Why does the slope of the term structure forecast excess returns?

*Gregory R. Duffee**

Haas School of Business
University of California – Berkeley

First version: January 2001
This Draft: May 10, 2001

ABSTRACT

In the U.S., the slope of the term structure is positively correlated with expected future excess returns to both stocks and long-maturity bonds. This paper empirically investigates possible explanations for this pattern. The results pose a significant challenge for representative agent, consumption-based asset-pricing models. I find that when the term structure is more steeply sloped than average, (1) future volatility of aggregate consumption growth is much lower than average; (2) future correlations between aggregate consumption growth and returns to both stocks and bonds are roughly zero; (3) the well-documented positive relation between excess stock and bond returns and subsequent growth in real variables (e.g., labor income and GDP) disappears. These facts are inconsistent with the hypothesis that a representative agent has standard power utility, recursive utility, or habit formation preferences.

*Voice 510-642-1435, fax 510-643-1420, email duffee@haas.berkeley.edu. Address correspondence to 545 Student Services Building #1900, Berkeley, CA 94720. I thank seminar participants at the FRB–San Francisco, the Stanford Fixed Income Research Conference, UBC, Barclays Global Investors, and the University of Rochester for helpful comments. The most recent version of this paper is at www.haas.berkeley.edu/∼duffee.

1 Introduction

When the term structure is more steeply sloped than average, future excess returns to stocks and long-term bonds (over short-term interest rates) tend to be higher than average. In this paper I attempt to interpret this well-known pattern in the context of representative-agent, consumption-based utility theory. The results pose a significant challenge to the theory.

Since Rubinstein (1976) and Lucas (1978), economists have attempted to explain the behavior of asset returns with equilibrium models in which a representative agent makes portfolio decisions to maximize utility defined over consumption. In these models, an asset's risk premium is proportional to the covariance between the asset's return and the change in the marginal utility of (aggregate per capita) consumption.

Because we cannot observe marginal utility directly, tests of consumption-based models rely on specifications of preferences that allow us to infer the dynamics of marginal utility from observables. The standard time-additive power utility framework used by Mehra and Prescott (1985) implies that an asset's expected excess return is proportional to the covariance between the asset's return and the growth of log consumption. The recursive utility framework of Epstein and Zin (1989) and Weil (1989) adds another explanatory variable, allowing expected excess returns to also depend on the covariance between the asset's return and the return to total wealth. The state-of-the-art perspective (e.g., Campbell (1999)) is that neither set of preferences can explain the observed time-variation in expected excess returns to stocks and bonds, at least in a representative-agent world. There is little evidence that either the volatility of consumption growth or covariances of asset returns with consumption growth vary through time. Covariances between asset returns and wealth (at least the portion of wealth that we can measure) do vary through time, but they do not appear to vary systematically with expected excess returns.

Habit formation frameworks, developed formally by Sundaresan (1989), Constantinides (1990) and Abel (1990), appear more promising. For example, the model of Campbell and Cochrane (1999) implies that an asset's expected excess return is the product of a time-varying price of consumption risk and the covariance between the asset's return and the growth of log consumption. The price of risk depends on how close current consumption is to a 'consumption habit.' Thus, expected excess returns can vary through time even if covariances do not. The conventional view is that expected excess returns to stocks and bonds positively covary with the slope of the term structure because the slope is countercyclical. When the slope is steep, economic output is relatively low, consumption is correspondingly low relative to habit, and the price of risk is correspondingly high.

However, this logic falters when we take a slightly closer look at the relation between the slope of the term structure and expected excess returns to long-term bonds. As noted by Fama and French (1993), the sign differs depending on whether the term-structure slope is steeper (positive expected returns) or flatter (negative expected returns) than usual. Therefore either the price of consumption risk or the covariance between long-term bond returns and the growth of log consumption has to change sign. Thus if asset prices (and in particular, bond prices) are consistent with habit formation preferences of a risk-averse representative agent, the slope of the term structure should have predictive power for covariances between the growth of log consumption and bond returns.

This intuition is part of the motivation for my paper. I investigate the link between the slope of the term structure and covariances between consumption growth and asset returns. As noted above, there is little existing evidence for predictable variations in such covariances. However, research to date has not looked at the predictive power of the slope. This is surprising, because the slope forecasts future consumption growth. In other words, we already know that the slope has predictive power for the first moment of consumption growth; it is worth exploring whether it has predictive power for second moments.

