

Luxury Goods and the Equity Premium

YACINE AÏT-SAHALIA, JONATHAN A. PARKER, and MOTOHIRO YOGO*

ABSTRACT

This paper evaluates the equity premium using novel data on the consumption of luxury goods. Specifying utility as a nonhomothetic function of both luxury and basic consumption goods, we derive pricing equations and evaluate the risk of holding equity. Household survey and national accounts data mostly reflect basic consumption, and therefore overstate the risk aversion necessary to match the observed equity premium. The risk aversion implied by the consumption of luxury goods is more than an order of magnitude less than that implied by national accounts data. For the very rich, the equity premium is much less of a puzzle.

AS DEMONSTRATED BY GROSSMAN AND SHILLER (1981), SHILLER (1982), Mehra and Prescott (1985), and the extensive literature that follows, the risk of the stock market as measured by its co-movement with aggregate consumption is insufficient to justify the extent to which its average return exceeds the return on short-term government debt. We propose a partial resolution to this equity premium puzzle by distinguishing between the consumption of basic goods and that of luxury goods. Intuitively, rich households that hold most equity are almost satiated in their consumption of basic goods; wealth shocks are reflected in the consumption of luxury goods, which is much more responsive to stock returns than the consumption of basic goods.

Specifically, we model utility as a function of the consumption of both a basic good, of which a certain amount is required in every period, and a luxury good, which has low marginal utility even at low consumption levels. With such preferences, households consume only basic goods at low levels of total expenditures, while the share of luxury goods in overall consumption rises with expenditures. Households display a high degree of risk aversion with respect to their consumption of basic goods, consistent with the *subsistence* aspect of

*Aït-Sahalia is with the Department of Economics and the Bendheim Center for Finance, Princeton University, and the NBER. Parker is with the Department of Economics and the Woodrow Wilson School of Public and International Affairs and the Bendheim Center for Finance, Princeton University, and the NBER. Yogo is with the Department of Economics, Harvard University. We thank the editor and an anonymous referee, Christopher Carroll, Angus Deaton, Karen Dynan, Gregory Mankiw, Masao Ogaki, Annette Vissing-Jørgensen, and participants at the NBER ME Meeting (November 2001) and the Wharton Conference on Household Portfolio-Choice and Financial Decision-Making (March 2002) for helpful comments and discussions. We thank Jonathan Miller at Miller Samuel Inc. for data on Manhattan real estate prices and Orley Ashenfelter and David Ashmore at Liquid Assets for data on wine prices. Aït-Sahalia and Parker thank the National Science Foundation (grants SBR-9996023 and SES-0096076, respectively) for financial support.

basic goods. Cutting down on basic goods is costly in utility terms. For wealthy households, the consumption of luxury goods responds to wealth shocks due to stock returns, consistent with the *discretionary* aspect of luxuries.

We derive the Euler equation associated with the consumption of each type of good. Our theory implies that households are more risk averse with respect to the consumption of basic goods, so that the equity premium puzzle—the high degree of risk aversion implied by the observed consumption of basic goods—is not inconsistent with our model. The real test of the model lies in the Euler equation for luxury goods, which evaluates the behavior and risk aversion of rich households. Does marginal utility measured by luxury consumption vary sufficiently with stock returns to rationalize the equity premium?

Since no extant datasets measure the consumption of high-end luxury goods, we construct our own data on luxury consumption. We depart from the typical approach of studying how much households spend and instead ask purveyors of luxury goods how much they sell. Household surveys typically contain few wealthy households and measure categories of consumption that do not distinguish between basic and luxury goods. The latter is also the case for national accounts data.

So, while we also evaluate readily available government statistics, we construct and analyze the following: U.S. imports data from a consortium of 70 French luxury good manufacturers (Comité Colbert); IRS data on charitable giving by households with adjusted gross income (AGI) over \$1 million; U.S. sales of imported luxury automobiles (BMW, Mercedes, Jaguar, and Porsche); and finally sales data for retailers of high-end luxury goods. For this final series, we define luxury retailers as those companies listed by Morgan Stanley and Merrill Lynch in their analysts' reports on the luxury goods retail sector. The series includes aggregate U.S. sales for seven luxury retailers—Bulgari, Gucci, Hermès, LVMH, Saks, Tiffany, and Waterford Wedgwood. Note that many of the luxury retailers whose names we do not list individually are owned by luxury powerhouses such as LVMH, for whom we have total U.S. sales data. As of 2000, LVMH owned 46 different luxury brand names, whose sales represented 15% of the \$68 billion global luxury-goods market, against 6% for Richemont, the next largest.

We find that the consumption of luxuries covaries significantly more with stock returns than does aggregate consumption. Our estimates of the coefficient of relative risk aversion are an order of magnitude lower than that found using data on Personal Consumption Expenditures (PCE) of nondurables and services from the National Income and Product Accounts (NIPA). Our main series on aggregated luxury retail sales yields a point estimate for risk aversion of 7, and similar results obtain for all of our luxury series. PCE nondurables and services yield point estimates ranging from 50 to 173, depending on data frequency. Given moderate sampling error, we cannot reject the hypothesis that households with completely reasonable levels of risk aversion generated the observed returns and consumption of luxury goods.

Figure 1 depicts this main result. Panel A is a scatter plot of luxury consumption growth, measured by the sales of luxury retailers, against excess returns of CRSP NYSE-AMEX value-weighted portfolio over 3-month T-bills. Panel B

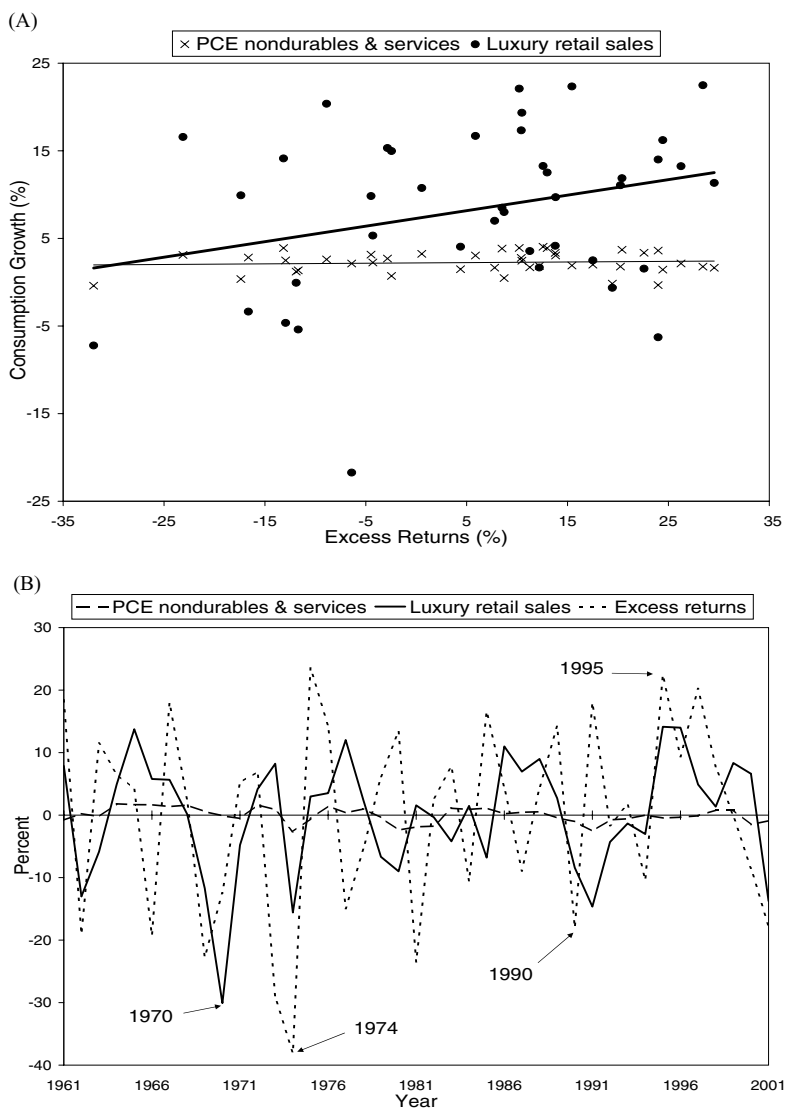


Figure 1. Response of basic and luxury consumption to stock returns. Panel A is a scatter plot of the growth rate for PCE nondurables and services and sales of luxury retailers against excess stock returns (CRSP NYSE-AMEX portfolio over 3-month T-bills). The thin (thick) line is the least squares regression line for PCE nondurables and services (sales of luxury retailers). Panel B is a time series plot of the growth rate for PCE nondurables and services, the growth rate for sales of luxury retailers, and excess stock returns. All series are normalized to have zero mean and are reported in percent.

is a time series plot of these series. For comparison, we include the growth rate of PCE nondurables and services in both plots. PCE is relatively smooth and almost nonresponsive to excess returns. By contrast, the consumption of luxuries is both more volatile and more correlated with excess returns. Luxury

consumption spikes down sharply in 1970 and 1974, both bad years for the stock market. Luxury consumption declines in the 1991 recession and experiences strong growth during the bull market of the late 1990s. At the end of the sample, it again dives as the stock market falls in 2001.

We present several extensions and robustness checks of our results. One potential concern is that luxury goods sales measure expenditures on durable goods, and so are more volatile than the correct measure, flow consumption. But our results are not driven by the volatility of expenditures rather than consumption. The increases in expenditure four years after an excess return implies even lower risk aversion than the contemporaneous movement in expenditures.

Next, we demonstrate that the *prices* of luxury goods in fixed supply reveal information about the equity premium. As an additional test of our theory, we construct time series on the prices of two high-end luxury goods whose supplies are highly inelastic: pre-war Manhattan coop apartments and Bordeaux wines from the finest châteaux and years. We calculate the equity premium implied by these price indexes. While the Manhattan coop prices covary strongly with returns and yield a point estimate for the equity premium as high as 7.8%, wine prices fail to rationalize the equity premium.

Finally, we also find that luxury consumption performs better than NIPA consumption in explaining the cross-section of stock returns. We examined the 25 Fama–French portfolios sorted by size and book-to-market equity, and 10 portfolios sorted by historical covariance with luxury consumption. Luxury consumption betas explain more of the variation in average returns across the portfolios than nondurable consumption betas (the usual CCAPM) or market betas (the CAPM), but not as much of the variation as the Fama–French three-factor model. We also formed stock deciles based on estimated luxury consumption betas. The results from this test are not robust and suffer from noise in the estimated luxury consumption betas at the individual stock level.

Our findings of reasonable levels of risk aversion based on luxury consumption data lead us to conclude that the single-good assumption that is embodied in most previous studies of asset prices and consumption leads to incorrect inference about the validity of the standard model, at least when applied to wealthy households. In particular, even within the basic power utility paradigm, there is no equity premium puzzle for the households that hold a large fraction of U.S. equity, or at least one not easily explained by sampling uncertainty. Put differently, we find no evidence that the risk faced by wealthy households does not justify the typical return on equity at reasonable levels of risk aversion.

The rest of the paper is organized as follows. Section I lays out our model with nonhomothetic utility and explains how the presence of luxury consumption changes asset-pricing equations. In Section II, we derive the testable implications of our model. The description of our dataset and the main empirical results are given in Section III. In Section IV, we extend our findings and test their robustness in four different directions. First, we consider the durability of luxury goods and argue that our findings are not driven by this issue. Second, we estimate the elasticity of intertemporal substitution with our luxury data.

Third, we estimate the equity premium using the price of luxury goods whose supply is inelastic. Fourth, we test whether our preferred measure of luxury consumption explains the cross-section of stock returns. Section V concludes. Derivations omitted from the main text are given in Appendix A, and a complete description of our dataset is given in Appendix B.

I. Luxury Goods, Basic Consumption, and Euler Equations

This section first lays the groundwork for studying the equity premium, then presents our modification of the canonical model to include multiple goods with nonlinear Engel curves. We explain the properties of this utility function and the implications for inference based on luxury goods. Finally, we derive the asset-pricing Euler equations and the implications for the covariance of aggregate consumption and returns.

A. The Equity Premium Puzzle

In the canonical model of investor behavior, households choose consumption expenditures (X_t) and the share of their saving invested in the stock market (ω_t) to maximize the expected present discounted value of utility flows for a given level of initial wealth A_0

$$\max_{X_t, \omega_t} E_0 \left[\sum_{t=0}^{\infty} \beta^t u(X_t) \right] \tag{1}$$

$$\text{s.t. } (\omega_t R_{t+1} + (1 - \omega_t) R_{t+1}^f)(A_t - X_t) = A_{t+1}, \tag{2}$$

$$X_t \geq 0, \tag{3}$$

where $u(\cdot)$ is the period utility function which is increasing, concave, and twice differentiable; $\beta > 0$ is the discount factor; A_t is household wealth at the beginning of period t ; R_{t+1} is the gross real return on stocks between time t and $t + 1$; and R_{t+1}^f is the gross real return on a conditionally risk-free asset between t and $t + 1$.

Note that for simplicity, and for consistency with the canonical model, we are assuming that households are infinitely lived, leisure is additively separable from consumption, and markets are complete so that labor income risk can be completely diversified. Complete markets also imply that the marginal utilities of all agents move together so that the moment condition can be estimated using aggregate consumption data. As is well known, this setup is easily extended to accommodate the choice of additional assets without changing the intertemporal conditions that we consider.

Assuming that the maximum of the objective is finite, we can rewrite the household optimization problem as a dynamic program

$$J(A_t | I_t) = \max_{\{X_t, \omega_t\}} \{u(X_t) + E_t[\beta J(\tilde{R}_{t+1}(A_t - X_t) | I_{t+1})]\}, \tag{4}$$

where J denotes the value function, I_t is the state of the economy at time t , $\tilde{R}_{t+1} = \omega_t R_{t+1} + (1 - \omega_t) R_{t+1}^f$ is the gross real return on wealth between time t and $t + 1$, and the program is subject to constraints (2) and (3). The first-order and envelope conditions imply the conditional moment restriction

$$E_t \left[\frac{\beta u'(X_{t+1})}{u'(X_t)} (R_{t+1} - R_{t+1}^f) \right] = 0, \quad (5)$$

which in turn implies the unconditional version of the same restriction.

Within this canonical model, the equity premium can only be explained by appealing to unappealingly high risk aversion. Given the observed joint stochastic process for the return on stocks, the return on bonds, and aggregate consumption, the coefficient of relative risk aversion implied by this model is implausibly high. Campbell (1999) surveys the last 15 years of research and shows that the puzzle is robust across countries and time.

One approach to solving the equity premium puzzle is to modify the canonical specification of the investor's marginal utility. The measured risk of equity can be raised substantially by altering the time-separable power utility framework so that marginal utility is more responsive to asset returns. Prominent examples of this approach include Abel (1990), Constantinides (1990), Epstein and Zin (1991), Bakshi and Chen (1996), and Campbell and Cochrane (1999).

A second solution, pioneered by Mankiw and Zeldes (1991), is to model markets so that only a subset of households hold equity and bear the aggregate risk of the market. Households that are not in the stock market—due to factors such as borrowing restrictions and fixed costs of investing in stocks—contaminate tests of the canonical theory that employ aggregate consumption data. Consumption data for households that hold equity directly confirm that these households do bear more market risk, although typically not enough to completely rationalize the high returns on equity (see Attanasio, Banks, and Tanner (2002), Brav, Constantinides, and Geczy (2002), Cogley (2002), Parker (2002), Vissing-Jørgensen (2002)).

B. Nonhomothetic Preferences

In this paper, we modify the period utility function, as in the first approach, and evaluate the risk of equity using the marginal utility of the wealthy only, as in the second approach. We modify the canonical utility function, but instead of dropping the assumption of time separability of utility, we drop the assumption that the period utility function is homothetic across goods. In other words, our point of departure from the canonical model is to drop the single-good assumption and model within-period utility as a function of two goods. We assume that households consume two types of goods: basic goods, C , and luxury goods, L . We conceptualize the former, which we treat as the numeraire in the economy, as the standard bundle of goods that most households in the United States regularly consume and that make up the bulk of the NIPA measure on consumption. The latter, luxury goods, are consumed only by the extremely rich, and that is where our hand-constructed data series play a role.

