

Time Variation in the Covariance between Stock Returns and Consumption Growth

GREGORY R. DUFFEE*

ABSTRACT

The conditional covariance between aggregate stock returns and aggregate consumption growth varies substantially over time. When stock market wealth is high relative to consumption, both the conditional covariance and correlation are high. This pattern is consistent with the “composition effect,” where agents’ consumption growth is more closely tied to stock returns when stock wealth is a larger share of total wealth. This variation can be used to test asset-pricing models in which the price of consumption risk varies. After accounting for variations in this price, the relation between expected excess stock returns and the conditional covariance is negative.

AFTER DECADES OF RESEARCH, financial economists have tentatively concluded that there are predictable variations in excess returns to the stock market. The uncertainty in this conclusion is partially driven by our inability to interpret comfortably these variations in a sensible model of asset pricing. In a world of rational investors who care about consumption, variations in expected excess returns should be driven by variations in either the amount of “consumption risk” in stock returns—the conditional covariance between stock returns and consumption—or the required compensation per unit of consumption risk. But attempts to link these potential explanations to the data have either failed empirically or have not been rigorously tested.

A major problem is that existing research has found little evidence for time variation in the conditional covariance between stock returns and aggregate consumption growth. Given this conclusion, there are only two interpretations we can place on time-varying expected excess returns in a consumption-based model. The first possibility is that investors are heterogeneous and there is time variation in the conditional covariance between stock returns and the consumption of marginal stockholders that is not matched by a similar variation at the aggregate level. The second possibility is the price of consumption risk varies through time. Neither possibility is easily tested empirically. Consumption data on individual investors is noisy and limited, while proxies for the price of consumption risk tend to be chosen based on the known behavior of expected stock returns, making the models close to tautological.

*Haas School of Business, University of California at Berkeley. Thanks to the many seminar participants who have commented on earlier versions of this paper, and also to an anonymous referee, John Campbell, Karen Dynan, Kris Jacobs, Martin Lettau, and Sydney Ludvigson.

In this paper I re-examine the relation between aggregate stock returns and aggregate consumption growth. As econometricians we do not observe their conditional covariance, thus we attempt to infer it by projecting the product of innovations to stock returns and consumption growth on a set of instruments. Therefore, conclusions about the time variation in the conditional covariance hinge on the econometrician's choice of instruments. I find that instruments related to the ratio of stock market wealth to consumption predict substantial time variation in the conditional covariance.

An important source of this predictive power is what I call the "composition effect," which refers to the composition of total wealth. The role of the composition effect in asset pricing was first noted by Santos and Veronesi (2003). Consumption moves with wealth. Wealth consists of both stock market wealth and non-stock-market wealth. At times when the share of stock market wealth in total wealth is relatively high, then (1) stock market wealth will be high relative to total consumption and (2) the sensitivity of consumption to changes in stock market wealth will be relatively high. Hence, both the conditional covariance and the conditional correlation between stock returns and consumption growth will be above average when the ratio of stock market wealth to consumption is above average. This simple theory is strongly supported in U.S. data from 1959 to 2001. For example, point estimates of the conditional correlation between monthly stock returns and monthly consumption growth over this period range from roughly 0 to about 0.6, with the highest values reached at the tail end of the bull market that ended in 2000.

An unfortunate consequence of this evidence is that it increases the difficulty of explaining time variation in expected excess stock returns.¹ Existing research documents that expected excess returns tend to be high when stock valuations are relatively low, which is also when conditional covariances are relatively low. This pattern demands dramatically more time variation in the price of consumption risk to explain time variation in expected stock returns. We can think of the ratio of expected excess returns to the conditional covariance as the price per unit of consumption risk. When I take into account the lagged response of consumption to stock returns, I estimate that over the past 40 years this ratio has an interquartile range from 28 to 360. Therefore, in order to fit these data, a representative-agent model needs to imply that 10-fold swings in the price of consumption risk are common.

A more important consequence of this evidence is that it allows us to better test models that produce time variation in the price of consumption risk. Such models typically imply that there is some observable variable that can be used as a proxy for investors' price of consumption risk. If the model is correct, the relation between conditional covariances and expected excess returns should depend on the level of this proxy. For example, when the price of risk is high, an increase in the conditional covariance should correspond to a larger increase in expected excess returns than when the price of risk is low. I test this hypothesis

¹ Santos and Veronesi (2003) arrive at a somewhat different conclusion, as I discuss at various points in the paper.

using various proxies for the price of consumption risk. The results are not encouraging. There is no statistically reliable evidence that, say, a higher level of surplus consumption (and therefore presumably a lower price of consumption risk) corresponds to a decreased sensitivity of expected excess stock returns to the conditional covariance.

The framework I use to produce these results is in the spirit of the instrumental variables approach of Campbell (1987) and Harvey (1989). The product of ordinary least-squares innovations to stock returns and consumption growth is treated as an ex post covariance estimate. The projection of this product on a set of instruments is the estimated conditional covariance. This setup imposes little structure on the dynamics of the conditional covariance and is sufficiently flexible to accommodate nonsynchronous dynamics between stock returns and consumption. I also verify the main results with a multivariate generalized autoregressive conditional heteroskedasticity (GARCH) model.

Section I presents the theoretical motivation behind the composition effect. Section II discusses the econometric approach and the data. Econometric details are presented in Section III. Section IV documents time variation in the conditional covariance. Section V links this time variation to expected excess stock returns. Section VI uses a GARCH model to examine the conditional correlation. Concluding comments are offered in Section VII.

I. The Theoretical Motivation

A. Consumption-Based Representative Agent Models

In consumption-based representative agent models, the stochastic discount factor is tied to the behavior of aggregate consumption. Therefore, conditional covariances between asset returns and aggregate consumption growth play a central role in determining expected excess returns. To fix ideas, denote the stochastic discount factor by $\exp(m_t)$ and the log excess return to some asset from $t - 1$ to t as r_t . In this paper I focus on the excess return to the aggregate stock market. If m_t and r_t are jointly conditionally normally distributed, then no arbitrage implies

$$E_{t-1}r_t + \frac{1}{2}\text{Var}_{t-1}(r_t) = -\text{Cov}_{t-1}(r_t, m_t). \quad (1)$$

In equation (1) expectations are conditioned on information available at the end of period $t - 1$. The Jensen's inequality term adjusts for the use of log returns over a discrete horizon.

Power utility, recursive utility, and habit formation models all express m_t as a function of aggregate consumption growth and perhaps other variables. For example, the habit formation model of Campbell and Cochrane (1999) implies

$$m_t = k(s_{t-1}) - (1 + \lambda(s_{t-1}))\Delta c_t, \quad (2)$$

where s_t is a measure of surplus consumption and Δc_t is log-differenced per capita real consumption. Power utility is a special case in which $k(s_{t-1}) = \log \delta$

(a constant) and $1 + \lambda(s_{t-1}) = \gamma$ (also constant). A model such as Campbell and Cochrane's habit formation implies

$$E_{t-1}r_t + \frac{1}{2}\text{Var}_{t-1}(r_t) = \gamma_{t-1}\text{Cov}_{t-1}(r_t, \Delta c_t), \quad (3)$$

where γ_t is the state-dependent sensitivity of expected returns to the conditional covariance between consumption growth and stock returns. In other words, expected excess returns to the stock market equal the product of the price per unit of consumption risk and the amount of consumption risk in stocks. In Campbell and Cochrane, $\gamma_t = 1 + \lambda(s_t)$.

An obvious way to test the time-series implications of (3) is to form conditional second moments (indicated with hats) and then estimate an equation such as

$$r_t + \frac{1}{2}\widehat{\text{Var}}_{t-1}(r_t) = b_0 + b_{t-1}\widehat{\text{Cov}}_{t-1}(r_t, \Delta c_t) + e_t, \quad (4)$$

where b_{t-1} is a parameterized observable proxy for the representative agent's state-dependent price of consumption risk. Given the strong interest among academics and market professionals in identifying and explaining fluctuations over time in the expected aggregate stock returns, we might think that the literature would be rife with conditional tests such as (4). Yet they are almost nonexistent in published work. Attanasio (1991) and Ferson and Harvey (1993) study restricted versions of (4) where $b_0 = 0$ and b_{t-1} is a constant. Yogo (2003) looks at the consumption of durables, while Li (2001) and Parker (2003) contain brief discussions of the variation in the conditional covariance. The main reason for this gap in the literature is that this research has identified only weak evidence of time variation in the conditional covariance.² This is surprising because there is a compelling reason to believe that the conditional covariance varies through time.

B. The Composition Effect

Time variation in the conditional covariance arises naturally in any model in which the stock market does not account for all of investors' wealth. A major source of this variation is variation in the conditional correlation between stock returns and consumption growth. As stock market wealth varies relative to the other determinants of aggregate consumption, so will the conditional correlation.

This point is first made clearly by Santos and Veronesi (2003) in an examination of the relative importance of labor income to consumption. They build a model of risk-averse agents who invest in assets that pay stochastic dividends, where the stochastic processes followed by dividend growth are fairly general. Here I use a more simplistic model to illustrate the point. Investors consume dividends from their total wealth. Wealth consists of stock market wealth and

² There is stronger evidence if consumption data are not seasonally adjusted. See Attanasio (1991) and Ferson and Harvey (1993).

all other assets such as other claims on firms' cash flows (e.g., corporate bonds), real estate, and human capital. Group these other assets together and denote real dividends to stocks and all other assets as $\delta_{s,t}$ and $\delta_{o,t}$ respectively. Aggregate real consumption equals total dividends

$$C_t = \delta_{s,t} + \delta_{o,t}. \tag{5}$$

For tractability I assume that dividend growth rates for both the stock market and all other wealth are random walks with constant variances and equal growth rates g . This growth rate is a simple net rate, as are the other rates used in this model. Formally,

$$\delta_{s,t+1}/\delta_{s,t} = 1 + g + \epsilon_{s,t+1}, \quad \delta_{o,t+1}/\delta_{o,t} = 1 + g + \epsilon_{o,t+1}. \tag{6}$$

Again for tractability, I assume that investors are risk neutral and the real discount rate is r_f . Risk neutrality is not important; it simply allows us to calculate easily the ex-dividend values of the stock market and all other wealth. These values are

$$W_{s,t} = \frac{1 + g}{r_f - g} \delta_{s,t}, \quad W_{o,t} = \frac{1 + g}{r_f - g} \delta_{o,t}. \tag{7}$$

Therefore, the gross returns to these two forms of wealth are

$$R_{k,t+1} = \frac{W_{k,t+1} + \delta_{k,t+1}}{W_{k,t+1}} = 1 + r_f + \frac{1 + r_f}{1 + g} \epsilon_{k,t+1}, \quad k = \{s, o\}. \tag{8}$$

The conditional variances of these returns are assumed to be constant over time. Denote them by σ_s^2 and σ_o^2 , respectively. Denote their correlation by ρ .

Substitution of (7) into (5) allows us to write the change in consumption from t to $t + 1$ as

$$\frac{C_{t+1}}{C_t} = \phi_t \frac{W_{s,t+1}}{W_{s,t}} + (1 - \phi_t) \frac{W_{o,t+1}}{W_{o,t}}, \tag{9}$$

where ϕ_t is the share of stock market wealth in total wealth

$$\phi_t \equiv \frac{W_{s,t}}{W_{s,t} + W_{o,t}}. \tag{10}$$

A similar derivation involving total wealth appears in Campbell (1996). Because his focus is on long-run behavior, he linearizes and replaces the time-varying coefficient ϕ_t with a constant coefficient. Lettau and Ludvigson (2001) and Jagannathan and Wang (1996) follow his lead. Here it is important to retain the time subscript in ϕ_t because the model implies that the conditional correlation between the return to the stock market and consumption growth increases in ϕ_t . For example, the conditional correlation is ρ when $\phi_t = 0$ and is 1 when $\phi_t = 1$.

With this model the conditional covariance between aggregate consumption growth and stock returns is

$$\text{Cov}_t \left(\frac{C_{t+1}}{C_t}, R_{w,t+1} \right) = \frac{1+g}{1+r_f} \sigma_s^2 \left[\phi_t + (1-\phi_t) \left(\rho \frac{\sigma_o}{\sigma_s} \right) \right]. \quad (11)$$

This covariance depends on $\rho\sigma_o/\sigma_s$. The available evidence indicates that stock returns are highly volatile relative to the other determinants of consumption. (In the United States, the standard deviation of aggregate stock returns is about 15 times the standard deviation of consumption growth.) Therefore, this ratio is much less than 1, which implies that the conditional covariance is an increasing function of ϕ_t . As ϕ_t increases, consumption growth becomes more volatile and more highly correlated with stock returns. Both effects raise the conditional covariance.

In practice, we do not observe the stock of nonfinancial wealth.³ However, we can infer ϕ_t from the ratio of stock market wealth to consumption

$$\phi_t = \left(\frac{r_f - g}{1 + g} \right) \frac{W_{s,t}}{C_t}. \quad (12)$$

Intuitively, when the ratio of stock market wealth to consumption is higher than usual, it is relatively more important in determining consumption.