The analysis, based on quarterly data from 1952 through 2000, produces two main results. The first result overturns the conventional wisdom that consumption volatility is largely unpredictable. There is a strong *negative* relation between the slope of the term structure and the volatility of consumption growth. When the term structure is flatter than usual at the end of quarter t, the standard deviations of growth in log consumption in quarters $t+1$, $t+2$, and $t+3$ are about 1.5 times the corresponding standard deviations when the term structure is steeper than usual. This is not good news for consumption-based models; it says that excess returns are higher when consumption volatility is lower.

Second, there is a strong *negative* relation between the slope of the term structure and correlations between consumption growth and aggregate stock returns. In fact, when the term structure is steeper than usual at the end of quarter t, the quarter $t+1$ correlation is essentially zero. Thus in either a power utility or habit formation framework, we should observe that expected excess stock returns are positive when the term structure is flatter than usual, and zero when the term structure is steeper than usual. In contrast to the results for stock returns, the correlation between consumption growth and long-term bond returns is weak–near zero–regardless of the shape of the term structure. In particular, there is no evidence that the sign of the correlation changes with the slope.

The most surprising of these findings is that stock returns and consumption growth are uncorrelated when the term structure is steeply sloped. I explore this issue further by considering the ability of stock and bond returns to forecast future growth in consumption, GDP, and labor income. I find that when the term structure slope is flatter than usual, future stock and bond returns lead growth in consumption, GDP, and labor income. But when the slope is steeper than usual, this forecasting power disappears.

The main message of these results is that the relation between the macroeconomy and asset (stock and bond) markets depends critically on the information impounded into the slope of the term structure. When the slope is steep, asset returns and the macroeconomy are largely decoupled; when the slope is flatter, the relation is tighter. A byproduct of this message is that representative agent, consumption-based asset-pricing models are inconsistent with the fact that expected excess returns to stocks and bonds are higher when the slope of the term structure is steep.

The next section reviews what we know about the slope of the term structure, expected excess returns, and macroeconomic growth. The third section discusses standard interpretations of this evidence in consumption-based asset-pricing frameworks, and points out some limitations of these interpretations. Section 4 presents most of the empirical results. Concluding comments are offered in Section 5. The Appendix contains a description of the data.

2 The evidence for predictability

In the U.S., the return to long-term Treasury bonds less the contemporaneous return to short-term Treasury bills is, on average, slightly positive. Future excess returns also vary predictably with the slope of the yield curve, a decades-old result that is equivalent to the failure of the expectations hypothesis of interest rates. A textbook summary of the evidence is in Chapter 10 of Campbell, Lo, and MacKinlay (1997). Fama and French (1993) note that the low mean and predictable variation together imply that when the slope of the term structure is steep, expected excess returns to bonds exceed zero, while the sign is reversed when the slope is relatively flat or inverted.

The relation between excess returns to the stock market (nominal stock returns less short-term Treasury bill returns) and the slope of the term structure is weaker and more recently discovered. Campbell (1987) first noted that information in the short end of the Treasury term structure could forecast excess stock returns. Fama and French (1989) found that the spread between long-term Aaa bond yields and short-term Treasury yields forecast excess stock market returns, although the statistical strength of the forecastability depended on both the sample period (stronger post-war) and the horizon over which forecasts were made (stronger at shorter horizons). Chen (1991) uses the slope of the Treasury term structure to forecast quarterly excess stock returns and finds statistically significant forecasting power for returns one and two quarters ahead.[1]

To get a concrete sense of the predictive power of the term-structure slope in the data set examined in this paper, Table 1 reports the results of regressions of quarterly excess returns on the slope of the term structure. One issued that needs to be addressed is the definition of "the" slope of the term structure. Earlier research used a variety of different measures. Alternative choices for the long end include yields on long-maturity, Aaa-rated corporate bonds, ten-year Treasury bonds, or five-year zero-coupon bond yields implied by coupon bonds. Choices for the short end include yields on one-month, three-month, or two-year Treasury securities. I follow Estrella and Hardouvelis (1991) in measuring the slope by the difference between a ten-year Treasury yield and a three-month Treasury yield. Because the results are, in some ways, sensitive to this choice, I also report in footnotes how different definitions affect the evidence of predictability.

