

NBER WORKING PAPER SERIES

LUXURY GOODS AND THE EQUITY PREMIUM

Yacine Aït-Sahalia
Jonathan A. Parker
Motohiro Yogo

Working Paper 8417
<http://www.nber.org/papers/w8417>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2001

We thank Angus Deaton for helpful discussions. We thank Jonathan Miller at Miller Samuel, Inc. for information on New York real estate prices and Orley Ashenfelter and David Ashmore at Liquid Assets for their raw data on wine prices. Preliminary versions of this paper were called “The Equity Premium is not So Puzzling for the Very Rich.” Aït-Sahalia and Parker thank the National Science Foundation (grants SBR-9996023 and SES-0096076 respectively) for financial support. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

©2001 by Yacine Aït-Sahalia, Jonathan A. Parker and Motohiro Yogo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Luxury Goods and the Equity Premium
Yacine Aït-Sahalia, Jonathan A. Parker and Motohiro Yogo
NBER Working Paper No. 8417
August 2001
JEL No. G12, E21, D12

ABSTRACT

This paper evaluates the return on equity using novel data on the consumption of luxury goods. Specifying household utility as a nonhomothetic function of the consumption of both a luxury good and a basic good, we derive and evaluate the riskiness of equity in such a world. Household survey and national accounts consumption data overstate the risk aversion necessary to match the observed equity premium because they contain basic consumption goods. The risk aversion implied by equity returns and the consumption of luxury goods is more than an order of magnitude less than found using national accounts consumption data. For the very rich, the equity premium is much less of a puzzle.

Yacine Aït-Sahalia
Department of Economics
Princeton University
Princeton, NJ 08544-1021
and NBER
yacine@princeton.edu

Jonathan A. Parker
Department of Economics
Princeton University
Princeton, NJ 08544-1021
and NBER
jparker@princeton.edu

Motohiro Yogo
Department of Economics
Harvard University
Cambridge, MA 02138
yogo@fas.harvard.edu

1 Introduction

Over the past century in the United States, the average return on the stock market has exceeded the return on short-term government bonds by over 6 percentage points at an annual rate. Economists have tried to understand this equity premium by appealing to the risks inherent in such an investment. However, as initially demonstrated by Mehra and Prescott (1985) and then by the extensive literature that followed, such a large premium has proven difficult to explain within the canonical consumption-based asset pricing framework. The source of the puzzle is clear: according to the model, asset returns should alter investors' marginal utility; their marginal utility is a function of their consumption; but aggregate consumption moves little with unpredictable returns. As a result, inordinately high degrees of risk aversion are necessary to reconcile the low variability of consumption with the high volatility of stock returns.

One possible solution to the puzzle is to modifying the specification of the investors' marginal utility. The canonical time-separable power utility function is abandoned in favor of specifications constructed so that marginal utility is more responsive to asset returns. Prominent examples of this approach include Constantinides (1990), Epstein and Zin (1991), Bakshi and Chen (1996) and Campbell and Cochrane (1999).

A second approach, initiated 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. The households that do not hold stock – due to uninsurable income risk or costs of investing in stocks – contaminate tests of the canonical theory that employ aggregate consumption data. In theory, the consumption of wealthy households who hold equity can be used to evaluate the canonical theory. Consumption data on wealthy households or stock owners confirms that these households do bear more market risk, but this additional risk is insufficient to save the model. Estimates of risk aversion, while lower than for the representative agent, are still unrealistically high.

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 model utility as a function of multiple goods and demonstrate that reasonable levels of risk aversion are consistent with the consumption of luxury goods and the equity premium. As in the first 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. Specifically, utility is a function of both the consumption of basic goods, of which a certain amount is required in every period, and the consumption of luxury goods, of which none are consumed when expenditure levels are not high. Given such nonhomothetic preferences, the share of luxury goods in overall consumption rises with total expenditures.

Households display a high degree of risk aversion with respect to their consumption of basic goods, consistent with a *subsistence* aspect of 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 a more *discretionary* aspect of luxuries. We derive the Euler equation associated with the consumption of each type of good. Our theory implies that households are less risk averse with respect to the consumption of luxury goods, so that the equity premium puzzle – the high degree of risk aversion implied by 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: is the covariance of returns and marginal utility measured by luxury consumption sufficient to justify the equity premium? By studying luxuries, we evaluate the behavior and risk aversion of rich households.

Since no extant datasets measure consumption of high-end luxury goods, we construct time series data on luxury consumption from sales data for luxury goods instead of household surveys of consumer spending or national accounts data. Household surveys typically contain few wealthy households and measure categories of consumption that do not distinguish between basic and luxury goods. While we also evaluate readily-available government statistics, we construct separate time series on sales of high-end luxury goods: luxury automobile purchases; U.S. sales of the following luxury retailers: Saks, Tiffany, Bulgari, Gucci, Hermès, Louis Vuitton, Moët & Chandon, Hennessy (the latter three part of the luxury conglomerate LVMH), and Waterford Wedgwood; aggregated U.S. imports from the Comité Colbert, a consortium of seventy French luxury good manufacturers (including Hermès and LVMH for whom we also obtained disaggregated U.S. sales data); and charitable contributions by high-income households. We define a luxury retailer as those companies listed by Morgan Stanley and Merrill Lynch in their analysts' reports on the luxury goods retail sector. Note that many of the luxury retailers whose names we do not list individually in our data appendix are owned by luxury powerhouses such as LVMH, for whom we have total U.S. sales data. At last count (in 2000), LVMH owns 46 different luxury brand names, whose sales represent 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 aggregate consumption does. Our estimates of the coefficient of relative risk aversion are an order of magnitude lower than those found using National Income and Product Accounts (NIPA) aggregate consumption data. For example, different series on high-end luxury retail sales yield point estimates for risk aversion ranging from 3.2 to 24.5 while consumption of nondurables and services in the NIPA yields a point estimate of 169. Given moderate sampling uncertainty, we cannot reject that completely reasonable levels of risk aversion generated the observed data on the consumption of luxury goods.

Figure 1 depicts this main result. Panel *A* plots the excess returns of U.S. stocks over Treasury Bills against consumption growth of nondurables and services from the U.S. NIPA data, and against consumption growth as measured by the U.S. sales of luxury retailers. Panel *B* plots the time series for these same three series. Aggregate consumption of basic goods is almost non-responsive to stock excess returns. By contrast, the consumption of luxuries is both more volatile than that of NIPA consumption and much more correlated with excess returns. This series is our flagship series, collected from U.S. sales data reported in the annual reports of luxury retailers. The estimated coefficient of relative risk aversion is 3.2. Some of our other measures of luxury consumption are not as highly correlated with stock excess returns, but all series on luxury consumption lead to strikingly lower estimates of risk aversion than NIPA data.

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. The increases in expenditure four years after an excess return implies even lower risk aversion than the contemporaneous movement in expenditures.

We analyze two additional features of luxury data. First, we estimate conditional Euler equations to measure intertemporal substitution and compare our results again with those from NIPA data. Estimated intertemporal elasticities of substitution are on average larger for the luxury consumption series than for NIPA consumption, but the variation across series is substantial. Second, we construct time series on the *prices* of three high-end luxury goods whose supplies are highly inelastic: pre-War Manhattan coop apartments, central London real estate and Bordeaux wines from the finest Chateaux and years. We calculate the equity premium from these price indexes under the assumption

that the supplies of these goods are fixed. The price series for luxury goods in fixed supply do not rationalize the equity premium as well as the quantity series for the luxury goods in elastic supply. Since this test relies on there being no close substitutes in more elastic supply, we think it inferior to the quantity indexes.

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 basic model. In particular, within the basic power utility paradigm, there is no equity premium puzzle for the households that hold most US equity, or at least one not easily explained by sampling uncertainty. Our estimates of risk aversion imply that the risk faced by wealthy households is sufficient to justify the typical return on equity.¹

The balance of the paper is laid out as follows. Section 2 lays out our assumptions about nonhomothetic utility and explains how the presence of multiple goods and luxury consumption changes inference. Section 3 presents the estimating equations. The main results are in Section 4, which describes the data that we gather to estimate risk aversion and the results of estimation on each series. In section 5, we push our findings in three different directions. First, we present results that deal with the issue of durability; second we ask what our data on luxury consumption imply for the intertemporal elasticity of substitution; third, we estimate the equity premium based on the price of those luxury goods whose supply is inelastic. The final section concludes. We include appendices with a complete description of our data series and derivations omitted from the main text.

2 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.

¹It could also well be that more sophisticated utility specifications such as those incorporating habit formation, but applied in a two-good world, could go even further towards reconciling the equity premium.

2.1 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_t :

$$\max_{X_t, \omega_t} E \left[\sum_{t=0}^{\infty} \beta^t u(X_t) \right] \quad (2.1)$$

$$s.t. \quad \sum_{t=0}^{\infty} (\omega_t R_{t+1} + (1 - \omega_t) R_{t+1}^f) (A_t - X_t) \geq 0 \quad (2.2)$$

$$X_t \geq 0 \quad (2.3)$$

where:

- $u()$ is the period utility function and is increasing, concave, and twice differentiable
- $\beta \in (0, 1)$ is the discount factor
- A_t is household wealth at the start of period t
- R_{t+1} is the gross real return on stocks between time t and $t + 1$
- R_{t+1}^f is the gross real return on Treasury bills between t and $t + 1$ (assumed to be conditionally risk free)

Note that for simplicity, and consistency with the canonical model, we are assuming that households are infinitely lived, that leisure is additively separable from consumption, and that markets are complete so that labor income risk can be completely diversified. 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

$$J(A_t | I_t) = \max_{\{X_t, \omega_t\}} \left\{ u(X_t) + E_t \left[\beta J(\tilde{R}_{t+1}(A_t - X_t) | I_{t+1}) \right] \right\} \quad (2.4)$$

where J denotes the value function, I_t 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 the constraints

(2.2)-(2.3). The envelope and first-order conditions imply the conditional moment restriction

$$E_t \left[\frac{\beta u'(X_{t+1})}{u'(X_t)} (R_{t+1} - R_{t+1}^f) \right] = 0. \quad (2.5)$$

To evaluate the model's predictions, the traditional approach is to assume that every agent's utility function is of the constant relative risk aversion class and that markets are complete. In this case, the consumption of all agents move together, and one can evaluate the implications of this moment condition using aggregate consumption data.

Within this canonical model, the equity premium can only be explained by appealing to unappealingly high risk aversion. As first shown by Mehra and Prescott (1985), 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 fifteen years of research and shows that the puzzle is robust across countries and time.

2.2 Nonhomothetic Preferences

Our point of departure from the canonical model is to drop the single-good assumption and model the within-period utility function in greater depth. We assume that households consume two types of goods: basic consumption, C , and luxury goods, L . We conceptualize the former, which we treat as the numeraire in the economy, as the standard goods that most households in the U.S. regularly consume and that make up the bulk of the National Income and Product Account (NIPA) measures of consumption. The latter, luxury goods, are consumed only by the extremely rich.

We adopt a two-good specification for the household within-period utility function that encompasses our concepts of "luxury" and "standard" goods. We reinterpret the previous statement of the problem as follows. X represents total expenditure in period t , 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 suppressed), and the direct utility function $v(C, L)$, which

we assume for simplicity is additively separable,² is

$$X = C + PL \tag{2.6}$$

$$v(C, L) = \frac{(\text{Max}\{0, C - a\})^{1-\phi}}{1 - \phi} + \frac{(L + b)^{1-\psi}}{1 - \psi} \tag{2.7}$$

where P is the relative price of luxury goods, a , b , γ , and ψ are positive constants, $\psi < \phi$, and we add the constraint that $C > 0$ (implying from (2.3) that $L > 0$).

This specification of utility captures two features of basic and luxury goods. First, luxury goods are not consumed by the poor and middle class: there exists $\underline{C} = a + b^{\frac{\psi}{\phi}} P^{\frac{1}{\phi}} > a$ such that $L = 0$ for all $C \leq \underline{C}$. That is, when the marginal utility of wealth is high, the agent chooses to consume none of the luxury. Note that the assumption $a, b > 0$ assures that the minimal amount of basic consumption at which a consumer starts consuming luxury goods is greater than the sustenance level.