Before we take this model to the data, we will have to make it a little more complicated. The complications are noted in Section II.B. Even with the complications, a robust theoretical conclusion is that the composition effect induces a procyclical conditional covariance. Because stock prices are highly volatile and rise as the economy booms, stocks account for a larger share of total wealth in booms than in recessions. We know from the empirical literature on stock return forecastability that expected excess returns are countercyclical. Even without taking into account procyclical conditional covariances, researchers have struggled to understand the observed time variation in expected returns. Once time variation in covariances is introduced, the difficulty of explaining the behavior of expected returns is magnified substantially.

The conclusion that the composition effect complicates our ability to explain expected stock returns seems straightforward. However, Santos and Veronesi (2003) conclude that empirically, the composition effect actually helps to explain time variation in expected returns. They assume that there are two assets: human capital and stocks. Given this assumption, the ratio of human capital to consumption contains the same information as does the ratio of stock market wealth to consumption. We do not observe human capital, but we observe labor income, which is the dividend to human capital. Given certain restrictions on the dynamics of this dividend, the ratio of labor income to consumption will be negatively correlated with the conditional covariance between stock returns and consumption growth.

³ In this simplistic model, we can calculate wealth if we observe dividends. But this relies on the assumed random walk of dividend growth rates.

Santos and Veronesi use this logic to justify regressions of stock returns on the lagged ratio of labor income to consumption. They find a negative relation, consistent with their theory. (They do not consider the ratio of stock market wealth to consumption, nor do they check whether the ratio of labor income to consumption has predictive power for the conditional covariance.) I discuss the behavior of the labor income/consumption ratio in Section IV. Here it is sufficient to note that notwithstanding their results, the empirical analysis in this paper strongly supports the conclusion that conditional covariances and expected stock returns move in opposite directions over time.

C. The Composition Effect and Heterogeneous Agents

In a world of heterogeneous agents, an equation such as (4) can provide only indirect evidence on time variation in the risk-return tradeoff of stockholders. Because stockholder consumption is more sensitive to stock returns than is nonstockholder consumption, the conditional covariance at the stockholder level will exceed the conditional covariance at the aggregate level. But the central issue here is whether time variation in the aggregate-level covariance is associated with time variation in the stockholder-level covariance.

To understand the relation between the aggregate-level and stockholder-level covariances, first note that with heterogeneous agents the composition effect works at both the aggregate level and at the stockholder level. In any reasonable model of the stock market and consumption, some part of aggregate consumption will be tied to stock market wealth. Therefore, the composition of the determinants of aggregate consumption will vary over time with the level of the stock market. When the aggregate value of stocks is high relative to other determinants of consumption, a larger share of aggregate consumption is sensitive to stock market wealth. The result is a relatively high conditional covariance between aggregate consumption and aggregate stock returns, just as in the representative-agent setting.

At the stockholder level, the magnitude of the composition effect depends on the stockholder's stock wealth as a fraction of her total wealth. The relation is not monotonic. To simplify this discussion, assume that stock returns are uncorrelated with returns to non-stock wealth. Then the counterpart of (11) for agent i simplifies to

$$\text{Cov}_t \left(\frac{C_{i,t+1}}{C_{it}}, R_{w,t+1} \right) = \frac{1 + g}{1 + r_f} \sigma_s^2 \phi_{it}, \tag{13}$$

where ϕ_{it} is the ratio of stock market wealth to total wealth for agent i . The variability of the conditional covariance for agent i depends on the variability of ϕ_{it} , which in turn depends on the level of ϕ_{it} . Consider, for example, the effect on ϕ_{it} of an increase in the value of the aggregate stock market. If agent i either has no wealth in stocks or has her entire wealth invested in the stock market, there is no effect on ϕ_{it} and the conditional covariance is constant. More generally, the derivative of ϕ_{it} with respect to the log of investor i 's stock market wealth is $\phi_{it}(1 - \phi_{it})$, which is plotted in Figure 1.

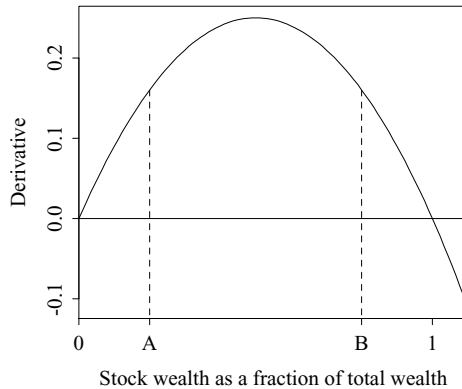


Figure 1. Derivative of (stock wealth/total wealth) with respect to log stock wealth. Total wealth is the sum of stock market wealth and non-stock wealth. The fraction held in the form of stock market wealth is ϕ . This figure plots the derivative of ϕ with respect to log stock market wealth, which is equivalent to the sensitivity of ϕ to the return to the stock market. The amount of non-stock wealth is held constant in computing this derivative.

What is a reasonable value of ϕ_{it} for a typical stockholder? As discussed in Jagannathan and Wang (1996), the contribution of the stock market to aggregate consumption is fairly low. (For example, aggregate dividends were less than 3% of total income from 1959 through 1992.) Thus, for a representative agent, the typical fraction of wealth in stocks is on the left side of the figure, say at point “A.” With heterogeneous agents, stock holdings are concentrated among investors who hold a larger fraction of their wealth in stocks than does the typical consumer. If that fraction lies between “A” and “B” in the figure, then the sensitivity of ϕ_{it} to the stock market is higher for stockholders than it is for the aggregate economy.

This argument suggests that variation in the marginal investor’s conditional covariance is a scaled-up version of variation in the aggregate conditional covariance. If so, estimating an aggregate-level equation such as (4) helps to evaluate heterogeneous-agent models as well as representative-agent models. More concrete conclusions must rely on specific heterogeneous-agent models. Unfortunately, the theory of asset markets with heterogeneous agents is not as well developed as is the corresponding theory with representative agents. In particular, the behavior of conditional correlations is hard to study because dynamic equilibrium models such as Basak and Cuoco (1998) and Chan and Kogan (2002) have a single random variable. (Shapiro (2002) is a notable exception.) We cannot even be sure that the sensitivity of the marginal investor’s ϕ_{it} to the value of the stock market is positive. In the model of Basak and Cuoco, stockholders have a leveraged position in stocks. (In Figure 1 this corresponds to a point to the right of “1” on the x axis.) An implication of their model is that stockholders’ conditional covariance is countercyclical because an increase in the stock market lowers stockholders’ leverage and hence lowers the volatility of their consumption. Thus, in the absence of either a formal model of the

composition effect in a world of heterogeneous agents or detailed information about the wealth of the marginal stockholder, any conclusions drawn from an equation such as (4) about the validity of heterogeneous-agent models should be tempered with caution.

II. The Econometric Approach

A. The Regressions

Period- t excess stock returns and consumption growth can be written as sums of one-step-ahead expectations and innovations

$$r_t = E_{t-1}r_t + \epsilon_{rt}, \quad \Delta c_t = E_{t-1}\Delta c_t + \epsilon_{ct}. \tag{14}$$

The product of the innovations, when projected on investors' information set at time $t - 1$, is the conditional covariance.

As econometricians, we do not observe either the true innovations or the entire set of conditioning information available to investors. A standard approach, which I follow here, is to use forecasting regressions to construct fitted residuals as proxies for true innovations. The product of the fitted residuals is then projected on a set of instruments identified by the econometrician. The forecasting regressions for returns and consumption growth can be written as

$$r_t = a_r' Y_{r,t-1} + e_{r,t-1,t-1,t}, \tag{15}$$

$$\Delta c_t = a_c' Y_{c,t-1} + e_{c,t-1,t-1,t}, \tag{16}$$

where a_r and a_c are parameter vectors and the vectors $Y_{r,t}$ and $Y_{c,t}$ are realized in period t or earlier. I refer to these regressions as “zero-stage” regressions, to distinguish them from the usual first-stage and second-stage instrumental variable regressions that are introduced below. The first time subscript on the residuals refers to the date of the instrument vector (i.e., the period at which the forecast is made). The second and third time subscripts are the starting and ending dates of the dependent variable. Here those subscripts are $t - 1$ and t ; the stock return is calculated from the end of period $t - 1$ to the end of period t , and consumption growth is the log change in consumption from $t - 1$ to t . Not all of these subscripts are necessary to uniquely denote residuals of (15) and (16), but they will be useful later when we consider n -period-ahead forecasts of longer horizon stock returns and consumption growth.

Denote the product of the fitted residuals as

$$\text{Cov}^*(r_t, \Delta c_t) \equiv \hat{e}_{r,t-1,t-1,t} \hat{e}_{c,t-1,t-1,t}. \tag{17}$$

The asterisk indicates an ex post estimate. The ex post estimate is projected on a set of instruments Z_{t-1} to produce a conditional covariance estimate:

$$\text{Cov}^*(r_t, \Delta c_t) = a_v' Z_{t-1} + \mu_t, \tag{18}$$

$$\widehat{\text{Cov}}(r_t, \Delta c_t) = \hat{a}_v' Z_{t-1}. \tag{19}$$

Equation (18) is the first-stage regression. I estimate it with OLS and test hypotheses about α_v with Wald statistics. I use a robust estimate of the variance-covariance matrix, which implies that the statistics have asymptotic χ^2 distributions.

The following second-stage regression is estimated with instrumental variables:

$$r_t + \frac{1}{2} \text{Var}^*(r_t) = b_0 + [b_1 + b_2 p_{t-1}] \text{Cov}^*(r_t, \Delta c_t) + w_t. \quad (20)$$

This regression is in the spirit of the instrumental regressions in Campbell (1987) and Harvey (1989). The second term on the left side is an ex post estimate of the variance of stock returns,

$$\text{Var}^*(r_t) \equiv \hat{e}_{r,t-1,t-1,t}^2. \quad (21)$$

The term in square brackets on the right-hand side of (20) is an observable proxy for the conditional price of consumption risk, which I also refer to as b_t :

$$b_t \equiv b_1 + b_2 p_{t-1}. \quad (22)$$

The base case of b_t is power utility, or $b_2 = 0$. A more general case is where the observable variable p_t picks up time variation in investors' willingness to bear consumption risk. For example, Campbell and Cochrane's model could be tested using $b_t = b_1 + b_2 \hat{s}_{t-1}$, where \hat{s}_{t-1} is a proxy for surplus consumption. The constant term b_0 forces identification of the relation between p_t and the price of risk to be picked up exclusively through a nonlinear relation between the conditional covariance and expected excess stock returns.

I use the generalized method of moments (GMM) technique of Hansen (1982) to estimate jointly the zero-stage regressions and second-stage instrumental variable regression. This allows the standard errors of the second-stage regression to incorporate the uncertainty in the fitted residuals owing to parameter uncertainty in the zero-stage regressions. In Section III.A, I evaluate the finite-sample properties of this method using Monte Carlo simulations.

B. Data Description and the Choice of Instruments

I use monthly data. Monthly consumption is measured by Bureau of Economic Analysis estimates of monthly per capita expenditures on nondurables and services, which are available beginning in January 1959. Expenditures are in 1996 dollars. Monthly consumption growth is defined as log-differenced consumption. I measure excess monthly returns to the aggregate stock market by the log return to the CRSP value-weighted NYSE/Amex/Nasdaq index less the continuously compounded yield on a 1-month Treasury bill as of the end of the previous month. The last observation of stock returns in my sample is December 2001.

The choice of instruments included in the vector Z_t used in the first-stage and second-stage regressions is motivated by the composition effect. The first

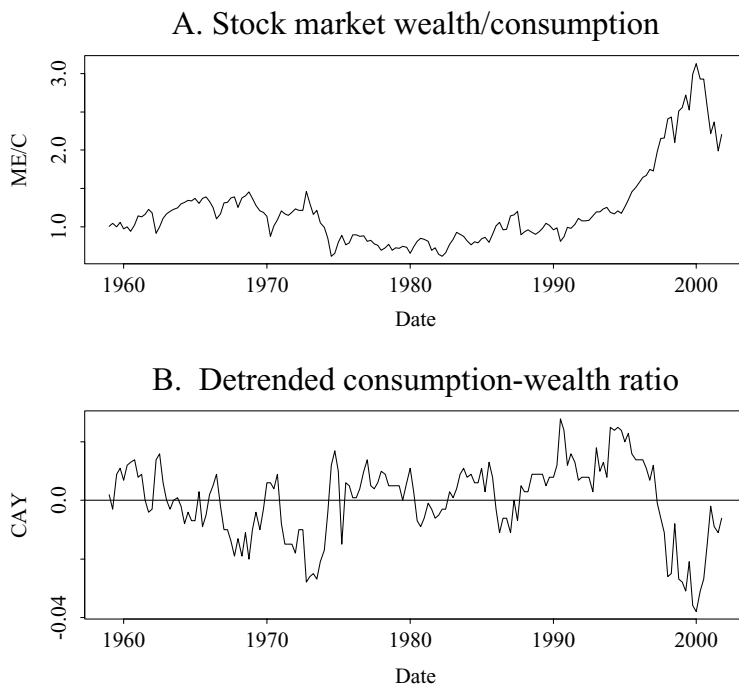


Figure 2. Ratios of wealth to consumption. The top panel plots the ratio of the market capitalization of publicly traded stocks to total consumption on nondurables and services. The bottom panel plots the detrended consumption–wealth ratio introduced in Lettau and Ludvigson (2001).

instrument is the ratio of stock market wealth to consumption. Stock market wealth is measured by the month-end market capitalization of the CRSP value-weighted index, expressed in real per capita terms for comparability to the consumption data. I denote this ratio by ME/C . Recall that one of the simplifying assumptions in the model of Section I.B is that the variance of stock returns is constant. However, stock return volatility temporarily falls (rises) after the market rises (falls).⁴ This means that an increase in ME/C is accompanied by a short-run decline in stock return volatility. This damps the role of the composition effect; the decline in volatility tends to lower the conditional covariance. To control for this effect, I also include lagged excess stock returns in the instrument vector. Denoting the sum $\sum_{i=0}^2 r_{t-i}$ by RET_t^q , I include RET_t^q and RET_{t-3}^q in Z_t .