The estimated regression is

$$r_t^k - r_t^f = b_0 + b_1 SL_{t-i} + e_t^k,\ k \in \{s, b\},\ i = 1, \ldots, 4, \tag{1}$$

where r_t^k is the log return in quarter t to either the stock market ($k = s$) or long-term Treasury bonds ($k = b$), r_t^f is the log return to short-maturity Treasury bills, and SL_t is the demeaned slope of the term structure at the end of quarter t. I demean the slope so that the constant term corresponds to the unconditional mean excess return in the sample. The construction of the return data is described in the Appendix. The sample period is $t \in$\{1952:1,2000:4\}. The beginning date is the first for which all of the data are available from the Center for Research in Security Prices (CRSP). The t-statistics are adjusted for generalized heteroskedasticity and one lag of moving average residuals using the technique of Newey and West (1987).

Table 1 illustrates that excess returns to long-term bonds are, on average, positive but statistically indistinguishable from zero, and positively associated with the slope. The statistical evidence for this association is strong; three of the four estimated coefficients are statistically different from zero at the 5% level.[2] A 100 basis point increase in the slope

[1] The statistical evidence in Chen may be overstated because the reported t-statistics appear to be based on OLS standard errors. I reestimated the regressions using heteroskedasticity-consistent standard errors and found statistically weak evidence of predictability.

[2] This forecast power is somewhat sensitive to the choice of maturity of the long bond used in constructing the slope. The evidence of forecastability is stronger when a five-year maturity is used and weaker when a thirty-year maturity is used. In the latter case, only one of the four reported regressions has a slope coefficient that is significant at the 5% level. The results also change when a one-month bill yield is used to measure the short end of the term structure, but there is no systematic pattern to the changes across the forecast horizons. If the short end of the slope is measured by a two-year yield, much of the forecast power of the slope disappears.

corresponds to an additional 2.2 percent excess return over the next four quarters. When these data are split into equal-sized two samples based on the magnitude of the quarter-t slope, the mean excess bond return in the "high-slope" sample is 0.68 percent per quarter. The mean excess bond return in the "low-slope" sample is -0.22 percent per quarter. (These results are not reported in the table.)

Excess stock returns are, on average, strongly positive. Their link with the slope is statistically weaker than for bonds.[3] A 100 basis point increase in the slope corresponds to an additional 3.5 percent excess return over the next four quarters. Although this evidence for stock return predictability might appear inconclusive, for our purposes a weak positive relation will turn out to be just as much of a puzzle as a strong positive relation; the real puzzle will be why the relation is not negative.

These regressions are subject to the "predictive regressions" finite-sample bias discussed in Stambaugh (1999). The coefficients on the slope are biased downwards because the contemporaneous correlations between the slope and excess returns are positive: 0.30 for bonds and 0.13 for stocks. Thus the statistical evidence in Table 1 understates the case for return predictability. In the case of excess returns to long-term bonds, this understatement is equivalent to the finite-sample bias in tests of the expectations hypothesis described in Bekaert, Hodrick, and Marshall (1997).

The slope of the term structure also forecasts economic growth. Chen (1991) and Estrella and Hardouvelis (1991) regress growth rates of real output and consumption on the slope of the yield curve and find that the slope is positively associated with real growth over the next one to two years.[4] Visual evidence is in Figure 1, which plots both the slope of the term structure and the log change in aggregate consumption. Consumption is measured by real expenditures on nondurables and services and is expressed per capita. More details are in the Appendix. The quarter-t slope is measured at quarter-end. The quarter-t change in consumption is measured by log consumption during quarter $t+1$ less log consumption during quarter t. This flow variable can be thought of as contemporaneous with the point-measured slope SL_t if all consumption in quarter t is determined at the beginning of the quarter. Campbell (1999) discusses consumption timing conventions in more detail.

We can see from the figure that the recessions of the mid-1970s, the early 1980s, and the early 1990s were all preceded by declines in the slope of the term structure. More formal

[3]The strength of the predictive relation is stronger if the long end of the slope is measured by either a Aaa-bond yield or a thirty-year Treasury yield. For both of these alternative measures, the slope coefficient in the one-quarter-ahead forecast equation has a t-statistic greater than 2.6. Replacing the three-month yield with a one-month yield has no qualitative effect, while replacing it with a two-year yield reduces the forecast power of the slope.

[4]Related evidence is in Harvey (1989).

evidence comes from regressions of log changes in consumption or GDP in quarter t on the slope of the term structure at the end of quarter $t-i$. The log change in consumption (GDP) from quarter t to quarter $t+1$ is denoted Δc_t (Δgdp_t).

$$\Delta x_t = b_{x,i,0} + b_{x,i,1} SL_{t-i} + e_{x,t,i}, \ x \in \{c, gdp\}, \ i = -8, \ldots, 8$$

The sample period is $t \in \{1952{:}1,\ 2000{:}3\}$ for $i \geq -1$. For $i \leq -2$, the sample period is shorter because the last observation of the term structure slope is 2000:4. Figure 2 plots estimated coefficients for each of the 34 regressions (17 each for consumption and GDP, from eight quarterly lags to eight quarterly leads). Also plotted are two-standard-error confidence bounds. The Newey-West asymptotic standard errors are adjusted for one lag of moving average residuals.