Second, the consumption of the rich is dominated by luxuries:

$$\begin{aligned} \lim_{X \rightarrow \infty} \frac{C}{X} &= 0 \\ \lim_{X \rightarrow \infty} \frac{PL}{X} &= 1 \end{aligned} \tag{2.8}$$

We prove this claim in Appendix A.1. The assumption that $\psi < \phi$ implies that as the marginal utility of wealth goes to zero, the budget share of the luxury approaches one.

An example of this expenditure behavior is presented in Figure 2. The limit behavior of the expenditure shares are governed by ψ and ϕ so that 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 and the fact that $-a < b$ delivers basic consumption in excess of luxury consumption at low expenditure levels.

²Modelling a nonseparability between C and L is unlikely to matter for our estimation. In order to overturn our argument and findings, the marginal utility of luxuries would have to decrease when the consumption of basic goods rises. Thus the stochastic discount factor would not covary much with market returns and the puzzle would remain even though luxury goods covary strongly with market returns. Such a nonseparability would have to be large, because basic consumption moves very little with market returns. Such conditions seem unlikely although they are maintained in the canonical model.

2.3 Euler Equations and Risk Aversion

We show in Appendix A.2 that the first-order condition and envelope condition from the dynamic program for the choice of C and L imply the following two sets of consumption conditional Euler equations:

$$\begin{aligned} E_t \left[\frac{\beta(C_{t+1} - a)^{-\phi}}{(C_t - a)^{-\phi}} R_{t+1} \right] &= E_t \left[\frac{\beta(C_{t+1} - a)^{-\phi}}{(C_t - a)^{-\phi}} R_{t+1}^f \right] = 1 \\ E_t \left[\frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}} \frac{P_t}{P_{t+1}} R_{t+1} \right] &= E_t \left[\frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}} \frac{P_t}{P_{t+1}} R_{t+1}^f \right] = 1 \end{aligned} \quad (2.9)$$

from which it follows that

$$E_t \left[\frac{\beta(C_{t+1} - a)^{-\phi}}{(C_t - a)^{-\phi}} (R_{t+1} - R_{t+1}^f) \right] = 0 \quad (2.10)$$

$$E_t \left[\frac{\beta(L_{t+1} + b)^{-\psi}}{(L_t + b)^{-\psi}} \frac{P_t}{P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] = 0. \quad (2.11)$$

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

The focus of the previous literature is on equation (2.5), or, if one takes the view that luxuries are not contained in NIPA nondurable consumption and services, then on equation (2.10). We instead focus on estimation and testing of equation (2.11). Equation (2.11) provides a test of whether the consumption Euler equation holds for wealthy households.

Our choice of utility function implies that the risk aversion of the rich is a lower bound for the risk aversion of the typical household.³ The coefficient of relative risk aversion with respect to gambles over C is $\gamma_C(C) = \phi \frac{C}{C-a}$ and so falls with C and $\lim_{C \rightarrow \infty} \gamma_C(C) = \phi$. Risk aversion with respect to gambles over L is $\gamma_L(L) = \psi L / (L + b)$ and so increases with L and $\lim_{L \rightarrow \infty} \gamma_L(L) = \psi$. Risk aversion with respect to gambles over X , which is the risk aversion coefficient that previous papers consider, is a weighted average of $\gamma_L(L)$ and $\gamma_C(C)$. $\gamma(X)$ declines with X and $\infty > \gamma(C) \geq \psi$. We estimate the level of risk aversion of the extremely rich, that is we estimate ψ . Thus our estimates provide a lower bound on the risk aversion over wealth gambles in the population.⁴

³To be clear, risk aversion refers to local risk aversion according to the Arrow-Pratt definition: $\frac{-Xu''(X)}{u'(X)}$. The utility function is continuous between (a, ∞) , but is not differentiable where $X = C$. Our construction does have the desirable property that the limit of risk aversion from above and below $X = C$ are both $\gamma_C(C)$.

⁴The inverse of this observation is that this utility function leads to intertemporal elasticities that vary with wealth.

It is worth noting that risk aversion that varies with wealth is an inherent feature of any non-homothetic intra-period utility function. There is no utility function that admits nonhomothetic Engel curves and delivers constant relative risk aversion (Hanoch (1977)).⁵

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

2.4 Risk Aversion from Basic Consumption and Luxury Consumption

It is important to address why an econometrician using the usual consumption Euler equation (2.10), to study the return on stocks and bonds would not accept the model since the equation holds in equilibrium. Further, why does the luxury consumption Euler equation provide the correct test?

Since there is some share of consumption that is necessary, the intra-period utility function does not exhibit constant relative risk aversion. For low levels of consumption, households are extremely unwilling to subject consumption to risk and so hold little equity and have stable consumption. Thus any test using the consumption of all households calculates risk aversion from an 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 weighted more heavily in this average than their weight in wealth. According to our theory then, inference based on aggregate consumption data should find high levels of risk aversion. This argument implies that limiting the sample studied to rich households provides a test of our theory, but there are further reasons to think basic consumption is inappropriate for such a test.

This may or may not provide an explanation for the fact that rich households are observed to consume a lower share of their total wealth than low income households.

⁵See also Stiglitz (1969) and the discussion of the intertemporal elasticity of substitution in Browning and Crossley (2000).

While not explicit in our above specification, it is reasonable that marginal utility from the consumption of basic goods is bounded from above or reaches 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 choose to have luxury consumption react to market returns. If we modified our function $u(C)$ to exhibit this feature, calculations based on basic consumption growth would find high risk aversion due to the unresponsive consumption of the rich as well as the unresponsive consumption of the poor.

In addition to the direct implications of nonhomothetic preferences, there are two classes of extant theories that imply in our two-good world that basic consumption is inappropriate for pricing 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 model the poor as not holding stocks due to fixed costs of participating in the equities 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, non-rich households are not be willing to incur this cost to invest and so their wealth is not directly affected by returns on equity and their consumption covaries less with the market.⁶ The theory positing incomplete markets argues that the poor do not hold stocks since households face uninsurable idiosyncratic endowment risk. According to this theory, households that face significant risk from labor income choose to invest less or not at all in stocks because their consumption is already (uninsurably) quite risky (Heaton and Lucas (1996), Constantinides and Duffie (1996) and Brav, Constantinides and Geczy (1999)). As with limited participation models, this theory predicts a low covariance between the consumption of the poor or middle-class and stock returns. 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. Since

⁶For some evidence on this theory see Vissing-Jorgenson (1998), Attanasio, Banks, and Tanner (1998). Guvenen (2000) calibrates a model in which only some households have access to the equity market and in which there are two types of agents, high-risk aversion agents and low risk aversion agents. The paper demonstrates that inference based on aggregate consumption in the canonical manner implies an implausibly high risk aversion, but it is also true that the model has to assume an equity premium an order of magnitude smaller than observed.

the consumption of luxury goods measures the marginal utility of the very rich, both theories imply that one should recover reasonable measures of risk aversion when using the consumption of luxury goods.⁷

Second, and less prominent, 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 are not separable from the consumption of goods that have high costs associated with adjusting the level of consumption. For example, items like transportation or fuels (subcategories of NIPA nondurable consumption) are in part determined by a household's consumption of housing and automobiles, items which are subject to large adjustment costs and are infrequently adjusted. Items like mobile phone service, health club memberships, and the like involve a degree of commitment over time. 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 extremely rich are unlikely to face labor income risk that is significant relative to their wealth; the rich are willing and able to pay any fixed costs for market participation. Moreover, very rich households hold most equity and most hold equity. While the first part of this 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 non-human 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

⁷The 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 discouraged from doing so.

pension funds, and 75% of all stock wealth.⁸ 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.⁹ Of the top 1% of households ranked by non-human wealth, 82% hold stock directly; of the top five percent of households, 78% hold stock directly. For the population as a whole, less than 50% hold stock directly.¹⁰

3 Estimating Equations

We seek to evaluate the risk aversion of the rich using equation (2.11) 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.¹¹

Linearizing the unconditional version of the Euler equation for luxury goods, as in Campbell (1999), risk aversion can be derived as a function of theoretical moments (details are contained in Appendix A.3). Rather than assume that time is discrete at the frequency of the observed data, we assume that the world operates in continuous time leading to an estimating equation

$$\psi = \frac{2}{3} \frac{E \left[R_{t+1} - R_{t+1}^f \right] - Cov \left[\Delta \ln(P_{t+1}), (R_{t+1} - R_{t+1}^f) \right]}{Cov \left[\Delta \ln(L_{t+1}), (R_{t+1} - R_{t+1}^f) \right]} \quad (3.1)$$

where the factor of $\frac{2}{3}$ is the adjustment for continuous time (see Grossman, Melino, and Shiller (1987)). 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. But higher demand translates into

⁸These numbers come from the 1998 Survey of Consumer Finances as calculated and reported in Poterba (2000), Table 2.

⁹Anecdotally, 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 has risen sharply and the total wealth controlled by households with assets of at least \$1 million grew 313% to approximately \$8.8 trillion (including Canada, source: Merrill Lynch-Cap Gemini's 2000 World Wealth Report). During the same period, the US stock market rose by 405%.

¹⁰These figures are from Heaton and Lucas (1999).

¹¹The 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.

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.¹² On the basis of equation (3.1), a positive correlation between luxury price inflation and positive returns reduces, other things equal, the coefficient of relative risk aversion required to reconcile the equity premium and luxury consumption data. And a relatively inelastic supply of luxury goods tend to make the denominator of equation (3.1) small and the covariance in the numerator large and positive.

The empirical counterparts to the theoretical moments on the right-hand side of equation (3.1) are used to estimate risk aversion. Appendix B describes the construction of sample moments and standard errors for each estimating equation. In particular, we note that we use all available data to calculate each moment of interest. That is, for our shorter series, we do not mistakenly assume that the recent high returns were largely expected but instead calculate $E[R_{t+1} - R_{t+1}^f]$ from returns from 1947 to 2000.

In addition to analyzing risk aversion, we estimate the intertemporal elasticity of substitution from a linear version of equation (2.11):

$$\Delta \ln(L_{t+1}) = -\frac{\rho}{\psi} + \frac{1}{\psi} r_{t+1}^L - \frac{1}{\psi} \varepsilon_{t+1} \quad (3.2)$$

where, as shown in Appendix A.3, $\rho \equiv -\ln(\beta)$ is the discount rate and $r_{t+1}^L \equiv \ln(R_{t+1}) - \Delta \ln(P_{t+1})$ is the real rate of return in terms of the luxury good. While $E_t[\varepsilon_{t+1}] = 0$, it is in general the case that $E[r_{t+1}^L \varepsilon_{t+1}] \neq 0$. Hence estimation treats r_{t+1}^L as endogenous and uses instruments dated $t - 1$, so that we are estimating

$$\Delta \ln(L_{t+1}) = -\frac{\rho}{\psi} + \frac{1}{\psi} E_t[r_{t+1}^L] - \frac{1}{\psi} \varepsilon_{t+1} \quad (3.3)$$

and $E_t[E_t[r_{t+1}^L] \varepsilon_{t+1}] = 0$.

The difficulty in evaluating the risk of equity using these equations is that good measures of the consumption of luxury goods do not exist.

¹²Some luxury goods have highly inelastic supply so that $\Delta \ln(L_{t+1}) = 0$; we analyze them separately.

4 Risk Aversion and the Consumption of Luxury Goods

This section describes the construction of measures of the consumption of luxury goods and presents risk aversion estimated from each series using equations (3.1) and (3.3). A complete description of the source and our use of each series of luxury consumption is contained in Appendix C. Excess returns are calculated as the return on the *S&P* 500 less the return on the 3-month treasury bill.¹³ We exactly align returns with each of our measures of consumption growth, whether it be monthly, quarterly, or annual, rather than lagging it. This maintains the comparability of our results to previous empirical work.

We construct new series on the consumption of luxury goods because NIPA consumption data is not classified into luxury and basic consumption. Moreover, available household survey data is not suited for this task. While there are a host of issues that arise with all household surveys, the main shortcomings of the most used surveys are as follows. The Panel Study of Income Dynamics (PSID) measures only the consumption of food and housing, has only infrequent measures of wealth, and undersamples the wealthy. The Survey of Consumer Finances (SCF), while oversampling the wealthy, does not collect consumption data beyond the stock of some consumer durable goods and has very little panel dimension and a small time dimension. The Consumer Expenditure Survey (CEX) covers very limited categories of wealth, has poor measurement of those that it does cover, and topcodes both consumption and wealth (wealth consists of four categories each topcoded at \$100,000 for most of the survey). The burden of detailing all consumption, as the CEX requires, is so large as to suggest very few high wealth households are in the survey or provide a full accounting of consumption.¹⁴

¹³We previously used the value-weighted NYSE portfolio available from CRSP instead of the S&P 500 and the results were very similar. We use the S&P 500 when extending our time series because more recent data was more readily available.