Not surprisingly, the ratio ME/C is related to the consumption–wealth ratio \bar{cay} introduced by Lettau and Ludvigson (2001). Their ratio is the trend deviation in the log consumption–wealth ratio, where wealth includes capitalized labor income. The two ratios are plotted in Figure 2. It is apparent from

⁴This pattern was discovered by Black (1976) and is the subject of a large literature. See, for example, French, Schwert, and Stambaugh (1987).

the figure that, although the two ratios move inversely (their correlation is about -0.58), ME/C is much more persistent than is \widehat{cay} . In quarterly data the first-order autocorrelation coefficient of ME/C is 0.96 , compared with 0.85 for \widehat{cay} . Therefore, \widehat{cay} is more closely related to the business cycle, as noted by Lettau and Ludvigson. The ratio ME/C is picking up lower frequency variation in the relative importance of the public stock market. Part of the low-frequency variation in ME/C may be variation in the proportion of stocks that are traded publicly versus privately. Small, less-well-known firms are probably better able to access capital through issuance of publicly traded stock today than they were in the 1960s.

The relevant ratio from the perspective of the composition effect is ME/C . Nonetheless, \widehat{cay} likely contains independent information about the conditional covariance, both because of the business cycle link and because it forecasts stock return volatility, as documented by Lettau and Ludvigson (2002a). (I look at predictions of stock return volatility in Section VI.) I therefore also include \widehat{cay} in the instrument vector.⁵ Since the consumption–wealth ratio is defined using quarter-end data, I assume that the observation available at the end of the first two months in each quarter is the value as of the end of the previous quarter.

Many variables in macroeconomics and finance exhibit persistent fluctuations in volatility. This empirical regularity suggests that lags of ex post covariance estimates should be included in Z_t . From an econometric perspective, it is more useful to use a slightly different proxy for the ex post covariance: the product of demeaned stock returns and demeaned consumption growth. The time series properties of this product are close to those of the ex post covariance, but unlike the ex post covariance this product does not depend on parameter estimates from the zero-stage regressions. To limit the number of explanatory variables, I use 3-month sums instead of the individual monthly products. These sums are denoted by

$$CV_t = \sum_{i=0}^2 (\Delta c_{t-i} - \overline{\Delta c})(r_{t-i} - \bar{r}). \quad (23)$$

I include four nonoverlapping lags of CV_t in Z_t . For completeness the set of instruments is reported in the following equation:

$$Z_t = \{1, (ME/C)_t, RET_t^q, RET_{t-3}^q, \widehat{cay}_t, CV_t, CV_{t-3}, CV_{t-6}, CV_{t-9}\}. \quad (24)$$

I do not include the ratio of labor income to consumption in this set of instruments because any information in this ratio should be subsumed in the ratio of stock market wealth to consumption. Nevertheless, to compare my results to those of Santos and Veronesi (2003), I report some results in which I use the ratio of labor income to consumption instead of the ratio of stock market wealth

⁵ The forecasting power of \widehat{cay} for stock returns is somewhat controversial. The look-ahead bias of the cointegrating relation is the subject of Brennan and Xia (2002), while Hahn and Lee (2001) argue the forecasting power is not stable. See also Lettau and Ludvigson (2002b).

to consumption. I compute this ratio using quarterly data following the definition of labor income in Lettau and Ludvigson (2001). As with \widehat{cay} , I assume that the value available at the end of the first 2 months in each quarter is the value as of the end of the previous quarter.

Two other instrument vectors are used in this paper. They are $Y_{c,t}$, used to construct fitted consumption growth residuals, and $Y_{r,t}$, used to construct fitted stock return residuals. Monthly consumption growth is autocorrelated, thus I include in $Y_{c,t}$ a constant term and monthly consumption growth for months t through $t - 2$. I experimented with a larger set of instruments that included stock returns. These experiments led to the conclusion that the results in this paper are unaffected by the precise composition of $Y_{c,t}$. The vector $Y_{r,t}$ consists of a constant and the consumption–wealth ratio \widehat{cay} . I experimented with also including the dividend/price ratio and the slope of the term structure. The use of this larger information set has only a minimal effect on the results, thus I use the smaller set here.

Higher frequency data allow us to produce more accurate estimates of second moments (see, e.g., Andersen et al. (2003)). I use the highest frequency consumption data available. Stock return data are available at higher frequencies, thus I could produce more accurate estimates of the ex post variance of stock returns to use as the Jensen’s inequality adjustment on the left-hand side of (20). In practice, nothing is gained by improving the accuracy of this variance estimate because it contributes little to the overall variation of the left-hand side of the equation. The standard deviation of the log excess return is around 25 to 30 times the standard deviation of the Jensen’s inequality term, regardless of how the variance is estimated. The ratio of standard deviations of 1-month-ahead forecasts of these components is similar. (These results are not shown in any table.) Therefore, I do not use higher frequency stock return data in this paper.

C. Specifications of the Price of Consumption Risk

I consider various specifications of the price of consumption risk $b_t = b_1 + b_2 p_{t-1}$. The first is power utility, or $b_2 = 0$. With $b_2 \neq 0$, I consider three choices of p_{t-1} . They are a measure of surplus consumption, the consumption–wealth ratio, and the dividend–price ratio. The motivation for surplus consumption is the implication of some habit formation models that the price of consumption risk should vary with the level of current consumption relative to past consumption. Wachter (2002) proposes a measure of past consumption growth that arises naturally in habit formation models. I follow her suggestion and define a proxy for surplus consumption at the quarterly frequency as

$$\hat{s}_t = \frac{1 - \Psi}{1 - \Psi^{40}} \sum_{j=0}^{39} \Psi^j \Delta c(t - j), \tag{25}$$

with the decay factor $\Psi = 0.96$. For consistency with Wachter, the notation in this equation differs from that used elsewhere in this paper. In (25), time is

measured in quarters and Δc_t refers to the log change in real per capita quarterly consumption on nondurables and services. My proxy for surplus consumption at the monthly frequency is the most recent quarterly measure of surplus consumption.

The use of the consumption–wealth and dividend–price ratios is motivated by the possibility that the price of consumption risk may vary over time for reasons unrelated to habit. Such variations are likely to show up in \widehat{cay} or D/P . A change in risk tolerance will alter the level of expected returns required by investors. Holding constant the time path of future cash flows, this change in expected returns will change stock valuations relative to dividends and to the long-run cointegrating relation among stock market wealth, consumption, and labor income, as discussed by Campbell and Shiller (1988) and Lettau and Ludvigson (2001). I denote the dividend price ratio at the end of the month t as $(D/P)_t$, where the numerator is the sum of dividends paid in the previous 12 months.

III. Details

In this section I discuss three issues. First, I report Monte Carlo evidence on the properties of the estimation procedure described in the previous section. Second, I consider how to choose the horizon over which consumption growth should be measured. The natural lumpiness of consumption expenditures, combined with the method used to measure such expenditures, implies that the “contemporaneous” relation between stock returns and consumption growth may appear as a relation between time- t stock returns and time- $(t + \tau)$ consumption growth. Third, I describe how I estimate the relation between expected stock returns and conditional covariances when both stock returns and consumption growth are measured over horizons longer than a month. This section is unavoidably tedious. For those readers who are uninterested in the details, the next two paragraphs summarize the results of this section.

The Monte Carlo evidence indicates that standard GMM statistical tests for the presence of time variation in the conditional covariance (the first-stage regression) are well behaved. Statistical GMM tests of the null hypothesis that the conditional covariance has no predictive power for returns (the second-stage regression) are reasonably well behaved (e.g., true critical values for relevant t -statistics are within 10–15% of asymptotic critical values) as long as there is a modest amount of true variation in the conditional covariance. Alternative estimation procedures, such as limited information maximum likelihood, do not have better finite-sample properties.

To address the slow adjustment of consumption, I use two measures of the covariance between stock returns and consumption growth. The first, which ignores slow adjustment, is the covariance between month t 's stock return and consumption growth from month $t - 1$ to month t . The other is the covariance between month t 's stock return and consumption growth from month $t - 1$ to month $t + 3$. I conclude that given the hypotheses of interest in this paper, there is no obvious reason to prefer one of these measures. Finally, I use sums

of monthly ex post covariance estimates as ex post estimates of covariances over horizons longer than a month.

A. Monte Carlo Evidence

I use Monte Carlo simulations to investigate the properties of the econometric methodology. There are two reasons why the usual GMM asymptotics may be inappropriate here. First, the instruments used to construct conditional covariances are weak. Second, there is an errors-in-variables problem created by the fact that we do not observe true residuals.

The properties of the instrumental variables regression (20) depend on the information in the instruments. Staiger and Stock (1997) suggest the instruments should be treated as weak if the F -statistic in the first-stage regression is small; say, less than 10. Kleibergen (2002) develops an instrumental variables statistical test that is robust to the presence of weak instruments. He uses limited information maximum likelihood (LIML) to estimate the second-stage regression and tests hypotheses using his K -statistic. The advantage of LIML and Kleibergen’s test is that they are robust to weak instruments.

The econometric setting in this paper is more complicated than that studied in Kleibergen because here there are “zero-stage” regressions used to construct the residuals to stock returns and consumption growth. Unfortunately, the analysis of more general GMM estimation in the presence of weak instruments is in its infancy.⁶ Stock, Wright, and Yogo (2002) suggest implementing GMM using the continuous-updating estimator of Hansen, Heaton, and Yaron (1996).

The errors-in-variables problem can be seen by expanding the product of fitted residuals into true residuals and errors in conditional expectations

$$\text{Cov}^*(r_t, \Delta c_t) = [\epsilon_{rt} + (\mathbf{E}_{t-1}r_t - \hat{\alpha}'_r Y_{r,t-1})][\epsilon_{ct} + (\mathbf{E}_{t-1}\Delta c_t - \hat{\alpha}'_c Y_{c,t-1})]. \quad (26)$$

Take the expectation of both sides with respect to the instrument vector Z_{t-1}

$$\begin{aligned} E(\text{Cov}^*(r_t, \Delta c_t) | Z_{t-1}) \\ = E(\epsilon_{rt}\epsilon_{ct} | Z_{t-1}) + E[(\mathbf{E}_{t-1}r_t - \hat{\alpha}'_r Y_{r,t-1})(\mathbf{E}_{t-1}\Delta c_t - \hat{\alpha}'_c Y_{c,t-1}) | Z_{t-1}]. \end{aligned} \quad (27)$$

If both the OLS forecasts of returns and the OLS forecasts of consumption growth differ from investors’ forecasts, then the second term on the right-hand side of (27) contaminates the proxy for the conditional expectation of $\epsilon_{rt}\epsilon_{ct}$. This errors-in-variables problem is basically the same as that identified by Pagan and Ullah (1988) in their discussion of regressions of stock returns on estimates of conditional variances. This problem has no effect under the null hypothesis that the price of consumption risk is zero. However, if this price differs from zero, the errors-in-variables problem results in inconsistent estimates of the parameters in (20).⁷

⁶ Stock and Wright (2000) discuss asymptotic distribution theory for GMM estimators in this case.

⁷ One way to circumvent this problem is to explicitly treat the conditional covariance as a latent variable, following the idea of Brandt and Kang (2004).