Figure 2 displays the standard results that a higher slope at the end of quarter t corresponds to higher consumption and output growth, both contemporaneously and in the future. There is also evidence that consumption and GDP growth slightly lead the slope of the term structure.[5] We now consider how this empirical evidence is interpreted in the context of consumption-based asset-pricing models.

3 Consumption-based interpretations of the evidence

Standard finance theory tells us that in a discrete-time framework, the gross nominal return to any asset i, denoted $1 + R_{i,t}$, satisfies

$$1 = E_t[(1 + R_{i,t+1})M_{t+1}]. \tag{2}$$

The random variable M_{t+1} is the stochastic discount factor. Equation (2) is a requirement of no arbitrage. In a utility-based framework, we can think of M_{t+1} as the ratio of the marginal utility of a dollar at time $t+1$ to the marginal utility of a dollar at time t. Denote $r_{i,t} \equiv \log(1 + R_{i,t})$, $m_{i,t} \equiv \log(M_{i,t})$, and the riskless nominal rate as r^f_{t+1}. If we assume conditional joint log-normality of $1 + R_{i,t+1}$ and M_{t+1}, (2) implies

$$E_t(r_{i,t+1}) - r^f_{t+1} + (1/2)Var_t(r_{i,t+1}) = -Cov_t(r_{i,t+1}, m_{t+1}). \tag{3}$$

[5] These results are not sensitive (at least qualitatively) to the slope measure. If the ten-year Treasury yield is replaced with either a thirty-year Treasury yield or a Aaa yield, the forecast power of consumption growth for future slopes increases, while the forecast power of the slope for future consumption growth decreases. This pattern is reversed if the ten-year yield is replaced with a five-year yield. On balance, it appears that lagged consumption growth is more closely tied to the very long end of the term structure, while future consumption growth is more closely tied to an intermediate range–maturities of five to ten years.

The left-hand-side of (3) is the expected excess return to the asset; the variance term adjusts for Jensen's inequality created by using log returns instead of returns. Asset-pricing models put additional content on the right-hand-side by restricting m_{t+1}. To understand why the slope of the term structure forecasts expected excess asset returns, we must understand why the covariance between returns and the stochastic discount factor varies systematically with the slope.

3.1 Power and recursive utility

Consumption-based asset-pricing models write m_{t+1} as a function of consumption. For example, time-separable power utility implies

$$m_{t+1} = \log(\delta) - \gamma \Delta c_{t+1}$$

where δ is the rate of time preference, γ is the coefficient of relative risk aversion, and Δc_{t+1} is the change in log consumption. This simple framework is nested in recursive utility, developed by Epstein and Zin (1989) and Weil (1989). Then m_{t+1} can be expressed as

$$m_{t+1} = \theta[\log(\delta) - (1/\Lambda)\Delta c_{t+1}] - (1-\theta) r_{w,t+1} \tag{4}$$

where the elasticity of intertemporal substitution is Λ, $\theta \equiv (1-\gamma)/(1-(1/\Lambda))$, and the log return to total tradeable wealth, including human capital, is $r_{w,t+1}$. Substituting (4) into (3) yields

$$E_t(r_{i,t+1}) - r^f_{t+1} + (1/2)Var_t(r_{i,t+1}) = (\theta/\Lambda)Cov_t(r_{i,t+1}, \Delta c_{t+1}) + (1-\theta)Cov_t(r_{i,t+1}, r_{w,t+1}) \tag{5}$$

An important limitation on our ability to test (5) is poor data. The first problem is that consumption is poorly measured. We have little individual-level data on consumption. In a representative agent model, individual consumption equals per capita aggregate consumption, but aggregate consumption is also measured with substantial noise. Second, consumption is measured infrequently (at best, monthly). This makes it difficult to observe whether there is any time-variation in $Cov_t(r_{i,t+1}, \Delta c_{t+1})$. Third, total wealth, which includes both financial and human wealth, is poorly measured. Financial wealth is measured reasonably well, but human wealth is unobserved. As noted by Roll (1977), even a slight mismeasurement in wealth can substantially distort tests of equations such as (5).