¹⁴The SCF and PSID take great pains to minimize the costs to participants. Since the CEX is designed to collect the distribution of expenditures across goods (although not classified by luxuriousness), complete compliance is quite costly to participants.

4.1 Results from NIPA, BEA and Trade data

To begin, we examine publicly-available government series which provide some evidence on the risk aversion implied by the consumption of luxuries. These series are not entirely satisfactory, so, as described in the next subsection, we construct better series for measuring the consumption of luxury goods.

We first estimate risk aversion from NIPA data on consumption of nondurable goods and services during the post-War period. The first two rows of Table 1 present the details of this exercise. The columns report the number of observations, the correlation between excess returns and the series, the standard deviation of the series, the point estimate of risk aversion, ψ , and the standard error of the estimate, which we compute from the empirical moments as described in appendix B. As is well known, annual or quarterly aggregate consumption data imply implausibly high estimates of risk aversion, 1096 and 169 respectively.

Two subcategories of NIPA personal consumption expenditures (PCE) capture luxury consumption to some extent: “jewelry and watches” and “boats and airplanes.” Unfortunately, both of these series are expenditures on durable goods rather than consumption, and both series contain some consumption of basic goods. The consumption of watches includes a significant amount of non-luxury consumption, while boats and planes also includes expenditures on sports and photographic equipment. The advantage of these series however is their long time dimension. Our main series are less than half this length. The second through fifth rows of Table 1 show that these series are significantly more volatile than the basic consumption series. Consumption of jewelry and watches however is negatively correlated with excess returns, leading to a negative estimate of risk aversion (a rejection of the model). While the NIPA series on boats and planes leads to an implausibly high or negative risk aversion depending on the frequency. These results are consistent with the findings of Poterba and Samwick (1995), but we report these results to be clear on our starting point.

The last three rows of table 1 display our results using government data on retail sales and imports of jewelry. Data on retail sales of jewelry are available from the Bureau of Economic Analysis (BEA) at a monthly frequency. Data on the imports of jewelry are collected by the US International Trade Commission. To help isolate luxury jewelry, we sum the imports of jewelry to the U.S. from France, Italy, and the U.K. These series have shorter time dimensions than the NIPA data, but probably

provide better measures of luxury goods.

The monthly retail sales of jewelry has a 0.57 correlation with excess returns, still below that of quarterly NIPA consumption data, but a significantly higher volatility. This leads to a point estimate for risk aversion of 69, below that of the NIPA data but still implausibly extreme. Focussing more on luxury goods, last two rows report the result using U.S. imports of jewelry. Here, we find still high but somewhat lower risk aversion. Given a fair amount of statistical uncertainty, we are unable to reject reasonable levels of risk aversion.

We now turn to the analysis of several series of luxury goods that we construct ourselves and that provide better evidence on the importance of nonhomothetic utility for understanding the risks of equity.

4.2 Results from Sales of High-End Luxury Goods

In this section we present the results of analyzing three categories of luxury goods: luxury automobiles, sales of high-end luxury retailers, and (tax adjusted) charitable contributions of households with incomes over one million dollars.

We begin by measuring luxury consumption as the sales of luxury automobiles from Ward's Automotive Reports. This series contains a large amount of what we view as basic consumption and not luxury consumption, but we are able to break out sales of Porsches and consider this category separately. While more focussed on luxury goods, automobile sales measure expenditures on a durable good rather than flow consumption, to which our model refers. We postpone directly tackling the issue of durability, and instead present results from our automobile purchases in rows three and four of table 2. To the extent that retail sales of luxury automobiles measure the consumption of luxury goods, risk aversion is significantly lower than estimated from NIPA data. Whether or not these expenditures measure flow consumption, evidence for the importance of luxuries is given by the fact that risk aversion declines, as it did in Table 1, as we move from less to more luxurious measures of expenditures on otherwise similar goods.

All of the measures considered so far are imperfect along two dimensions. First, the measures include basic goods purchased by middle class households and do not focus purely on the super-rich, as our theory requires. Second, as noted, due to durability, these series may have a weak mapping between consumption expenditures and marginal utility. We partly deal with this second issue by

checking that our findings are not driven by the volatile response of expenditures to wealth for any of the series: in section 5.1, we redo our analysis with the change in expenditures over the four periods following the innovation to the market. As we discuss, our main conclusions stand, but this approach is not a perfect fix.

Thus we turn to sales from the extremely high-end market for luxury consumption goods directly. By doing so we are measuring by definition consumption of very expensive luxuries. Durability is also likely to be less of an issue as fashion is fickle: an Hermès tie, Gucci bag or designer dress lasts only one season for those who can afford them. We collect retail sales data from luxury retailers. We do this in two ways.

First, an organization of French high-end luxury producers collects information on the exports of 70 French producers of luxury goods to the United States. This group, Comité Colbert, shared with us the total annual exports to the United States from 1984 to 1998. This series covaries significantly with US excess returns. Figure 3 displays the time series of excess returns and many of our consumption series. Panel *A* shows the low covariance between consumption and excess returns that is at the heart of the equity premium puzzle. NIPA consumption varies little and covaries less. Panel *B* shows the series for Porsche sales and the significantly higher covariance of these sales with excess returns. Finally, panel *C* shows the strong relationship between excess returns and our series on the imports of luxury goods from the Comité Colbert. As reported in the fifth row of Table 2, consumption of high-end French luxury goods imply a risk aversion of the wealthiest households of only 12.7. While 12.7 is still implausibly large, the estimate is more than an order of magnitude lower than that estimated on NIPA data and could easily be biased upward by the slow adjustment of the flow of international goods.

Second, we collect data on U.S. sales from the annual reports of high-end luxury retailers. Of the 32 companies so defined by two major investment banks, we are able to collect data for two major U.S. retailers – Saks and Tiffany – and five European retailers – Bulgari, Gucci, Hermès, LVMH, and Waterford Wedgwood. We aggregate sales growth across these seven retailers weighting by market share to create a common series from an unbalanced panel of firm sales. The averaging lessens the measurement error which arises from the fact that individual companies may misprice products, produce poor products, and so forth, and so suffer sales movements not indicative of total consumption of luxury goods. The construction of this series requires assuming that the one-month

differences in fiscal years among companies listed on different exchanges does not affect our results.

Panel *D* of Figure 3 displays the close relationship over time between excess returns and sales data from these seven luxury retailers. The sixth row of Table 2, labelled Luxury Retail Sales presents the correlation and risk aversion estimated from this series. The results are striking: estimated risk aversion is 3.6, an estimate $1/300th$ and $1/47th$ respectively of the estimates from NIPA annual and quarterly consumption series.

To check that this series is not contaminated by the aggregation of sales data across slightly different reporting periods, we construct a consumption of luxuries series from only the sales of Gucci, Saks, and Tiffany, which trade on the NYSE and so report on a common fiscal year ending in January. This series, Luxury Retail Sales (US Retailers), has a correlation with excess returns of 0.356 and implies an estimate of risk aversion of 3.2. The data that generate such a plausible estimate of risk aversion are those presented in Figure 1 in the introductory section of the paper.

As our last measures of high-end luxury goods, we consider sales data for Tiffany by itself since it is the company for which we have the longest and most consistent 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. In addition, besides studying annual sales, we have been able to gather data on sales at the New York city flagship store, and data on US sales at the quarterly frequency. We are somewhat concerned about the shorter time horizon since by looking only at the last 14 years for the N.Y. series and 8 years for the quarterly data, some of the unexpected comovement of the two series, caused by high returns over the most recent period, may be treated as expected. To the extent this is true, we underestimate the covariance and overestimate risk aversion. Nevertheless, as shown in the next three rows of Table 2, these series, while suggesting higher risk aversion than the series for all luxury goods, still suggest a much lower risk aversion than typically found. Panel *E* of Figure 3 displays the shortest Tiffany time series, the quarterly sales data.

Our final measure of the consumption of luxuries captures more purely expenditures that give utility to the extremely wealthy. We construct a measure of luxury consumption from the charitable contributions of households with (adjusted gross) incomes over a million dollars. This is not a standard measure of consumption, but it is the leading theory to explain charitable bequests to model gifts of medical care to the suffering, chairs in economics, art to a museum, and so forth

provide “warm glow” utility to donors.¹⁵ One strength of this series is its length. The data is available from the Internal Revenue Service biannually from 1952 through 1972 and annually since 1973. The price of charitable giving is the tax-price of charitable giving and varies with the average marginal tax rate of the households in this category.

The final Panel of Figure 3 and the final row of table 2 shows that luxury consumption of charitable donations are highly correlated with returns, implying a level of risk aversion of 4.7. While slightly higher than most economists would believe, a coefficient of relative risk aversion of 2 is within one standard deviation of our point estimate.

In sum, our estimates based on high-end luxury consumption suggest an entirely different picture of the risk of equity. We turn now to robustness and extensions.

5 Extensions

This section first present results that suggest that durability is not driving the findings of the previous section. Second, our data on luxury consumption is used to estimate the conditional Euler equation and so provide estimates of the intertemporal elasticity of substitution of the very rich. The final subsection derives a method for estimating the equity premium implied by the price of luxury goods when these goods are in fixed supply. This method is applied to the price of pre-War Manhattan co-op apartments, central London real estate, and high-end Bordeaux wines.

5.1 Durability

As noted, many of the publicly available series measure the expenditures on 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 2.4. One might be concerned that this problem is present even in our measures of sales of high-end luxury goods. In this section we provide some evidence that durability is not driving our results.

¹⁵Indeed treating the provision of public goods as a commodity generating a “warm glow” passes revealed-preference tests (Andreoni and Miller (forthcoming)).

Suppose that utility comes from the service flow from the stock of a durable good, K_t . The durable good is related to expenditures as

$$K_{t+1} = (1 - \delta) K_t + L_t \quad (5.1)$$

where δ is the rate at which the durable depreciates. If there are no adjustment costs on the stock, then expenditures (L_t) 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 since risk aversion decreases with the covariance of expenditure growth and excess returns.

If the growth rate of consumption is stationary, equation (5.1) implies that the stock of the good, and so its service flow, is cointegrated with expenditures. This suggests that we look not at the instantaneous change in expenditures following an innovation to the market, but rather at a longer-run increases in expenditures. To the extent that a large positive return leads to an upwards revision in the stock of a durable, this is still apparent a few years later in higher expenditures. Thus we do exactly this. 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 alternative estimates of risk aversion from the equation¹⁶

$$\psi = \frac{E \left[R_{t+1} - R_{t+1}^f \right] - Cov \left[\ln \left(\frac{P_{t+4}}{P_t} \right), (R_{t+1} - R_{t+1}^f) \right]}{Cov \left[\ln \left(\frac{L_{t+4}}{L_t} \right), (R_{t+1} - R_{t+1}^f) \right]}. \quad (5.2)$$

Table 3 provides evidence that our main findings are not driven by durability. Panel *A* displays the results for the government statistics of Table 1 and shows that, if anything, the estimated coefficients of risk aversion are more reasonable. The absolute value of the coefficients of relative risk aversion tend to be lower, but the only estimate that is at all plausible is that of annual imports of jewelry, which is 5.

Turning to our main series, displayed in Panel *B* of Table 3, for every series except the sales of all luxury automobiles, the coefficients of risk aversion estimated from these long-differences are lower than those in the Table 2 and more plausible. The one estimate that seems driven by durability is the coefficient of relative risk aversion estimated from luxury automobile sales, which becomes negative. The remainder of Panel *B* shows that estimates of risk aversion are more reasonable than the baseline

¹⁶Note that we correctly make no continuous time adjustment.

results in Table 2. Despite their robustness to durability, these estimates are not necessarily better than those in Table 2. By taking a long difference we have reduced the effective time dimension of our series.

Our main results are not simply driven by the volatility of expenditures, but by a high covariance of long-lasting movements in sales and excess returns. We conclude that the consumption of luxury goods implies that much lower estimates of risk aversion are required to rationalize the premium on equity.