To implement the Monte Carlo simulations, I model the evolution of the relevant variables with a vector autoregression (VAR). The variables in the VAR are (demeaned) monthly excess stock returns, log consumption growth, the detrended consumption–wealth ratio, and the ratio of stock market wealth to consumption growth. (The ratio ME/C cannot be recovered from the combination of stock returns and consumption growth because of flows in and out of the stock market.) I do not impose any cointegrating relations among the variables. Shocks are drawn from a multivariate normal distribution. Because I construct residuals to consumption growth with a third-order autoregression (AR(3)), I also include two lags of consumption growth in the VAR. Stack these variables in a vector x_t where the first two variables are stock returns and consumption growth and the final two variables are lags 2 and 3 of consumption growth. Then the VAR is

$$x_t = Ax_{t-1} + \mu_t, \quad E_{t-1}(\mu_t) = 0, \quad E_{t-1}(\mu_t \mu_t') = \begin{pmatrix} \Sigma_t & 0_{(2 \times 2)} \\ 0_{(2 \times 4)} & 0_{(2 \times 2)} \end{pmatrix}. \quad (28)$$

I set the first four rows of the parameter matrix A to the parameters from OLS estimation of the VAR over the sample period examined in this paper. (Exceptions to this are noted below.) The remaining two rows are determined by the companion form. I use two versions of Σ_t . One is a constant Σ , set equal to the sample variance–covariance matrix of the residuals from OLS estimation of the VAR. The other version is identical except for the covariance between innovations in stock returns and consumption growth. I assume the conditional covariance varies over time because the conditional correlation changes over time. The setup is

$$\Sigma_{12t} = \sqrt{\Sigma_{(11)} \Sigma_{(22)}} \rho_t, \quad (29)$$

where

$$\rho_t = \bar{\rho} + \alpha(ME/C)_{t-1}. \quad (30)$$

I set α to 0.35. This choice generates variability in the conditional covariance similar to that observed in the sample. The mean $\bar{\rho}$ is set to the corresponding sample correlation. The conditional correlation ρ_t must be bounded in order to ensure the invertibility of Σ_t . If the above equation results in a correlation less than -0.36 or greater than 0.56 , ρ_t is set to the relevant bound.

I produce three sets of Monte Carlo simulations that differ in the dynamics of returns and/or covariances. All of them satisfy the null hypothesis that the conditional covariance has no true forecast power for stock returns:

1. Neither conditional expectations of stock returns nor conditional covariances vary over time. For this hypothesis, I replace the first row of the A matrix with zeros.
2. Conditional expectations of stock returns are constant and conditional covariances vary over time. Here the first row of A is also zero, and covariances are produced with (29) and (30).

3. Conditional expectations of stock returns vary over time, and conditional covariances are constant. Here the first row of A is given by the OLS sample estimates.

A single simulation proceeds as follows. A panel of 515 monthly observations is generated with the specified process. (This is the length of the actual data sample.) Given this simulated data, I estimate the zero-stage regressions and use the results to construct ex post estimates of the covariance between stock returns and consumption growth. I then estimate the first-stage and second-stage regressions (18) and (20). In the latter regression, I set the parameter b_2 to 0.

I estimate the second-stage regression with both asymptotically efficient GMM (jointly with the zero-stage regressions) and LIML, following Kleibergen’s methodology. The weighting matrix for GMM estimation is calculated using residuals from the OLS estimation of both the zero-stage and second-stage regressions. Experiments with the continuous-updating GMM estimator of Hansen et al. (1996) revealed that in this setting it has worse properties than LIML. Thus I do not discuss the results for this alternative estimator.

Table I summarizes the results of the first-stage regressions and GMM estimation of the second-stage regressions. The results for the first-stage regression are straightforward. The test of the hypothesis that the conditional covariance is predictable has reasonable size properties under the null of no predictability. The relevant results are in rows 1 and 3 of the table. Differences between

Table I
Finite-Sample Properties of the Estimation Procedure

This table summarizes the empirical distribution of asymptotic p -values and test statistics calculated from 1,000 Monte Carlo simulations. A vector autoregression is used to simulate 515 monthly observations of excess stock returns, consumption growth, the ratio of stock market wealth to total consumption ME/C , and the consumption–wealth ratio $\widehat{c\bar{w}}$. Given a simulated dataset, innovations to stock returns and consumption growth are constructed with OLS (zero-stage regressions). The product of these innovations is an ex post covariance estimate. In the first-stage regression, the ex post covariance estimate is regressed on a set of instruments. In the second-stage regression, excess stock returns are regressed on the ex post covariance estimate using the same set of instruments. The second-stage regression is estimated using asymptotically efficient GMM (incorporating the zero-stage regressions).

	First Stage Regression			Second Stage Regression					
	Mean F -Stat	Freq. that χ^2 -Stat > Asymp. p -Value		Mean t -Stat	Median t -Stat	Empirical Critical Values for the t -Statistic			
		10%	5%			0.025	0.05	0.95	0.975
Null									
Constant mean return, constant covariance	1.00	0.127	0.058	0.00	-0.03	-1.87	-1.61	1.62	1.75
Constant mean return, stochastic covariance	2.65	0.754	0.672	-0.33	-0.37	-2.16	-1.91	1.35	1.64
Stochastic mean return, constant covariance	1.00	0.117	0.059	0.03	0.12	-2.45	-2.26	2.33	2.55

the empirical and asymptotic values are not large. For example, if neither returns nor conditional covariances are forecastable, the null is rejected in less than 6% of the simulations when the 5% critical value is used. The power of the first-stage regression under the alternative hypothesis (29) is reasonably high for the chosen value of α . The relevant results are in the second row. The null hypothesis is rejected about 70% of the time at the 5% significance level. The typical F -statistic is about 2.6, which is close to the empirical F -statistics reported in Section IV.

The results for the second-stage regression depend on the specification of the null. First examine the case of unpredictable stock returns and predictable covariances. These results are in the second row of the table. I report summary information about t -statistics of the estimates of b_1 in (20) rather than information about the estimates themselves because the distribution of the t -statistics is more robust to variations in the process for generating conditional covariances.⁸ When covariances are predictable and returns are not, the t -statistics are negatively biased. The explanation for the bias is consistent with the logic of Stambaugh (1999). If the value of ME/C at month t is higher than its sample mean, then stock returns prior to t are likely higher than their sample mean, and stock returns after t are thus likely lower than their sample mean. This induces a spurious negative relation between ME/C at month t and future stock returns. Since the conditional covariance is positively associated with ME/C, there is also a spurious negative relation between the conditional covariance and future stock returns. Therefore, to reject at the 5% level the null of no predictability in returns in favor of the hypothesis that the conditional covariance predicts returns, the t -statistic on b_1 must be less than -2.16 or greater than 1.64 .

The magnitude of this bias is, of course, lower when the true relation between ME/C and the conditional covariance is weaker. At the extreme of no relation (the first row in the table), the bias disappears and the standard critical values for the second stage regression are excessively conservative for both positive and negative parameter estimates. These results lead to a somewhat counterintuitive conclusion: when we are more confident about the presence of predictability in the conditional covariance, we must use more stringent critical values in the second-stage regression.

Given the existing research documenting predictability in stock returns, perhaps the most relevant case is when returns are predictable and conditional covariances are constant. To reject this null, we should rely on the first-stage regression. The third row of the table reports that the true critical values for the second-stage regression are much larger (in absolute value) than standard critical values. The problem in this case is that because we are using the same variables that predict returns to predict the conditional covariance, any spurious predictive power for the conditional covariance will necessarily show up as predictive power for expected returns as well.

⁸ This conclusion is based on experimentation with alternative processes for the conditional covariance.

The table does not report the results of LIML estimation because the Monte Carlo results indicate that LIML offers no advantages to GMM. The finite-sample properties of the K -statistic are substantially different from the corresponding asymptotic properties. The statistic is a measure of the distance of the parameter estimate from zero and has an asymptotic χ^2 distribution. Like the GMM estimate, the LIML estimate is biased. Under the null of constant mean returns and stochastic covariances, the empirical 5% critical value for the K -statistic is 5.25 conditional on a negative LIML estimate. Conditional on a positive LIML estimate, this critical value is only 3.15. (The asymptotic critical value is 3.84 and does not depend on the sign of the estimate.) Viewing the same issue from a slightly different perspective, 9.2% of the negative estimates of b_1 have χ^2 values that exceeded the asymptotic 5% critical value.

There are three messages to take from these simulations. First, if the first-stage regression does not provide strong statistical evidence of time-varying conditional covariances, then we do not know what critical values to use in the second-stage regression. Second, if we can confidently reject the null of constant conditional covariances, the magnitude of this variation does not need to be large in order to produce reasonable results from the second-stage regression. An F -statistic of about 2.5, although weak in the Staiger/Stock sense, appears sufficient. Third, there is no standard estimation procedure that dominates GMM.

B. The Horizon Used to Measure Consumption Growth

A well-known empirical result is that a period- t stock return shock is positively correlated with measured aggregate consumption growth in both period t and future periods. Grossman, Melino, and Shiller (1987) show that a positive correlation between period- t returns and period- $(t + 1)$ aggregate consumption growth is created by the time averaging that is built into measured consumption.⁹ Lynch (1996) and Gabaix and Laibson (2001) note that transactions costs that delay adjustments to consumption can produce nonzero correlations over longer intervals. Thus the purely contemporaneous covariance between returns and consumption growth underestimates the true covariance. For example, a power utility model written in terms of measured aggregate consumption instead of the instantaneous consumption of a consumer that faces no transaction costs implies (see Gabaix and Laibson (2001))

$$E_{t-1}r_t + \frac{1}{2}\text{Var}_{t-1}(r_t) = \gamma\text{Cov}_{t-1}(r_t, c_{t+K^*} - c_{t-1}), \tag{31}$$

where K^* is sufficiently large. This intuition suggests that we should replace the OLS forecasting regression for consumption growth (16) with a multiperiod counterpart

$$c_{t+k} - c_{t-1} = a_c Y_{c,t-1} + e_{c,t-1,t-1,t+k}. \tag{32}$$

⁹The research on the effects of time averaging began with Working (1960).

The form of regression (20) is unchanged except for the replacement of the single-period fitted residual $\hat{e}_{c,t-1,t-1,t}$ with the multiperiod residual $\hat{e}_{c,t-1,t-1,t+k}$. Denote the corresponding ex post covariance by

$$\text{Cov}^*(r_t, c_{t+k} - c_{t-1}) \equiv \hat{e}_{r,t-1,t-1,t} \hat{e}_{c,t-1,t-1,t+k}. \tag{33}$$

Then the instrumental variables regression is

$$r_t + \frac{1}{2} \text{Var}^*(r_t) = b_0 + [b_1 + b_2 p_{t-1}] \text{Cov}^*(r_t, c_{t+k} - c_{t-1}) + w_t. \tag{34}$$

We can set $k = 0$ in (32) and (34) to recover the regressions (16) and (20).

Notwithstanding the logic underlying (31), for the purposes of this paper it is not clear that an empirical implementation that uses $k > 0$ is superior to an implementation that uses $k = 0$. Gabaix and Laibson emphasize the effect of k on the ratio of mean excess returns to the unconditional covariance between stock returns and consumption growth. In (34) these unconditional moments are effectively picked up by b_0 , which is not of direct interest. Under the null hypothesis that the price of consumption risk is zero, the choice of k has no effect. Under an alternative hypothesis the critical issue is how the choice of k affects the power of the statistical test.

To understand the effect of k on the power of the test, first note that the left-hand side of (32) is the sum of k single-period changes in log consumption. Therefore, we can write the multiperiod residual as the sum of k single-period residuals. These are the residuals from forecasts made at $t - 1$ of $\Delta c_t, \Delta c_{t+1}$, and so on. Similarly, we can write the product of the fitted residuals to stock returns and multiperiod consumption growth as the sum of k products of fitted residuals to stock returns and single-period consumption growth:

$$\text{Cov}^*(r_t, c_{t+k} - c_{t-1}) = \sum_{j=0}^k \hat{e}_{r,t-1,t-1,t} \hat{e}_{c,t-1,t-1+j,t+j}. \tag{35}$$

Each of the ex post covariance estimates on the right-hand side of (35) is the sum of its conditional expectation and a residual:

$$\hat{e}_{r,t-1,t-1,t} \hat{e}_{c,t-1,t-1+j,t+j} = \text{Cov}(r_t, \Delta c_{t+j} | Z_{t-1}) + n_{t-1,t-1+j,j}. \tag{36}$$

The first subscript on the noise term refers to the date of the instrument vectors used to construct the fitted residuals. The second and third are the beginning and ending dates used to define the change in log consumption. We can think of these ex post covariances as the sum of a signal (the conditional covariance) and noise ($n_{t-1,t-1+j,j}$). Increasing k , or in other words, adding additional ex post covariance estimates to (35), increases both the total signal and the total noise. From the perspective of hypothesis testing, the net effect depends on the change in the total signal to total noise.

An example makes this clear. Assume the true model is power utility plus transaction costs, so that expected excess returns are proportional to the conditional covariance between returns and consumption growth over K^* periods.

Also assume that investors and econometricians are working with the same information set, so we can ignore any errors-in-variables problem. Denote the conditional covariance between returns and long-horizon consumption growth as B_{t-1} :

$$\text{Cov}(r_t, c_{t+K^*} - c_{t-1} | Z_{t-1}) \equiv B_{t-1}. \tag{37}$$

Then excess returns are the sum of a term proportional to B_{t-1} plus a residual

$$r_t + \frac{1}{2} \text{Var}(r_t | Z_{t-1}) = \gamma B_{t-1} + \epsilon_{r,t}. \tag{38}$$

To simplify the analysis, assume that for $k \leq K^*$, the conditional covariance between returns and the j -period-ahead consumption growth is proportional to B_{t-1} :

$$\text{Cov}_{t-1}(r_t, \Delta c_{t+j}) = \rho_j B_{t-1}, \quad \rho_j \geq 0, \quad \sum_{j=0}^{K^*} \rho_j = 1. \tag{39}$$

In other words, these conditional covariances all move together.

