The literature on consumption-based asset pricing neglected this observation until recently because, according to standard theory as set forth in Robert Hall's seminal 1978 paper, consumption should follow a random walk.⁴ In particular, the stock market should not be correlated with future consumption changes, so the contemporaneous risk and the medium-term risk (as measured in Parker's table 1) should be the same. Yet this theoretical prediction has never been fully confirmed by the data. In that same 1978 paper, Hall tested the theory and found that the random-walk hypothesis worked well, with a single exception: the stock market predicted future consumption growth. He hypothesized that some part of consumption takes time to respond to changes in wealth. This conjecture

3. Gabaix and Laibson (forthcoming).

4. Hall (1978).

foreshadowed by two decades this paper by Parker and the parallel one by Gabaix and Laibson.

Although Hall initially proposed this delayed-adjustment hypothesis, he never took it very seriously, and for good reason. If you add adjustment costs to the standard permanent income model, you are likely to get strong positive autocorrelation in consumption growth, because people will slowly respond to news about permanent income. By contrast, as Hall showed, the actual univariate process for consumption is close to a random walk, which means that consumption growth is close to white noise. Gabaix and Laibson salvage the model from this critique by assuming that, although there is lagged adjustment to stock market news, consumers respond immediately to all other news, such as news about labor income. Thus delayed adjustment is important for explaining the equity premium, but all other fluctuations influence consumption as in the Hall random-walk model. This set of assumptions makes the model work, but it seems like a *deus ex machina*, descending out of the sky to save us when we need it and assuming itself back to heaven when it would prove inconvenient. Parker avoids such criticism by skirting the question of why the short-run and the medium-run dynamics are so different. This is an important topic for future work.

One place where I part company with Parker is over the magnitude of this time horizon effect. I think the paper overstates the case because it emphasizes the difference between medium-term risk and contemporaneous risk measured using quarterly data. Much, although not all, of the previous literature, including my paper with Zeldes, has used annual data. Parker's figures 1 and 2 show that most of the gain in this exercise comes in the first few quarters. Had he started with data at an annual frequency, he would have seen a much smaller improvement from going beyond the contemporaneous consumption risk.

Despite this complaint, I am convinced by the more fundamental point: Consumption risk measured over the medium term exceeds that measured contemporaneously. I do not believe that this observation changes the estimates by a factor of ten, but it could well be a factor of two or three.

In addition to limited participation and time horizon, Parker suggests a third piece to the equity premium puzzle: aggregation. He claims that moving from National Income and Product Accounts data to a correctly aggregated consumption series from the Consumer Expenditure Survey increases measured consumption risk substantially. This inconsistency

between the NIPA and the CEX data is interesting and potentially important, but the explanation for it is left as a bit of a puzzle in its own right. If the NIPA data are artificially smoother than they should be, this fact would have major ramifications for the literature on consumption and asset pricing. But it is also possible that the CEX data are flawed in some way. The paper shows that the two data sets lead to different results, but it does not make a compelling case for preferring one data set over the other. The CEX data may simply be noisy: the large standard errors in Parker's table 6 certainly suggest this possibility.

Besides the explanations for the equity premium that this paper explores, another deserves to be on the table: sheer luck. Parker emphasizes that it would be a mistake to judge the equity premium using data from the 1980s and 1990s alone, because those decades saw extraordinarily high returns, a fact that someone at the beginning of the period could not have foreseen. I agree with that logic. But the same argument could be made for the entire twentieth century. From the standpoint of humankind's longer history, the economic progress of the past hundred years has been exceptional. Naturally, those who held an equity stake in this progress were well rewarded. Yet this outcome was probably not what most people living at the beginning of the twentieth century expected. In other words, the *ex ante* equity premium may have been much smaller than the *ex post* equity premium of 6 percent.

One person who did envisage this remarkable growth was John Maynard Keynes. In a famous essay, "Economic Possibilities for Our Grandchildren,"⁵ Keynes suggested that humankind was embarking on an economic miracle of a sort that previous generations could not have imagined. It is worth taking a moment to reflect on Keynes's prognostications and their implications for the equity premium. Keynes predicted that, a century after he was writing, incomes would be four to eight times higher than they were in 1930. This translates into growth in real income per capita between 1.4 and 2.1 percent a year, which has turned out to be remarkably accurate. Keynes then went on to suggest that if this growth occurred, "the economic problem may be solved." In other words, humankind would become satiated. The marginal utility of consumption would hit zero.