5.2 Results from the Conditional Euler Equation

We turn from measuring risk aversion from unexpected returns to estimating its inverse from the relationship between the growth in the consumption of luxury goods and predictable returns using the conditional Euler equation (3.3). As is well-known, the coefficient on the real interest rate in the conditional Euler equation estimates the intertemporal elasticity of substitution (IES). In preferences more general than those we specify, it is possible that the IES differs from the inverse of the coefficient of relative risk aversion.

Table 4 presents results from estimation of the conditional Euler equation on the same series of measures of luxury goods analyzed in the previous two sections. To predict returns, we use the second lags of: NIPA consumption growth, the return on equity, the return on the Treasury bill, and the log price-dividend ratio. The first column of results gives the estimated IES when the key independent variable is the predictable return on equities, and the second column gives the associated specification tests based on the orthogonality of the residuals and the instruments. The last two columns present the IES and specification test when the independent variable is the predictable variation in the return on Treasury bills. Given the concerns about durability, there is little information in the rows with the results based on the NIPA subcategories (“Jewelry . . .” and “Boats . . .”) as well as the results from jewelry and automobile sales. We present these results for completeness and focus on our high-end luxury sales measures and charitable contributions.¹⁷

The results based on high-end luxury retail sales typically find significantly higher IES than the NIPA data, with the exceptions of Tiffany NY series and the shorter quarterly series. Although we

¹⁷We also experimented with nonlinear GMM estimation. Estimates were unstable and the discount factor tended to be estimated greater than one. We suspect that this is due to the short time series available for many of our measures.

do not have long time series, we are able to obtain somewhat tight estimates for the intertemporal elasticity of substitution for many of our series. These results imply that wealthy households have higher IES than the typical household.¹⁸

While not the main focus of our paper, these results suggest that low estimates of the IES in aggregate data may in part be due to the use of basic consumption rather than luxuries. Indeed, there is significant evidence from household survey data that the IES rises with the level of consumption of the household (Attanasio and Browning (1995) , Vissing-Jorgensen (2001)). Economists have been less concerned by the low IES estimated on aggregate data than with the equity premium implied by the covariance of returns and consumption. But for many applications, it is likely to be important that wealthy households have higher IES than the typical household.

5.3 Results from the Prices of Luxury Goods Assumed in Fixed Supply

In addition to using data on sales, we can also 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. In this section, we use the covariance of excess returns and the prices of luxury goods in fixed supply to construct the implied equity premium.

Letting $L_t = L_{t+1} = L$, the stochastic discount factor for these goods is $\beta (M_{t+1}^L = \beta)$, so that the Euler equation (2.11) becomes

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

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

$$E[R_{t+1} - R_{t+1}^f] = - \frac{Cov \left[\frac{P_t}{P_{t+1}}, R_{t+1} - R_{t+1}^f \right]}{E \left[\frac{P_t}{P_{t+1}} \right]}. \quad (5.4)$$

Note that equation (5.4) does not give information about risk aversion. Instead, our data on the prices of luxury goods directly imply a given premium on equity that is independent of the specific

¹⁸See Vissing-Jorgenson (2001) and the discussion therein.

choice of within-period utility function.¹⁹

We construct price series from three sets of goods that are in close to fixed supply. Our first good is pre-War Manhattan Co-ops, which represent the high end of the Manhattan apartment market. Our data are price series for pre-War coops in Manhattan at an annual frequency from 1989 on. We consider four average price series: 1) all pre-War coops in Manhattan, 2) all pre-War coops in Manhattan with four or more bedrooms, 3) all pre-War luxury-location (Central Park West, Park Avenue, Fifth Avenue) coops, and 4) all pre-War luxury-location coops with four or more bedrooms.

The second price series we use measures real estate prices in central London. We employ returns on the London stock exchange's FTSE index to calculate the implied equity premium.

The final data on the price of luxury goods is the price of great Bordeaux wines. We create indexes from raw data provided by Ardmore and Ashenfelter based on cases of wine sold at US auctions from 1989 to 1997 at a quarterly frequency. The index includes only the best vintages: 1961, 1966, 1970, 1975, 1978, 1982, 1983, 1985, and 1986. The "fine" index includes the nine best château's Lafite, Latour, Margaux, Mouton, Cheval Blanc, Ducru Beaucaillou, Leoville Lascasses, Palmer, and Pichon Lalande. The "finest" series includes only the first five of these château's. The "great" index includes only Lafite and Latour. To give an idea of the quality, in 1997, the average price of a case of wine in the fine index is over \$2,200 and over \$2,600 for the finest and great index. More detailed information on each index is contained in Appendix C.

The weakness of each measure is that there are close substitutes to each product. A wealthy Manhattanite can live on Park Avenue West or choose to live in a mansion in Greenwich, CT or Chappaqua, NY, instead... There are fine California wines and new vintages that are close to those that we single out in quality, as well as other high-end alcoholic beverages. To the extent that when wealth levels rise, the price increase of the items that we study is limited by the increase in the supply of substitutes, our method underestimates the equity premium implied by these price movements.

Figure 4 plots the time series of these prices of luxury goods and excess returns. While the high frequency fluctuations in the series are not matched by returns, there is a fair amount of lower frequency correlation. That said, it is insufficient to rationalize the equity premium. Table 5 reports our findings for prices of pre-War Manhattan coops, London real estate, and fine wines. All price series on luxury real estate imply equity premia that are quite small, although still significantly larger

¹⁹Note also that we do not make a continuous time adjustment.

than would be implied by NIPA consumption data with risk aversion of unity. A concern with the real estate data is that there are significant adjustment costs for households changing their stock of housing, and this could reduce the estimated premium by reducing the correlation between price and returns.²⁰ Of course, the effect of adjustment costs are mitigated by the fact that housing is an asset so that its price should reflect future demand.

While the premia implied by real estate prices are small, the price series for wine are inconsistent with a positive premium on equity. This rejects some combination of our assumption of fixed supply or our model. Indeed, since much good wine is ultimately drunk, the assumption of fixed supply for this series may fail not only due to the presence of substitutes but due to consumption of wine.

Our findings using data on the prices of luxuries in fixed supply are inconsistent with the marginal utility of the rich rationalizing the equity premium.²¹ We suspect that our assumption of fixed supply does not properly account for close substitutes, but we present the results of our investigation and let the reader decide.

6 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 luxuries. We find that the marginal utility of the rich, who hold most US equity, moves significantly with the return on equity. The covariance of luxury goods and excess returns implies coefficients of relative risk aversion more than an order of magnitude lower

²⁰Based on information from the firm constructing the price indexes for NY co-ops, we lag the price series two quarters to best capture when the sale, rather than the closing takes place. While this may in general bias down our estimates, slightly different choices or working with annual data do not change our conclusions. The London real-estate series is lagged only one period and different choices lead to slightly lower equity premia.

²¹For high-end luxury goods in fixed supply, our assumptions imply a unit elasticity between the price change of these goods and the predictable variations in return in the stock market:

$$\Delta \ln(P_{t+1}) = -\rho + E_t [R_{t+1}] - \varepsilon_{t+1}. \quad (5.5)$$

Thus we can run a conditional Euler equation and test whether the coefficient on the returns is unity. Not surprisingly given the results on risk aversion, the coefficients are all below unity, and some are even negative.

than the covariance of NIPA consumption and excess returns. 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.

It is quite possible that our estimates are biased upward since our main series on luxury goods are significantly shorter than data typically used to estimate excess returns. During the second half of the 1990's, excess returns were unusually large and our estimate of the covariance of returns and luxury consumption might as a result be biased down as some of what is unexpected is measured as expected.²² Consistently with the empirical evidence on the faster increase of the income and wealth of top 1% relative to the rest of the population, the growth rate of the luxury goods market has far exceeded that of aggregate consumption over the past decade.

While the marginal utility of the rich moves nearly enough with the market to justify the equity premium, we do not completely answer why the basic consumption equation implies such high risk aversion for the bottom two thirds of the consumption distribution. Is it enough that necessities imply that these households have high risk aversion? Or are there sufficient costs to adjusting consumption to rationalize the lack of consumption movement? Alternatively, are background risks or costs to participating in the equity market substantial enough to explain the consumption and investment behavior of these households?

²²As detailed in the appendix, we take steps to minimize this.

References

- ANDERSON, T. W. (1971): *The Statistical Analysis of Time Series*, Wiley, New York.
- ANDREONI, JAMES AND JOHN MILLER (forthcoming): "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism," *Econometrica*.
- ATTANASIO, ORAZIO, AND MARTIN BROWNING (1995): "Consumption Over the Life Cycle and Over the Business Cycle," *American Economic Review*, 85(5), 1118-1137.
- ATTANASIO, ORAZIO, JAMES BANKS AND SARAH TANNER (1998): "Asset Holding and Consumption Volatility," NBER Working Paper No. 6567.
- BAKSHI, G. S., AND Z. CHEN (1996): "The Spirit of Capitalism and Stock-Market Prices," *American Economic Review*, 86(1), 133-57.
- BRAV, A., G. M. CONSTANTINIDES, AND C. GECZY (1999): "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," NBER Working Paper No. W7406.
- BROWNING, M., AND T. F. CROSSLEY (2000): "Luxuries are Easier to Postpone: A Proof," *Journal of Political Economy*, 108(5), 1022-26.
- CAMPBELL, J. Y. (1999): "Asset Prices, Consumption, and The Business Cycle," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1C, chap. 19, pp. 1231-1303. Elsevier Press.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205-51.
- CONSTANTINIDES, G. (1990): "Habit Formation: A Resolution of the Equity Premium Puzzle" *Journal of Political Economy*, 98(3), 519-43.
- CONSTANTINIDES, G. M., AND D. DUFFIE (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104(2), 219-40.
- DYNAN, K., J. SKINNER, AND S. 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(2), 263-86.
- GROSSMAN, SANFORD J., ANGELO MELINO, AND ROBERT SHILLER (1987): "Estimating the Continuous-Time Consumption-Based Asset-Pricing Model," *Journal of Business and Economic Statistics*, 5(3), 315-27.
- GUVENEN, MUHAMMET F. (2000): "Mismeasurement of the Elasticity of Intertemporal Substitution: The Role of Limited Stock Market Participation," working paper, Carnegie Mellon University.
- HANOCH, G. (1977): "Risk Aversion and Consumer Preferences," *Econometrica*, 45(2), 413-26.
- HEATON, J., AND D. LUCAS (1996): "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104(3), 443-87.
- HEATON, J., AND D. LUCAS (1999): "Stock Prices and Fundamentals," in *N.B.E.R. Macroeconomics Annual*, ed. by B. Bernanke, and J. Rotemberg, pp. 213-242. Cambridge: MIT Press.
- HUGGETT, M., AND G. VENTURA (2000): "Understanding Why High Income Households Save More Than Low Income Households," *Journal of Monetary Economics*, 45(2), 361-97.
- MANKIW, N. G., AND S. ZELDES (1991): "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics*, 29(1), 97-112.
- MEHRA, R., AND E. C. PRESCOTT (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15(2), 145-61.
- POTERBA, J. M. (2000): "Stock Market Wealth and Consumption," *Journal of Economic Perspectives*, 14(2), 99-118.

POTERBA, J. M., AND A. SAMWICK (1995): “Stock Ownership Patterns, Stock Market Fluctuations, and Consumption,” *Brookings Papers on Economic Activity*, (2), 295–357.

STIGLITZ, J. E. (1969): “Behavior Towards Risk with Many Commodities,” *Econometrica*, 4(37), 660–67.

VISSING-JORGENSEN, A. (1998): “Limited Stock Market Participation,” working paper, MIT.

VISSING-JORGENSEN, A. (2001): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution” working paper, University of Chicago.

Appendix

A Proofs and Derivations

A.1 Limiting Consumption Shares

We claim

$$\begin{aligned}\lim_{E \rightarrow \infty} \frac{C}{E} &= 0 \\ \lim_{E \rightarrow \infty} \frac{PL}{E} &= 1\end{aligned}$$

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

$$\frac{\tilde{C}}{\tilde{E}} = \frac{1}{1 + P^{1-1/\psi} \tilde{C}^{\phi/\psi-1}} \quad (\text{A.1})$$

$$\frac{P\tilde{L}}{\tilde{E}} = \frac{1}{1 + P^{1/\phi-1} \tilde{L}^{\psi/\phi-1}} \quad (\text{A.2})$$

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

$$0 = \lim_{\tilde{E} \rightarrow \infty} \frac{\tilde{C}}{\tilde{E}} = \lim_{E \rightarrow \infty} \frac{C}{E} + \frac{a}{E} = \lim_{E \rightarrow \infty} \frac{C}{E}.$$

Similar logic demonstrates the second claim in equation (2.8).