We can combine this equation with (38) to show that for arbitrary k ,

$$r_t + \frac{1}{2} \hat{\epsilon}_{r,t-1,t-1,t}^2 = \frac{\gamma}{\sum_{j=0}^k \rho_j} \text{Cov}^*(r_t, c_{t+k} - c_{t-1}) + \left[\epsilon_{r,t} - \frac{\gamma}{\sum_{j=0}^k \rho_j} \sum_{j=0}^k n_{t-1,t-1+j,t+j} + \frac{1}{2} (\hat{\epsilon}_{r,t-1,t-1,t}^2 - \text{Var}(r_t | Z_{t-1})) \right]. \tag{40}$$

The term in square brackets is unforecastable with the instruments Z_{t-1} .

Now consider the instrumental variables regression (34) with $b_2 = 0$. The dependent variable in the regression is the left-hand side of (40). The explanatory variable is the ex post covariance on the right-hand side of (40). Therefore, the estimate of b_1 should approach the coefficient in front of the ex post covariance. More precisely, for a given k , the instrumental variable estimate of b_1 , denoted $\hat{b}_{1\{k\}}$, satisfies

$$\text{plim } \hat{b}_{1\{k\}} = \frac{\gamma}{\sum_{j=0}^k \rho_j}. \tag{41}$$

For $k < K^*$, this plim exceeds the true coefficient of relative risk aversion, a result consistent with the intuition of Gabaix and Laibson. Therefore, an increase

in k tends to push the parameter estimate toward the true value of γ . However, the asymptotic t -statistic of $\hat{b}_{1(k)}$ does not necessarily increase with k . If we ignore the dependence of $\hat{e}_{r,t-1,t-1,t}\hat{e}_{c,t-1,t-1,t+k}$ on first-stage parameter estimates, it is straightforward to show that the asymptotic t -statistic is inversely proportional to the standard deviation of the term in square brackets in (40). Therefore, from the perspective of the power of the test that $b_1 = 0$, adding an additional lead of consumption growth has two offsetting effects. The first is that an additional noise term is added to the sum $\sum_{j=0}^k n_{t-1,t-1+j,t+j}$. The second is that $\sum_{j=1}^k \rho_j$ increases, which dampens the effect of the sum of noise terms.

It is hard to say more about the effects of k without assuming something about the joint distribution of the error terms ϵ_{rt} , $n_{t-1,t-1+j,t+j}$ and $\hat{e}_{r,t-1,t-1,t}^2 - \text{Var}(r_t | Z_{t-1})$. If they are jointly uncorrelated, the effect of k depends on the ratio

$$\frac{\sqrt{\sum_{j=0}^k \text{Var}(n_{t-1,t-1+j,t+j})}}{\sum_{j=0}^k \rho_j}. \quad (42)$$

A higher value of k results in a less powerful test if it increases this ratio. The extreme case is if excess stock returns r_t and future consumption growth Δc_{t+j} are uncorrelated. Then their product is pure noise, so including the j^{th} lead of consumption results in lower power.

In summary, the choice of k has no effect under the null and has an ambiguous effect on power under alternative hypotheses. The best choice of k likely depends on the precise alternative hypothesis of interest. I adopt an ad hoc procedure for choosing k based on the unconditional sample covariances between fitted residuals to stock returns and leads of consumption growth.

Table II reports the unconditional covariance between innovations in stock returns and consumption growth, where the innovations are residuals from the first-stage regressions (15) and (16). The sample period is February 1960 through December 2001. The relation between stock returns and consumption growth is the strongest for contemporaneous consumption growth. The covariances drop off to zero beyond $j = 3$. I therefore, focus on two different measures of 'contemporaneous' consumption growth. They are the growth in consumption from $t - 1$ to t ($k = 0$) and from $t - 1$ to $t + 3$ ($k = 3$), respectively.

C. The Stock Return Horizon

Most of the data used in this paper are available at a monthly frequency, which leads me to focus on the predictability of monthly returns. For comparability with other research, I also predict returns over longer horizons using conditional covariances appropriate for those horizons.

Table II
The Relation between Aggregate Stock Returns and Real Consumption Growth

The table reports sample covariances and correlations between the innovation in month t 's log excess aggregate stock return and innovations in contemporaneous and future log changes in real per capita consumption on nondurables and services. Innovations are residuals from OLS regressions on a set of instruments realized prior to the beginning of month t . The sample is 503 observations from February 1960 through December 2001.

Consumption Measure	Residuals	
	Covariance ($\times 10^{-4}$)	Correlation
Δc_t	0.308	0.182
Δc_{t+1}	0.116	0.066
Δc_{t+2}	0.116	0.067
Δc_{t+3}	0.130	0.074
Δc_{t+4}	0.026	0.015
Δc_{t+5}	-0.001	-0.001

Consider the n -month excess log stock return from month $t - (n - 1)$ to month t , $\sum_{i=0}^{n-1} r_{t-i}$. An ex post estimate of the covariance of this return with contemporaneous consumption growth is

$$\text{Cov}^* \left(\sum_{i=0}^{n-1} r_{t-i}, c_{t+k} - c_{t-n} \right) \equiv \sum_{i=0}^{n-1} \hat{e}_{r,t-i-1,t-i-1,t-i} \hat{e}_{c,t-i-1,t-i-1,t-i+k}. \quad (43)$$

If k is positive in (43), then consumption growth is measured over a longer horizon than returns, as discussed in the previous subsection.

This measure is not simply the product of the residuals of the n -month stock return and the $(n + k)$ -month log change in consumption. This latter product is also an ex post covariance estimate, but it does not take advantage of higher frequency data. An example will help clarify the difference between the measure in (43) and the product of multi-month residuals. Consider the case of quarterly stock returns, or $n = 3$. Then (43) is an ex post estimate of the quarterly covariance between returns and consumption growth that is constructed with monthly data.¹⁰ For example, if $k = 0$, this estimate is the sum of the monthly ex post contemporaneous covariance estimates for each month in the quarter. The alternative ex post quarterly covariance estimate $(\hat{e}_{r,t-3,t-3,t})(\hat{e}_{c,t-3,t-3,t})$ is effectively the sum of nine monthly ex post covariance estimates; the residual for each month's stock return multiplied by the residual for each month's consumption growth. Since the noncontemporaneous covariances should be zero with $k = 0$, including them in the covariance estimate simply adds noise.

¹⁰ Quarterly consumption growth here is defined as the log change from month t to month $t + 3$. Because of time averaging this differs from the log change in average consumption during quarter τ to quarter $\tau + 1$.

Multi-month estimates of the ex post variance of stock returns are defined analogously to the covariance estimate:

$$\text{Var}^* \left(\sum_{i=0}^{n-1} r_{t-i} \right) \equiv \sum_{i=0}^{n-1} \hat{\sigma}_{r,t-i-1,t-i-1,t-i}^2. \quad (44)$$

The excess stock return from $t - n + 1$ to t (adjusted for Jensen's inequality) is regressed on (43) using the instrument vector Z_{t-n} :

$$\begin{aligned} & \sum_{i=0}^{n-1} r_{t-i} + \frac{1}{2} \text{Var}^* \left(\sum_{i=0}^{n-1} r_{t-i} \right) \\ &= b_0 + [b_1 + b_2 p_{t-n}] \text{Cov}^* \left(\sum_{i=0}^{n-1} r_{t-i}, c_{t+k} - c_{t-n} \right) + w_t. \end{aligned} \quad (45)$$

If $n = 1$ in (45), we recover the single-period instrumental regression (20), or with a longer horizon for measuring consumption growth, (34).

In this paper I report results for $n = 1$ (monthly returns) and $n = 3$ (quarterly returns). In order to avoid overlapping observations with $n = 3$, I estimate (45) on every third monthly observation. I choose the last month of each quarter. Since GMM estimation requires that this equation be estimated jointly with the zero-stage regressions, which are defined at a monthly frequency, I adopt the following procedure. The zero-stage regressions used to construct monthly residuals are estimated using just the last month of each quarter. Since construction of quarterly ex post variances and covariances requires residuals for each month in the quarter, the parameter estimates from these zero-stage regressions are then used to construct fitted residuals for all months in the sample.

Regressions of stock returns on higher frequency estimates of second moments have a long history in finance, although applications have focused on estimates of stock return volatility. In spirit, the combination of (43) and (45) is similar to the procedure adopted by Whitelaw (1994). He uses the sum of squared daily aggregate stock returns to construct ex post estimates of longer horizon aggregate stock return volatility. He then uses an instrumental variables setup to regress longer horizon returns on these volatility estimates.

IV. The Predictability of the Conditional Covariance

In this section I summarize the evidence from first-stage regressions in which ex post covariance estimates are projected on a set of instruments. Figure 3 displays the time series of the ex post estimates. For Panel A, consumption growth is measured by the change in log consumption from $t - 1$ to t . For Panel B, it is measured by the change from $t - 1$ to $t + 3$. The slow adjustment of consumption means that the latter measure produces ex post estimates that are both larger and more volatile than those produced with the former measure. The figure also shows that the outliers are more extreme when consumption growth is measured over the longer horizon. In particular, the observation for

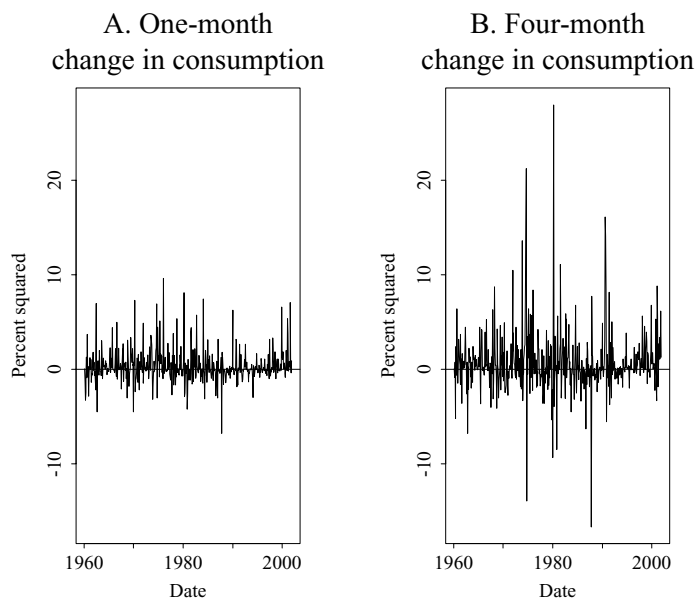


Figure 3. Ex post estimates of the covariance between aggregate stock returns and consumption growth. The innovations to month- t 's aggregate excess stock return and per capita real consumption growth are constructed with OLS forecasting regressions. Consumption growth is measured with either the change from month $t - 1$ to month t (Panel A) or from month $t - 1$ to month $t + 3$ (Panel B). The product of these innovations is an ex post estimate of the month t conditional covariance. The sample period is February 1960 through December 2001.

March 1980 is more than 8 standard deviations away from 0. This observation was driven by a particular event: the credit controls announced by the Carter administration in March. The stock market fell 12% in the month, while aggregate consumption fell dramatically through May before leveling off in June.¹¹ The credit controls probably distorted investors' first-order conditions. Conveniently, this single observation generally has little effect on any of the results reported below. When there is an exception to this rule I note it.

Results from the first-stage regressions are displayed in Table III. In Panel A of the table, the horizon over which the covariance is measured is noted in the first column. Consumption growth is measured as indicated in the second column. The first four rows of Panel A report results for the entire set of instruments itemized in (24). There is overwhelming statistical evidence against the null hypothesis of constant conditional covariances. For each horizon and consumption measure, the null hypothesis is rejected at less than the 1% level. The F -statistics are similar in size to those used in the Monte Carlo simulation of Section III.A. The fitted values for the monthly regressions are displayed in Figure 4. In relative terms, there is substantial variation in the conditional

¹¹ The controls were weakened in late May and completely removed at the beginning of July.

Table III
Predicting Ex Post Covariances between Stock Returns and Consumption Growth

The month t ex post covariance is the product of the innovations in month t 's aggregate stock return and contemporaneous consumption growth (defined by "measure of change in monthly consumption"). Quarterly ex post covariances are sums of three monthly ex post covariances. Ex post covariances are regressed on instruments realized prior to month t . They are the ratio of stock market wealth to consumption ME/C, the consumption-wealth ratio \widehat{cay} , two lags of quarterly excess stock returns, and four lags of quarterly ex post covariances. "Ex covariance lags" refers to the first four instruments and "only covariance lags" refers to the final four instruments. In Panel A the p -values correspond to tests of the hypothesis that the coefficients are all zero. The covariance lags are not included in the regressions reported in Panel B. In this panel, p -values are in square brackets. Regressions that do not include lagged ex post covariances are adjusted for 12 months of moving average residuals. The monthly sample is 503 observations from February 1960 through December 2001. The quarterly sample is 167 observations from 1960Q2 through 2001Q4.