These data problems make it difficult to interpret the empirical evidence that (5) cannot explain the observed time-variation in expected excess returns to stocks and bonds. There is substantial evidence of time-variation in the volatility of returns to financial assets, but this variation appears unrelated to variation in expected excess returns.[6] There is no evidence of substantial time-variation in either the volatility of consumption growth or in covariances of asset returns with consumption growth.[7] We could chalk up these results to poor data, but an alternative explanation is that (4) does not correctly describe preferences.

3.2 Habit formation

A plausible interpretation of the slope's forecast power for expected excess asset returns is that it is driven by the effects that recessions have on investors' attitudes toward risk. The slope of the term structure is one of a number of financial market measures that predicts both asset returns and future macroeconomic growth. If investors are more risk-averse in recessions, variables that are correlated with business cycles will also be correlated with expected excess returns. This interpretation is formalized and defended by Campbell and Cochrane (1999) in a habit formation model. Although I follow their intuition, it is more convenient here to use the model of Wachter (2001), which extends Campbell and Cochrane's work by introducing predictable consumption growth. Details and references are suppressed here and can be found in the original papers.

Aggregate log consumption growth follows the process

$$\Delta c_{t+1} = z_t + v_{t+1},$$

$$z_{t+1} = (1-\psi)g + \psi z_t + \mu_{t+1},$$

where v_t and μ_t are jointly normally distributed. The parameter $\psi > 0$ creates persistence in consumption growth, which is consistent with U.S. data. The representative agent maximizes

$$E_t \sum_{i=t+1}^{\infty} \delta^{i-t} \frac{(C_i - X_i)^{1-\gamma} - 1}{1-\gamma}$$

where X_t is the agent's time-t habit. (Capital letters represent levels and lower case letters represent logs.) Power utility is a special case of habit formation with $X_t \equiv 0$. The surplus consumption ratio is defined as

$$S_t = \frac{C_t - X_t}{C_t}.$$

[6]See Campbell (1999) and Whitelaw (2000) for discussions and references.
[7]See, e.g., Boudoukh (1993).

The log stochastic discount factor is therefore

$$m_{t+1} = \log(\delta) - \gamma(\Delta s_{t+1} + \Delta c_{t+1}).$$

To capture the notion that risk aversion increases in recessions, Wachter follows Campbell and Cochrane in writing the dynamics of surplus consumption as

$$s_{t+1} = (1-\phi)\overline{s} + \phi s_t + \lambda(s_t)v_{t+1} \tag{6}$$

where $\lambda(s_t)$ is a decreasing function. Thus when past shocks to consumption growth have been negative, log surplus s_t is both low and volatile. Expected excess asset returns satisfy

$$E_t(r_{i,t+1}) - r^f_{t+1} + (1/2)Var_t(r_{i,t+1}) = \gamma(1+\lambda(s_t))Cov_t(r_{i,t+1}, \Delta c_{t+1}). \tag{7}$$

Variations in expected excess returns are produced by both variations in the covariance of the asset's return with consumption and variations in surplus consumption. With power utility, only the former channel operates. Wachter shows that around the steady state, deviations of log surplus from its mean are approximately[8]

$$s_t - \overline{s} \approx \ln(1-\overline{S}) + \frac{g}{1-\phi} + (1-\phi)\sum_{j=0}^{\infty}\phi^j \Delta c_{t-j}. \tag{8}$$

Although this model generates nontrivial joint dynamics between expected excess returns and the real term structure, it does not include inflation and thus says little about the nominal term structure. However, inflation is easily introduced as long as we are willing to treat the relation between consumption and inflation as exogenous.[9] For example, we can follow Pennacchi (1991) and Boudoukh (1993) by assuming a vector autoregression joint process for consumption growth and inflation. More generally, we could write the log growth in the price level as

$$\Delta p_t = f(v_t, \mu_{t-1}, \ldots, \mu_t, \mu_{t-1}, \ldots, \epsilon_t, \epsilon_{t-1}, \ldots)$$

where the ϵ's are shocks to inflation that are independent of consumption growth. The important restriction that allows us to patch an inflation process on to the model is that the

[8]For tractability, she does not allow shocks to z_t enter into (6), thus (8) is perhaps better interpreted as a weighted average of past changes in log consumption around the stochastic means of these changes. Of course, when z_t is at its mean, there is no difference between these interpretations.