A.2 First Order and Envelope Conditions in the Presence of Two Types of Consumption Goods

With the period utility function $v(C, L)$ in (2.7) written as $v(C, L) = \varpi(C) + v(L)$, the value function $J_t(W_t)$ satisfies

$$J_t(W_t) = \max_{\{C_t, L_t, \omega_t\}} \{\varpi(C_t) + v(L_t) + E_t[\beta J_{t+1}(W_{t+1})]\} \quad (\text{A.3})$$

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 $\{C_t, L_t, \omega_t\}$

$$\begin{cases} \varpi'(C_t) - \beta E_t \left[J'_{t+1}(W_{t+1}) \tilde{R}_{t+1} \right] = 0 \\ v'(L_t) - \beta E_t \left[J'_{t+1}(W_{t+1}) P_t \tilde{R}_{t+1} \right] = 0 \\ E_t \left[J'_{t+1}(W_{t+1}) (W_t - C_t - P_t L_t) (R_{t+1} - R_{t+1}^f) \right] = 0 \end{cases} \quad (\text{A.4})$$

Replacing the optimal controls in (A.3), yields

$$J_t(W_t) = \varpi(C_t^*(W_t)) + v(L_t^*(W_t)) + E_t[\beta J_{t+1}(W_{t+1}^*(W_t))] \quad (\text{A.5})$$

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 (A.5) with respect to the state variable W_t then yields

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

which after simplification using (A.4) and the fact that all variables subscripted with t are contained in the information set at t , reduces to the envelope conditions

$$J'_t(W_t) = \beta E_t \left[J'_{t+1}(W_{t+1}^*(W_t)) \tilde{R}_{t+1}^*(W_t) \right] \quad (\text{A.6})$$

$$= \beta \varpi'(C_t^*(W_t)) \quad (\text{A.7})$$

$$= \beta v'(L_t^*(W_t)) / P_t \quad (\text{A.8})$$

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 (A.4) becomes

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

which yields the set of conditional Euler equations (2.9). From this it follows that

$$\begin{cases} E_t \left[\beta \frac{\varpi'(C_{t+1})}{\varpi'(C_t)} (R_{t+1} - R_{t+1}^f) \right] = 0 \\ E_t \left[\beta \frac{v'(L_{t+1}) P_t}{v'(L_t) P_{t+1}} (R_{t+1} - R_{t+1}^f) \right] = 0 \end{cases} \quad (\text{A.9})$$

i.e., equations (2.10)-(2.11).

A.3 Estimating Equations

Let $M_{t+1}^L \equiv \beta v'(L_{t+1}) / v'(L_t)$ denote the marginal rate of substitution for luxury consumption, which is the stochastic discount factor in this case. With $v(L) \equiv (L + b)^{(1-\psi)} / (1-\psi)$, it follows that $M_{t+1}^L \equiv \beta (L_{t+1} + b)^{-\psi} / (L_t + b)^{-\psi}$. From (2.9), we have that

$$1 = E_t \left[M_{t+1}^L \frac{P_t}{P_{t+1}} R_{t+1}^f \right] = E_t \left[M_{t+1}^L \frac{P_t}{P_{t+1}} \right] R_{t+1}^f \quad (\text{A.10})$$

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

$$E \left[M_{t+1}^L \frac{P_t}{P_{t+1}} \right] = E \left[\frac{1}{R_{t+1}^f} \right]. \quad (\text{A.11})$$

Then the unconditional version of (2.11) is

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

We then employ a fairly standard linearization (see e.g., Campbell (1999)). With our choice of utility function for luxury goods consumption,

$$M_{t+1}^L \frac{P_t}{P_{t+1}} = \beta \frac{v'(L_{t+1})}{v'(L_t)} \frac{P_t}{P_{t+1}} = \beta \exp(-\psi \Delta \ln(L_{t+1} + b) - \Delta \ln(P_{t+1})) \quad (\text{A.12})$$

with the notation $\Delta \ln(X_{t+1}) \equiv \ln(X_{t+1}) - \ln(X_t)$ for any positive X . Under the under the reasonable approximation that $b/L_t \approx 0$, equation (A.12) implies that

$$M_{t+1}^L \frac{P_t}{P_{t+1}} \approx \beta \exp(-\psi \Delta \ln(L_{t+1}) - \Delta \ln(P_{t+1})) \quad (\text{A.13})$$

$$\approx \beta [1 - \psi \Delta \ln(L_{t+1}) - \Delta \ln(P_{t+1})]. \quad (\text{A.14})$$

where the linearization approximation consists of $\psi \Delta \ln(L_{t+1}) + \Delta \ln(P_{t+1}) \ll 1$. This type of approximation is true more generally; it follows from the Taylor expansion $v'(L_{t+1}) \approx v'(L_t) + v''(L_t)(L_{t+1} - L_t)$ which in our case gives

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

Therefore we have from (A.14) that

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

Finally, we assume that $\beta E \left[1/R_{t+1}^f \right]^{-1} \approx 1$ and make the adjustment as if our model were in continuous time (see Grossman, Melino, and Shiller (1987)), yielding equation (3.1).

The conditional Euler equation (2.9) 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 (A.13) 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 $|\varepsilon_{t+1}| \ll 1$) 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 (3.2).

B Empirical Moments and Inference

B.1 The Unconditional Euler Equation

The unconditional Euler equation (2) for quantities involves estimating $E[R_{t+1} - R_{t+1}^f]$, $\text{Cov} [\Delta \ln(L_{t+1}), (R_{t+1} - R_{t+1}^f)]$ and $\text{Cov} [\Delta \ln(P_{t+1}), (R_{t+1} - R_{t+1}^f)]$. In practice, to estimate these moments we use all available post-war data to construct each sample moment. The last few years have seen unusually high excess returns; by using all available data we can minimize the small-sample bias associated with this unusual occurrence.

Regarding inference, define the vector $X_t \equiv \left(R_t - R_t^f, \Delta \ln(L_t), \Delta \ln(P_t) \right)'$ and assume that $X \equiv \{X_t\}_{t=2, \dots, T+1}$ is covariance-stationary. With x_{it} denoting the i^{th} component of the vector X_{it} , $i = 1, 2, 3$, it is clear that the

asymptotic properties of the three moment estimators we are considering, and ultimately the estimator of ψ based on equation (2), can be deduced from those of the vector

$$M_T \equiv T^{-1} \left(\sum x_{1t}, \sum x_{2t}, \sum x_{3t}, \sum x_{1t}x_{2t}, \sum x_{1t}x_{3t} \right)' \quad (\text{B.1})$$

by judicious use of the delta method. Let us define $x_{4t} \equiv x_{1t}x_{2t}$ and $x_{5t} \equiv x_{1t}x_{3t}$ and the extended vector $\tilde{X}_t \equiv (x_{1t}, x_{2t}, x_{3t}, x_{4t}, x_{5t})'$. Further, define the vector and the matrix

$$\mu \equiv E \left[\tilde{X}_t \right] \quad (\text{B.2})$$

$$\Gamma_j \equiv E \left[\left(\tilde{X}_t - \mu \right) \left(\tilde{X}_{t-j} - \mu \right)' \right] \quad (\text{B.3})$$

and assume that the autocovariance matrices are absolutely summable, i.e., $\sum_{j=0}^{+\infty} |\Gamma_j| < \infty$. The natural estimators of these quantities are the corresponding sample moments. Under a number of reasonable assumptions on the time series dependence of the process $\{\tilde{X}_t\}_{t=2, \dots, T+1}$, such as for example $\tilde{X}_t = \mu + \sum_{j=0}^{+\infty} \alpha_j \varepsilon_{t-j}$ with ε_t i.i.d., $E[\varepsilon_t^2] < \infty$, $\sum_{j=0}^{+\infty} |\alpha_j| < \infty$, the Central Limit Theorem (see e.g., Anderson (1971,)) implies that

$$\sqrt{T}(M_T - \mu) \longrightarrow N \left(0, \sum_{j=-\infty}^{+\infty} \Gamma_j \right). \quad (\text{B.4})$$

The quantities of interest are

$$\begin{aligned} \widehat{E}[R_{t+1} - R_{t+1}^f] &= M_{1T} \\ \widehat{Cov} \left[\Delta \ln(L_{t+1}), (R_{t+1} - R_{t+1}^f) \right] &= M_{4T} - M_{1T}M_{2T} \\ \widehat{Cov} \left[\Delta \ln(P_{t+1}), (R_{t+1} - R_{t+1}^f) \right] &= M_{5T} - M_{1T}M_{3T} \end{aligned}$$

and the delta method yields the asymptotic distribution of our unconditional estimator of ψ given by the empirical counterpart of equation (2), that is:

$$\hat{\psi}_{UNC} \equiv \hat{\psi}(M_T) = \frac{2}{3} \frac{\widehat{E}[R_{t+1} - R_{t+1}^f] - \widehat{Cov} \left[\Delta \ln(P_{t+1}), (R_{t+1} - R_{t+1}^f) \right]}{\widehat{Cov} \left[\Delta \ln(L_{t+1}), (R_{t+1} - R_{t+1}^f) \right]} = \frac{2}{3} \frac{M_{1T} - (M_{5T} - M_{1T}M_{3T})}{M_{4T} - M_{1T}M_{2T}} \quad (\text{B.5})$$

from which it follows that

$$\sqrt{T}(\hat{\psi}_{UNC} - \psi) \longrightarrow N \left(0, \vec{\nabla}(\mu)' \left(\sum_{j=-\infty}^{+\infty} \Gamma_j \right) \vec{\nabla}(\mu) \right) \quad (\text{B.6})$$

where $\vec{\nabla}(\mu) \equiv \left(\partial \hat{\psi}(M_T) / \partial M_T \right)_{M_T = \mu}$ can be estimated consistently by $\vec{\nabla}(M_T)$.

As noted, we use all available data for our point estimates. Let T be the number of periods of available data on returns and the price of luxuries, and N be the number of periods of available data on luxury consumption.

Note that for most series $T > N$. Our point estimates of $M_{1,t}$ and $M_{4,T}$ are

$$\begin{aligned}
M_{1,T} &= \frac{1}{T} \sum_{t=1}^T (R_t - R_t^f) \\
M_{2,T} &= \frac{1}{N} \sum_{t=1}^N \Delta \ln(L_t) \\
M_{3,T} &= \frac{1}{N} \sum_{t=1}^N \Delta \ln(P_t) \\
M_{4,T} &= \frac{1}{N} \sum_{t=1}^N \Delta \ln(L_t) (R_t - R_t^f) \\
M_{5,T} &= \frac{1}{N} \sum_{t=1}^N \Delta \ln(P_t) (R_t - R_t^f).
\end{aligned}$$

As with our point estimates, we use all available data to construct an estimate of each sample element of the variance covariance matrix. The variance covariance matrix is constructed using Newey-West triangular weighting with 1 lag.

For prices of luxury goods in fixed supply, the unconditional Euler equation (5.4) requires estimating moments of P_t/P_{t+1} and $R_{t+1} - R_{t+1}^f$ which we do using T periods. The asymptotic properties of the implied estimator of the risk premium follow using the same method as above.

B.2 The Conditional Euler Equation

With respect to the linearized conditional Euler equations, we estimate the first-stage with as much data as is available and then the second-stage with the smaller amount of luxury consumption data. Let $\mathbf{Z}_T = (Z_0, \dots, Z_{T-1})'$ be the $T \times k$ matrix where Z_0 is the $k \times 1$ column vector of instruments known at time period 0 appropriate for use to instrument R_t^i where i denotes risk-free or equity return. Let $\mathbf{R}_T^i = (R_1^i, \dots, R_T^i)'$ be the $T \times 1$ column vector of the available returns and let $\mathbf{X}_T = (\mathbf{1}', \mathbf{R}_T^i)'$ an $T \times 2$ matrix where $\mathbf{1}$ is an $T \times 1$ vector of ones. Let \mathbf{Z}_N and \mathbf{X}_N be the same matrixes truncated to the N periods in which we have data on luxuries. Finally, let $\mathbf{L}_N = (\Delta \ln(L_1), \dots, \Delta \ln(L_N))'$ be the vector of data on the growth of the consumption of luxuries.