Panel A. Multivariate Regression Results					
Horizon	Measure of Change in Monthly Consumption	Instruments	F -Statistic	χ^2 -Statistic	p -Value
Monthly	$c_t - c_{t-1}$	All	2.89	22.66	0.004
Monthly	$c_{t+3} - c_{t-1}$	All	2.06	20.28	0.009
Quarterly	$c_t - c_{t-1}$	All	3.33	29.62	0.000
Quarterly	$c_{t+3} - c_{t-1}$	All	2.89	25.26	0.001
Monthly	$c_t - c_{t-1}$	Only covariance lags	3.18	11.24	0.024
Monthly	$c_{t+3} - c_{t-1}$	Only covariance lags	2.57	12.48	0.014
Quarterly	$c_t - c_{t-1}$	Only covariance lags	4.19	19.01	0.001
Quarterly	$c_{t+3} - c_{t-1}$	Only covariance lags	4.79	13.89	0.008
Monthly	$c_t - c_{t-1}$	Ex covariance lags	2.29	18.82	0.001
Monthly	$c_{t+3} - c_{t-1}$	Ex covariance lags	1.73	8.89	0.064
Quarterly	$c_t - c_{t-1}$	Ex covariance lags	2.49	18.14	0.001
Quarterly	$c_{t+3} - c_{t-1}$	Ex covariance lags	1.66	11.16	0.025

Panel B. Detailed Results for $c_t - c_{t-1}$ at the Monthly Horizon					
Coefficient ($\times 10^4$) on:					
ME/C	\widehat{cay}	Quarterly Stock Return (1st Lag)	Quarterly Stock Return (2nd Lag)	F -Statistic	χ^2
0.267	—	−0.900	−1.996	2.60	10.37
(2.30)		(−0.96)	(−1.91)		[0.016]
—	−11.259	−0.672	−2.008	2.81	17.24
	(−3.12)	(−0.73)	(−1.93)		[0.001]

covariance. As a fraction of the mean conditional covariance, the interquartile range of the fitted values is about 0.3–1.6.

The table's remaining results examine the predictive power of subsets of these instruments. The second four rows of Panel A report results of regressions of ex post covariances on four quarterly lags of ex post covariances.¹² We see that

¹² Recall that the explanatory variables are calculated from the product of demeaned returns and consumption growth rather than from the product of OLS regression residuals. This difference has little effect on the results.

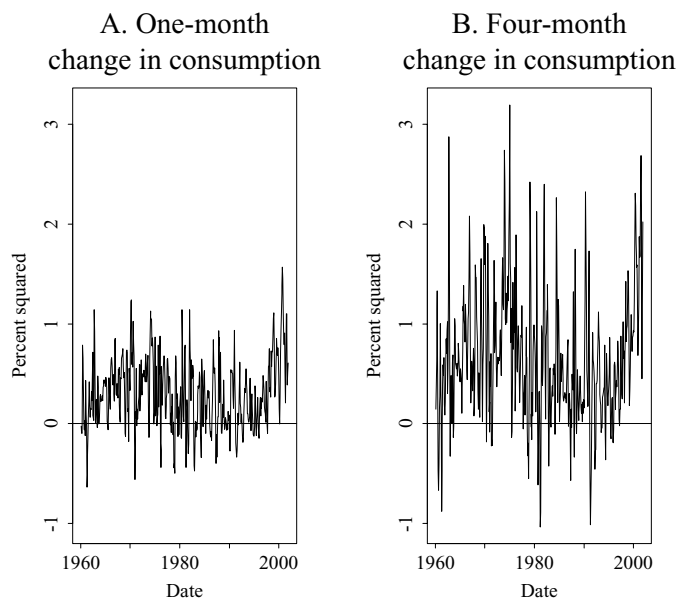


Figure 4. Fitted conditional covariances between aggregate stock returns and consumption growth. The monthly ex post covariance estimates displayed in Figure 3 are regressed on instruments realized prior to month t . The fitted values are displayed in this figure. The sample period is February 1960 through December 2001.

ex post covariances are autocorrelated. For all specifications the null hypothesis of no predictability of the conditional covariance is rejected at the 3% level. However, the signs of these autocorrelations are not typical of second moments in financial market data. The autocorrelations are negative for the first few lags and are positive thereafter. (These coefficients are not reported in any table.) These autocorrelations are responsible for the high-frequency fluctuations in conditional covariances that are evident in Figure 4. This pattern is not inherited from the autocorrelation properties of the volatility of either stock returns or consumption growth. The nonmonotonicity of these autocorrelations raises a number of questions, but addressing them takes us too far afield of the main purpose of this paper.

The final four rows of Panel A examine the joint explanatory power of the stock market wealth–consumption ratio, the consumption–wealth ratio, and the lags of stock returns. Because of the serial correlation in ex post covariances discussed in the preceding paragraph, the test statistics are adjusted for 12 months (four quarters) of moving average residuals using the technique of Newey and West (1987). The pattern of results supports the conclusion that these instruments have predictive power for the conditional covariance. In three of the four regressions, the hypothesis that the coefficients are jointly zero is rejected at the 3% level. The regression that does not exhibit strong statistical significance is somewhat sensitive to the inclusion of the March 1980 observation. When it

is excluded, the regression's F -statistic rises from 1.73 to 2.07 and the p -value of the χ^2 -statistic falls from 0.064 to 0.056.

One potential criticism of the predictive power of ME/C and \widehat{cay} is that investors do not observe them in real time. To address this criticism, I replace the month- t values of ME/C and \widehat{cay} in the instrument vector Z_t with their one-quarter-earlier values. I do not report the results in detail, because they are easily summarized. Lagging these instruments by a quarter typically increases their predictive power for the conditional covariance. For example, the F -statistics of the final four regressions in Panel A increase to 2.55, 2.72, 3.42, and 3.46, respectively.

To conserve space I do not report individual regression coefficients for any of the regressions in Panel A. However, one characteristic of these coefficients deserves mention. Because ME/C and \widehat{cay} are so highly (negatively) correlated, the individual t -statistics associated with these variables are generally too small to reject the hypothesis that any individual coefficient differs from zero. Thus in Panel B I take a closer look at the relative predictive power of ME/C and \widehat{cay} . Again to conserve space, I focus on monthly covariances defined using $c_t - c_{t-1}$. The first regression includes ME/C and the two lagged stock returns. The second regression replaces ME/C with \widehat{cay} . Asymptotic t -statistics are in parentheses and the p -value of a $\chi^2(3)$ test that the coefficients are jointly zero is in brackets. Standard errors are adjusted for 12 lags of moving-average residuals.

We see in Panel B that, as implied by the composition effect, higher values of ME/C correspond to higher conditional covariances. Higher lagged stock returns correspond to lower conditional covariances, consistent with their inverse relation to stock return volatility. Higher values of \widehat{cay} correspond to lower conditional covariances. The test statistics indicate that \widehat{cay} is relatively more important than ME/C in capturing variation in the conditional covariance. This is probably because \widehat{cay} not only picks up much of the composition effect, but also has independent forecasting power for stock return volatility. We will see in Section VI that ME/C is unrelated to stock return volatility.

The model of Santos and Veronesi (2003) implies that the ratio of labor income to consumption is inversely related to the conditional covariance. To test their model, I replace ME/C with their ratio and repeat the regression in Panel B. The results, which are not detailed in any table, do not support the model. The estimated coefficient for the ratio is actually positive, although it is statistically almost identical to zero. The t -statistic is 0.20. The F -statistic of the regression (which includes the two lags of quarterly stock returns) is only 1.51. Given the logic of Santos and Veronesi, this result is surprising because labor income/consumption should move in lockstep with (the inverse of) ME/C. However, over the sample period examined in this paper, their correlation is only -0.68 . The wedge between these two ratios suggests that other assets are also important to consumption, breaking the tight link between labor income's contribution and the stock market's contribution. Hence when Santos and Veronesi regress stock returns on the lagged labor income/consumption ratio, their results are uninformative about the effect on expected stock returns of variations in the conditional covariance.

The main conclusions from these results are that the conditional covariance is strongly predictable and positively associated with stock market wealth. Given this evidence, why does earlier research find only weak variability in the conditional covariance? The likely reason is that the instruments used in this earlier work do not pick up the composition effect. For example, Li (2001) uses lagged stock returns, D/P, and term structure variables. Although there are probably many potential instruments other than ME/C and \widehat{cay} that are reasonable proxies for time variation in the composition of wealth, researchers did not happen to choose them.

V. Predicting Stock Returns with the Conditional Covariance

This section reports results from estimation of (20) and two related regressions: equation (31), which is the version that accounts for slow adjustment in consumption, and equation (45), which examines longer horizon covariances. In interpreting the results, it is helpful to recall that mean excess stock returns are very high relative to the unconditional covariance between shocks to stock returns and shocks to consumption growth. Over the sample period examined in this paper, the mean excess stock return (adjusted for Jensen's inequality) is 0.49% per month. As reported in Table II, the contemporaneous monthly covariance is about 3.1×10^{-5} , or 6.7×10^{-5} if three additional leads of consumption growth are included. The implied coefficients of relative risk aversion are 160 and 75, respectively. Consumption-based asset-pricing models with representative agents have difficulty explaining why investors are so sensitive to consumption risk. Qualitatively, the same pattern is exhibited in the parameter estimates reported below: small changes in conditional covariances correspond to large changes in expected excess returns.

A. The Linear Case

I begin by considering the linear instrumental variables regression of excess stock returns on ex post covariances. The results are displayed in Table IV. The main message of these results is that there is a strong inverse relation between conditional covariances and expected excess returns. Consider, for example, the first row in the table. It reports results at the monthly horizon using the entire set of instruments. When consumption growth is measured from month $t - 1$ to month t , the point estimate is -134 . The t -statistic of -2.32 allows us to reject the null hypothesis of no predictability at the 5% level. (This rejection is based on the finite-sample distribution of the statistic under the relevant null, from the Monte Carlo results in Table I.) When consumption growth is measured from $t - 1$ to t , the estimated coefficient on the conditional covariance is -84 with a t -statistic of -2.47 . The point estimate is closer to zero because conditional covariances are larger and more volatile when consumption growth is measured over the longer period. The estimates (regardless of the measure of consumption growth) imply an interquartile range of expected excess returns from about 3% per year to about 11% per year.

Table IV
Regressions of Stock Returns on the Conditional Covariance

Excess stock returns measured over the given horizon are regressed, using instrumental variables, on ex post estimates of the contemporaneous covariance. The table reports the parameter estimates and asymptotic t -statistics. The “contemporaneous” change in month t 's consumption growth is measured by either $c_t - c_{t-1}$ or $c_{t+3} - c_{t-1}$. The instruments used in the regression are the ratio of stock market wealth to consumption ME/C, the consumption–wealth ratio \widehat{cay} , two lags of quarterly excess stock returns, and four lags of quarterly ex post covariances. “Ex covariance lags” refers to the first four instruments and “only covariance lags” refers to the final four instruments. Estimation is with GMM. The monthly sample is 503 observations from February 1960 through December 2001. The quarterly sample is 167 observations from 1960Q2 through 2001Q4.

Horizon	Instruments	Measure of Change in Monthly Consumption	
		$c_t - c_{t-1}$	$c_{t+3} - c_{t-1}$
Monthly	All	-134.48 (-2.32)	-84.07 (-2.47)
Quarterly	All	-168.07 (-3.19)	-94.92 (-3.73)
Monthly	Ex covariance lags	-237.14 (-2.29)	-144.05 (-2.55)
Quarterly	Ex covariance lags	-180.95 (-1.98)	-133.91 (-2.63)
Monthly	Only covariance lags	-27.52 (-0.34)	-31.01 (-0.71)
Quarterly	Only covariance lags	-101.58 (-1.37)	-52.93 (-1.66)

One way to interpret the economic significance of this predictability is to compute the ratio of expected excess returns to the conditional covariance. When consumption growth is measured from $t - 1$ to t , the point estimates imply that the interquartile range of this ratio—intuitively, the sensitivity of expected returns to consumption risk—is from 52 to 1600. When consumption growth is measured from $t - 1$ to $t + 3$, the corresponding interquartile range is from 28 to 360. Of course, time variation in conditional relative risk aversion is not news. Since the numerator of the ratio varies through time, the standard assumption of a constant denominator is sufficient to conclude that relative risk aversion varies. However, the magnitude of variability documented here is dramatically wider than is implied by the assumption of constant conditional covariances.

The rest of the table fleshes out the nature of the inverse relation. There are two other points to take from these additional results. First, the strength of the relation is largely unaffected by the horizon—monthly or quarterly—over which returns and covariances are measured. Second, only part of the conditional covariance is related to expected returns. Recall from Table III that the stock price variables (ME/C, \widehat{cay} , and the two lags of quarterly stock returns) and the lagged ex post covariances have independent predictive power for conditional covariances. The results in this table tell us that only the first set of instruments produce a statistically strong relation between expected returns and conditional

covariances. When only the lagged covariances are used as instruments in the regression of stock returns on ex post covariances, the estimated coefficients are uniformly statistically insignificant. In fact, in results not detailed here, I find the statistical strength relies on the inclusion of \widehat{cay} in the instrument set. Without \widehat{cay} there is not enough business cycle variation in the conditional covariance to pick up the business cycle variation in expected returns. (But if only \widehat{cay} is used in the instrument set, there is no statistically significant predictability either; \widehat{cay} alone does not generate sufficient variability in the conditional covariance.)