[9]An endogenous relation, using a cash-in-advance constraint, is proposed for term-structure models by Backus, Gregory, and Zin (1989) and implemented by Labadie (1994).

shocks to inflation do not affect investors' habit.[10] For the purposes of the current paper, specifying a particular inflation process and solving for bond prices will take us too far afield. It will be sufficient to note that the predictive power of the slope for expected excess returns is consistent with a model in which inflation is temporarily low in recessions. Longer-term bond yields reflect expected future increases in inflation, thus the term structure slopes up in bad economic times, when expected excess returns are high.

3.3 Conceptual difficulties with habit formation

This model has two channels through which the slope can forecast expected excess asset returns. Either the slope is associated with covariances between consumption growth and asset returns or it is associated with changes in surplus consumption. The conventional wisdom is that there is no evidence for the first channel, thus the second channel must be the important one. The main goal of this paper is to revisit the relation between the slope and covariances, but it is also important to explore why the second channel may not be consistent with the data.

Assume, for now, that the only reason for the slope's forecast power is that the slope is associated with surplus consumption. Then past consumption growth should capture part of the forecast power of the slope. There is no doubt that consumption data are measured poorly, especially in contrast to the measurements of financial instrument prices. This noise dampens the forecasting power of measured consumption for expected asset returns–the empirical counterpart to (8) is a noisy measure of true surplus consumption. But the term structure is also a noisy measure of surplus consumption. The term structure reflects information about past consumption growth ($v_{t-i}, i > 0$), expected future consumption growth (z_t), and also has a component that varies independently of the process driving consumption ($e_{t-i}, i > 0$). Because the noise in measured consumption is probably unrelated to the extraneous information in the slope, both lagged consumption growth and the slope should help forecast future excess returns.

To investigate this issue, we require an empirical proxy for surplus consumption. I follow Wachter and truncate the infinite sum in (8) at 40 quarters, with the decay factor $\phi = 0.969$. The proxy is

$$\hat{s}_t = \frac{1-\phi}{1-\phi^{40}} \sum_{j=0}^{39} \phi^j \Delta c_{t-j}. \tag{9}$$

I then estimate a modification of (1):

[10] Brandt and Wang (1999) allow inflation to affect investors' habit.

$$r_t^k - r_t^f = b_0 + b_1 SL_{t-i} + b_2 \hat{s}_{t-i-1} + e_t^k,\ k \in \{s, b\},\ i = 1, \ldots, 4. \tag{10}$$

I use lagged surplus, $\hat{s}_{t-i-1}$, instead of $\hat{s}_{t-i}$ because of a potential problem introduced by consumption timing. Recall that Δc_t is defined as the change in consumption from t to $t+1$. Thus $\hat{s}_{t-i}$ depends, in part, on consumption during quarter $t-i+1$. As long as consumption in a quarter is determined at the beginning of the quarter, $\hat{s}_{t-1}$ will be uncorrelated with the error term e_t^k in (10). But to prevent any possible spurious correlation, I use lagged surplus.[11] The empirical results are displayed in Table 2. Because ten years of consumption data are needed to construct the measure of surplus, the sample period is 1957:2 through 2000:4. Newey-West asymptotic t-statistics are adjusted for one lag of moving average residuals.

At first glance, the results seem to provide partial support for this habit formation model. The measure of surplus is a strong predictor of excess stock returns. In three of the four stock-return regressions, surplus is statistically significant at the 5% level (and is significant at the 10% level in the fourth regression). Including surplus in the regression eliminates the explanatory power of the slope. By contrast, surplus has very limited predictive power for excess bond returns. None of the coefficients on surplus is statistically different from zero at conventional levels, while the slope remains a strong predictor. However, in combination these two sets of results are damning evidence against the model. The model is consistent with the bond-return regressions only if the noise in measured consumption growth drowns out the explanatory power of true surplus consumption. But this is contradicted by the predictive power of surplus consumption for stock returns.

The results for stock returns are somewhat sensitive to the rate of decay used in (9). The value $\phi = 0.969$ implies a half-life of about 22 quarters. I recalculated surplus assuming a decay of $\phi = 0.93303$, which weights more heavily relatively recent growth in consumption. (This value produces a half-life of 10 quarters.) I then reestimated the regressions with this measure. The results are not reported in any table. Of the four stock-return regressions, only one has a coefficient on surplus that is significant at the 5% level. A look at Figure 2 explains why the results are weaker when recent consumption growth is emphasized. Increases in consumption growth lead the slope of the term structure by a couple of quarters. In other words, when the slope is steep in quarter t, consumption tends to be high relative to where it was a few quarters ago. Thus a consumption habit based on very recent consumption would forecast lower excess asset returns, not higher returns.