An optimal two-step instrument variables estimator is

$$\hat{\gamma} = \left(\widehat{-\frac{\delta}{\psi}}, \widehat{\frac{1}{\psi}} \right)' = \left(\mathbf{X}_T' \mathbf{Z}_T \hat{V}^{-1} \mathbf{Z}_T' \mathbf{X}_T \right)^{-1} \mathbf{X}_T' \mathbf{Z}_T \hat{V}^{-1} \mathbf{Z}_N' \mathbf{L}_N$$

where

$$\hat{V} \xrightarrow{p} V = E[\mathbf{Z}_t \varepsilon_t \varepsilon_t' \mathbf{Z}_t'] + \sum_{j=1}^{\infty} E[\mathbf{Z}_t \varepsilon_t \varepsilon_{t+j}' \mathbf{Z}_{t+j}'] + E[\mathbf{Z}_{t+j} \varepsilon_{t+j} \varepsilon_t' \mathbf{Z}_t']$$

and

$$\varepsilon_t = \mathbf{L}_t - \gamma \mathbf{X}_t.$$

To construct this estimator, we first take $\hat{V} = \mathbf{Z}_T' \mathbf{Z}_T$ and construct a preliminary estimate of γ . We then use this estimate to construct a vector $\hat{\varepsilon}_N = \mathbf{L}_N - \hat{\gamma} \mathbf{X}_N$ and then \hat{V} is calculated using a Newey-West estimator as for the unconditional inference. With \hat{V} in hand, we apply the above formulae to calculate $\left(\widehat{-\frac{\delta}{\psi}}, \widehat{\frac{1}{\psi}} \right)$ and its variance-covariance matrix.

To test the model based on the degree of overidentification, the statistic $\frac{1}{N} \varepsilon_N' \mathbf{Z}_N \hat{V}^{-1} \mathbf{Z}_N' \varepsilon_N$ is distributed chi-squared with degrees of freedom equal to the degree of overidentification under the null that the model is correct.

C Data

A detailed description of the source and our use of each series on luxury goods follows. We seasonally adjust the monthly/quarterly data by regressing consumption growth on indicator variables for month/quarter and using the residuals in place of consumption growth. In every case, using the raw, unadjusted data leads to the same substantive conclusions.

Equity returns compared to consumption growth between $t - 1$ and t are the price at the end of period t plus dividends during period t both divided by the price at the start of the period t . The nominal quarterly risk-free rate of return is the return on 3 month treasury bills reported for the first month of the quarter. We use post-war data only, returns from 1947 to 2000 on 2001Q1.

C.1 Quantity Data for Luxury Goods in Elastic Supply

C.1.1 BEA Data

We make use of real PCE of nondurables and services, jewelry and watches, and boats and planes.²³ Nondurables and services is available at both an annual and quarterly frequency since 1946, and the latter series are available since 1959. Unfortunately, PCE of jewelry and watches includes many non-luxury items. For instance, most of the PCE-measured consumption of watches is unlikely to represent high-end luxury. Hence, we also use the real retail sales of jewelry published by the BEA as a measure of aggregate consumption of jewelry. This series is available at a monthly frequency since 1967.

C.1.2 U.S. Imports of Jewelry

Another measure of luxury consumption at an aggregate level is U.S. imports for consumption of jewelry (SITC 897), made available through the U.S. International Trade Commission. To isolate the luxury items, we only consider imports from France, Italy, and U.K. Our choice of these European countries is motivated by our list of foreign luxury retailers, which is described below. In 2000, France accounted for 1.3% of total U.S. jewelry imports, Italy 23.4%, and U.K. 0.9%. Hence this series that we construct is mainly driven by imports from Italy. To deflate the nominal value of imports, we use BEA's price index for retail sales of jewelry stores.

C.1.3 Luxury Cars Sales

We have obtained data on total U.S. sales of luxury vehicles from Ward's Automotive. Luxury vehicles includes cars and light trucks with a list price over \$24,800, SUV's over \$26,500, and minivans and vans over \$26,000. We also have data on sales of Porsches. Both series are available since 1980. For price deflation, we use BEA's price index for retail sales of automotive dealers.

C.1.4 Comité Colbert Data: French Luxury Exports to the U.S.

Comité Colbert is a consortium of seventy French companies that specialize in luxury products. We collected data on their total U.S. sales from 1984 through 1998. Among the Comité Colbert members, the sixty 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, Christoffe, 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, Lenôtre, 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 & C. Parlon, Rochas, Champagne Ruinart, Cristal Saint-Louis, Souleïado, S.T. Dupont and Champagne Veuve Clicquot Ponsardin.

²³Specifically, the boats and planes category includes "wheel goods, sports and photographic equipment, boats and pleasure aircraft."

C.1.5 U.S. Sales of Luxury Retailers

We initially targeted sales data for a list of 7 U.S. 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 (1983) and Tiffany (1983). The European retailers are Bulgari (1992), Gucci (1992), Hermès (1984), LVMH (1993), and Waterford Wedgwood (1992). The years in parentheses indicate the first year in which we were able to obtain sales data. 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 in order to reliably measure correlations. Nine of the companies are yet to respond to our request for information. The remaining company is Neiman Marcus whose sales data we have since 1983. However, we do not use Neiman Marcus because their fiscal year ends in July rather than December or January for all the other retailers. Hence, we cannot reliably aggregate their sales data with our other retailers.

We construct two measures of aggregate sales of luxury retailers. The first, which we call Sales of All Luxury Retailers, is constructed from the sales data for the seven luxury retailers we mention above. The second, which we call Sales of US Luxury Retailers, is constructed from Gucci, Saks, and Tiffany. These three companies trade in the NYSE and report on a fiscal year which ends in January. The remaining retailers trade in European stock exchanges and report on a fiscal year which ends in December. So our second measure of sales is a robustness check for our first measure which we construct on the assumption that the difference of a month does not influence our results. To give a flavor for our cross section of data, in 1993 Saks accounted for 16% of total sales, Tiffany 22%, Bulgari 2%, Gucci 5%, Hermès 5%, LVMH 34%, and Waterford Wedgwood 15%. In 1999, it was Saks 72%, Tiffany 8%, Bulgari 1%, Gucci 4%, Hermès 1%, LVMH 9%, and Waterford Wedgwood 5%. In 1997, Proffitt’s merged with Saks Holdings to become the current Saks Incorporated. Hence, although individual firms may go through changes, we take care to compute growth rates in sales over the same firms in adjacent periods to assure that our series is as consistent as possible.

We take sales data from the annual reports for all companies. For Tiffany, we obtained the quarterly data from the 10-Q filings. For all European retailers, we isolate sales in the U.S. In the initial annual report, a retailer typically reports sales for five or so years before the company became public, but do not report the share of sales attributed to the U.S. In those cases, we compute the average percentage of sales in the U.S. for years in which that figure is available to estimate U.S. sales. Sales data reported in foreign currencies are converted to U.S. dollars using the average exchange rate over the fiscal year. Since jewelry is the main line of business for many of the companies on our list, we use the price index for retail sales of jewelry stores to deflate nominal sales.

C.2 Charitable Contributions by the Very Rich

As a proxy for charitable contributions of the wealthy, we use the average contributions data for households with adjusted gross income (AGI) over \$1 million, which is taken from the IRS publication Individual Income Tax Returns. The data is available biannually from 1952 through 1972 and annually since 1973. The nominal values are deflated by the CPI of All Urban Consumers. The price of a charitable contribution is its tax price. That is, the relative price of luxuries in this case is $1 - \tau_t$. We compute the marginal tax rate for households with AGI over \$1 million as

$$\tau_t = \frac{Tax^{1m} - Tax^{0.5m}}{AGI^{1m} - AGI^{0.5m}} \quad (C.1)$$

where Tax^{1m} ($Tax^{0.5m}$) are the average taxes per capita for households with AGI over \$1 million (with AGI from \$0.5–1 million) and AGI^{1m} and $AGI^{0.5m}$ are the correspondingly defined average AGI’s for each group. The tax adjustment makes little difference to the results.

C.3 Price Data for Luxury Goods in Fixed Supply

C.3.1 Manhattan Pre-War Coop Apartments

We have obtained the average closing price of pre-war coops in Manhattan at a quarterly frequency since 1989. Pre-war coops are especially appropriate for our analysis since they are in fixed supply. We consider four average price series: 1) all pre-war coops in Manhattan, 2) all pre-war coops in Manhattan with four or more bedrooms, 3) all pre-war luxury (Central Park West, Park Avenue, Fifth Avenue) coops, and 4) all pre-war luxury coops with four or more bedrooms. Since the price data that we have is the price recorded at the time of closing, there is delay between the time a sale price is negotiated and the time the price is recorded. Hence, in our calculations in Table 5 lag the price series 6 months which is the time frame recommended by our data provider. 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.

C.3.2 London Real Estate

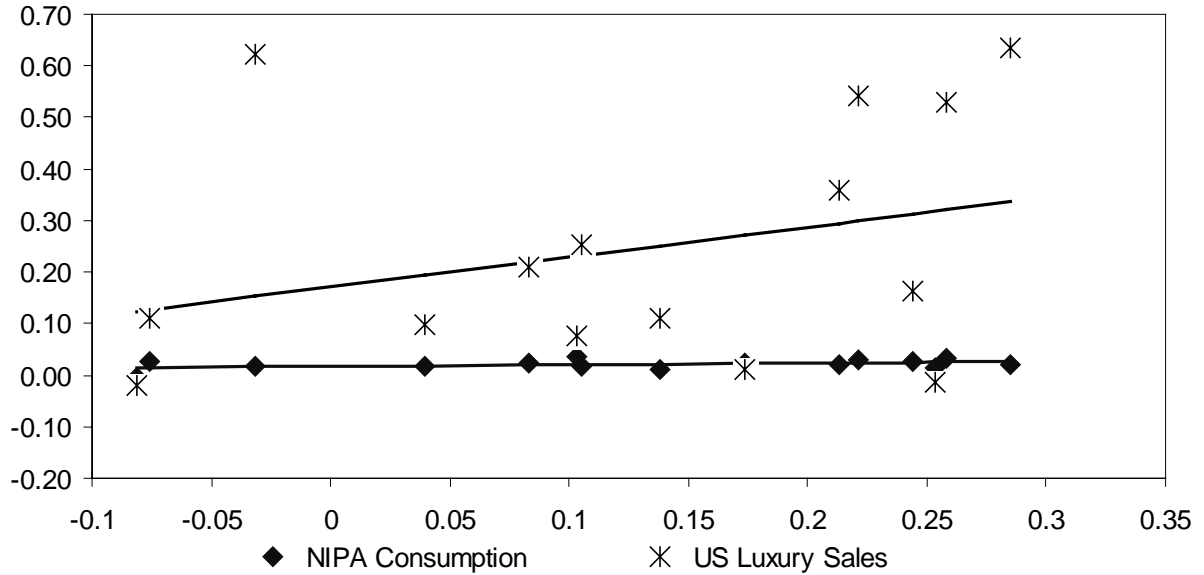
Real estate prices for central London are produced by the Nationwide Building Society in London. The price series is seasonally adjusted, and is available at the quarterly frequency since 1973:4. For stock returns, we use the FTSE all-share returns.