5. Keynes (1931).

If Keynes's view on the latter point were right, buying equities at the beginning of the twentieth century would not have been a good investment. If the economy experienced the growth that Keynes forecast, as it has, equities would pay off, as they have. But so what? If we are all going to be satiated, the higher return of stocks is of no value, and therefore there is no point in taking the risk. If the marginal utility of consumption is heading rapidly toward zero, people will exhibit very high risk aversion and will require a very large extra return to take on the risks of holding equities.

With the benefit of hindsight, we can see that Keynes was right about the growth, but wrong about the satiation. Although the twentieth century did generate a high return on equities, it also generated a surprising increase in the number of ways for investors to spend their marginal wealth. Perhaps this is why today we look back and view the equity premium as such a puzzle.

Jonathan Parker has moved us closer to resolving the puzzle. I am convinced that he is getting at a large part of the truth. Limited participation in the stock market and a longer time horizon can surely help explain the large observed equity premium. And the discrepancy between the NIPA and the CEX data is an issue that deserves a closer look. But part of the explanation for the equity premium may also lie in the fact that few people expected the rapid growth and high returns that we have experienced over the past hundred years, and some of those who did thought we would end up satiated as a result. In other words, what we take for granted at the beginning of the twenty-first century was not at all obvious a century ago.

Paul Willen: In this paper Jonathan Parker calculates the risk aversion of consumers who invest in the stock market. He exploits the first-order conditions of the individual utility maximization problem, which relate expected asset returns and the covariance of marginal utility of consumption with asset returns. Parker innovates on the existing literature by measuring that covariance as the covariance of consumption with lagged asset returns, rather than with contemporaneous asset returns as previous researchers have done. He finds that consumers who invest nonzero amounts in the stock market have risk aversion coefficients in the range that most researchers would consider reasonable. Thus one might be tempted to argue that the equity premium puzzle has been solved.

But has it been? Strictly speaking, no. Although Parker shows that the equity premium is consistent with reasonable values of risk aversion for those holding stocks, the puzzle remains why so many consumers hold no stocks. One natural explanation of Parker's finding is that some consumers do not participate because they are highly risk averse. If we average together the participants, with their reasonable risk aversion coefficients, and the nonparticipants, with their extremely high coefficients, we get a very high average level of risk aversion (and a puzzlingly large amount of dispersion in risk aversion). Thus the equity premium puzzle is alive and well.

Alternatively, suppose we assume that there is some reason *other* than high risk aversion why people fail to participate in equity markets. Then we can conjecture that these nonparticipants' risk aversion is comparable to that of participants and claim a solution to the equity premium puzzle. Thus Parker's analysis solves one puzzle, the equity premium, by invoking another, the participation puzzle: if nonparticipants have reasonable risk aversion coefficients, why don't they own equity?

Recall two basic stylized facts about equity ownership:¹

—Most people own little or no equity, even when equity ownership is defined very broadly to include equity held in defined-contribution pension plans and similar vehicles.

—Even many people with significant liquid financial assets own no equity, again even when equity ownership is broadly defined.

Why do these facts puzzle economists? Consider the individual investment decision: all potential investors face roughly the same opportunity to invest in risky assets—a better-than-fair bet. Theory tells us that no matter how risk averse you are, you should invest a strictly positive amount in such a bet. In the real world, indeed, equity is not just a better-than-fair bet, but a much, much better than fair bet.

But in the real world certain circumstances still might lead to nonparticipation. Three of these are trading costs, borrowing constraints, and labor income risk. In reviewing these, one should keep in mind that a plausible explanation must generate the nonparticipation of a large fraction of the population *and* predict positive equity holdings for the consumers that Parker focuses on. In other words, if something reduces demand for

1. See Poterba and Samwick (1995) for details.

equity for everyone, it might explain the equity premium but will not explain the participation puzzle.

Trading costs. If investors can borrow unlimited amounts at the risk-free interest rate, standard models predict that they will take enormous positions in equity, by borrowing against future labor income. For example, Steven Davis and I show that a plumber aged thirty will invest more than \$700,000 in risky assets.² Such investments allow for dramatically higher lifetime consumption. And as a result, no reasonable trading cost can dissuade investors from participating in risky asset markets.