Another problem with the view that surplus drives variations in expected asset returns is that it is inconsistent with variations over time in the sign of expected excess returns to

[11] The empirical results are not sensitive to this choice.

bonds. To generate sign changes without changes in the sign of $Cov_t(r^b_{t+1}, \Delta c_{t+1})$, the sign of $(1 + \lambda(s_t))$ must change. This means that consumers switch from avoiding consumption gambles to seeking consumption gambles. This is both intuitively implausible and hard to reconcile with the fact that expected excess returns to other assets, such as stocks, do not appear to exhibit sign changes.

One way to interpret the sign-switching in expected excess bond returns is that it is an artifact of the data, and not a robust feature of bond markets. When these quarterly data are split into two samples based on the magnitude of the quarter-t slope, the mean quarter $t + 1$ excess bond return in the low-slope sample is negative, but not statistically different from zero. Yet there is no *a priori* reason to reject the hypothesis that the sign of $Cov_t(r^b_{t+1}, \Delta c_{t+1})$ depends on the slope of the term structure. If, say, the central bank accomodates macroeconomic shocks in bad times, but does not in good times, this time-varying reaction function should show up in the sign of $Cov_t(r^b_{t+1}, \Delta c_{t+1})$.

This habit formation model can be reconciled with the behavior of both stock and bond returns if covariances between consumption growth and bond returns change sign with the term structure. This would explain why surplus consumption forecasts stock returns, but not bond returns; the effect of changing covariances on bond risk premia can offset the effect of changing surplus on these risk premia. However, as noted in Section 3.1, the existing literature finds no evidence for the link between consumption covariances and risk premia.

This lack of evidence indicates that either our models of representative-agent, consumption-based asset prices are wrong, consumption data is too poor to allow us to observe time-varying conditional second moments, or researchers simply have not used the right conditioning information. Existing research has not examined the predictive power of the term-structure slope for conditional variances and covariances involving consumption growth. Given the close link between the term structure and predictable variations in future aggregate output, this is a surprising gap. In the next section, I take a detailed look at the ability of the slope to forecast standard deviations, correlations, and covariances involving consumption growth.

4 The empirical evidence

The main goal of this section is to understand the empirical relation between the slope of the term structure and subsequent covariances between consumption growth and asset returns. It will be illuminating to decompose covariances into standard deviations and correlations. Therefore the first issue addressed here is the predictive power of the term structure slope for subsequent volatilities of consumption growth and asset returns.

4.1 Split sample results

Perhaps the most intuitive way to see the predictive power of the slope for the volatility of consumption growth is to split the data sample in two, based on whether the slope of the term structure at the end of quarter t is greater or less than its median. Then compute sample standard deviations ($sdev$) of Δc_{t+i} for each subsample:

$$sdev(\Delta c_{t+i}|SL_t > median(SL_t)),\ sdev(\Delta c_{t+i}|SL_t < median(SL_t)), i > 0. \tag{11}$$

It is worth repeating two definitions from Section 2: SL_t is the (demeaned) slope at the end of quarter t and Δc_{t+i} is the change in log consumption from quarter $t+i$ to quarter $t+i+1$. (This log change is also multiplied by 100%.) Therefore Δc_{t+1} does not depend on consumption realized before the observation of the term structure.

Table 3 reports these conditional standard deviations for $i = 1, 2, 3$. Three sample periods are considered. The first is the full sample $t \in \{1952{:}1\text{-}2000{:}3\}$. (The ending date is earlier for consumption data and, for $i > 1$, for asset return data. The last observation of consumption growth is 2000:3 and the last observation of asset returns is 2000:4.) Because observations of both the slope and consumption growth are most volatile during the Fed monetarist experiment during 1979 through 1982, I also report results for pre-experiment (1952:1 through 1977:4) and post-experiment (1983:2 through 2000:3) periods.

The results demonstrate that the volatility of consumption growth is substantially higher when the slope of the term structure is flatter than usual. In the full sample, the ratios of low-slope standard deviations to high-slope standard deviations range from 1.44 to 1.50. F-tests overwhelmingly reject the hypothesis of constant variances. These tests should be interpreted with caution because the assumptions of normality and independence are unlikely to be appropriate. Alternative statistical tests are discussed in Section 4.2. This volatility pattern holds across the smaller samples. In the non-experiment subperiods, the ratios range from 1.11 to 1.49. The subperiod results should not be emphasized too heavily because there are relatively few observations. In the early subperiod, there are 52 observations in each of the two slope-dependent samples. In the latter subperiod, this number falls to 34.