C.3.3 Fine French Wine, US Auction Prices

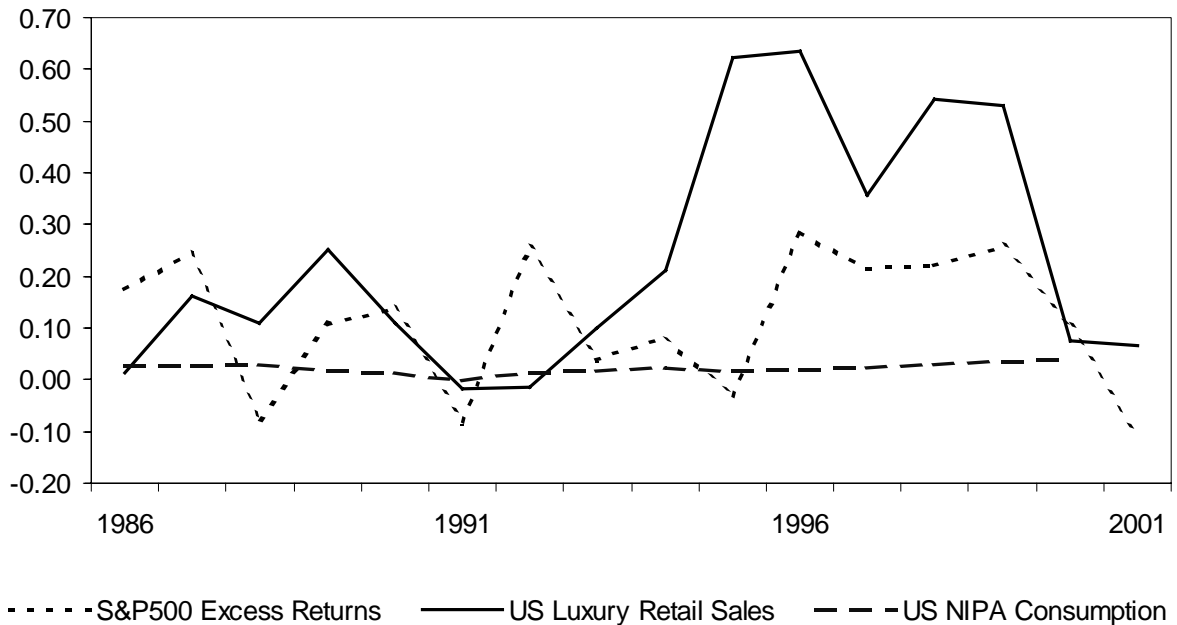
The fine and finest are the Ardmore-Ashenfelter indexes for the finer and finest wines, reconstructed so as to reflect hammer price per dozen 750ml bottles at US auctions only. For more information on the raw data see: www.liquidassets.com. The finest index covers wines from the Château's Lafite, Latour,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

Figure 1: Growth in luxury Consumption, Growth in Basic Consumption, and Excess Returns, 1986-2000

Panel A: Consumption growth against Excess Returns



Panel B: Consumption Growth and Excess Returns over Time



Note: All data are annual growth rates. Both figures display returns at the frequency of annual reports, from February 1985 to January 2001.

Figure 2: Expenditure Shares

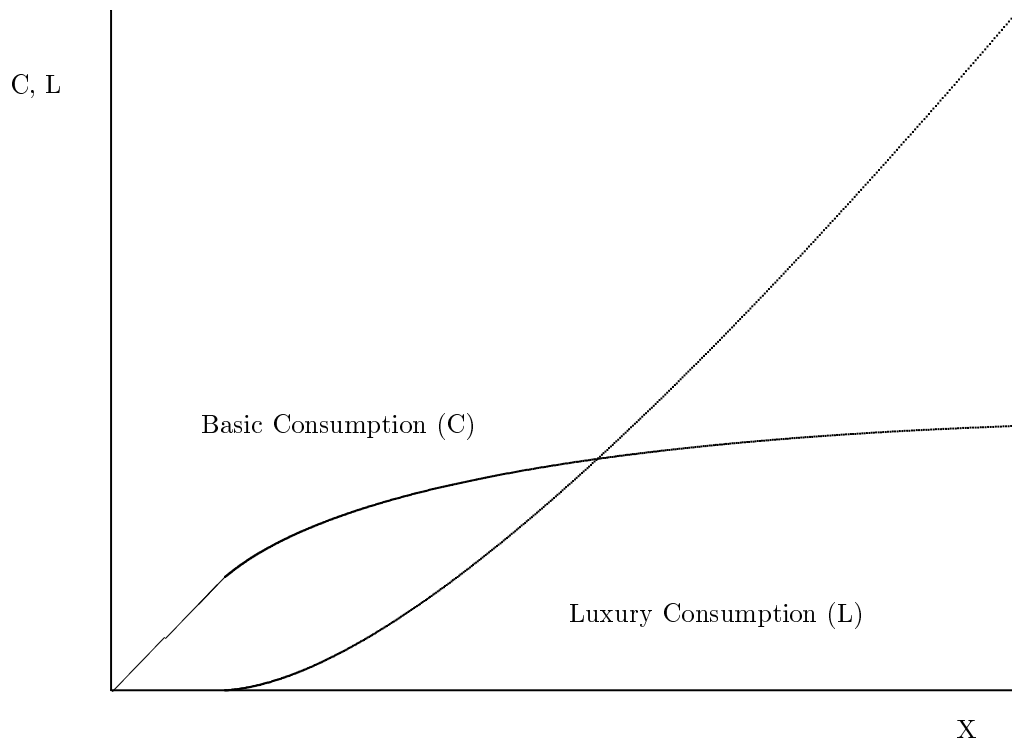
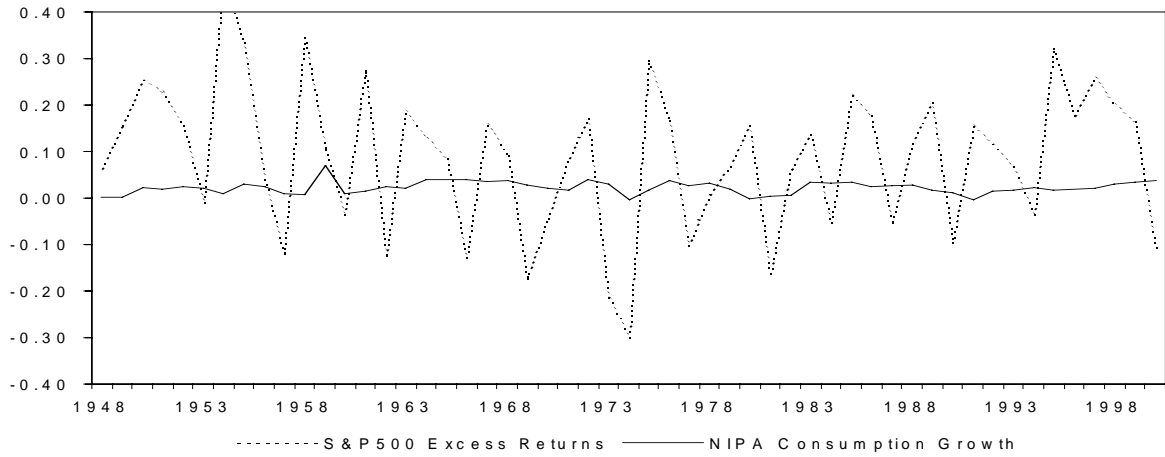


Figure 3: Luxury Consumption Growth and Excess Returns

Panel A: NIPA Consumption



Panel B: U.S. Sales of Luxury Automobiles, Porsche

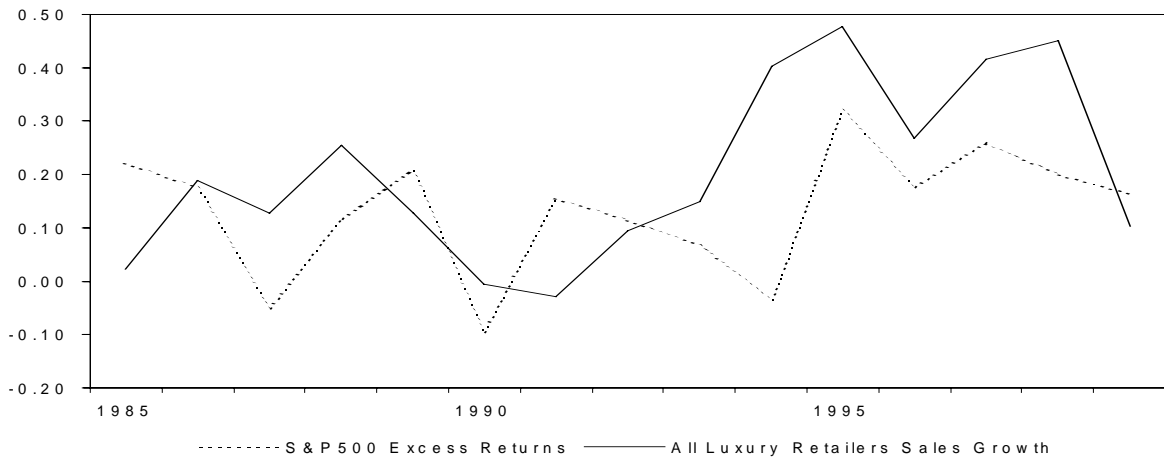


Panel C: U.S. Sales of French Luxury Group, Comité Colbert




Figure 3 (Continued)

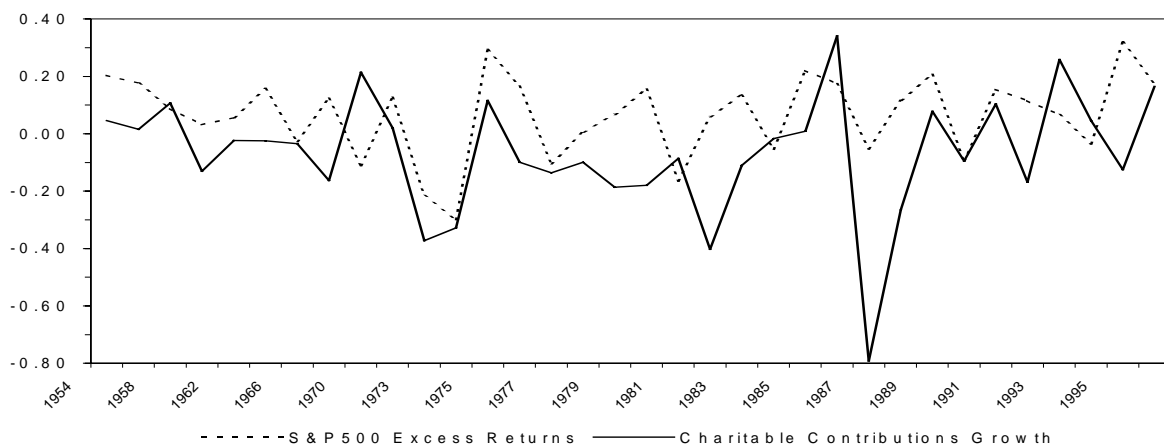
Panel D: U.S. Sales from All Luxury Retailers



Panel E: U.S. Sales from Tiffany (Q)