These results are not particularly surprising given the evidence from the last section that conditional covariances are procyclical. The more interesting application of the time variation in conditional covariances is to test models in which investors exhibit time-varying sensitivity to consumption risk.

B. The Nonlinear Case

An inverse relation between conditional covariances and expected excess stock returns is not necessarily inconsistent with consumption-based asset pricing. If the price of consumption risk tends to be low when conditional covariances are high and vice versa, we may observe an inverse unconditional relation between conditional covariances and expected excess returns. We cannot immediately rule out this possibility because intuition suggests that the price of consumption risk is countercyclical, and thus inversely related to observed conditional covariances.

If we have an observable proxy for time variation in the price of consumption risk, then we can test whether time variation in the price of risk explains the empirical results documented above. The most obvious test is to use instrumental variables to regress excess stock returns on both the ex post covariance and the proxy for the price of risk. Then we could interpret the coefficient on the ex post covariance as the relation between expected returns and conditional covariances, holding constant the price of risk.

But we have a more powerful test at our disposal. For a fixed price of risk, the relation between the conditional covariance and expected excess returns should not only be positive, but it should be larger when the price of risk is high than when it is low. Thus there should be a nonlinear relation between expected returns and conditional covariances that is picked up by the price of risk. This is the logic behind my use of the more general functional form $b_1 + b_2 p_{t-1}$ for the sensitivity of expected returns to the conditional covariance, where p_t is a proxy for the price of risk. If, say, p_t is positively related to the price of risk, then b_2 should be positive.

To fix ideas, assume the true relation between expected excess returns and the conditional covariance is

$$E_{t-1}r_t + \frac{1}{2}\text{Var}_{t-1}(r_t) = f(p_{t-1})\text{Cov}_{t-1}(r_t, \Delta c_t), \tag{46}$$

where $f(\cdot)$ is some monotonic function of the proxy for the price of risk. In the empirical test I use $b_1 + b_2 p_{t-1}$ in place of $f(p_{t-1})$, allowing us to interpret it as a linear approximation to $f(p_{t-1})$ and in particular to interpret b_2 as an average value of $f'(p_{t-1})$. Under the null that $f(\cdot)$ is a constant (i.e., p_{t-1} is unrelated to the price of risk), then $b_2 = 0$. We can test this hypothesis formally. An informal test is to check that the fitted values of $b_1 + b_2 p_{t-1}$ are positive.

The first proxy for the price of risk I examine is Wachter's measure of surplus consumption. High surplus should correspond to a low price of risk, so we expect b_2 to be less than 0. In order for the negative relation between expected returns and conditional covariances to be explained by variations in surplus consumption, surplus and conditional covariances must be positively correlated. For the data examined in this paper, the sample correlations between surplus consumption and the fitted conditional covariances (based on all the instruments) are around 0.2. The second proxy is the consumption–wealth ratio \widehat{cay} . A high consumption–wealth ratio suggests a high price of risk (low valuations of wealth). The sample correlations between \widehat{cay} and the fitted conditional covariances are about -0.3 and -0.4 at the monthly and quarterly horizons, respectively. The third proxy I examine is the dividend–price ratio. Its relation to the conditional covariance is weaker. The sample correlations between D/P and the fitted conditional covariances are about -0.2 and -0.1 at the monthly and quarterly horizons, respectively.

I now turn to estimates of (20) and its generalizations (31) and (45). To simplify interpretation of the parameter estimates, I transform the proxies for the price of risk to have a mean of 0 and a standard deviation of 100.¹³ The results are displayed in Table V. The results offer only minimal support for the view that any of the proxies is associated with the price of consumption risk. The strongest conclusion supported by the results is that all but one of the point estimates of b_2 have the correct sign: negative for surplus consumption and positive for \widehat{cay} and D/P. However, only one of the 12 point estimates has an asymptotic t -statistic that exceeds 2 in absolute value.

Moreover, because the estimated coefficients are so close to zero in an economic sense, the fitted values of $b_1 + b_2 p_t$ are typically negative. Consider, for example, the evidence in the table that provides the strongest support for a role for surplus consumption: monthly covariances and consumption growth measured from $t - 1$ to $t + 3$. The interquartile range of the fitted values of $b_1 + b_2 \hat{s}_{t-1}$ is -4 to -88 . Thus even when we fix surplus consumption to a relatively low value (and thus presumably fix the price of consumption risk to be relatively high), there is an inverse relation between expected excess returns and conditional covariances. The strongest evidence in the table for a time-varying sensitivity of expected returns to the conditional covariance is with \widehat{cay} . Again using monthly covariances and consumption growth measured from $t - 1$ to $t + 3$, the interquartile range of the fitted values of $b_1 + b_2 \widehat{cay}_{t-1}$ is 1 to

¹³ In addition, because surplus consumption and D/P are not elements of Z_t , I add the relevant variable to the set of instruments used to estimate the regression. They both have little explanatory power for the conditional covariance.

Table V
Nonlinear Regressions of Stock Returns on the Conditional Covariance

Excess stock returns measured over the given horizon are regressed, using instrumental variables, on ex post estimates of the contemporaneous covariance. The coefficient, denoted b_t , is allowed to depend on the “price of risk variable”: $b_t = b_1 + b_2 p_{t-1}$, where p_t is the specified variable. This variable is transformed to have a mean of 0 and standard deviation of 100. The table reports the parameter estimates and t -statistics. The “contemporaneous” change in month t ’s consumption growth is measured by either $c_t - c_{t-1}$ or $c_{t+3} - c_{t-1}$. The instruments used in the regression are the ratio of stock market wealth to consumption ME/C, the consumption–wealth ratio \widehat{cay} , two lags of quarterly excess stock returns, and four lags of quarterly ex post covariances. Estimation is with GMM. The monthly sample is 503 observations from February 1960 through December 2001. The quarterly sample is 167 observations from 1960Q2 through 2001Q4.

Horizon	Price of Risk Variable	Measure of Change in Monthly Consumption			
		$c_t - c_{t-1}$		$c_{t+3} - c_{t-1}$	
		b_1	b_2	b_1	b_2
Monthly	Surplus consumption	-132.11 (-1.98)	-0.01 (-0.02)	-55.91 (-1.48)	-0.54 (-1.72)
Quarterly	Surplus consumption	-152.24 (-2.82)	-0.50 (-0.95)	-77.73 (-2.54)	-0.38 (-1.16)
Monthly	\widehat{cay}	-110.80 (-1.65)	0.48 (0.82)	-48.01 (-1.17)	0.71 (2.26)
Quarterly	\widehat{cay}	-161.67 (-2.84)	0.04 (0.06)	-67.05 (-2.25)	0.38 (1.03)
Monthly	D/P	-123.58 (-2.18)	0.25 (0.68)	-84.18 (-2.37)	0.04 (0.18)
Quarterly	D/P	-168.43 (-3.20)	-0.24 (-0.63)	-102.59 (-3.46)	0.02 (0.09)

-91. As with surplus consumption, the relation between expected returns and conditional covariances is generally negative even when holding constant the proxy for the price of consumption risk.

The main conclusions to take from this section are easy to summarize. First, there is a strong inverse relation between conditional covariances and expected excess stock returns. Second, there is little evidence that the sensitivity of this relation varies with standard proxies for the price of consumption risk. The latter conclusion is the one that is more troubling from the perspective of consumption-based asset-pricing models. If investors are particularly risk averse in bad economic times, then why are expected stock returns not more sensitive to fluctuations in conditional covariances at such times?

VI. GARCH Estimates of the Conditional Correlation

In this section I focus on the conditional correlation between stock returns and consumption growth. Although conditional covariances play a central role

in asset pricing, conditional correlations are more intuitive. In addition, the composition effect implies that the conditional correlation depends on the ratio of stock market wealth to consumption.¹⁴ The tool that I use to study conditional correlations is a multivariate GARCH model.

A multivariate GARCH framework is a natural choice for modeling persistent covariance dynamics. The conditional covariances produced by the instrumental variables methodology, displayed in Figure 4, exhibit negative autocorrelation over short horizons. This feature of the data creates specification problems for a GARCH model, which is why I use the IV approach throughout most of this paper. However, conditional variances of both stock returns and consumption growth both exhibit highly persistent dynamics that fit nicely into a GARCH setting. Therefore, from the perspective of modeling conditional correlations, the GARCH framework is a reasonable compromise.

To jointly model conditional variances of stock returns and consumption growth as well as their conditional correlation, I use the Dynamic Conditional Correlation (DCC) model of Engle (2002). Expected stock returns are assumed to depend on the conditional covariance, while consumption growth follows an AR(3).

$$r_t = a_{r0} + a_{r1}H_{t(12)} + \epsilon_{1t}, \quad (47)$$

$$\Delta c_t = a_{c0} + \sum_{i=1}^3 a_{ci} \Delta c_{t-i} + \epsilon_{2t}, \quad (48)$$

$$\epsilon_t \equiv (\epsilon_{1t} \epsilon_{2t})', \quad (49)$$

$$\epsilon_t \epsilon_t' \sim N(0, H_t). \quad (50)$$

The conditional variance–covariance matrix H_t is

$$H_t = D_t R_t D_t, \quad (51)$$

where D_t is a diagonal matrix with $\sqrt{h_{it}}$ on the diagonal and R_t is the conditional correlation matrix. The conditional correlation matrix is written as

$$R_t = Q_t^{*-1} Q_t Q_t^{*-1}, \quad (52)$$

where Q_t^* is a diagonal matrix with the square root of the diagonal elements of Q_t on the diagonal. Standardized shocks to stock returns and consumption growth are denoted by

$$z_t = D_t^{-1} \epsilon_t. \quad (53)$$

¹⁴ We could also look at conditional betas. However, betas are more relevant for cross-sectional comparisons than for time-series comparisons. If, say, the consumption beta of the stock market increases from t to $t + 1$, the implications for expected excess returns are ambiguous: The beta could have risen because the conditional variance of consumption growth declined, while the conditional covariance held constant.

The conditional variances of stock returns and consumption growth are modeled as univariate GARCH(1,1) processes that can also depend on predetermined variables. The predetermined variables are the same variables used earlier: the ratio of stock market wealth to consumption, \widehat{cay} , and two lags of quarterly excess stock returns, that is,

$$h_{it} = v_{it} + \alpha_i(\epsilon_{it-1}^2 - v_{it-1}) + \beta_i(h_{it-1} - v_{it-1}), \tag{54}$$

$$v_{it} = \omega_{0i} + \omega_{1i}\widehat{cay}_{t-1} + \omega_{2i}(ME/C)_{t-1} + \omega_{3i}RET_{t-1}^q + \omega_{4i}RET_{t-4}^q, \quad i = 1, 2. \tag{55}$$

The conditional correlation is modeled similarly

$$Q_t = \begin{pmatrix} 1 & q_t \\ q_t & 1 \end{pmatrix} + \alpha_3 \left[z_{t-1}z'_{t-1} - \begin{pmatrix} 1 & q_{t-1} \\ q_{t-1} & 1 \end{pmatrix} \right] + \beta_3 \left[Q_{t-1} - \begin{pmatrix} 1 & q_{t-1} \\ q_{t-1} & 1 \end{pmatrix} \right], \tag{56}$$

$$q_t = \omega_{03} + \omega_{13}\widehat{cay}_{t-1} + \omega_{23}(ME/C)_{t-1} + \omega_{33}RET_{t-1}^q + \omega_{43}RET_{t-4}^q. \tag{57}$$

I use this functional form so that the effect on H_t of the predetermined variables does not feed back through the GARCH terms. To simplify the interpretation of the coefficients on ME/C, \widehat{cay} , and the lagged quarterly returns, I normalize these variables to have mean 0 and variance 1. I estimate this model with maximum likelihood over the sample July 1959 through December 2001. The initial H_1 is

$$H_1 = \begin{pmatrix} v_{11} & q_1\sqrt{v_{11}v_{21}} \\ q_1\sqrt{v_{11}v_{21}} & v_{21} \end{pmatrix}, \tag{58}$$

where the values of the predetermined variables as of month-end June 1959 are used to construct v_{11} , v_{21} , and q_1 .

The parameter estimates are displayed in Table VI. The fitted values of the conditional standard deviations and correlation are displayed in Figure 5. I discuss the results for stock returns, consumption growth, and their correlation in turn.

Expected excess stock returns are negatively associated with the conditional covariance. The estimated coefficient on the conditional covariance, from Panel A of the table, is -377 . This estimate is about twice as large as the corresponding estimate in Table IV. We see in Panel B that the conditional variance of stock returns is negatively associated with \widehat{cay} and past stock returns, consistent with the previous literature. The more interesting result here is that unlike \widehat{cay} , ME/C has no predictive power for stock return volatility.