We reinterpret the previous statement of the problem as follows. The variable $X = C + PL$ represents total consumption expenditure per period, measured in terms of the numeraire (C) and optimally allocated between C and L . The utility function $u(X)$ represents an indirect utility function (with the relative price P of luxury goods suppressed), and the direct utility function $v(C, L)$, which we assume for simplicity is additively separable, is

$$v(C, L) = \frac{(C - a)^{1-\phi}}{1 - \phi} + \frac{(L + b)^{1-\psi}}{1 - \psi}, \tag{6}$$

where a, b, ϕ , and ψ are positive constants with $\phi > \psi$.¹ This implies that the subsistence level (a) is positive for basic goods, and negative ($-b$) for luxury goods.² Constraint (3) is replaced by $C > 0$ and $L > 0$.

This specification of utility captures two features of basic and luxury goods. First, luxury goods are not consumed by the “poor”: there exists $\bar{C} = a + b^{\psi/\phi} P^{1/\phi} > a$ such that $L = 0$ for all $C \leq \bar{C}$. That is, when the marginal utility of wealth is high, the agent chooses to consume none of the luxury goods. Second, the consumption of the “rich” is dominated by luxuries

$$\lim_{X \rightarrow \infty} \frac{C}{X} = 0, \tag{7}$$

$$\lim_{X \rightarrow \infty} \frac{PL}{X} = 1. \tag{8}$$

The assumption that $\phi > \psi$ implies that as the marginal utility of wealth goes to zero, the budget share of the luxury good approaches one. We prove this claim in Section A of Appendix A.

This expenditure behavior is illustrated in Figure 2. The limit behavior of the expenditure shares as X gets large are governed by ψ and ϕ ; the assumption $\psi < \phi$ delivers luxury consumption in excess of basic consumption at large expenditure levels. The local-to-zero behavior of expenditure shares are governed by a and b ; the assumption that $-a < b$ delivers basic consumption in excess of luxury consumption at low expenditure levels.

C. Euler Equations and Risk Aversion

We show in Section B of Appendix A that the first-order and envelope conditions from the dynamic program for the choice of C and L imply the following

¹ We assume that if $C \leq a$ the agent has infinite negative utility and purchases only basic goods, so $C = X$.

² The assumption that utility is separable in C and L is strong, but modeling nonseparability is unlikely to matter for our estimation. In order to overturn our results, the nonseparability must cause the marginal utility of luxuries to rise when the consumption of luxuries rises. In that case, the observed high covariance of luxury consumption and stock returns would not resolve the puzzle. Such nonseparability would have to be large because basic consumption is very smooth and moves so little with market returns. Houthakker (1960) and Atkeson and Ogaki (1996) use addilog utility functions that share this separability feature.

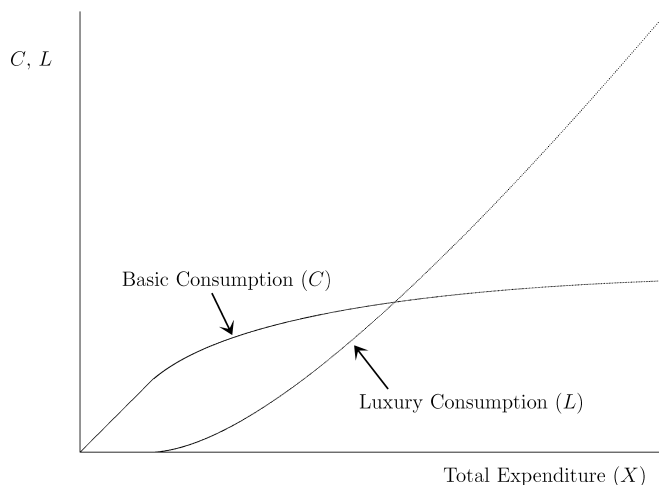


Figure 2. Consumption of basic and luxury goods under nonhomothetic utility. The figure plots the consumption of basic and luxury goods as a function of total expenditure.

two conditional Euler equations,

$$E_t \left[\frac{\beta(C_{t+1} - a)^{-\phi}}{(C_t - a)^{-\phi}} (R_{t+1} - R_{t+1}^f) \right] = 0, \quad (9)$$

$$E_t \left[\frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}} \frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] = 0. \quad (10)$$

The law of iterated expectations implies the unconditional versions of these equations.

The focus of the previous literature is on the unconditional version of equation (5), or if one takes the view that luxuries are not contained in NIPA nondurables consumption, of equation (9). We instead focus on the estimation and testing of equation (10). Equation (10) provides a test of whether the consumption Euler equation holds for wealthy households.

Our choice of utility function implies that the relevant curvature parameter that determines a household's attitude toward risk depends on the level of its total expenditures X .³ Consider the Arrow–Pratt definition of relative risk aversion $\gamma(X) = -Xu''(X)/u'(X)$. The coefficient of relative risk aversion with respect to gambles over C is $\gamma_C(C) = \phi C/(C - a)$, which falls with C and asymptotically approaches ϕ . Hence, for households with sufficiently low levels of X that only consume C , $\gamma(X) = \gamma_C(C)$, so ϕ is the curvature parameter that controls risk aversion. Risk aversion with respect to gambles over L is

³ Risk aversion that varies with wealth is an inherent feature of any nonhomothetic intra-period utility function. There is no utility function that admits nonhomothetic Engel curves and delivers constant relative risk aversion (see Stiglitz (1969), Hanoch (1977), and the discussion of the elasticity of intertemporal substitution in Browning and Crossley (2000)).

$\gamma_L(L) = \psi L / (L + b)$, which increases with L and asymptotically approaches ψ . Hence, for households at high levels of X that only consume L on the margin, ψ is the curvature parameter that is relevant for risk aversion. In general, for households that consume both C and L , $\gamma(X)$ is a weighted sum of ϕ and ψ that approaches ψ as X becomes large. Since we estimate ψ , our estimates provide a lower bound on the risk aversion over wealth gambles in the population at large.

This specification has two additional desirable features. First, since risk aversion declines with wealth, our model predicts that the wealthy hold a larger share of their wealth in equity, which is consistent with observed behavior (detailed in the next section). Second, the consumption of basic goods is a smaller share of expenditures for the rich. The distribution of basic consumption, as measured in the Panel Study of Income Dynamics (PSID) or the Consumer Expenditure Survey (CEX), is more equally distributed across households than the distribution of permanent income or wealth (see Dynan, Skinner, and Zeldes (2000) and Huggett and Ventura (2000)). Thus, the consumption of the poor and middle class remains a significant share of aggregate consumption despite the skewness in the wealth distribution.

D. Why Is There an Equity Premium Puzzle for Basic Consumption?

It is important to address why an econometrician using the usual consumption Euler equation (9) to study the return on stocks may not accept the model although the equation holds in equilibrium. Moreover, why does our Euler equation for luxury consumption (10) lead to the correct test of the theory? We present four answers. The first two are direct implications of our model; the second two are extant theories that complement our model.

The assumed intra-period utility function does not exhibit constant relative risk aversion, since some share of consumption is necessary for subsistence. Thus, for low levels of consumption, households are extremely unwilling to subject consumption to risk, so they hold little equity and have stable consumption. Thus, any test using an aggregate measure of consumption and assuming $a \approx 0$ calculates risk aversion from a weighted average of this nonresponsive consumption and the consumption of higher wealth households. Since the budget share of basic consumption declines with wealth, poor households are more heavily weighted in this average than their weight in wealth. According to our theory then, inference based on NIPA nondurables consumption data should find high levels of risk aversion. For this explanation to rationalize the low covariance between returns and NIPA consumption would still require extremely large risk aversion for some households, most plausibly by assuming that a and b are in fact large. There is a tension between our assumption for estimation that b is "small" and this explanation for previous findings.

Second, while not explicit in our model, it is reasonable for marginal utility from the consumption of basic goods to be bounded from above or reach zero (satiation), as in the cases of constant absolute risk aversion utility and quadratic utility, respectively. In either case, the coefficient of relative risk aversion for

basic consumption goes to infinity as wealth rises and marginal utility falls. Thus, high-wealth households maintain relatively stable basic consumption and react to market returns by changing luxury consumption. If we modified our utility function $v(C, L)$ to exhibit this feature, calculations based on basic consumption growth would find high risk aversion due to the unresponsive basic consumption of the rich as well as the poor.

In addition to the direct implications of nonhomothetic preferences, there are two extant theories that suggest that basic consumption may be inappropriate for measuring asset risk while luxury consumption provides the correct measure. Our study provides a test of both of these classes of theories, which both pass.

The first class of theories models the poor as not holding stocks due to fixed costs of participating in the stock market or due to uninsurable labor income risk. The theory of limited participation (Mankiw and Zeldes (1991)) posits that households must pay a fixed cost in order to invest in the stock market. In this case, nonrich households are not willing to incur this cost to invest, so their wealth is not directly affected by returns on equity and their consumption covaries less with the market.⁴ The theory positing incomplete markets argues that the nonrich do not hold stocks since households face uninsurable idiosyncratic endowment or income risk (see Heaton and Lucas (1996) and Brav et al. (2002)).⁵ As with limited participation models, this theory predicts a low covariance between the consumption of the poor or middle-class and stock returns, hence rationalizing the equity premium puzzle. Aggregate consumption includes the consumption of these households, and so the theory predicts an equity premium puzzle with respect to aggregate consumption and stock returns. On the other hand, the consumption of luxury goods is dominated by the rich, who are actually investing in both equity and risk-free assets.

Second, the basic consumption Euler equation may fail because there are costs to adjusting either basic consumption or an item that is nonseparable with basic consumption. Some items in basic consumption require commitment or are subject to direct or indirect adjustment costs associated with changing consumption. Similarly, the marginal utility of some items is not separable from the consumption of goods that have high costs associated with adjusting the level of consumption. For instance, items like transportation or fuels (subcategories of NIPA nondurables consumption) are in part determined by a household's consumption of housing and automobiles, which are subject to large adjustment costs and are hence infrequently adjusted. Items like mobile

⁴ For additional evidence on this theory, see Attanasio et al. (2002), Parker (2002), and Vissing-Jørgensen (2002). Guvenen (2000) calibrates a model with two types of agents, high risk aversion and low, in which only some agents have access to the stock market. The paper demonstrates that inference based on aggregate consumption in the canonical manner implies an implausibly high risk aversion, but the model has to assume an equity premium an order of magnitude smaller than that observed.

⁵ The one caveat we must note is that some rich households receive some share of labor income as stock options. While the idiosyncratic component of such risk is easy to unwind, the employee is often discouraged from doing so.

phone service, health club memberships and the like involve a degree of commitment over time.⁶ In contrast, such costs and nonseparabilities apply less to the consumption of luxury goods.

By studying the behavior of luxury goods, we test the central predictions of asset pricing in a way that is robust to these deviations from the canonical theory. Conversely, by testing a prediction of asset pricing that is consistent with these modified theories, we provide a test of these theories that is consistent with the presence of luxury goods.

The difficulties just discussed in using basic consumption to study the equity premium do not apply to the consumption of luxury goods. The consumption of necessities by the poor and rich does not contaminate a luxury-based measure of marginal utility; luxury goods are “discretionary.” The rich are willing and are able to pay any fixed costs for market participation. Moreover, rich households hold most equity and most hold equity. While the latter statement is to some extent tautological, the wealth distribution is so highly skewed that the concentration is extreme. The top 1% of households ranked by nonhuman wealth own over one-third of all privately held wealth, over half of stock wealth not held in pension funds, and 47% of all stock wealth. The top 5% of households own over half of all privately held wealth, over 80% of stock wealth not held in pension funds, and 75% of all stock wealth.⁷ It is also the case that most of the very rich own some stock, and investable wealth is a larger share of wealth for the rich than for the typical household, again, almost tautologically.⁸ Of the top 1% of households ranked by nonhuman wealth, 82% hold stock directly; of the top 5% of households, 78% hold stock directly. For the population as a whole, less than 50% hold stock directly (see Heaton and Lucas (2000)).

Finally, we do not require that all rich households consume luxuries or that no middle class households consume luxuries. If some rich households do not consume luxuries (the recipe for becoming a millionaire according to Stanley and Danko (1998)), then our test prices equity using only the subset of wealthy households that do purchase high-end luxury goods. The remaining households may save for bequests or have a “capitalist spirit” (Bakshi and Chen (1996)) but do not contaminate our main result. Second, if some middle-class households occasionally consume high-end luxury goods and if these households do not hold equity, this implies that our estimates are an upper bound on the risk aversion of the truly rich.

⁶ This is consistent with Parker (2002), who measures the risk of the stock market using a method that is robust to some of these issues. He finds that, in aggregate data, the ultimate movement of consumption following a return implies nearly an order of magnitude more consumption risk of equity than the contemporaneous movement.

⁷ These numbers are from the 1998 Survey of Consumer Finances as calculated and reported in Poterba (2000, Table II).

⁸ Anecdotally, Bill Gates saw his wealth drop from \$85 billion to \$63 billion between 1999 and 2000, a percentage decrease that closely mirrors that of Microsoft stock. Between 1986 and 2000, the number of millionaires rose sharply, and the total wealth controlled by households with assets of at least \$1 million grew 313% to approximately \$8.8 trillion (including Canada, as reported in the Merrill Lynch–Cap Gemini’s 2000 World Wealth Report). During the same period, the U.S. stock market rose by 405%.

II. Estimating Equations

We seek to evaluate the risk aversion of the rich using equation (10) and observations on high-end luxury goods. We assume that expenditures on any category of luxury goods move in proportion to those on all luxury goods. Thus, we can use observations on a subset of luxury goods to evaluate the model.⁹

Linearizing the unconditional version of the Euler equation for luxury goods, as in Campbell (1999), risk aversion can be derived as a function of population moments,

$$\psi = \frac{E[(R_{t+1} - R_{t+1}^f)P_t/P_{t+1}]}{\text{Cov}[\Delta l_{t+1}, (R_{t+1} - R_{t+1}^f)P_t/P_{t+1}]}, \quad (11)$$

where $l_{t+1} \equiv \ln L_{t+1}$. Throughout the rest of the paper, lowercase letters will be used to denote the logs of the corresponding uppercase variables. Details of this derivation are contained in Section C of Appendix A. The relative price of luxuries is present in the equation because returns are defined in terms of the price of basic consumption.

There is ample anecdotal evidence that the sales of luxury goods have benefited greatly from the bull market of the last decade, and that they have fallen after the stock market bubble burst in 2000–2001. But higher demand translates into higher consumption of luxury goods only to the extent that the supply is elastic enough so that price inflation does not completely crowd out the increase in nominal consumption of luxury goods.

We estimate the risk aversion coefficient using the method of moments by replacing the population moments in (11) with sample counterparts. The standard error is estimated by the delta method.

Two practical issues arise in the estimation of ψ through equation (11). The first is the timing convention used to convert the time average of expenditures to consumption flows. In equation (11), $R_{t+1} - R_{t+1}^f$ and Δl_{t+1} are both measured from the end of period t to the end of period $t + 1$. In practice, what we observe is total expenditures on the luxury good during period $t + 1$, which we denote as \tilde{l}_{t+1} . To translate from measured data to the model, we use the “end of the period” timing convention, where $\Delta l_{t+1} \approx \Delta \tilde{l}_{t+1}$. We make an exception to this rule in the analysis of imports of luxury goods. That is, for U.S. sales of Comité Colbert, sales of BMW and Mercedes, and sales of Jaguar and Porsche, we use the “beginning of the period” convention, where $\Delta l_{t+1} \approx \Delta \tilde{l}_{t+2}$. This exception is motivated by the shortness of these series and the possibility of shipping delays (in fact, beginning of period timing shows more reaction to stock returns).