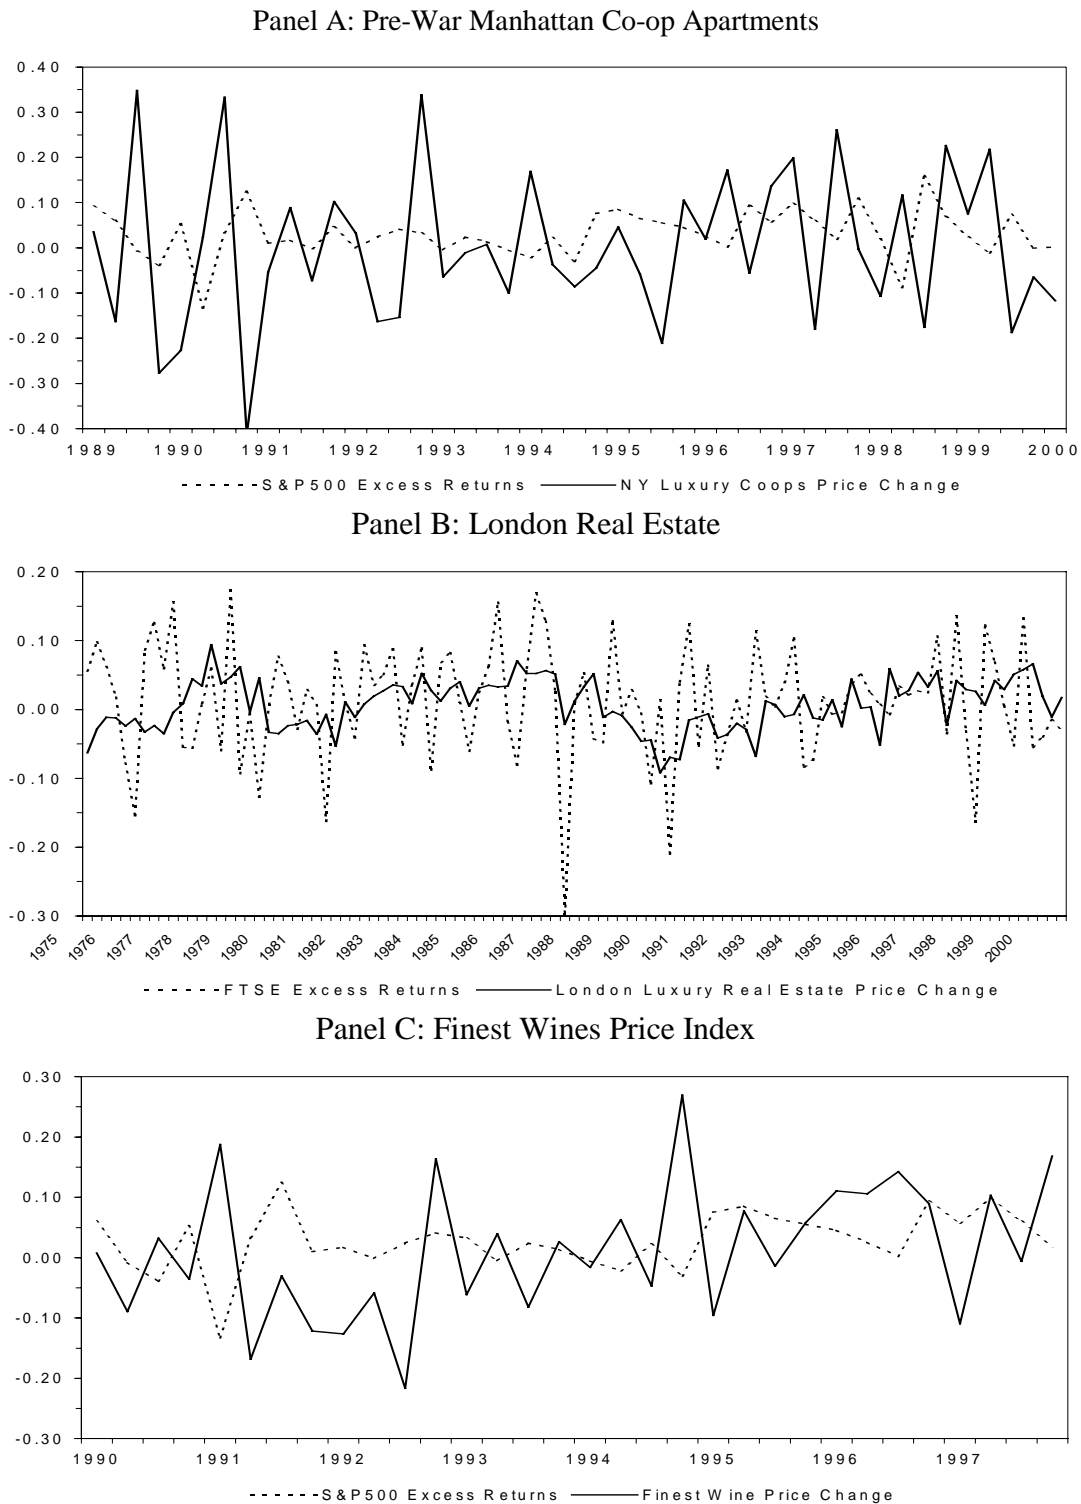


Panel F: Charitable Contributions of the Rich



Note: Figures present time series of returns and aligned sales/consumption. See text and appendix for details on the series. Quarterly data is seasonally adjusted by removing the component of the series correlated with quarter indicator variables.

Figure 4: The Prices of Luxury Goods in Fixed Supply



Note: Figures present time series of returns and aligned price changes (Tables lag price change of closing prices as described in the text). Quarterly data is seasonally adjusted by removing the component of the series correlated with quarter indicator variables. See text and appendix for details on the series.

Series	Number of Obs.	Corr. with Returns	Standard Deviation	Risk Aversion	Standard Error
NIPA PCE Nondur. & Serv.	54	0.000	0.014	1095.8	4945.8
NIPA PCE Nondur. & Serv. (Q)	215	0.177	0.013	169.0	75.5
NIPA PCE Jewelry & Watches	39	-0.125	0.066	-26.7	27.1
NIPA PCE Jewelry & Watches (Q)	161	-0.019	0.082	-133.5	254.6
NIPA PCE Boats and Planes	39	0.039	0.074	-93.4	217.5
NIPA PCE Boats and Planes (Q)	161	0.099	0.070	70.5	79.2
Retail Sales of Jewelry (M)	408	0.057	0.118	69.0	61.7
Imports of Jewelry	11	0.336	0.084	13.3	12.1
Imports of Jewelry (Q)	47	0.038	0.154	50.5	137.5

Table 1: Risk Aversion Implied by NIPA Data and Jewelry Imports

Note: Based on estimation of the unconditional Euler equation for luxury goods. (Q) denotes quarterly and (M) denotes monthly frequency. Series are at an annual frequency unless otherwise noted. Standard deviation of series is reported at an annual rate. (For instance, by multiplying the quarterly rate by 2.) See text for further description of series and estimating equations.

Series	Number of Obs.	Corr. with Returns	Standard Deviation	Risk Aversion	Standard Error
NIPA PCE Nondur. & Serv.	54	0.000	0.014	1095.8	4945.8
NIPA PCE Nondur. & Serv. (Q)	215	0.177	0.013	169.0	75.5
Luxury Automobile Sales	19	0.240	0.085	15.4	15.0
Porsche Sales	19	0.134	0.343	9.9	15.8
US Sales of French					
Luxury Group (Com. Colbert)	14	0.184	0.116	12.7	12.9
Luxury Retail Sales	15	0.363	0.168	3.6	2.8
Luxury Retail Sales (US Retailers)	16	0.356	0.230	3.2	2.2
Tiffany	16	0.241	0.085	8.7	6.6
Tiffany NY	13	0.063	0.085	24.5	32.5
Tiffany (Q)	31	0.512	0.098	5.5	4.3
Charitable Contrib. of Rich	34	0.370	0.211	4.7	3.3

Table 2: Risk Aversion Implied by the Consumption of Luxury Goods

Note: Based on estimation of the unconditional Euler equation for luxury goods. (Q) denotes quarterly frequency. Series are at an annual frequency unless otherwise noted. Standard deviation of series is reported at an annual rate. (For instance, by multiplying the quarterly rate by 2.) See text for further description of series and estimating equations.

Series	Number of Obs.	Corr. with Returns	Standard Deviation	Risk Aversion
Panel A				
Government Series on Luxury Goods				
NIPA PCE Jewelry & Watches	36	-0.047	0.091	-11.6
NIPA PCE Jewelry & Watches (Q)	158	0.076	0.079	112.0
NIPA PCE Boats & Planes	36	0.047	0.092	-15.3
NIPA PCE Boats & Planes (Q)	158	0.150	0.085	32.8
Jewelry Retail Sales (M)	405	0.131	0.092	29.4
Imports of Jewelry	8	0.790	0.055	5.0
Imports of Jewelry (Q)	44	0.084	0.106	16.5
Panel B				
Constructed Series on Luxury Goods				
Luxury Automobile Sales	16	0.015	0.117	-260.0
Porsche Sales	16	0.047	0.492	14.0
US Sales of French				
Luxury Group (Com. Colbert)	11	0.285	0.071	9.7
Luxury Retail Sales	12	0.197	0.255	2.5
Luxury Retail Sales (US Retailers)	13	0.146	0.355	1.6
Tiffany	13	0.403	0.129	2.7
Tiffany NY	10	0.320	0.120	5.4
Tiffany (Q)	25	0.336	0.063	3.1
Charitable Contrib. of Rich	31	0.282	0.178	3.8

Table 3: Risk Aversion Implied by Change in Sales Luxury Goods Over Four Periods

Note: Based on estimation of the unconditional Euler equation for luxury goods without adjustment from continuous time. (Q) denotes quarterly and (M) denotes monthly frequency. Series are at an annual frequency unless otherwise noted. Standard deviation of series is reported at an annual rate. See text for a complete description of estimation and data.

Series	Number of Obs.	NYSE		T-Bill	
		Inverse of Risk Aversion	Test of Overident.	Inverse of Risk Aversion	Test of Overident.
NIPA PCE Nondur. & Serv (A)	51	0.073 (0.037)	4.45 (0.22)	0.352 (0.210)	2.82 (0.42)
NIPA PCE Nondur. & Serv (Q)	213	-0.066 (0.052)	2.74 (0.43)	0.143 (0.100)	7.98 (0.05)
NIPA PCE Jewelry & Watches (A)	39	0.403 (0.211)	0.80 (0.85)	0.544 (0.541)	5.71 (0.13)
NIPA PCE Jewelry & Watches (Q)	161	-0.035 (0.133)	3.03 (0.39)	0.021 (0.423)	3.19 (0.36)
NIPA PCE Boats & Planes (A)	39	0.295 (0.180)	2.82 (0.42)	0.376 (0.409)	4.54 (0.21)
NIPA PCE Boats & Planes (Q)	161	0.012 (0.165)	12.34 (0.01)	0.844 (0.452)	12.43 (0.01)
Jewelry Retail Sales (M)	393	-0.082 (0.262)	2.50 (0.47)	0.012 (0.910)	2.59 (0.46)
Imports of Jewelry (A)	11	0.020 (0.034)	5.23 (0.16)	0.167 (0.202)	4.99 (0.17)
Imports of Jewelry (Q)	44	-0.051 (0.125)	5.30 (0.15)	0.286 (0.419)	4.72 (0.19)
Luxury Automobile Sales	19	0.162 (0.131)	5.76 (0.12)	0.126 (0.894)	6.36 (0.10)
Porsche Sales	19	0.400 (0.390)	2.01 (0.57)	-2.596 (3.209)	1.40 (0.71)
US Sales of French Luxury Group (Com. Colbert)	14	0.146 (0.066)	5.30 (0.15)	0.117 (0.195)	6.46 (0.09)
Luxury Retail Sales	15	0.202 (0.076)	5.35 (0.15)	1.032 (0.356)	5.23 (0.16)
Luxury Retail Sales (US Retailers)	16	0.265 (0.063)	3.39 (0.34)	1.776 (0.409)	3.64 (0.30)
Tiffany (A)	16	0.174 (0.080)	5.17 (0.16)	0.709 (0.354)	5.43 (0.14)
Tiffany NY (A)	13	0.020 (0.031)	3.15 (0.37)	0.174 (0.197)	3.21 (0.36)
Tiffany (Q)	24	-0.073 (0.099)	12.61 (0.01)	-0.176 (0.215)	12.59 (0.01)
Charitable Contrib. of Rich	34	0.496 (0.365)	1.14 (0.77)	1.442 (1.014)	0.43 (0.93)

Table 4: Estimates from the Conditional Euler Equation

Note: Based on estimation of the conditional Euler equation for luxury goods. Standard errors are reported in parentheses below coefficients and p-values are reported in parentheses below tests of overidentification based on the $\chi^2(3)$ which has 95 percent critical value 7.82. Instruments include second lags of NIPA consumption growth, T-Bill return, stock return, and the log price-dividend ratio. See text for complete description of series and estimating equations.

Series	Number of Obs.	Corr. with Returns	Standard Deviation	Equity Premium	Standard Error
NIPA C & $\psi = 1$	215	0.177	0.013	0.000	0.000
Pre-War Manhattan Coops	45	0.190	0.253	0.005	0.002
NY Coops (4+ bedrms)	45	0.340	0.396	0.015	0.004
NY Luxury Coops	45	0.296	0.340	0.011	0.004
NY Luxury Coops (4+ bedrms)	45	0.322	0.558	0.024	0.008
London Real Estate	102	0.275	0.073	0.003	0.001
Fine Wine	32	-0.206	0.225	-0.004	0.002
Finest Wine	32	-0.227	0.225	-0.004	0.002
Great Wine	32	-0.119	0.214	-0.002	0.002

Table 5: The Equity Premium Implied by Prices of Luxury Goods in Fixed Supply

Note: Based on estimation of the unconditional price relation for luxury goods with fixed quantities. All series are at a quarterly frequency. The standard deviation of series, estimated equity premium, and standard error of the estimate are reported at an annual rate. (For instance, by multiplying the equity premium by 4 and standard error by 2.) See text for further description of series and estimating equations.