A glance at Panel B of Figure 5 reveals that consumption growth volatility has declined steadily over the sample period. This negative drift is consistent with the evidence of Romer (1999) that a variety of macroeconomic measures were substantially more volatile in the first half of the postwar period than in the

Table VI

A GARCH-in-Mean Model of Stock Returns and Consumption Growth

The joint evolution of monthly log excess aggregate stock returns r_t and consumption growth Δc_t is

$$r_t = a_{r0} + a_{r1}\sqrt{h_{1t}h_{2t}}\rho_t + e_{1t}, \quad \Delta c_t = a_{c0} + \sum_{i=1}^3 a_{ci}\Delta c_{t-i} + e_{2t}.$$

The conditional variance of ϵ_{it} is

$$h_{it} = v_{it} + \alpha_i(\epsilon_{it-1}^2 - v_{it-1}) + \beta_i(h_{it-1} - v_{it-1}),$$

$$v_{it} = \omega_{0i} + \omega_{1i}\widehat{cay}_{t-1} + \omega_{2i}(ME/C)_{t-1} + \omega_{3i}RET_{t-1}^q + \omega_{4i}RET_{t-4}^q, \quad i = 1, 2.$$

The conditional correlation ρ_t has the same functional form. Predetermined variables in the variance equations are standardized to have zero means and unit variances. Estimation is with maximum likelihood. The coefficients on the predetermined variables for the variance equations are scaled up by 10^4 . t -statistics are in parentheses. The sample period is July 1959 through December 2001.

Panel A. Conditional Mean Parameters					
a_{r0}	a_{r1}	a_{c0}	a_{c1}	a_{c2}	a_{c3}
0.015	-377.66	0.002	-0.278	0.017	0.180
(2.83)	(-2.33)	(9.00)	(-6.22)	(0.32)	(3.90)

Panel B. Conditional Variance Parameters							
Equation	α	β	Constant	\widehat{cay}_{t-1}	$(ME/C)_{t-1}$	RET_{t-1}^q	RET_{t-4}^q
Stock returns	0.022	0.868	19.211	-5.743	-1.131	-3.261	-3.261
	(0.67)	(4.77)	(13.39)	(-4.80)	(-0.83)	(-2.55)	(-3.28)
Consumption growth	0.009	0.994	0.240	-	-	-	-
	(1.68)	(152.19)	(5.45)				
Correlation	-	-	0.184	-0.085	0.058	0.006	-0.014
			(3.81)	(-1.76)	(1.85)	(0.24)	(-0.53)

second. I exclude the predetermined variables from the variance of consumption growth. I exclude lagged returns and \widehat{cay} because they have no significant explanatory power. The ratio ME/C has modest explanatory power, but this explanatory power is driven by the joint drift in consumption-growth volatility and ME/C. If a linear time trend is included in the variance equation, the significance of ME/C disappears. A similar result holds if the aggregate dividend/price ratio is included in the variance equation. As noted by Bansal and Yaron (2004), consumption growth volatility and the dividend/price ratio are positively associated. Both variables drift down in the postwar period. If the dividend/price ratio and a linear time trend are included in the variance equation, the coefficient on the dividend/price ratio is statistically insignificant.¹⁵

¹⁵ In fact, if the linear time trend is included, the GARCH terms are also statistically insignificant.

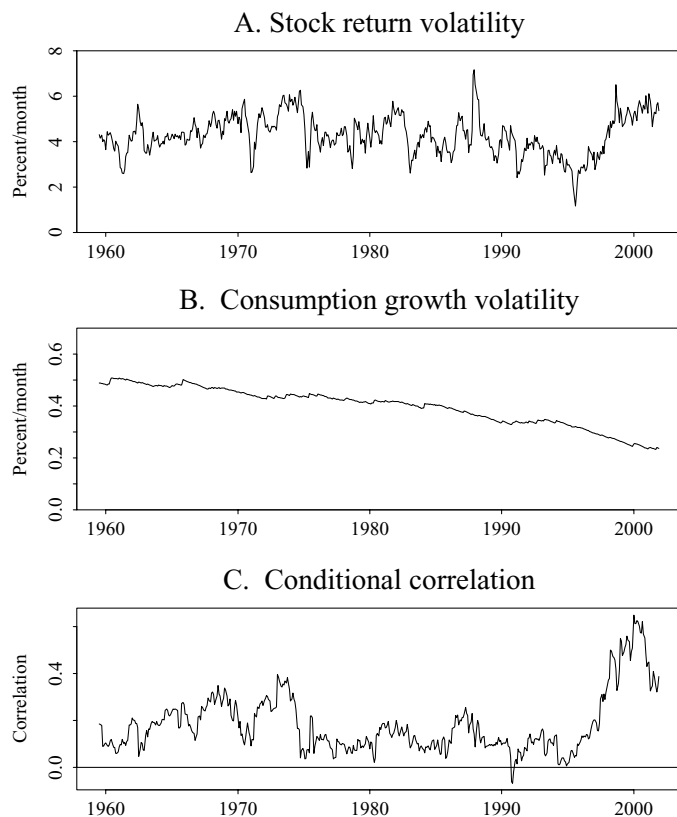


Figure 5. GARCH estimates of the conditional variance–covariance matrix of monthly stock returns and consumption growth. Fitted values from a dynamic conditional correlation representation of a two-dimensional GARCH(1,1) model are displayed. The sample period is July 1959 through December 2001.

Based on this evidence, I am unwilling to draw any conclusions about the explanation for the drift in consumption-growth volatility. Therefore, I exclude ME/C (as well as a linear time trend and the dividend/price ratio) from this equation to avoid forcing an implicit interpretation on this drift.

The equation for the conditional correlation does not include GARCH effects. When the GARCH parameters are restricted to be nonnegative, their maximum likelihood estimates are on the boundary of the parameter space. The problem is basically the one illustrated in Figure 4. Conditional on the predetermined variables, conditional correlations are not positively autocorrelated. Adding additional ARCH terms (while restricting them to be nonnegative) does not help. However, when α is allowed to be negative, the model has too much flexibility. When the GARCH parameters are unrestricted, the estimate of α is negative. One of the fitted conditional correlations is negative 1, producing a log-likelihood of infinity.

The parameter estimates for the conditional correlation equation tell us that the conditional correlation is positively related to ME/C, as implied by the composition effect. It is also negatively related to \widehat{cay} . Because the variables are so strongly correlated, the test statistics do not allow us to distinguish between the roles of ME/C and \widehat{cay} in predicting the conditional correlation. When all variables except for ME/C are dropped from the conditional correlation equation, the coefficient on ME/C is 0.096 with a t -statistic of 2.99.

Lettau, Ludvigson, and Wachter (2004) use a regime-switching model to argue that the high stock valuations in the late 1990s can be explained by the decline in the volatility of consumption growth. Once investors realize they are in a low-volatility regime, their required risk premium on stocks falls. The conditional covariance between stock returns and consumption growth is low because in their model dividend growth is perfectly correlated with consumption growth. Although this story has a nice intuitive flavor, the evidence here casts some doubt on it. During the past 40 years, the conditional correlation between stock returns and consumption growth varied widely around its sample mean of 0.18. It reached a low of roughly 0 in late 1990 and a high of 0.65 in early 2000. Combining these dynamics with those of the conditional variances reveals that the conditional covariance peaks in both mid 1973 and in early 2000. (Similar evidence is displayed in Figure 4.) Hence, although the volatility of consumption growth is low in the late 1990s, the amount of consumption risk in stocks is nevertheless at an historical high because of the high conditional correlation and the high volatility of stock returns.

VII. Conclusions

The conditional covariance between aggregate excess stock returns and aggregate consumption growth is procyclical, a pattern that is consistent with the composition effect. Regardless of the source of this time variation, we can use it to test consumption-based asset-pricing models. The conditional covariance is inversely related to the expected excess return to the market. This relation holds even when we allow the price of consumption risk to vary with measures such as surplus consumption. It will be quite difficult to shoehorn these results into a model with representative agents who care about consumption. Models with heterogeneous agents may fare better, but this is more of a hope than an expectation. As long as the marginal stockholder holds a nontrivial amount of her wealth in a form other than stocks, the composition effect will still imply that the amount of her consumption risk in stocks is procyclical.

REFERENCES

- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 579–626.
- Attanasio, Orazio P., 1991, Risk, time-varying second moments and market efficiency, *Review of Economic Studies* 58, 479–494.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of the equity premium puzzle, *Journal of Finance* 59, 1481–1509.

- Basak, Suleyman, and Domenico Cuoco, 1998, An equilibrium model with restricted stock market participation, *Review of Financial Studies* 11, 309–341.
- Black, Fischer, 1976, Studies of stock price volatility changes, *Proceedings of the 1976 Meetings of the Business and Economics Statistics Section*, 177–181.
- Brandt, Michael W., and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economics* 72, 217–257.
- Brennan, Michael J., and Yihong Xia, 2002, Tay's as good as cay, Working paper, Wharton School, University of Pennsylvania.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.
- Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.
- Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Chan, Yeung Lewis, and Leonid Kogan, 2002, Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices, *Journal of Political Economy* 110, 1255–1285.
- Engle, Robert F., 2002, Dynamic conditional correlation—A simple class of multivariate GARCH models, *Journal of Business and Economic Statistics* 20, 339–350.
- Ferson, Wayne E., and Campbell R. Harvey, 1993, Seasonality and heteroskedasticity in consumption-based asset pricing: An analysis of linear models, *Research in Finance* 11, 1–35.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- Gabaix, Xavier, and David Laibson, 2001, The 6D bias and the equity premium puzzle, *NBER Macroeconomics Annual* 16, 257–312.
- Grossman, Sanford J., Angelo Melino, and Robert J. Shiller, 1987, Estimating the continuous-time consumption-based asset-pricing model, *Journal of Business and Economic Statistics* 5, 315–327.
- Hahn, Jaehoon, and Hangyong Lee, 2001, On the estimation of the consumption-wealth ratio: cointegrating parameter instability and its implications for stock return forecasting, Working paper, University of Washington.
- Hansen, Lars P., 1982, Large sample properties of generalized method of moment estimators, *Econometrica* 50, 1029–1053.
- Hansen, Lars Peter, John Heaton, and Amir Yaron, 1996, Finite-sample properties of some alternative GMM estimators, *Journal of Business and Economic Statistics* 14, 262–280.
- Harvey, Campbell R., 1989, Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics* 24, 289–317.
- Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3–53.
- Kleibergen, Frank, 2002, Pivotal statistics for testing structural parameters in instrumental variables regression, *Econometrica* 70, 1781–1803.
- Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 55, 815–849.
- Lettau, Martin, and Sydney Ludvigson, 2002a, Measuring and modeling variation in the risk-return tradeoff, Working paper, NYU.
- Lettau, Martin, and Sydney Ludvigson, 2002b, Tay's as good as cay: Reply, Working paper, NYU.
- Lettau, Martin, Sydney Ludvigson, and Jessica Wachter, 2004, The declining risk premium: What role does macroeconomic risk play?, Working paper, NYU.
- Li, Yuming, 2001, Expected returns and habit persistence, *Review of Financial Studies* 14, 861–899.
- Lynch, Anthony W., 1996, Decision frequency and synchronization across agents: Implications for aggregate consumption and equity return, *Journal of Finance* 51, 1479–1497.

- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Pagan, Adrian, and Aman Ullah, 1988, The econometric analysis of models with risk terms, *Journal of Applied Econometrics* 3, 87–105.
- Parker, Jonathan A., 2003, Consumption risk and expected stock returns, *American Economic Review Papers and Proceedings* 93, 376–382.
- Romer, Christina D., 1999, Changes in business cycles: Evidence and explanations, *Journal of Economic Perspectives* 13, 23–44.
- Santos, Tano, and Pietro Veronesi, 2003, Labor income and predictable stock returns, Working paper, University of Chicago Graduate School of Business.
- Shapiro, Alexander, 2002, The investor recognition hypothesis in a dynamic general equilibrium: Theory and evidence, *Review of Financial Studies* 15, 97–141.
- Staiger, Douglas, and James H. Stock, 1997, Instrumental variables regression with weak instruments, *Econometrica* 65, 557–586.
- Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.
- Stock, James H., and Jonathan H. Wright, 2002, GMM with weak identification, *Econometrica* 68, 1055–1096.
- Stock, James H., Jonathan H. Wright, and Motohiro Yogo, 2002, A survey of weak instruments and weak identification in generalized method of moments, *Journal of Business and Economic Statistics* 20, 518–529.
- Wachter, Jessica A., 2002, Habit formation and returns on bonds and stocks, Working paper FIN-01-024, Stern School of Business, NYU.
- Whitelaw, Robert F., 1994, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Finance* 49, 515–541.
- Working, Holbrook, 1960, Note on the correlation of first differences of averages in a random chain, *Econometrica* 28, 916–918.
- Yogo, Motohiro, 2003, A consumption-based explanation of expected stock returns, Working paper, Harvard University.