The second issue that arises in estimation of the population covariance in the denominator of (11) is time aggregation in consumption data. Given the timing assumptions above, we are using consumption data that are averages over quarters or years in place of the desired instantaneous flow at a point

⁹ The primitive assumptions needed to ensure this are the same aggregation results across goods implicitly assumed to employ aggregate consumption data and imply homothetic Engel curves among luxury goods.

in time. As shown by Breeden, Gibbons, and Litzenberger (1989), using time-aggregated consumption data to estimate the covariance can bias the estimated covariance downward by a factor of $1/2$, which biases the estimated risk aversion upward by a factor of 2. Thus, we report time-aggregation corrected estimates of the risk aversion in addition to the conventional one that does not correct for time aggregation.

III. Risk Aversion and the Consumption of Luxury Goods

This section describes our data on the consumption of luxury goods and presents the associated estimates of risk aversion based on equation (11). A complete description of the source and construction of each series is contained in Appendix B.

Stock returns are measured as the return on value-weighted NYSE-AMEX portfolio from the Center for Research in Security Prices (CRSP). The risk-free return is the yield on the 3-month T-bill from CRSP's Fama Risk-Free Rates File. The excess return is the difference of the returns. Real returns, in units of basic consumption, are then computed using the implicit price deflator for PCE nondurables and services. We make use of the annual, quarterly, and monthly CRSP data to match the frequency of data on consumption.

We construct new series on the consumption of luxury goods because NIPA consumption data are not classified into luxury and basic consumption. Moreover, available household survey data are not suited for this task. While there is a host of issues that arise with all household surveys, the main shortcomings of the commonly used surveys are as follows. The PSID measures only the consumption of food and housing, has only infrequent measures of wealth, and under-samples the wealthy. The Survey of Consumer Finances (SCF), while over-sampling the wealthy, does not collect consumption data beyond the stock of some consumer durable goods and has very small panel dimension as well as a short time dimension. The CEX covers very limited categories of wealth, has poor measurement of those that are covered, and topcodes both consumption and wealth. The burden of detailing all consumption, which the CEX requires, is so large that very few high wealth households are in the survey and provide a full accounting of consumption.

A. Results from Government Aggregate Data

To begin, we examine publicly available government series that provide some evidence on the risk aversion implied by the consumption of luxury goods. These series are not entirely satisfactory, so we construct better measures of the consumption of luxury goods, as discussed below.

A.1. NIPA Personal Consumption Expenditures

We first estimate risk aversion from NIPA data on PCE nondurables and services. The first three rows of Table I present the results from this exercise. The columns, from left to right, report: (1) the sample period and size; (2) the

Table I
Risk Aversion Implied by Government Data

The table reports the descriptive statistics for consumption growth using government data. Data frequency is annual (A), quarterly (Q), or monthly (M). Correlation is with respect to excess stock returns deflated by the corresponding price deflator. Standard deviation (SD) is annualized (for instance, by multiplying the quarterly value by 2). The last two columns report the coefficient of relative risk aversion estimated by the method of moments. The first estimate is the conventional one that uses the sample covariance, and the second estimate corrects for time aggregation in consumption data by a factor of 1/2. Standard error is in parentheses.

Series	Period (Obs.)	Correlation	SD	Risk Aversion	
				Sample	Corrected
PCE nondurables and services (A)	1930–2001 (72)	0.173	0.023	100.029 (118.072)	50.015 (59.036)
PCE nondurables and services (Q)	1947–2001 (219)	0.137	0.011	346.556 (185.747)	173.278 (92.873)
PCE nondurables and services (M)	1959–2001 (515)	0.147	0.013	210.597 (104.842)	105.298 (52.421)
PCE jewelry and watches (A)	1960–2001 (42)	–0.147	0.063	–41.813 (56.729)	–20.906 (28.364)
PCE jewelry and watches (Q)	1959–2001 (171)	–0.029	0.079	–161.359 (419.576)	–80.680 (209.788)
PCE jewelry and watches (M)	1959–2001 (515)	0.085	0.102	47.352 (33.909)	23.676 (16.955)
PCE boats and aircraft (A)	1960–2001 (42)	0.027	0.138	106.084 (715.485)	53.042 (357.742)
PCE boats and aircraft (Q)	1959–2001 (171)	0.076	0.194	25.330 (30.852)	12.665 (15.426)
PCE boats and aircraft (M)	1959–2001 (515)	–0.053	0.363	–21.312 (20.939)	–10.656 (10.470)
Jewelry retail sales	1967–2001 (411)	0.109	0.118	31.811 (20.508)	15.906 (10.254)
Jewelry imports	1981–1999 (19)	–0.021	0.167	–206.961 (2768.212)	–103.480 (1384.106)

correlation between excess returns (deflated by the appropriate price deflator) and consumption growth; (3) the annualized standard deviation of the series; (4) the point estimate and standard error of risk aversion; and (5) the estimated risk aversion corrected for time aggregation by a factor of 1/2. As is well known, the risk aversion implied by PCE nondurables and services is implausibly high, with the point estimate ranging from 50 (annual) to 173 (quarterly) even after correcting for time aggregation. This is the consequence of the fact that nondurables consumption has low correlation with returns and low volatility.

Two subcategories of NIPA PCE capture luxury consumption to some extent: PCE jewelry and watches and PCE boats and aircraft. Unfortunately, both of these series are expenditures on durable goods rather than consumption, and they contain some consumption that should be categorized under basic goods. The consumption of watches includes a significant amount of nonluxury consumption, while PCE boats and aircraft also include expenditures on “durable

toys” and sports equipment. Table I shows that these series are significantly more volatile than PCE nondurables and services. However, only three of the six series are positively correlated with returns. These results are broadly consistent with those of Poterba and Samwick (1995), who find weak stock market wealth effects for government series that to some extent focus on luxury expenditures.

A.2. Retail Sales and Imports of Jewelry

The last two rows of Table I report our results using government data on retail sales of jewelry and imports of jewelry. Data on retail sales of jewelry are available from the Bureau of Economic Analysis (BEA) at monthly frequency since 1967. Data on U.S. imports of jewelry are from the World Trade Analyzer. To try to focus on higher-end jewelry, we aggregate imports from only France, Italy, and the United Kingdom. These series have shorter time dimensions than the PCE data, but are likely to measure luxury consumption better.

Monthly retail sales of jewelry has correlation of only 0.11 with excess returns, but has a high level of volatility. This leads to a point estimate for risk aversion of 16 when corrected for time aggregation. The last row of Table I reports the estimates for the series on U.S. imports of jewelry. The estimated risk aversion is negative due to the negative correlation between excess returns and the growth rate of imports.

B. Results from High-end Luxury Goods

We now turn to the analysis of measures of luxury good consumption that we have constructed ourselves. These provide strong evidence that nonhomothetic utility is important for understanding the risk of equity.

B.1. Sales of Luxury Automobiles

We begin by measuring luxury consumption as the sales of luxury automobiles from *Ward's Automotive Yearbook*. We construct two series on the sales of luxury automobiles: (1) BMW and Mercedes and (2) Jaguar and Porsche. The former is available since 1970, and the latter since 1962. While luxurious, these series are not ideal. Automobile sales measure expenditures on a durable good rather than flow consumption, to which our model refers. Thus, sales data on luxury goods capture expenditures rather than the service flow from the stock of durable goods, and as such should be more volatile than the stock (see, e.g., Mankiw (1982)).

We deal more formally with the issue of durability in Section IV.A. There, we show how to estimate risk aversion in a manner robust to some degree of durability by using the change in expenditures over several periods following an innovation to the stock market. In brief, with this adjustment, our main conclusions stand. For now, we treat sales as consumption and present results from our automobile purchases in rows one and two of Table II.

Table II
Risk Aversion Implied by Consumption of Luxury Goods

The table reports the descriptive statistics for consumption growth using luxury goods data. Data frequency is annual (A) or quarterly (Q). Correlation is with respect to excess stock returns deflated by the corresponding price deflator. Standard deviation (SD) is annualized (for instance, by multiplying the quarterly value by 2). The last two columns report the coefficient of relative risk aversion estimated by the method of moments. The first estimate is the conventional one that uses the sample covariance, and the second estimate corrects for time aggregation in consumption data by a factor of 1/2. Standard error is in parentheses.

Series	Period (Obs.)	Correlation	SD	Risk Aversion	
				Sample	Corrected
BMW and Mercedes sales	1970–1999 (30)	0.216	0.108	20.016 (18.683)	10.008 (9.342)
Jaguar and Porsche sales	1962–1999 (38)	0.036	0.236	49.877 (258.541)	24.939 (129.271)
Luxury retail sales (A)	1961–2001 (41)	0.299	0.095	13.984 (11.151)	6.992 (5.575)
Tiffany sales (A)	1961–2001 (41)	0.288	0.098	14.072 (11.550)	7.036 (5.775)
Luxury retail sales (Q)	1987–2001 (60)	0.199	0.196	14.339 (13.685)	7.170 (6.842)
Tiffany sales (Q)	1987–2001 (60)	0.221	0.217	11.677 (10.078)	5.838 (5.039)
Comité Colbert	1984–1997 (14)	0.587	0.110	13.558 (8.838)	6.779 (4.419)
Charitable contributions	1954–1999 (37)	0.339	0.204	7.388 (5.208)	3.694 (2.604)

To the extent that retail sales of luxury automobiles measure the consumption of luxury goods, risk aversion is significantly lower than that estimated from NIPA data. The estimated coefficient of relative risk aversion implied by sales of luxury automobiles is 10 and 25 for BMW and Mercedes and Jaguar and Porsche, respectively.

Figure 3A is a time series plot of the growth rate in sales of BMW and Mercedes along with excess returns. There appears to be strong covariation with large negative growth rates in 1987 and 1990, which were bad years for the stock market. More recently, sales of luxury automobiles were strong in the 1990s during the market boom. Figure 3B is a time series plot for the sales of Jaguar and Porsche. Again, there is evidence for covariation with the stock market with large slumps in 1987 and 1990.

B.2. Sales of Luxury Retailers

All of the measures considered so far are imperfect along two dimensions. First, the measures include basic goods purchased by the middle-class households and do not focus purely on the rich. Second, as noted, due to durability, these series may provide a weak mapping between consumption expenditures and marginal utility.

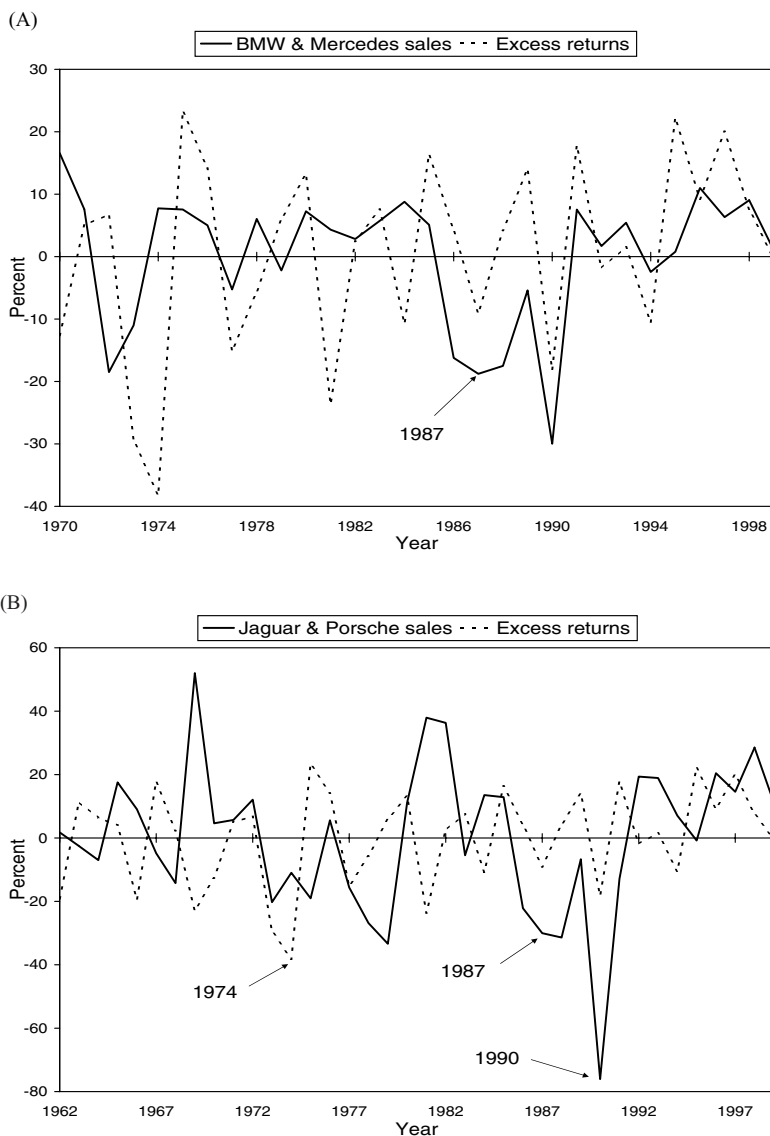


Figure 3. Response of luxury consumption to stock returns. Shown here is a time series plot of excess stock returns (CRSP NYSE-AMEX portfolio over 3-month T-bills) and the growth rate for (A) U.S. sales of imported BMW and Mercedes, (B) U.S. sales of imported Jaguar and Porsche, (C) annual sales of Tiffany, (D) quarterly sales of luxury retailers, (E) quarterly sales of Tiffany, (F) U.S. sales of C omite Colbert, and (G) charitable contributions by households with AGI over \$1 million. All series are normalized to have zero mean and are reported in percent.

Thus, we turn to sales from the high-end market for luxury consumption goods directly. By doing so, we are by definition measuring consumption of very expensive luxuries, which answers the first concern. As far as durability is concerned, it is likely to be less of an issue for high-end luxury goods sold to

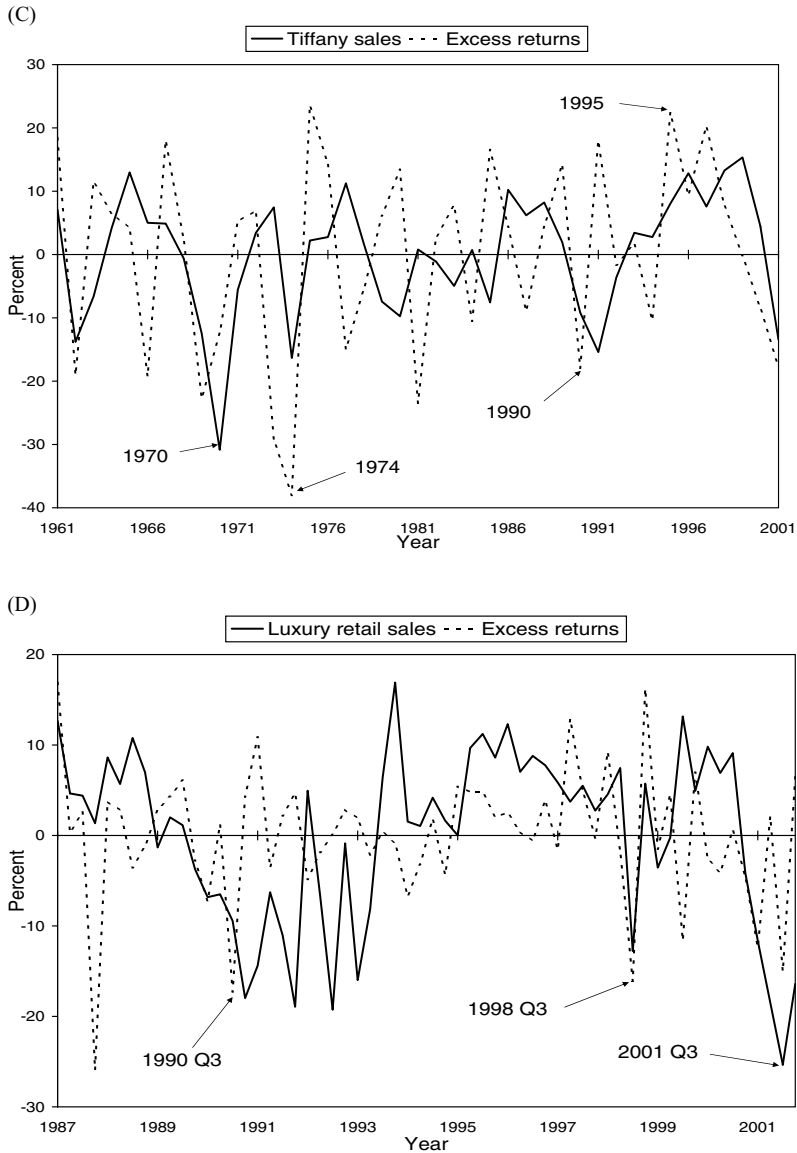


Figure 3.—*Continued*

the super rich by these retailers since fashion is fickle. A Hermès tie, Prada handbag, or designer dress lasts only one season (if not one social event) for those who can afford them. As for our other series, in Section III.A, we show that our results are not driven by the volatility of expenditures.