Borrowing constraints. Under borrowing constraints, consumers cannot take leveraged positions. But reasonably calibrated examples show that they will still invest all (or almost all) their liquid wealth in risky assets. Life-cycle models with borrowing constraints predict low levels of liquid wealth until age forty.³ Thus, even if a younger consumer invests all his or her wealth in the stock market, the consumption and utility benefits of equity ownership are small. Relatively small trading costs could then explain nonparticipation among these younger consumers. So borrowing constraints, coupled with reasonable trading costs, could in principle explain the first stylized fact cited above, but they cannot explain the second: that of nonparticipation among consumers with considerable liquid wealth.

Labor income risk. Can the presence of risk to labor income reduce the benefits of stock ownership? Yes, in principle. First, if labor income is positively correlated with stock returns, then labor income acts as an implicit holding of equity and reduces the benefits of additional equity ownership.⁴ However, empirical evidence shows that the correlation of equity and labor income is highest for men with the most education and for the self-employed,⁵ two groups that are highly likely to own stock; the correlation is negative for least educated men, the group *least* likely to own stock. Second, even if labor income risk is uncorrelated with equity, it can depress demand for equity by making consumers more averse to all risks. Simulation evidence suggests that labor income risk uncorrelated

2. Davis and Willen (2000a).

3. Coco, Gomes, and Maenhout (2001); Constantinides, Donaldson, and Mehra (forthcoming).

4. See Davis and Willen (2000b) for an explanation.

5. For evidence on the first group see Davis and Willen (2000a); for evidence on the self-employed see Heaton and Lucas (2000). See also Bertaut and Haliassos (1995).

with equity does not have a significant effect on portfolio choice.⁶ Gregory Mankiw argues that if the variance of labor income shocks is negatively correlated with return shocks, demand for equity will fall.⁷ Simulation evidence on the Mankiw effect is mixed.⁸ Empirical evidence suggests that the effects of uncorrelated labor income risk are present but small.⁹

Other researchers have argued that habit formation and other modifications to preferences can explain why consumers are much more averse to holding equity than standard measures suggest they should be. Although such explanations could, in principle, explain the equity premium puzzle, they cannot explain the cross section of portfolio holding. It is not clear, for example, why habit formation would affect one part of the population and not another.

Widespread nonparticipation is, in many ways, an even more puzzling phenomenon for equilibrium theory than it is for optimization theory. General-equilibrium models with incomplete markets predict that marketable risks will be shared across individuals.¹⁰ Equity risk is marketable, yet we see that only a small fraction of the population owns equity. If frictions such as borrowing constraints account for the poor distribution of equity across the population, then the puzzle is why society has not created institutions to share equity risk more broadly.

In conclusion, the failure of large sections of the population to hold stock is a puzzle for economic theory. Although many things could explain this nonparticipation, researchers have yet to find a reasonable model that matches the observed distribution of participation. But if Parker has not solved the participation puzzle, he has at least shown that a solution to the participation puzzle will take us a long way toward solving the equity premium puzzle.

General discussion: Robert Gordon commented on the fact that the CES data that Parker uses are more volatile than the NIPA consumption data. He noted that for many reasons—sample survey error, imperfect recall of purchases, and lumping together of large purchases that are strung out over

6. Heaton and Lucas (1997); Coco, Gomes, and Maenhout (2001).

7. Mankiw (1986).

8. Heaton and Lucas (1996, 1997) find small effects; Storesletten, Telmer, and Yaron (2001) find larger effects.

9. Vissing-Jørgensen (2000).

10. See Willen (2001), for example.

time—the CES survey data might exhibit a lot of spurious volatility. This could give the impression of greater variability in the marginal utility of consumption, and hence of greater risk, than actually exists. Gregory Mankiw noted, however, that spurious volatility per se is not the problem, but rather spurious volatility that is correlated with stock market returns, so that it is not clear that these errors bias Parker's results.