We collect data on U.S. sales of luxury retailers, defined as any company listed by Morgan Stanley and Merrill Lynch in their analysts' reports on the

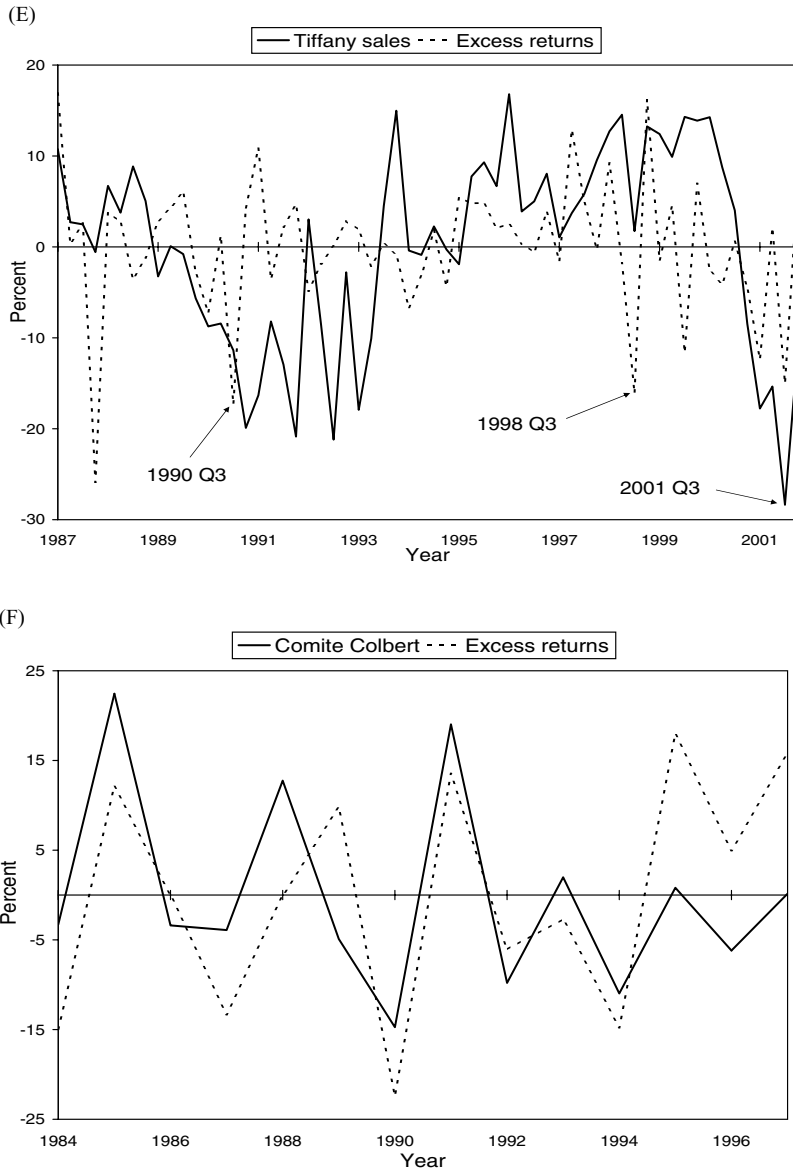


Figure 3.—Continued

luxury goods retail sector. Of these 32 companies, we use sales data for two major U.S. retailers (Saks and Tiffany) and five European retailers (Bulgari, Gucci, Hermès, LVMH, and Waterford Wedgwood); see Section B.1 of Appendix B for details. We aggregate sales across these seven retailers to create the total sales of luxury goods. Since the goods sold by each retailer are close substitutes, we sum sales across retailers to reduce the volatility of each series arising from

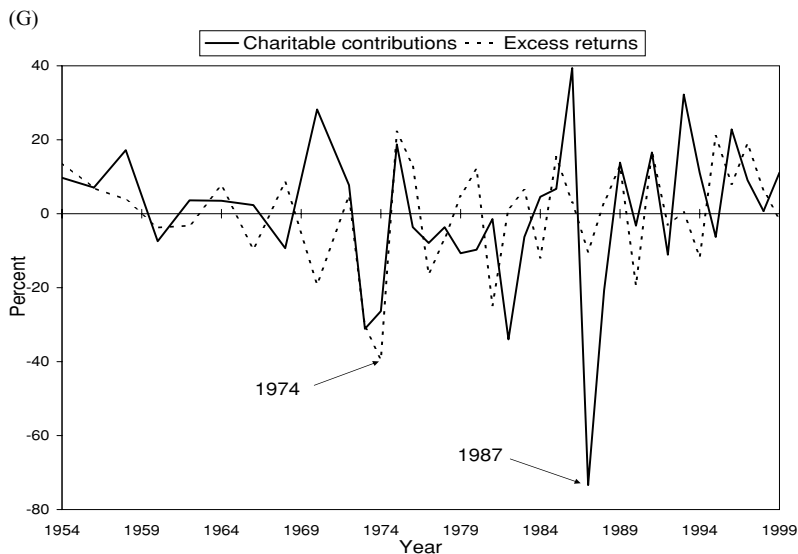


Figure 3.—Continued

idiosyncratic sales shocks. That is, any one retailer may misprice products, produce poor (e.g., unfashionable) products in any given season, and so forth, and hence suffer sales movements that are not indicative of total consumption of luxury goods.¹⁰

The length of time for which sales are observed differs by retailer. The longest series that we have is Tiffany, with sales data going back to 1960. For the other retailers, the series begin mostly in the early 1990s. To avoid artificial increases in sales as firms enter our dataset, we compute growth rates in sales using the same set of retailers at date t and $t + 1$. We also report results using only Tiffany sales since this is the retailer for which we have the longest and most consistent time series. It is consistent in the sense that the nature of the business for the company has not changed significantly over time, which is not the case for a company like Saks, which has gone through numerous mergers and acquisitions.

As Figures 1A and 1B show, there is a close relationship over time between excess stock returns and sales growth for luxury retailers. As reported in Table II, the correlation between excess returns and luxury retail sales is 0.3, about twice as large as that for PCE nondurables and services. Luxury retail sales is also quite volatile with a standard deviation of about 10%. This level of correlation and volatility implies that the risk aversion of the rich is about 7 (14 without the correction for time aggregation). That is, because stock market risk for the rich is high, when measured with movements in marginal utility

¹⁰ While not an issue for consistency, such noise can artificially inflate a measured covariance in finite samples.

of luxury consumption, the high equity premium can be justified with a low degree of risk aversion. This estimate of risk aversion is an order of magnitude less than that implied by PCE nondurables and services used in most previous consumption-based asset-pricing models.

When focusing our analysis on the single retailer for which we have a long time series of sales, the results are highly similar. Figure 3C shows the high covariance of excess returns and the growth rate of Tiffany sales. In years that experienced sharp declines in the stock market (1969, 1974, 1990, and 2001), Tiffany sales decline sharply. There appears to be some delay of consumption adjustment for the market shocks in 1969 and 1990, since sales growth bottom out in the subsequent years, 1970 and 1991, respectively. During the bull market of the 1990s, the sales for Tiffany experienced rapid growth, thus confirming anecdotal evidence of strong luxury consumption that was plentiful in the news during that period. This significant covariance implies a risk aversion estimate of about 7, shown in the fourth row of Table II.

For three of the luxury retailers (Gucci, Saks, and Tiffany) whose equity trade in the United States, we are also able to obtain quarterly sales data since 1986 (see Section B.1. of Appendix B for details). As Figure 3D shows, this series also covaries significantly with returns. Sales growth experienced large negative shocks in the third quarter of 1998 and 2001, which were bad quarters for the stock market. Starting around the third quarter of 1990, luxury retail sales experienced several quarters of negative growth rates, which appears to be a reaction to the 1991 recession. As shown in Table II, the quarterly sales of luxury retailers has correlation with excess returns of 0.2 with annualized standard deviation of almost 20%. This results in a point estimate for risk aversion of about 7, which agrees with our results using annual data. As shown in row 6 of Table II and Figure 3E, the results for quarterly sales of Tiffany are similar to those of all three retailers.

As a final measure of sales on high-end luxury goods, we have obtained data from Comité Colbert, a consortium of 70 French companies that specialize in high-end luxury products. Comité Colbert shared with us their total annual exports to the United States from 1984 to 1998. Figure 3F plots the growth rate of this series along with excess returns. Although the time series is short, there is significant covariation of U.S. luxury import growth and excess returns. As shown in Table II, the correlation with excess returns is almost 0.6, and the standard deviation is 11%. This leads to a point estimate for risk aversion of 6.8, when corrected for time aggregation. These results suggest that the lack of evidence for strong covariation using aggregate government data on U.S. imports of jewelry, discussed in Section III.A.2, may be due to the contamination of imports data by nonluxury items. It appears that movements in the consumption of *high-end* luxury items are very different from movements in broader measures of consumption that include *some* luxury goods.

B.3. Charitable Contributions

Despite our attempts to isolate the consumption of the very rich, it is possible

that some of the consumption we measure is sales to the middle class. While the middle class probably does not typically consume designer clothing, and while to the extent they do our measured risk aversion is biased *upwards*, we consider an alternative measure of luxury consumption that is even less subject to this criticism. Using the IRS publication *Individual Income Tax Returns*, we obtain data on the charitable contributions of households with AGI over \$1 million. One strength of this series is its length; the data is available biannually from 1952 to 1972 and annually since 1973.

Treating charitable giving as consumption is not standard in finance, but it is the leading theory explaining the phenomenon. According to these explanations, donating for medical aid to the suffering, endowing chairs in finance departments, donating art to a museum, and so forth provide so-called “warm glow” utility to donors. Andreoni and Miller (2002) show that the warm glow theory of charitable giving passes revealed-preference tests. Furthermore, Carroll (2000) argues that bequest giving enters utility as a luxury good.

In our framework, the price of charitable giving, relative to basic goods, is the tax price of charitable giving and varies with the marginal tax rate of the household; see Section B.3 of Appendix B for details.

Figure 3G is a time series plot of the growth in charitable contributions along with excess stock returns. Charitable contributions by high-income households track the stock market closely with large negative growth rates in 1973 and 1974, and even more significantly in 1987, the year of the “Black Monday” stock market crash. The excess stock return is clearly negative for 1987, but from a historical perspective on market returns, it is not unusually low.

As reported in the final row of Table II, the correlation of charitable contributions and excess returns is high at about 0.34, and the standard deviation is over 20%. This high covariance with excess returns leads to a point estimate of risk aversion of 3.7—a number completely consistent with most economists’ views on a plausible level of risk aversion. According to this measure, the equilibrium risk of equity for the marginal utility of the rich is sufficient to rationalize the observed equity premium.

In sum, our estimates based on high-end luxury consumption suggest an entirely different picture of the risk of equity than nondurable consumption in the NIPA.

IV. Robustness and Extensions

We now turn to robustness checks and extensions of our results. We first argue that durability is not driving the findings of the previous section. Second, we use the conditional Euler equation on our data on luxury consumption to estimate the elasticity of intertemporal substitution of the very rich. Third, we derive a method for estimating the equity premium implied by the price of luxury goods that are in inelastic supply. We apply this method to the rental prices of luxury pre-war Manhattan coop apartments and the auction prices of high-end Bordeaux wines. Fourth, we present the cross-sectional asset-pricing results using various factor models including our measure of luxury consump-

tion. While consumption betas calculated using luxury consumption do not provide a completely satisfactory explanation for the cross-section of returns, they do explain more of the variation in average returns across portfolios than betas constructed from either NIPA nondurable consumption (the usual CCAPM) or total market returns (the CAPM).

A. Dealing with Durability

As we already noted, many of the publicly available series measure the consumption expenditures that include some durable goods rather than being entirely flow consumption. To some extent, this criticism contaminates all empirical work in this area, as even NIPA consumption of nondurable goods and services contains items like shoes, financial services, health care, and items that may not be easily adjusted as discussed in Section I.D. We are concerned that this problem might be present in our measures of sales of high-end luxury goods. In this section, we provide some evidence that durability is not driving our results.

Suppose that utility comes from the service flow from the stock of a durable good, K_{t+1} . The stock is related to expenditures L_{t+1} by the equation

$$K_{t+1} = (1 - \delta)K_t + L_{t+1}, \quad (12)$$

where δ is the rate at which the durable good depreciates. If there are no adjustment costs, then expenditures are volatile as they increase or decrease to adjust the stock, while the stock is relatively stable. If this were the case, we would underestimate risk aversion using expenditures data, since risk aversion is decreasing in the covariance of expenditure growth and excess returns.

If the growth rate of consumption is stationary, equation (12) implies that the ratio of expenditures to stock, L_{t+1}/K_t , is stationary. In other words, the stock of the good, and hence its service flow, is cointegrated with expenditures. To the extent that a large positive return at $t + 1$ leads to an upwards revision in the stock of a durable good K_{t+1} , this should still be observed a few periods later as a higher level of expenditures, which are proportional to K_{t+1} .

Thus, we measure the covariance of returns with changes in the stock of durable goods using long-run increases in expenditures. In practice, our exercise is limited by the length of our sample, so we choose to look at the increase in expenditures from immediately before the excess return to four periods out. That is, we provide an alternative estimate of risk aversion with the equation

$$\psi = \frac{E[(R_{t+1} - R_{t+1}^f)P_t/P_{t+1}]}{\text{Cov}[l_{t+4} - l_t, (R_{t+1} - R_{t+1}^f)P_t/P_{t+1}]}. \quad (13)$$

Parker and Julliard (2003) motivate this measure of risk aversion using both the same assumptions needed to derive equation (11) and a variety of other deviations from the canonical model.

Table III provides evidence that our main findings are not driven by dura-

Table III
Risk Aversion Implied by Consumption Growth over Four Periods

The table reports the descriptive statistics for consumption growth measured over four periods. Data frequency is annual (A), quarterly (Q), or monthly (M). Correlation is with respect to excess stock returns deflated by the corresponding price deflator. The last two columns report the point estimate and Newey–West standard error (SE) for the coefficient of relative risk aversion estimated by the method of moments.

Series	Correlation	Risk Aversion	
		Point Estimate	SE
Panel A: Government Data			
PCE nondurables and services (A)	0.409	18.850	17.348
PCE nondurables and services (Q)	0.228	87.264	36.640
PCE nondurables and services (M)	0.180	97.120	48.327
PCE jewelry and watches (A)	-0.034	-74.126	344.789
PCE jewelry and watches (Q)	0.091	27.276	23.045
PCE jewelry and watches (M)	0.058	39.535	34.273
PCE boats and aircraft (A)	0.039	38.583	118.711
PCE boats and aircraft (Q)	0.105	9.822	8.413
PCE boats and aircraft (M)	0.021	39.697	92.692
Jewelry retail sales	0.113	20.554	14.228
Jewelry imports	-0.166	-11.782	23.756
Panel B: Consumption of Luxury Goods			
BMW and Mercedes sales	-0.035	-41.776	155.306
Jaguar and Porsche sales	0.021	28.489	156.072
Luxury retail sales (A)	0.206	10.033	6.474
Tiffany sales (A)	0.268	6.917	3.814
Luxury retail sales (Q)	0.200	5.078	4.207
Tiffany sales (Q)	0.195	4.310	3.348
Comité Colbert	0.345	12.818	9.369
Charitable contributions	0.227	4.346	3.191

bility. Panel A presents the results for the government data series of Table I and shows that, if anything, the estimated coefficients of relative risk aversion are more reasonable. For instance, the point estimate using annual PCE nondurables and services is about 19, which is much lower than the estimate of 100 that we obtained using contemporaneous consumption growth in Table I. The absolute value of the coefficients tend to be lower compared to Table I, but the only point estimate that is below 10 is for quarterly PCE boats and aircraft.

For our series on luxury consumption (Panel B of Table III), the coefficients of relative risk aversion estimated from long-run changes in expenditures are lower than those computed in Table II using contemporaneous changes. The point estimates are strikingly similar to those reported in Table II using the correction for time aggregation. The one exception is the estimate of the coefficient of relative risk aversion based on sales of BMW and Mercedes, which becomes negative. However, we obtain small and reasonable estimates of risk aversion for all the other series.

Thus, our results are not driven by high volatility of expenditures while service flows and marginal utility are relatively stable. In fact, the low risk aversion implied by the covariance of luxury goods and returns is driven by long-lasting movements in expenditures following excess returns. We conclude that the consumption of luxury goods implies that much lower values of risk aversion are consistent with the premium on equity.

B. Elasticity of Intertemporal Substitution

In this section, we turn from the task of measuring risk aversion to estimating its inverse, the elasticity of intertemporal substitution (EIS).¹¹ We linearize the conditional Euler equation (10)

$$E_t \left[\frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}} \frac{P_t}{P_{t+1}} R_{t+1} \right] = 1$$

to obtain

$$\Delta l_{t+1} = -\frac{\rho}{\psi} + \frac{1}{\psi} r_{t+1}^L - \frac{1}{\psi} \varepsilon_{t+1}, \quad (14)$$

where $\rho \equiv -\ln(\beta)$ is the discount rate and $r_{t+1}^L \equiv r_{t+1} - \Delta p_{t+1}$ is the real rate of return in terms of the price of luxuries (see Section E of Appendix A for the derivation). Since in general, $E_t[r_{t+1}^L \varepsilon_{t+1}] \neq 0$, we instrument r_{t+1}^L with instruments, z_t , uncorrelated with ε_{t+1} and estimate the model using two-stage least squares (TSLS) identified by the conditional moment restriction $E_t[z_t \varepsilon_{t+1}] = 0$. We choose instruments that are known to predict stock returns (see Campbell (1987), Fama and French (1988)): the second lags of the 3-month T-bill return, the yield spread, the log dividend-price ratio, and the luxury price inflation (see Section C of Appendix B). Since we use instruments dated at $t - 1$ for z_t , estimates of the EIS are unaffected by time aggregation in consumption data (see Hall (1988)).

First-order asymptotic inference can yield a poor approximation of the true statistical uncertainty when the instruments are weakly correlated with the endogenous regressor (see Nelson and Startz (1990), Staiger and Stock (1997)). This is a problem in estimating the EIS since both consumption growth and asset returns are difficult to predict (Stock and Wright (2000), Neely, Roy, and Whiteman (2001), Yogo (2004)). As suggested by Stock, Wright, and Yogo (2002) in a recent survey on weak identification in GMM, we report the first-stage F -statistic to assess whether weak instruments are a problem. We also report confidence intervals based on Moreira's (2003) conditional likelihood ratio (LR) test, which is robust to weak instruments.

The first three rows of Panel A in Table IV confirm that the estimated EIS is small when measured using aggregate consumption data (see Hall (1988)). The first three columns report results when the regressor in equation (14) is stock returns, and the next three columns are for the T-bill rate. For each of

¹¹ In preferences more general than the one in this paper (e.g., Epstein and Zin (1991)), the EIS may not be the inverse of the coefficient of relative risk aversion. However, the coefficient estimated by equation (14) can still be interpreted as the EIS.

Table IV
Estimates of the Elasticity of Intertemporal Substitution

The table reports the elasticity of intertemporal substitution (EIS), estimated through the conditional Euler equation. The test assets are real stock return (CRSP NYSE-AMEX portfolio) and the 3-month T-bill rate. From left to right, the table reports the first-stage F -statistic, the EIS estimated by two-stage least squares (TSLS) with standard error in parentheses, and the 95% confidence interval for the EIS implied by the conditional likelihood ratio (LR) test. Instruments include second lags of the 3-month T-bill rate, yield spread, log dividend-price ratio, and change in the log price deflator.

Series	Stock Return			T-Bill Rate		
	F	TSLS	Conditional LR	F	TSLS	Conditional LR
Panel A: Government Data						
PCE nondurables and services (A)	1.019	-0.084 (0.071)	$[-\infty, \infty]$	8.591	-0.145 (0.111)	$[-0.373, 0.140]$
PCE nondurables and services (Q)	4.669	0.015 (0.016)	$[-0.032, 0.076]$	39.481	0.013 (0.093)	$[-0.191, 0.196]$
PCE nondurables and services (M)	2.702	0.058 (0.030)	$[0.011, 0.603]$	38.983	-0.112 (0.166)	$[-0.474, 0.226]$
PCE jewelry and watches (A)	0.378	-0.382 (0.437)	$[-\infty, \infty]$	1.926	-0.375 (0.629)	$[-17.098, 0.418]$
PCE jewelry and watches (Q)	1.851	0.298 (0.209)	$[-\infty, \infty]$	10.566	0.387 (0.444)	$[-0.863, 1.237]$
PCE jewelry and watches (M)	2.175	0.203 (0.227)	$[-\infty, \infty]$	3.326	0.276 (0.752)	$[-4.349, 2.208]$
PCE boats and aircraft (A)	2.207	0.000 (0.305)	$[-\infty, \infty]$	1.377	-0.849 (1.469)	$[-\infty, \infty]$
PCE boats and aircraft (Q)	3.364	0.016 (0.331)	$[-1.311, 1.050]$	6.166	-0.471 (1.403)	$[-4.343, 2.602]$
PCE boats and aircraft (M)	3.023	-0.040 (0.685)	$[-2.385, 2.295]$	2.570	-1.327 (2.927)	$[-17.764, 6.816]$
Jewelry retail sales	2.190	0.360 (0.262)	$[-0.491, 57.503]$	5.792	0.486 (0.679)	$[-1.582, 2.073]$
Jewelry imports	0.760	-0.147 (0.715)	$[-\infty, \infty]$	1.019	1.610 (0.957)	$[-\infty, \infty]$
Panel B: Consumption of Luxury Goods						
BMW and Mercedes sales	1.008	0.335 (0.357)	$[-\infty, \infty]$	2.109	-0.745 (1.268)	$[-\infty, \infty]$
Jaguar and Porsche sales	0.843	1.445 (1.105)	$[-\infty, \infty]$	1.162	-0.535 (4.010)	$[-\infty, \infty]$
Luxury retail sales (A)	0.629	-0.112 (0.419)	$[-\infty, \infty]$	2.438	-0.579 (0.825)	$[-8.459, 0.857]$
Tiffany (A)	0.629	-0.084 (0.424)	$[-\infty, \infty]$	2.438	-0.371 (0.824)	$[-11.713, 1.087]$
Luxury retail sales (Q)	1.381	0.100 (0.548)	$[-\infty, \infty]$	2.509	0.133 (1.652)	$[-7.315, 7.366]$
Tiffany (Q)	1.381	-0.028 (0.617)	$[-\infty, \infty]$	2.509	1.782 (1.869)	$[-4.068, 23.704]$
Comité Colbert	1.740	0.092 (0.349)	$[-31.988, 0.693]$	0.226	-0.128 (2.819)	$[-\infty, \infty]$
Charitable contributions	0.151	2.040 (2.105)	$[-\infty, \infty]$	1.626	1.476 (1.025)	$[-\infty, \infty]$

these assets, the first column reports the first-stage F -statistic, the second reports the point estimate and standard error of the EIS using TSLS, and the third reports the 95% confidence interval constructed from the conditional LR test (see Yogo (2004) for details on the implementation of robust confidence intervals for the EIS).

In Panel B, we report the estimates of the EIS for consumption of luxury goods. First, note that the first-stage F -statistic is always less than 10, which indicates that inference based on TSLS is unreliable. This leads to uninformative confidence intervals when the EIS is estimated from stock returns. Using the T-bill rate, which is somewhat more predictable than stock returns, we are able to obtain informative confidence intervals for our series on luxury retail sales. The confidence intervals include rather large values of the EIS; for instance, the upper end of the confidence interval using the quarterly series on luxury retail sales is 7.4. It is possible that the EIS is larger for the consumption of luxuries. While not strong, this evidence complements that of Attanasio and Browning (1995) and Vissing-Jørgensen (2002), who show that the estimated EIS rises with the level of consumption or wealth of the household.

C. Equity Premium Implied by the Prices of Luxury Goods in Fixed Supply

In addition to using data on sales, we use the price movements of high-end luxury goods that are in perfectly inelastic supply to evaluate the equity premium. Intuitively, when a luxury good is in fixed supply, its price rises when excess returns are positive as household demand for the goods increases. When there is no increase in supply, this price change can be used as a measure of the change in marginal utility. We show how to use the covariance of excess returns and the prices of luxury goods in fixed supply to construct the implied equity premium. We assume that the utility of each series is additively separable from that of all other consumption, as we did for luxury and basic consumption in previous results.

When $L_t = L_{t+1} = L$, the stochastic discount factor for these goods is β (i.e., $M_{t+1}^L = \beta$), so that the Euler equation (10) becomes

$$E \left[\frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] = 0.$$

Using the definition of covariance and rearranging, this implies that the equity premium is given by

$$E[R_{t+1} - R_{t+1}^f] = - \frac{\text{Cov}[P_t/P_{t+1}, R_{t+1} - R_{t+1}^f]}{E[P_t/P_{t+1}]}. \quad (15)$$

Note that equation (15) does not give information about risk aversion. Instead, our data on the prices of luxury goods directly imply an equity premium independent of preference parameters.

C.1. Manhattan Pre-war Coops

We construct price series from two types of goods that are plausibly in fixed supply. The first is quarterly observations on the implied rental prices of

Manhattan pre-war coops, which we have obtained from Miller Samuel Inc., a real estate appraisal company in Manhattan. Rents represent the price of flow consumption derived from real estate. Manhattan pre-war coops represent a close-to-ideal market for our analysis, since these apartments are in fixed supply due to their “pre-war” nature and enjoy clear luxury status in New York City, where they represent the high end of the real estate market.

We focus on four data series (dollar values in parentheses represent the average value of the apartments sold in the fourth quarter of 1999): (1) all pre-war coops in Manhattan (\$630,356); (2) all pre-war coops in Manhattan with four or more bedrooms (\$3,393,750); (3) all pre-war luxury (defined as Central Park West, Park Avenue, and Fifth Avenue) coops (\$2,256,618); and (4) all pre-war luxury coops with four or more bedrooms (\$4,431,250). See Section D.1 of Appendix B for further details.

With real estate data, one might be concerned that there are significant adjustment costs for households to change their stock of housing, and this could reduce the estimated premium by reducing the correlation between price and returns. Furthermore, due to the length of typical rental leases and the presence of price controls in a limited segment of the New York market, rental prices might not immediately adjust to innovations in marginal utility. However, the effect of adjustment costs is mitigated by the fact that housing is an asset, so its price should reflect expected future demand. One could address this concern by inferring rents from sales prices.

Figure 4A plots the time series of the growth rate in the price of luxury pre-war coops along with excess returns. Although there is a significant amount of high frequency noise, there is significant covariation of price and stock returns. For instance, the poor stock market performance in the third quarter of 1990 is matched by a sharp decrease in the price of coops that quarter. As reported in Table V, the correlation of the price growth rate and excess returns is about 0.33.

Table V reports the equity premium estimated by using equation (15). For the series on luxury coops with four or more bedrooms, which is the most luxurious market, the implied equity premium is 3.9%. We also report the estimated equity premium that has been corrected for time aggregation (as determined in Section II) in the last column of Table V. After the adjustment, the equity premium using the price of luxury coops with four or more bedrooms is 7.8%. This is roughly in line with the historical equity premium.

For comparison, the canonical consumption CAPM calibrated to aggregate consumption data with a conservative relative risk aversion $\phi = 10$ implies an equity premium of a mere 0.2%, or 0.4% after adjusting for time aggregation, as reported in the first row of Table V.

C.2. U.S. Auction Prices of Fine Bordeaux Wine

The second series we use for the prices of luxury goods is the price of fine Bordeaux wines. We create quarterly indexes from raw data, provided by Orley Ashenfelter and David Ashmore at Liquid Assets, on cases of wine sold in U.S.

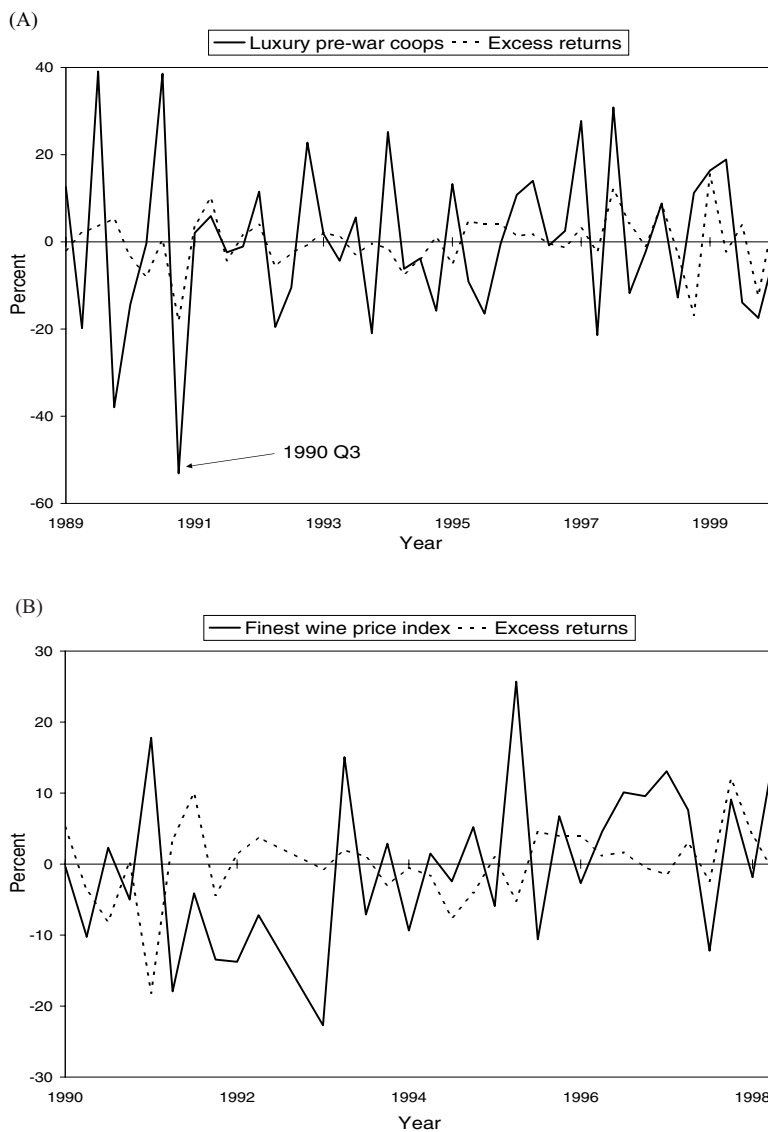


Figure 4. Response of the price of luxury goods in inelastic supply to stock returns. Shown here is a time series plot of excess stock returns (CRSP NYSE-AMEX portfolio over 3-month T-bills) and the growth rate for (A) price of luxury (Central Park West, Park Avenue, and Fifth Avenue) Manhattan pre-war coop apartments and (B) finest Bordeaux wine price index. All series are normalized to have zero mean and are reported in percent.

auctions from 1989 to 1997. We construct three indexes: fine, finest, and great. The fine index contains wines from the nine best *château*. The finest index contains only the best five of the nine *château*, and the great index only contains the best two. (For details, see Section D.2 of Appendix B.) To give an idea of the

Table V
Equity Premium Implied by the Price of Luxury Goods
in Inelastic Supply

The table reports the descriptive statistics for the price inflation of luxury goods in inelastic supply. Correlation is with respect to excess stock returns. Standard deviation (SD) is annualized by multiplying the quarterly value by 2. The last two columns report the equity premium (percent per year) estimated by the method of moments. The first estimate is the conventional one that uses the sample covariance, and the second estimate corrects for time aggregation in price data by a factor of 2. Standard error is in parentheses.