William Nordhaus observed that the paper assumes agents know the parameters of the distribution of stock market returns estimated from the entire sample period. He wondered whether estimates of risk aversion would be noticeably different if they were based only on past observations at each point in time. Elaborating on this point, Gordon compared the extraordinarily high returns from 1982 to 2000 with the low returns from 1965 to 1982 and cautioned against taking the last twenty years of stock market data as a representative sample of returns. He viewed the higher stock prices of the last twenty years as the result of a series of positive supply shocks that unwound the earlier adverse supply shocks that had depressed prices. Gordon suggested that it would be imprudent for investors today to draw conclusions based on either series alone. However, he thought that the risk of the stock market might be substantially less now than it had been perceived to be earlier in the twentieth century. Stabilization policy in the postwar period has reduced the exposure of individual investors to the greater part of the risk actually experienced, in both the stock market and consumption, in the period before 1950. In addition, mutual funds now protect more people from the risk associated with holding individual stocks. He concluded that there are substantial reasons why the perceived risk of stocks, as opposed to their return, may have fallen significantly. Although the expected return has probably gone up, the risk has probably gone down.

The discussion turned to the paper's implications for Social Security. Alberto Alesina questioned the relevance of the model to anyone with a paternalistic view of Social Security. Because Parker's is a model in which investors are rational, any constraint imposed by the Social Security system would appear to decrease utility. Peter Orszag noted the importance of distinguishing between a policy of investing Social Security funds in the stock market and a policy that increases saving, either by prefunding the existing plan or by increasing compulsory saving. He argued that although the paper supports proposals to diversify the Social Security trust fund, it does not provide evidence in favor of increased saving. He noted that a

footnote in the paper suggests that the increase in estimated risk from lengthening the horizon is actually greater for bonds than for stocks, which would increase the desired portfolio share in stocks. Elaborating on this point, William Nordhaus noted that, for liquidity-constrained households, forced increases in saving could be utility decreasing, even if equities are extremely attractive relative to risk-free investments. A household that finds itself borrowing on its credit card at 16 percent interest will not find it advantageous to borrow more to invest in stocks. He noted further that if a parentalistic view with respect to Social Security reflects the belief that consumers are not well informed with respect to either financial investing or calculating future savings, it may be undesirable to give individuals choice over the portfolio allocation of their Social Security savings.

William Brainard suggested it would be interesting to examine whether characteristics of households that vary over the life cycle are useful in explaining risk taking or the response of consumption to unexpected gains or losses. Even for a given measured wealth-income ratio, young households have a great deal more human wealth relative to financial wealth than do old households; unexpected returns on stocks should therefore have a smaller percentage effect on the consumption of younger households. Similarly, it would be interesting to compare homeowners with renters, the college educated with the less educated, and households whose earned income has a high correlation with market returns with those for whom the correlation is low—stockbrokers and college professors, for example.

References

- Ait-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo. 2001. "Luxury Goods and the Equity Premium." Working Paper 8417. Cambridge, Mass.: National Bureau of Economic Research (August).
- Allen, Franklin, and Douglas Gale. 1994. "Limited Market Participation and Volatility of Asset Prices." *American Economic Review* 84(4): 933–55.
- Attanasio, Orazio, James Banks, and Sarah Tanner. 1998. "Asset Holding and Consumption Volatility." Working Paper 6567. Cambridge, Mass.: National Bureau of Economic Research (May).
- Attanasio, Orazio P., and Guglielmo Weber. 1993. "Consumption Growth, the Interest Rate and Aggregation." *Review of Economic Studies* 60 (3): 631–49.
- . 1995. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy* 103(6): 1121–57.
- Basu, Susanto, and Miles S. Kimball. 2000. "Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption." Unpublished paper. University of Michigan (December).
- Baxter, Marianne, and Urban J. Jermann. 1999. "Household Production and the Excess Sensitivity of Consumption to Current Income." *American Economic Review* 89(4): 902–20.
- Bertaut, Carol C., and Michael Haliassos. 1995. "Why Do So Few Hold Stocks?" *Economic Journal* 105(432): 1110–29.
- Brainard, William C., William R. Nelson, and Matthew D. Shapiro. 1991. "The Consumption Beta Explains Expected Returns at Long Horizons." Unpublished paper. Yale University (June).
- Brav, Alon, George M. Constantinides, and Christopher C. Geczy. 1999. "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." Working Paper 7406. Cambridge, Mass.: National Bureau of Economic Research (October).
- Browning, Martin, and Annamaria Lusardi. 1996. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature* 34(4): 1797–1855.
- Caballero, Ricardo J. 1990. "Consumption Puzzles and Precautionary Savings." *Journal of Monetary Economics* 25(1): 113–36.
- . 1993. "Durable Goods: An Explanation for Their Slow Adjustment." *Journal of Political Economy* 101(2): 351–84.
- . 1995. "Near-Rationality, Heterogeneity, and Aggregate Consumption." *Journal of Money, Credit, and Banking* 27(1): 29–48.
- Campbell, John Y. 1987. "Stock Returns and the Term Structure." *Journal of Financial Economics* 18(2): 373–99.
- . 1999. "Asset Prices, Consumption, and the Business Cycle." In *Handbook of Macroeconomics*, volume 1C, edited by John B. Taylor and Michael Woodford. Elsevier Press.

- Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107(2): 205–51.
- Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112(1): 1–55.
- Cochrane, John H., and Lars Peter Hansen. 1992. "Asset Pricing Explorations for Macroeconomics." In *NBER Macroeconomics Annual*, edited by Olivier Blanchard and Stanley Fischer. MIT Press.
- Coco, João F., Francisco J. Gomes, and Pascal J. Maenhout. 2001. "Consumption and Portfolio Choice over the Life-Cycle." Working paper. London Business School (May).
- Constantinides, George M. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98(3): 519–43.
- Constantinides, George M., John B. Donaldson, and Rajnish Mehra. Forthcoming. "Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle." *Quarterly Journal of Economics*.
- Constantinides, George M., and Darrell Duffie. 1996. "Asset Pricing with Heterogeneous Consumers." *Journal of Political Economy* 104(2): 219–40.
- Daniel, Kent D., and David Marshall. 1997. "The Equity Premium Puzzle and the Risk-Free Rate Puzzle at Long Horizons." *Macroeconomic Dynamics* 1(2): 452–84.
- Davis, Steven J., Jeremy Nalewaik, and Paul Willen. 2000. "On the Gains to International Trade in Risky Financial Assets." Working Paper 7796. Cambridge, Mass.: National Bureau of Economic Research (July).
- Davis, Steven J., and Paul Willen. 2000a. "Using Financial Assets to Hedge Labor Income Risks: Estimating the Benefits." Unpublished paper. University of Chicago Graduate School of Business (March).
- . 2000b. "Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice." Working Paper 7905. Cambridge, Mass.: National Bureau of Economic Research (September).
- Deaton, Angus. 1992. *Understanding Consumption*. Oxford, England: Clarendon Press.
- Dynan, Karen E. 2000. "Habit Formation in Consumer Preferences: Evidence from Panel Data." *American Economic Review* 90(3): 391–406.
- Dynan, Karen E., and Dean M. Maki. 2001. "Does the Stock Market Matter for Consumption?" Finance and Economics Discussion Working Paper 2001-23. Washington: Federal Reserve Board (May).
- Flavin, Marjorie A. 1981. "The Adjustment of Consumption to Changing Expectations about Future Income." *Journal of Political Economy* 89(5): 974–1009.
- Gabaix, Xavier, and David Laibson. Forthcoming. "The 6D Bias and the Equity Premium Puzzle." In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Kenneth Rogoff. MIT Press.

- Glassman, James K., and Kevin A. Hassett. 1999. *Dow 36,000: The New Strategy for Profiting from the Coming Rise in the Stock Market*. New York: Three Rivers Press.
- Gourinchas, Pierre-Olivier, and Jonathan A. Parker. Forthcoming. "Consumption Over the Life Cycle." *Econometrica*.
- Grossman, Sanford J., and Guy Laroque. 1990. "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica* 58(1): 25–51.
- Grossman, Sanford J., Angelo Melino, and Robert J. Shiller. 1987. "Estimating the Continuous-Time Consumption-Based Asset-Pricing Model." *Journal of Business and Economic Statistics* 5(3): 315–27.
- Grossman, Sanford J., and Robert J. Shiller. 1981. "The Determinants of the Variability of Stock Market Prices." *American Economic Review Papers and Proceedings* 71(2): 222–27.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86(6): 971–87.
- Hall, Robert E., and Frederic S. Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica* 50(2): 461–81.
- Hansen, Lars Peter, and Kenneth J. Singleton. 1983. "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy* 91(2): 249–65.
- Heaton, John, and Deborah Lucas. 1996. "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing." *Journal of Political Economy* 104(3): 443–87.
- . 1997. "Market Frictions, Savings Behavior, and Portfolio Choice." *Macroeconomic Dynamics* 1(1): 76–101.
- . 1999. "Stock Prices and Fundamentals." In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Julio J. Rotemberg. MIT Press.
- . 2000. "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk." *Journal of Finance* 55(3): 1163–98.
- Keynes, John Maynard. 1931. "Economic Possibilities for Our Grandchildren." In *Essays in Persuasion*. London: Macmillan.
- Kocherlakota, Narayana R. 1996. "The Equity Premium: It's Still a Puzzle." *Journal of Economic Literature* 34(1): 42–71.
- Lamont, Owen. 1998. "Earnings and Expected Returns." *Journal of Finance* 53(5): 1563–87.
- Lettau, Martin, and Sydney Ludvigson. 2001. "Consumption, Aggregate Wealth, and Expected Stock Returns." *Journal of Finance* 56(3): 815–49.
- Ludvigson, Sydney, and Charles Steindel. 1999. "How Important is the Stock Market Effect on Consumption?" Federal Reserve Bank of New York *Economic Policy Review* 5(2): 29–51.