Series	Period (Obs.)	Correlation	SD	Equity Premium (%)	
				Sample	Corrected
PCE nondurables and services ($\phi = 10$)	1947–2001 (219)	0.137	0.108	0.218 (0.109)	0.437 (0.219)
Manhattan coops	1988–1999 (45)	0.095	0.294	0.370 (0.540)	0.740 (1.080)
Manhattan coops (4+ bed)	1988–1999 (45)	0.265	0.405	1.554 (1.129)	3.108 (2.257)
Manhattan luxury coops	1988–1999 (45)	0.329	0.373	1.679 (1.152)	3.358 (2.303)
Manhattan luxury coops (4+ bed)	1988–1999 (45)	0.333	0.635	3.920 (2.301)	7.841 (4.603)
Fine wine	1989–1997 (32)	-0.249	0.225	-0.573 (0.521)	-1.147 (1.042)
Finest wine	1989–1997 (32)	-0.210	0.225	-0.456 (0.473)	-0.912 (0.946)
Great wine	1989–1997 (32)	-0.157	0.214	-0.342 (0.367)	-0.684 (0.734)

quality of these wines, the average price of a case of wine in the fine index was over \$2,200 in 1997. For the finest and great indexes, it was over \$2,600.

Figure 4B plots the growth rate for the finest wine price index along with excess returns. The price of these wines is quite noisy, but there does not appear to be much covariation with returns. As reported in Table V, the correlation is actually -0.21 . This implies a negative equity premium. For the fine and great wine price indexes, we similarly estimate a negative equity premium. This indicates a rejection of some combination of our assumptions necessary for identifying the equity premium.

Some factors may explain why data on top wines does not imply a significant equity premium. One assumption that may be violated for fine wines is fixed supply since at least some top wine is ultimately consumed. In addition, there are other fine wines and alcoholic beverages that are close substitutes for the Bordeaux vintages we consider. These close substitutes may be subject to production shocks and be more responsive to demand shifts. To the extent that when wealth levels rise, the price increase of the good is limited by the increase in the supply of close substitutes, our method may underestimate the equity premium implied by the price movements of luxury goods.

D. Luxury Consumption and the Cross-section of Expected Returns

In this section, we evaluate the ability of luxury consumption to price the average returns on different portfolios of stocks instead of just the aggregate excess market return. Stocks that have a high (low) covariance with the consumption of luxury goods should have a high (low) average return. If luxury consumption is a better measure of the marginal utility of stockholders than nondurable consumption, then it should do a better job of explaining observed differences in average returns across portfolios.

We first consider the ability of our luxury retail sales series to explain the pattern of average returns across the 25 Fama–French portfolios. The Fama–French portfolios are the intersections of 5 quintiles of stocks based on size (market equity, ME) and 5 quintiles based on the ratio of book equity to market equity (B/M) (see Fama and French (1992) and (1993)). These portfolios have large and persistent differences in returns over time.

We estimate a linear factor model in which the average excess return on a given portfolio i is a linear function of the beta for each portfolio

$$E[R_{i,t+1}] = \beta_i \lambda, \quad (16)$$

where

$$\beta_i = \frac{\text{Cov}[f_{t+1}, R_{i,t+1}]}{\text{Var}[f_{t+1}]} \quad (17)$$

and f_{t+1} represents the factor(s) that we are using to price the portfolios.¹² Keeping the same timing convention, the excess return on a portfolio, $R_{i,t+1}$, represents the excess return on portfolio i during the quarter $t + 1$.

Our main factor of interest is growth in luxury consumption, $f_{t+1} = \Delta l_{t+1}$, and we focus on luxury retail sales, our longest measure of luxury consumption sales. We also compare these results to those using the aggregate market return, as suggested by the CAPM, and nondurable consumption, $f_{t+1} = \Delta c_{t+1}$, as suggested by the typical application of the CCAPM. Finally, as a benchmark, we consider the Fama–French three-factor model. Fama and French (1993) construct three factors to explain the cross-sectional variation in returns in the 25 portfolios: the overall market return (denoted R^m), the difference between returns on small and large stocks (“small minus big,” denoted SMB), and the difference between returns on high and low B/M stocks (“high minus low,” denoted HML). Although this model fits the returns well, it is not a structural model of behavior. Section D of Appendix A describes our estimation methodology.

Table VI reports the estimated λ , standard errors, and fit of equation (16) in terms of pricing errors. The second column of results for each model tests the model by estimating equation (16) with an intercept. Since this estimate should equal zero, its statistical and economic significance provides a test of the model.¹³

¹² This formula follows from equation (10) for the excess return on each portfolio.

¹³ The intercept is denominated in terms of percent returns and can be interpreted as the degree to which the risk-free rate predicted by the model exceeds the mean of the T-bill rate.

Table VI
Cross-Sectional Asset-Pricing Test

The table reports the cross-sectional tests of the CAPM, the Fama–French three-factor model, the CCAPM, and the luxury CCAPM. Luxury consumption is measured by luxury retail sales. The test assets are the 25 Fama–French portfolios sorted by size and book-to-market equity at (A) annual and (B) quarterly frequency. The first five rows report the factor risk premia estimated from a cross-sectional regression of average returns onto the estimated betas. The sixth row reports the intercept, if included in the regression. Standard error is in parentheses. The last two rows report the mean-absolute-pricing error (MAE) and the root-mean-squared-pricing error (RMSE).

Factor	CAPM		Three-Factor		CCAPM		Luxury	
Panel A: Annual 1961–2001								
Market	0.082 (0.031)	−0.041 (0.064)	0.058 (0.027)	0.002 (0.060)				
SMB			0.032 (0.023)	0.033 (0.023)				
HML			0.062 (0.024)	0.063 (0.024)				
Basic consumption					0.048 (0.070)	0.011 (0.012)		
Luxury consumption							0.152 (0.098)	0.155 (0.173)
Intercept (%)		13.110 (5.609)		5.410 (5.410)		7.289 (3.382)		−0.169 (7.872)
MAE (%)	3.177	2.716	1.131	1.128	4.054	2.467	2.160	2.152
RMSE (%)	3.826	3.016	1.421	1.387	4.735	2.974	2.735	2.735
Panel B: Quarterly 1987–2001								
Market	0.021 (0.012)	−0.014 (0.018)	0.020 (0.011)	−0.036 (0.019)				
SMB			0.002 (0.008)	−0.003 (0.008)				
HML			0.010 (0.010)	0.008 (0.010)				
Basic consumption					−0.002 (0.003)	0.001 (0.002)		
Luxury consumption							0.138 (0.130)	−0.053 (0.110)
Intercept (%)		3.762 (1.536)		5.789 (1.492)		2.368 (1.237)		3.161 (1.557)
MAE (%)	0.853	0.407	0.486	0.340	2.314	0.441	0.676	0.428
RMSE (%)	1.086	0.513	0.611	0.472	2.401	0.607	0.922	0.608

Panel A demonstrates that in annual data, the luxury CCAPM does a better job at pricing the returns on the Fama–French portfolios than the CAPM or the nondurable CCAPM. The relationship between luxury consumption betas and average returns has greater statistical significance than either the relationship between nondurable consumption betas and returns or the relationship between market betas and returns. Furthermore, the luxury CCAPM has lower pricing errors. Not surprisingly, all three one-factor models are dominated by

the Fama–French three-factor model. However the estimate of the intercept from the luxury CCAPM is both statistically and economically the closest to zero, as required by the theory.

Importantly, these estimates imply reasonable levels of risk aversion. The $\hat{\lambda}$ for the CCAPM is 0.048 and from Table I, the standard deviation of basic consumption is 0.023, implying an estimate of risk aversion of 91. On the other hand, for luxury consumption, $\hat{\gamma} = 0.152/(0.095)^2 = 17$, which is more reasonable and not far above our estimates of risk aversion reported in Table II. However, we caution the reader that this estimate has a large amount of statistical uncertainty, since $\hat{\lambda}$ is not even statistically different from zero.

Panel B reports results using quarterly data. The luxury consumption factor has smaller pricing errors than the CAPM or nondurable CCAPM, but the statistical significance of the CAPM factor is slightly higher and the intercept estimates are similar across models. Using quarterly data, we cover the period 1987 to 2001, which is not only short but is also an unusual period for the stock market. It is also the case that the effect of the slow consumption adjustment is more significant at the quarterly frequency (a wealth shock may not translate into an adjustment of consumption immediately), which is not an issue for the CAPM (see Parker (2003) for evidence on returns and the slow adjustment of consumption).

Figures 5 and 6 visually demonstrate these results. In annual data, the luxury CCAPM prices returns better than nondurable consumption growth or total returns. These figures plot the average returns of each portfolio (vertical axis) against the return predicted by the fitted model (horizontal axis) for each model and period. If the models were true, asymptotically, all portfolios would lie on the 45° line. The pricing errors are the horizontal difference between this line and each portfolio.

In addition to working with the Fama–French portfolios, we constructed and analyzed luxury beta portfolios following the methodology that Pastor and Stambaugh (2003) use to construct liquidity portfolios. That is, using rolling windows, we calculated consumption betas using luxury consumption for all stocks, grouped stocks into deciles based on these estimated betas, and computed value-weighted portfolio returns on these portfolios. We experimented with several choices of window length: 8 years, 10 years, and assuming a single grouping for the entire sample. Unfortunately, we found that the results were not robust across these choices. Nor were the results on balance supportive of the luxury CCAPM model, although the nondurable CCAPM fared no better. Either there is insufficient temporal stability in stock-specific covariances with luxury consumption for this exercise to be informative or these results provide some evidence against our model. Supportive of the former interpretation, in one case the portfolio formed from the highest beta stocks had a lower luxury consumption beta than the portfolio formed from the lowest beta stocks, and in no case did the order of the luxury consumption betas match those of the grouping. This suggests that the grouping is quite random and that the variation in returns across portfolios is not systematically related to actual luxury consumption betas.

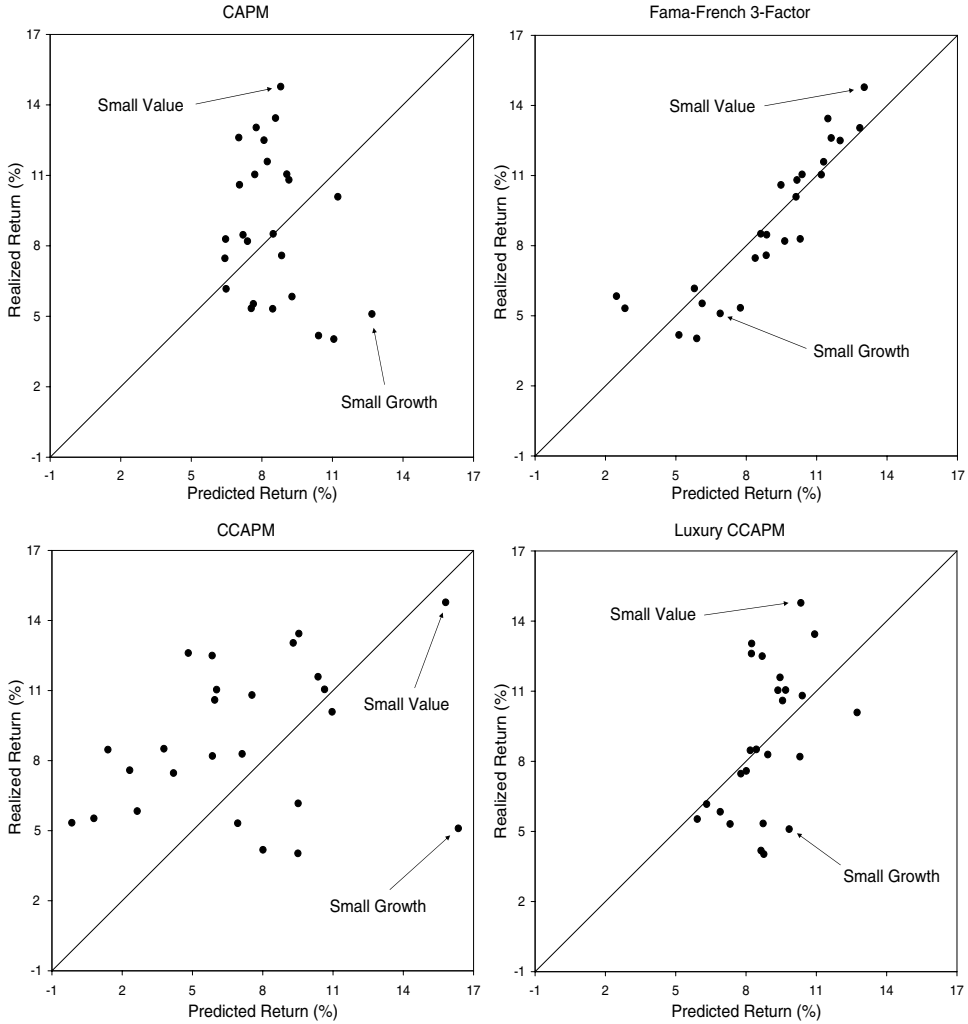


Figure 5. Realized versus predicted excess returns (without intercept). Realized versus predicted excess returns for the CAPM, the Fama–French three-factor model, the CCAPM, and the luxury CCAPM. Luxury consumption is measured by luxury retail sales. Black dots represent annual returns (1961–2001) on 25 Fama–French portfolios sorted by size and book-to-market equity. Pricing errors are generated by a cross-sectional regression of average returns onto the estimated betas.

We conclude that luxury consumption betas explain more of the variation in average returns across the Fama–French portfolios than betas constructed either from NIPA nondurable consumption (the usual CCAPM) or from total market returns (the CAPM). This provides further evidence that luxury consumption measures the marginal utility of rich households, which bear market risk, better than aggregate nondurable consumption. The results from the

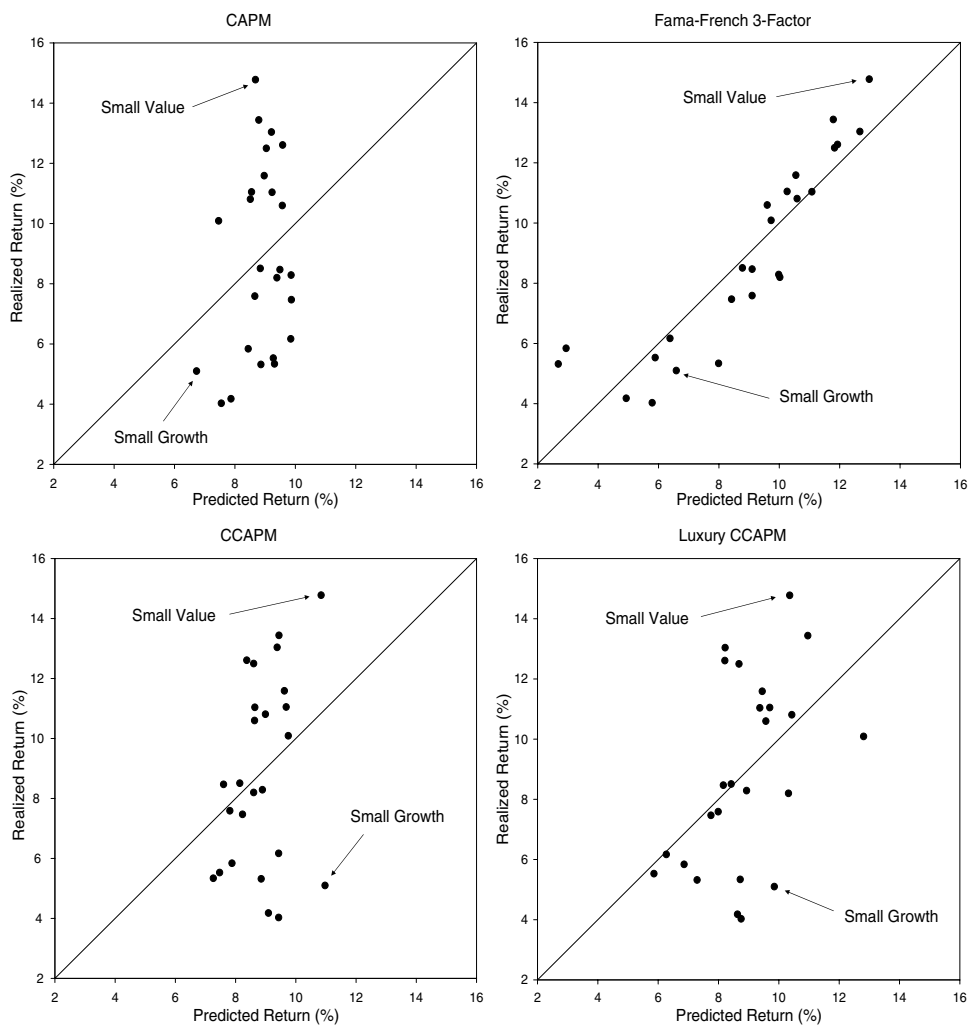


Figure 6. Realized versus predicted excess returns (with intercept). See notes to Figure 5. Pricing errors are generated by a cross-sectional regression of average returns onto the estimated betas and a constant.

construction of luxury beta portfolios, however, suffer from noise in the estimated consumption betas, which makes them unstable over time.

V. Conclusions

Evaluating the risk of equity for a given household requires measuring the marginal utility of that household. We argue that aggregate consumption fails to measure the marginal utility of the representative agent because the poor are quite risk averse, and the rich do not vary their consumption of basic goods,

only their consumption of luxury goods. Furthermore, many U.S. households, particularly those with low net worth, do not participate in the stock market. The consumption of these households should not vary much with stock returns, and aggregate consumption, therefore, does not measure the marginal risk of investing in the stock market.

By postulating a nonhomothetic utility function, and constructing series on the consumption of high-end luxury goods, we show how small modifications to the basic paradigm can go a long way towards reconciling the observed equity premium and the marginal utility of the very rich.

We find that their marginal utility moves significantly with the return on equity. The covariance of luxury goods and excess returns implies a coefficient of relative risk aversion more than an order of magnitude lower than that implied by NIPA consumption. Our main point estimates suggest a level of risk aversion only slightly higher than most economists would believe plausible. Confidence intervals contain plausible estimates of risk aversion.

While this paper shows that the marginal utility of the rich moves enough with the market to justify the large equity premium, the fact remains that the covariance of consumption and stock returns is low for the typical household that only consumes basic goods. If these households are stock owners, our Euler equation for basic goods implies that the risk aversion for these households is implausibly high, leading to the equity premium puzzle. If these households are not investing in stocks, then nonparticipation becomes a puzzle. For a household whose marginal utility moves little with returns, stocks are nearly riskless assets that deliver high returns.

Appendix A: Derivations

A. Limiting Consumption Shares

We claim that

$$\lim_{\tilde{X} \rightarrow \infty} \frac{C}{\tilde{X}} = 0,$$

$$\lim_{\tilde{X} \rightarrow \infty} \frac{PL}{\tilde{X}} = 1.$$

Let $\tilde{C} = C - a$, $\tilde{L} = L + b$, and $\tilde{X} = X - a + Pb$. Then the intratemporal first-order condition is $\tilde{C}^{-\phi} = \tilde{L}^{-\psi}/P$, and the budget constraint is $\tilde{X} = \tilde{C} + P\tilde{L}$. It follows that

$$\frac{\tilde{C}}{\tilde{X}} = \frac{1}{1 + P^{1-1/\psi} \tilde{C}^{\phi/\psi-1}}, \quad (\text{A1})$$

$$\frac{P\tilde{L}}{\tilde{X}} = \frac{1}{1 + P^{1/\phi-1} \tilde{L}^{\psi/\phi-1}}. \quad (\text{A2})$$

As $\bar{X} \rightarrow \infty$, either $\tilde{C} \rightarrow \infty$, in which case equation (A1) implies that $\frac{\tilde{C}}{\bar{X}} \rightarrow 0$, or \tilde{C} is bounded and so $\frac{\tilde{C}}{\bar{X}} \rightarrow 0$. Finally, since $\lim_{\bar{X} \rightarrow \infty} \frac{X}{\bar{X}} = 1$ and $0 < \tilde{C} < C$,

$$0 = \lim_{\bar{X} \rightarrow \infty} \frac{\tilde{C}}{\bar{X}} = \lim_{\bar{X} \rightarrow \infty} \frac{X}{\bar{X}} \frac{C}{\bar{X}} - \frac{X}{\bar{X}} \frac{a}{\bar{X}} = \lim_{\bar{X} \rightarrow \infty} \frac{C}{\bar{X}}.$$

Similar arguments demonstrate the claim in equation (8).

B. First-Order and Envelope Conditions in the Presence of Two Types of Consumption Goods

With the period utility function $v(C, L)$ in (6) written as $v(C, L) = \varpi(C) + \nu(L)$, the value function $J_t(W_t)$ satisfies

$$J_t(W_t) = \max_{\{C_t, L_t, \omega_t\}} \{ \varpi(C_t) + \nu(L_t) + E_t[\beta J_{t+1}(W_{t+1})] \}, \tag{A3}$$

where $W_{t+1} = (W_t - C_t - P_t L_t) \tilde{R}_{t+1}$. The optimal controls $C_t^*(W_t), L_t^*(W_t)$, and $\omega_t^*(W_t)$ are the solutions of the three first-order conditions with respect to the three controls $\{C_t, L_t, \omega_t\}$

$$\begin{cases} \varpi'(C_t) - \beta E_t[J'_{t+1}(W_{t+1}) \tilde{R}_{t+1}] = 0 \\ \nu'(L_t) - \beta E_t[J'_{t+1}(W_{t+1}) P_t \tilde{R}_{t+1}] = 0 \\ E_t[J'_{t+1}(W_{t+1})(W_t - C_t - P_t L_t)(R_{t+1} - R_{t+1}^f)] = 0. \end{cases} \tag{A4}$$

Replacing the optimal controls in (A3) yields

$$J_t(W_t) = \varpi(C_t^*(W_t)) + \nu(L_t^*(W_t)) + E_t[\beta J_{t+1}(W_{t+1}^*(W_t))], \tag{A5}$$

where

$$\begin{aligned} W_{t+1}^*(W_t) &\equiv (W_t - C_t^*(W_t) - P_t L_t^*(W_t)) \tilde{R}_{t+1}^*(W_t), \\ \tilde{R}_{t+1}^*(W_t) &\equiv R_{t+1}^f + (R_{t+1} - R_{t+1}^f) \omega_t^*(W_t). \end{aligned}$$

Differentiating (A5) with respect to the state variable W_t then yields

$$\begin{aligned} J'_t(W_t) &= \varpi'(C_t^*(W_t)) \frac{\partial C_t^*}{\partial W_t} + \nu'(L_t^*(W_t)) \frac{\partial L_t^*}{\partial W_t} \\ &+ E_t \left[\beta J'_{t+1}(W_{t+1}^*(W_t)) \left\{ \left(1 - \frac{\partial C_t^*}{\partial W_t} - P_t \frac{\partial L_t^*}{\partial W_t} \right) \tilde{R}_{t+1}^*(W_t) \right. \right. \\ &\left. \left. + (W_t - C_t^*(W_t) - P_t L_t^*(W_t)) (R_{t+1} - R_{t+1}^f) \frac{\partial \omega_t^*}{\partial W_t} \right\} \right], \end{aligned}$$

which after simplification using (A4) and the fact that all variables subscripted with t are contained in the information set at t , reduces to the envelope

conditions

$$\begin{aligned} J'_t(W_t) &= E_t[\beta J'_{t+1}(W_{t+1}^*)\tilde{R}_{t+1}^*(W_t)] \\ &= \beta \varpi'(C_t^*(W_t)) \\ &= \beta v'(L_t^*(W_t))/P_t. \end{aligned}$$

Evaluating the expressions for $J'_t(W_t)$ given by the envelope conditions at $t + 1$, and suppressing the superscript $*$ and the dependence of the optimal policies on current wealth, the system of first-order conditions (A4) becomes

$$\begin{cases} \varpi'(C_t) - \beta E_t[\varpi'(C_{t+1})R_{t+1}^f] = 0 \\ v'(L_t) - \beta E_t[v'(L_{t+1})P_t R_{t+1}^f] = 0 \\ E_t[J'_{t+1}(W_{t+1})(R_{t+1} - R_{t+1}^f)] = 0. \end{cases}$$

From this follow two sets of consumption conditional Euler equations:

$$\begin{aligned} E_t \left[\frac{\beta(C_{t+1} - a)^{-\phi}}{(C_t - a)^{-\phi}} R_{t+1} \right] &= E_t \left[\frac{\beta(C_{t+1} - a)^{-\phi}}{(C_t - a)^{-\phi}} R_{t+1}^f \right] = 1, \\ E_t \left[\frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}} \frac{P_t}{P_{t+1}} R_{t+1} \right] &= E_t \left[\frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}} \frac{P_t}{P_{t+1}} R_{t+1}^f \right] = 1, \end{aligned} \tag{A6}$$

which deliver equations (9) and (10).

C. Unconditional Euler Equation

Let

$$M_{t+1}^L \equiv \frac{\beta v'(L_{t+1})}{v'(L_t)} = \frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}}$$

denote the marginal rate of substitution for luxury consumption, which is the stochastic discount factor in this case. From (A6), we have that

$$1 = E_t \left[M_{t+1}^L \frac{P_t}{P_{t+1}} R_{t+1}^f \right] = E_t \left[M_{t+1}^L \frac{P_t}{P_{t+1}} \right] R_{t+1}^f,$$

since R_{t+1}^f is known at t . Thus,

$$E \left[M_{t+1}^L \frac{P_t}{P_{t+1}} \right] = E \left[\frac{1}{R_{t+1}^f} \right].$$

Then the unconditional version of (10) is

$$\begin{aligned} 0 &= E \left[M_{t+1}^L \frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] \\ &= E[M_{t+1}^L] E \left[\frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] + \text{Cov} \left[M_{t+1}^L, \frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right]. \end{aligned} \tag{A7}$$

It follows from the Taylor expansion

$$v'(L_{t+1}) \approx v'(L_t) + v''(L_t)(L_{t+1} - L_t)$$

that

$$M_{t+1}^L = \beta \frac{v'(L_{t+1})}{v'(L_t)} \approx \beta \left[1 + \frac{v''(L_t)}{v'(L_t)}(L_{t+1} - L_t) \right].$$

With our choice of utility function for luxury goods consumption, under the reasonable approximation that $b/L_t \ll 1$, it follows that

$$\frac{v''(L_t)}{v'(L_t)} = -\psi \frac{(L_t + b)^{-\psi-1}}{(L_t + b)^{-\psi}} = -\frac{\psi}{(L_t + b)} \approx -\frac{\psi}{L_t}.$$

Thus,

$$M_{t+1}^L \approx \beta \left[1 - \psi \frac{(L_{t+1} - L_t)}{L_t} \right] \approx \beta [1 - \psi \Delta \ln(L_{t+1})]. \tag{A8}$$

Substituting (A8) in (A7) yields

$$0 = E [M_{t+1}^L] E \left[\frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] - \beta \psi \text{Cov} \left[\Delta \ln(L_{t+1}), \frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right]$$

or

$$E \left[\frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] = \frac{\beta}{E[M_{t+1}^L]} \psi \text{Cov} \left[\Delta \ln(L_{t+1}), \frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right].$$

Finally, under $\beta/E [M_{t+1}^L] \approx 1$, we obtain equation (11).

D. Estimation of Cross-sectional Returns

Following the methodology of Fama and MacBeth (1973) and Fama and French (1992), we estimate equations (16) and (17) in two steps. First, each β_i is estimated directly from the empirical counterparts to the population moments in equation (17). This is simple to implement since $\hat{\beta}_i$ is the slope coefficient from a time series regression of return i onto a constant and the growth in the luxury goods sales. Tests and inference are constructed using a moment approach as described in Cochrane (2000, pp. 241–242). The moment vector is given by

$$g_t(b) = \begin{bmatrix} R_t - a - \beta f_t \\ f_t \otimes (R_t - a - \beta f_t) \\ R_t - \alpha I - \beta \lambda \end{bmatrix},$$

where

- $R_t = N \times 1$ vector of excess returns,
- $f_t = F \times 1$ vector of factors,
- $I = N \times 1$ vector of ones,
- $\alpha = N \times 1$ vector of parameters,
- $\beta = N \times F$ matrix of betas,
- $\lambda = F \times 1$ vector of factor risk premia,
- $\alpha =$ scalar parameter,

and

$$b = \begin{bmatrix} \alpha \\ \text{vec}(\beta) \\ \lambda \\ \alpha \end{bmatrix}.$$

E. Conditional Euler Equation

The conditional Euler equation (A6) is

$$1 + \varepsilon_{t+1} = M_{t+1}^L \frac{P_t}{P_{t+1}} R_{t+1}, \quad E_t [\varepsilon_{t+1}] = 0.$$

Taking logs, it follows from (A8) that

$$\ln(1 + \varepsilon_{t+1}) \approx \ln \beta - \psi \Delta \ln(L_{t+1}) - \Delta \ln(P_{t+1}) + \ln(R_{t+1}),$$

which itself is approximated (under $|\varepsilon_{t+1}| \ll 1$, we have $\ln(1 + \varepsilon_{t+1}) \approx \varepsilon_{t+1}$) by

$$\varepsilon_{t+1} \approx \ln \beta - \psi \Delta \ln(L_{t+1}) - \Delta \ln(P_{t+1}) + \ln(R_{t+1}).$$

Rearranging gives equation (14), with $r_{t+1}^L = \ln(R_{t+1}) - \Delta \ln(P_{t+1})$.

Appendix B: Data

A detailed description of the source and our use of each data series follows. All standard data series were downloaded from DRI Webabstract unless otherwise noted.

A. Government Data

A.1. NIPA Data

The series on PCE nondurables and services is the sum of real (chained) “PCE nondurables” and “PCE services,” divided by the “population used to calculate per capita income.” Using the current dollar value for PCE nondurables and services, we back out the implicit price deflator. The series is available at an annual frequency since 1929, quarterly since 1946, and monthly since 1959. The

series on real “PCE jewelry and watches” and PCE boats and aircraft (“PCE pleasure boats and aircraft”) are available at annual, quarterly, and monthly frequency since 1959. We divide these series by the population to compute the per-capita consumption levels and use their price deflators for the relative price of luxuries.

Unfortunately, PCE jewelry and watches includes many nonluxury items. For instance, most of the consumption of watches is unlikely to represent high-end luxury. Hence, we also use the real “retail sales of jewelry stores” published by the BEA as a measure of aggregate consumption of jewelry. This series and its implicit price deflator, used to compute the relative price, are available at monthly frequency since 1967.