- Lynch, Anthony W. 1996. "Decision Frequency and Synchronization across Agents: Implications for Aggregate Consumption and Equity Return." *Journal of Finance* 51(4): 1479–97.
- Mankiw, N. Gregory. 1986. "The Equity Premium and the Concentration of Aggregate Shocks." *Journal of Financial Economics* 17(1): 211–19.
- Mankiw, N. Gregory, and Stephen Zeldes. 1991. "The Consumption of Stockholders and Nonstockholders." *Journal of Financial Economics* 29(1): 97–112.
- Marshall, David A., and Nayan G. Parekh. 1999. "Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles?" *Journal of Finance* 54(2): 623–54.
- Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15(2): 145–61.
- Ogaki, Masao, and Carmen M. Reinhart. 1998. "Measuring Intertemporal Substitution: The Role of Durable Goods." *Journal of Political Economy* 106(5): 1078–98.
- Otrok, Christopher, B. Ravikumar, and Charles H. Whiteman. 2001. "Habit Formation: A Resolution of the Equity Premium Puzzle." University of Iowa Working Paper 98-04 (May).
- Parker, Jonathan A. 1998. "The Consumption Function Re-Estimated." Unpublished paper. University of Wisconsin (December).
- . 1999a. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *American Economic Review* 89(4): 959–73.
- . 1999b. "Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate." In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Julio J. Rotemberg. MIT Press.
- Piazzesi, Monika. Forthcoming. "Comment on Gabaix and Laibson, 'The 6D Bias and the Equity Premium Puzzle.'" In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Kenneth Rogoff. MIT Press.
- Poterba, James, and Andrew Samwick. 1995. "Stock Ownership Patterns, Stock Market Fluctuations and Consumption." *BPEA*, 2:1995, 295–372.
- Shiller, Robert J. 1982. "Consumption, Asset Markets, and Macroeconomic Fluctuations." *Carnegie Rochester Conference Series on Public Policy* 17: 203–38.
- Sims, Christopher A. 2001. "Implications of Rational Inattention." Unpublished paper. Princeton University (June).
- Souleles, Nicholas S. Forthcoming. "Consumer Sentiment: Its Rationality and Usefulness in Forecasting Expenditure—Evidence from the Michigan Micro Data." *Journal of Money, Credit, and Banking*.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2001. "Asset Pricing with Idiosyncratic Risk and Overlapping Generations." Working paper. Carnegie Mellon University (May).
- Vissing-Jørgensen, Annette. 1998. "Limited Stock Market Participation." Unpublished paper. Massachusetts Institute of Technology (May).

- . 2000. “Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures.” Working paper. University of Chicago (September).
- . Forthcoming. “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution.” *Journal of Political Economy*.
- Weil, Phillippe. 1989. “The Equity Premium Puzzle and the Risk-Free Rate Puzzle.” *Journal of Monetary Economics* 24(3): 401–21.
- Wilcox, David W. 1992. “The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work.” *American Economic Review* 82(4): 922–41.
- Willen, Paul. 2001. “New Financial Markets: Who Gains and Who Loses.” Unpublished paper. University of Chicago (September).
- Zeldes, Stephen P. 1989. “Consumption and Liquidity Constraints: An Empirical Investigation.” *Journal of Political Economy* 97(2): 305–46.