A.2. U.S. Imports of Jewelry

Another measure of luxury consumption at an aggregate level is U.S. imports of jewelry (SITC 897, “jewelry, goldsmiths and other art of precious metals”), taken from the World Trade Analyzer CD. This data is reported to the United Nations Statistical Office and compiled by Statistics Canada. To isolate the luxury items, we only aggregate imports from France, Italy, and the United Kingdom. In 1999, France accounted for 3%, Italy for 88%, and the United Kingdom for 9% of the total imports from these three countries. Hence, the series that we construct is mainly driven by imports from Italy. Our choice of these European countries is motivated by our list of foreign luxury retailers, which is described below. To deflate the nominal value of imports, we use the U.S. import price index for SITC 897 (SITC 89 before 1985) from the Bureau of Labor Statistics. We also use this price index as the relative price of luxuries.

B. Quantity Data on Luxury Goods in Elastic Supply

B.1. U.S. Sales of Luxury Retailers

We initially targeted sales data for a list of seven United States and 25 European luxury retailers based on the list of luxury retailers contained in Morgan Stanley’s “Luxury Goods Weekly” (June 9, 2000) and Merrill Lynch’s report “Luxury Goods” (June 16, 2000). Of the 32 companies in our list, we consider the sales data for two U.S. retailers and five European retailers. The U.S. retailers are Saks (1991) and Tiffany (1960). The European retailers are Bulgari (1992), Gucci (1991), Hermès (1992), LVMH (1993), and Waterford Wedgwood (1994). The years in parentheses indicate the first year for which we are able to obtain U.S. sales data for these retailers. Nine of the companies in our list are not public and hence do not disclose sales information. Six of the companies have been public for less than five years, and hence we do not have enough observations to reliably measure correlations. Nine of the companies did not respond to our (repeated) requests for information. The remaining company is Neiman Marcus, whose sales data we have since 1984. However, we do not use

Neiman Marcus because their fiscal year ends in July rather than in December or January like all the other retailers. Hence, we cannot reliably aggregate their sales data with our other retailers.

The sales data for these companies were for the most part collected from the annual reports. For European retailers that report sales in foreign currencies, they are converted to U.S. dollars using the average exchange rate over the fiscal year. For the three companies whose equity trade in U.S. exchanges—Gucci (GUC), Saks (SKS.Z 1991 to 1997), and Tiffany (TIFF 1960–1977, TIF since 1986)—we were able to obtain sales data from COMPUSTAT. For Saks, we only take the part of sales attributed to Saks Fifth Avenue Stores, which is the only luxury component of the parent company Saks Inc. Before Saks Fifth Avenue Stores was acquired by Saks Inc. (formerly Proffitt's) in 1998, it was part of Saks Holdings Inc. In the late 1970s and early 1980s, Tiffany was owned by Avon Products and hence did not report sales separately. However, we were able to obtain sales data for the period 1978–1982 from *The New York Times* (1983). For the three companies that trade in the United States—Gucci (1996), Saks (1995), and Tiffany (1986)—we were able to construct a series on quarterly U.S. sales with data from quarterly reports and COMPUSTAT. Again, the years in parentheses indicate the first year in which data are available.

The three companies that trade in the U.S. report on a fiscal year that ends in January, whereas the remaining retailers that trade in European stock exchanges report on a fiscal year that ends in December. Since the bulk of sales occur during the holiday season in December, the 1-month difference in report dates can be ignored in aggregation across retailers. (We obtained essentially the same results isolating sales from only the three companies that trade in the United States.) To give a flavor of the cross-section of our dataset, in 2001 Saks accounted for 49% of total sales, Tiffany for 16%, LVMH for 12%, Gucci for 10%, Waterford Wedgwood for 8%, Hermès for 3%, and Bulgari for 2%.

Since sales data start in different years for different retailers, we compute growth rates in sales over the same set of retailers to assure that our series is as consistent as possible. For the quarterly data series, seasonality is a dominant feature of the data, mostly due to holiday purchases. Hence, we obtain deseasonalized growth rates by computing growth rate with respect to the same quarter in the previous year. This is how quarterly growth rate of sales is typically computed in the retail industry. Since jewelry is the main line of business for many of the companies on our list, we use the implicit price deflator for retail sales of jewelry stores as the deflator for nominal sales since 1967. Before then, we use the price index for PCE jewelry and watches due to data availability. These price series are also used to compute the relative price of luxuries.

B.2. Comité Colbert: U.S. Imports from French Luxury Retailers

Comité Colbert is a consortium of 70 French companies that specialize in luxury products. We collected data on their total sales as U.S. exports from 1984 to 1998. Among the Comité Colbert members, the 60 companies

with U.S. sales are Baccarat, Bernardaud, Champagne Bollinger, Boucheron, Breguet, Bussière, Caron, Céline, Chanel, Parfums Chanel, Château Cheval Blanc, Château Lafite-Rothschild, Château d'Yquem, Christian Dior, Parfums Christian Dior, Christofle, D. Porthault, Daum, Ercuis, Faïenceries de Gien, Flammarion Beaux Livres, Givenchy, Parfums Givenchy, Guerlain, Guy Laroche, Hédiard, Hermès, Parfums Hermès, Jean Patou, Parfums Jean Patou, Jean-Louis Scherrer, Jeanne Lanvin, John Lobb, Champagne Krug, La Chemise Lacoste, Lalique, Lancôme, Parfums Lanvin, Champagne Laurent-Perrier, Lenôte, Léonard, Champagne Louis Roederer, Louis Vuitton, La Maison du Chocolat, Mauboussin, Mellerio dits Meller, Nina Ricci, Parfums Nina Ricci, Pierre Balmain, Pierre Frey, Puiforcat, Rémy Martin, Revillon, Robert Haviland and C. Parlon, Rochas, Champagne Ruinart, Cristal Saint-Louis, Souleïado, S.T. Dupont, and Champagne Veuve Clicquot Ponsardin. The sales data in French francs were converted to U.S. dollars using the average exchange rate. The implicit price deflator of jewelry retail stores was used to deflate nominal sales and to compute the relative price of luxuries.

B.3. Charitable Contributions by the Wealthy

As a proxy for charitable contributions by the wealthy, we use data on the average contributions for households with AGI over \$1 million, which is taken from the IRS publication *Individual Income Tax Returns*. The data is available biannually from 1952 to 1972 and annually since 1973. The nominal values are deflated by the price index for PCE nondurables and services. The relative price for charitable contributions is its tax price, that is $1 - \tau$, where τ is the marginal tax rate. For each year, we compute the marginal tax rate for households with AGI over \$1 million as

$$\tau = \frac{\text{Tax}_1 - \text{Tax}_{0.5}}{\text{AGI}_1 - \text{AGI}_{0.5}},$$

where Tax_1 ($\text{Tax}_{0.5}$) is the average tax paid per capita for households with AGI over \$1 million (AGI \$0.5 to \$1 million) and AGI_1 and $\text{AGI}_{0.5}$ are the correspondingly defined average AGIs for each group. Since stock returns are noisy, the tax adjustment makes little difference for the results.

B.4. Luxury Automobile Sales

We have obtained data on total U.S. sales of imported luxury automobiles from *Ward's Automotive Yearbook*. We aggregate the quantities sold for BMW and Mercedes since 1970 and for Jaguar and Porsche since 1962. We have aggregated quantities sold in this way under the assumption that these brands are close substitutes. To compute the relative price of luxuries, we use BEA's implicit price deflator for retail sales of automotive dealers.

C. Instruments for Expected Returns

The asset returns that we use are real returns on stocks (NYSE-AMEX portfolio) and the risk-free rate (3-month T-bill). Both are deflated using the

relevant price deflator for luxuries. As for instruments, the yield spread is the difference between Moody's Seasoned Aaa Corporate Bond Yield and the 1-month T-bill rate from CRSP's Fama Risk-Free Rates File. The log dividend-price ratio is computed from CRSP data on returns including and excluding distributions.

D. Price Data on Luxury Goods in Inelastic Supply

D.1. Manhattan Pre-war Coops

We start with the closing prices of pre-war coops in Manhattan at a quarterly frequency since 1989. We consider four series: (1) all pre-war coops in Manhattan, (2) all pre-war coops in Manhattan with four or more bedrooms, (3) all luxury (Central Park West, Park Avenue, and Fifth Avenue) pre-war coops, and (4) all luxury pre-war coops with four or more bedrooms. Since the price data that we have is the price recorded at the time of closing, there is a delay between the time a sales price is negotiated and the time the price is recorded. Hence, we lag the price series two quarters, which is the time frame recommended by our data provider Miller Samuel Inc. In other words, we assume that the recorded closing price in the third quarter of 1999 is the effective price of real estate in the first quarter of 1999.

While direct observations of rents in Manhattan are available, the rental market there generally covers the entry to mid-level apartments, rather than the luxury market that we target. Since prices of apartments are closely related to rents, we use the gross rent monthly multiplier, assumed to be constant during our period, to map changes in the sales price into changes in the rental price.

The correction for time aggregation reflects the fact that in equation (15), the ratio P_t/P_{t+1} is the ratio of spot prices, whereas we observe prices that have been averaged over apartments sold each quarter. The adjustment factor compensates for the fact that time aggregation biases the estimated equity premium downwards by a factor of two.

D.2. U.S. Auction Prices of Fine Bordeaux Wine

The fine and finest are the Ashmore-Ashenfelter indexes for the fine and finest wines, reconstructed so as to reflect hammer price per dozen 750 ml bottles at U.S. auctions only. For more information on the raw data, we refer the reader to www.liquidassets.com. The "finest" index covers wines from Château Lafite, Latour, Margaux, Mouton, and Cheval Blanc. The "fine" index also Leoville Lascasses, Palmer, and Pichon Lalande. The "great" index covers only the top two: Lafite and Latour. All three indexes use wines from quite good vintages only: 1961, 1966, 1970, 1975, 1978, 1982, 1983, 1985, and 1986. The price index is constructed from regressions of log price on year, month, vintage, and château dummies. The series are log of nominal price of a constant-quality basket of wines.

REFERENCES

- Abel, Andrew, 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* 80, 38–42.
- Andreoni, James, and John Miller, 2002, Giving according to GARP: An experimental test of the consistency of preferences for altruism, *Econometrica* 70, 737–753.
- Atkeson, Andrew, and Masao Ogaki, 1996, Wealth-varying intertemporal elasticities of substitution: Evidence from panel and aggregate data, *Journal of Monetary Economics* 38, 507–534.
- Attanasio, Orazio, James Banks, and Sarah Tanner, 2002, Asset holding and consumption volatility, *Journal of Political Economy* 110, 771–792.
- Attanasio, Orazio, and Martin Browning, 1995, Consumption over the life cycle and over the business cycle, *American Economic Review* 85, 1118–1137.
- Bakshi, Gurdip S., and Zhiwu Chen, 1996, The spirit of capitalism and stock-market prices, *American Economic Review* 86, 133–157.
- Brav, Alon, George M. Constantinides, and Christopher C. Geczy, 2002, Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, *Journal of Political Economy* 110, 793–824.
- Breeden, Douglas T., Michael R. Gibbons, and Robert H. Litzenberger, 1989, Empirical test of the consumption-oriented CAPM, *Journal of Finance* 44, 231–262.
- Browning, Martin, and Thomas F. Crossley, 2000, Luxuries are easier to postpone: A proof, *Journal of Political Economy* 108, 1022–1026.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.
- Campbell, John Y., 1999, Asset prices, consumption, and the business cycle, in John B. Taylor, and Michael D. Woodford, eds.: *Handbook of Macroeconomics* (Elsevier Science Publishers, New York).
- 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, 205–251.
- Carroll, Christopher D., 2000, Why do the rich save so much?, in Joel B. Slemrod, ed.: *Does Atlas Shrug? The Economic Consequences of Taxing the Rich* (Harvard University Press, Cambridge).
- Cochrane, John H., 2000, *Asset Pricing* (Princeton University Press, Princeton).
- Cogley, Timothy, 2002, Idiosyncratic risk and the equity premium: Evidence from the Consumer Expenditure Survey, *Journal of Monetary Economics* 49, 309–334.
- Constantinides, George M., 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.
- Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes, 2000, Do the rich save more? Working paper, Columbia University.
- Epstein, Larry G., and Stanley E. Zin, 1991, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis, *Journal of Political Economy* 99, 263–286.
- Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.
- Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene F., and James MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.
- Grossman, Sanford J., and Robert J. Shiller, 1981, The determinants of the variability of stock market prices, *American Economic Review* 71, 222–227.
- Güvenen, Muhammet F., 2000, Mismeasurement of the elasticity of intertemporal substitution: The role of limited stock market participation, Working paper, Carnegie Mellon University.
- Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339–357.
- Hanoch, Giora, 1977, Risk aversion and consumer preferences, *Econometrica* 45, 413–426.

- Heaton, John, and Deborah J. Lucas, 1996, Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.
- Heaton, John, and Deborah J. Lucas, 2000, Stock prices and fundamentals, in Ben S. Bernanke, and Julio J. Rotemberg, eds.: *NBER Macroeconomics Annual 1999* (MIT Press, Cambridge).
- Houthakker, Hendrik S., 1960, Additive preferences, *Econometrica* 28, 244–257.
- Huggett, Mark, and Gustavo Ventura, 2000, Understanding why high income households save more than low income households, *Journal of Monetary Economics* 45, 361–397.
- Mankiw, N. Gregory, 1982, Hall's consumption hypothesis and durable goods, *Journal of Monetary Economics* 10, 417–425.
- Mankiw, N. Gregory, and Stephen P. Zeldes, 1991, The consumption of stockholders and nonstockholders, *Journal of Financial Economics* 29, 97–112.
- Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.
- Moreira, Marcelo J., 2003, A conditional likelihood ratio test for structural models, *Econometrica* 71, 1027–1048.
- Neely, Christopher J., Amlan Roy, and Charles H. Whiteman, 2001, Risk aversion versus intertemporal substitution: A case study of identification failure in the intertemporal consumption capital asset pricing model, *Journal of Business and Economic Statistics* 19, 395–403.
- Nelson, Charles R., and Richard Startz, 1990, Some further results on the exact small sample properties of the instrumental variable estimator, *Econometrica* 58, 967–976.
- The New York Times*, 1983, At Tiffany, a troubled transition, October 16, 1.
- Parker, Jonathan A., 2002, The consumption risk of the stock market, *Brookings Papers on Economic Activity* 2, 279–348.
- Parker, Jonathan A., 2003, Consumption risk and expected stock returns, *American Economic Review Papers and Proceedings* 93, 376–382.
- Parker, Jonathan A., and Christian Julliard, 2003, Consumption risk and cross-sectional returns, NBER Working paper 9538.
- Pastor, Lubos, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.
- Poterba, James M., 2000, Stock market wealth and consumption, *Journal of Economic Perspectives* 14, 99–118.
- Poterba, James M., and Andrew A. Samwick, 1995, Stock ownership patterns, stock market fluctuations, and consumption, *Brookings Papers on Economic Activity* 2, 295–357.
- Shiller, Robert J., 1982, Consumption, asset markets, and macroeconomic fluctuations, *Carnegie Mellon Conference Series on Public Policy* 17, 203–238.
- Staiger, Douglas, and James H. Stock, 1997, Instrumental variables regression with weak instruments, *Econometrica* 65, 557–586.
- Stanley, Thomas J., and William D. Danko, 1998, *The Millionaire Next Door: The Surprising Secrets of America's Wealthy* (Simon and Schuster, New York).
- Stiglitz, Joseph E., 1969, Behavior towards risk with many commodities, *Econometrica* 4, 660–667.
- Stock, James H., and Jonathan H. Wright, 2000, GMM with weak identification, *Econometrica* 68, 1055–1096.
- Stock, James H., Jonathan H. Wright, and Motohiro Yogo, 2002, A survey of weak instruments and weak identification in generalized method of moments, *Journal of Business and Economic Statistics* 20, 518–529.
- Vissing-Jørgensen, Annette, 2002, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825–853.
- Yogo, Motohiro, 2004, Estimating the elasticity of intertemporal substitution when instruments are weak, *Review of Economics and Statistics* 86, 797–810.