This evidence suggests that consumption growth is far from *i.i.d.* Instead, a steep term structure slope forecasts high-mean, low-volatility consumption growth. Although in this paper our primary focus is on covariances of returns with consumption growth, these results on volatility have implications for general equilibrium models of asset pricing. A common approach to exploring issues such as the equity premium puzzle is to model the stock market as a claim on the consumption process (e.g., Whitelaw (2000)). Then the variance of consumption growth is used to measure the risk exposure associated with stocks. The evidence here indicates that this risk exposure is higher when the term structure is flatter. In power utility or recursive utility setups, this implies higher expected excess stock returns when the term structure is flatter, which is counterfactual.

A natural question is whether this pattern in consumption volatility carries over to the volatility of asset returns. The remainder of Table 3 reports versions of (11) for excess returns to the aggregate stock market (denoted er_t^s) and long-term bonds (er_t^b). The returns are expressed in percent. The evidence indicates that the predictive power of the slope for asset return volatility is economically weaker than that for consumption growth volatility. In the full sample, ratios of low-slope standard deviations to high-slope standard deviations are about 1.20 for stock returns and 1.08 for bond returns. The F-tests typically cannot reject the hypothesis of constant variances. This might seem surprising, given the well-known heteroskedasticity in asset returns, but the tests here have very low power compared with the usual tests that rely on high-frequency data.

We now turn to correlations. Again, we split the data sample based on the slope of the term structure at the end of quarter t and compute correlation matrices for each subsample. The variables of interest are Δc_{t+1}, er_{t+1}^s, er_{t+1}^b and the log change in real per capita labor income, Δli_{t+1}. The latter variable is included because, as argued by Lettau and Ludvigson (2001), it is a proxy for the return to human capital. The definition of labor income follows Lettau and Ludvigson.

Table 4 reports these conditional correlation matrices for the three sample periods examined in Table 3. Three features of this table are striking. First, the correlation between the stock market and consumption growth strongly depends on the slope. When the slope is less steep than usual, the contemporaneous correlation between future stock returns and consumption growth is around 0.4 in all sample periods. But when the slope is steeper than usual, the correlation is roughly zero. In the full sample the correlation is 0.01 and ranges from 0.09 to −0.21 in the subperiods. I test whether the correlation is constant across the two slope-sorted groups using Fisher's z. (These statistical tests are not reported in the table.) Over the entire sample period, we can reject at the one percent level the hypothesis that the correlation is constant. tatistical rejections are weaker in the pre-experiment and post-experiment periods, with significance levels of seven percent and two percent, respectively.

Second, the correlation between the stock market and labor income growth also strongly depends on the slope. When the slope is less steep, the correlation between future stock returns and future labor income is 0.27 in the full sample. When the slope is steeper, the correlation is −0.08. We can reject the hypothesis of constant correlations over the entire sample period at the two percent significance level. Statistical significance in the subperiods is weaker.

Third, there is no clear relation between the slope and the magnitude of the correlation between long-term bond returns and consumption growth. When the slope is less steep than usual, the correlation is close to zero and of uncertain sign, ranging from 0.08 to −0.20 in the samples. When the slope is steeper than usual, the correlations remain small, ranging from 0.21 to 0.05. Although none are negative, suggesting that correlations are higher when the slope is steeper, in the full sample the correlation is 0.08 regardless of the slope.

All three observations are bad news for consumption-based, representative-agent asset-pricing models. Power utility and the habit formation model described here have difficulty with the first observation. These models, combined with the first observation, imply that expected excess stock returns should be roughly zero when the slope is steeper than usual. From (7), stocks with consumption betas of zero should earn no risk premium, regardless of the level of surplus consumption.

The second observation suggests that with recursive utility, expected excess stock returns should also be lower when the slope is steep. From (3), expected excess returns vary with both covariances with consumption and covariances with the return to total wealth, including human capital. Labor income is an important component to the return to human capital. Therefore a lower correlation between the stock market and labor income suggests a lower correlation between the stock market and total wealth. Hence with recursive utility, both sources of time-variation in expected excess returns lead to a counterfactual conclusion. The third observation indicates that neither power utility nor habit formation can explain the sign-switching in expected excess returns to bonds. The relation between excess returns to long-term bonds and consumption growth is weak regardless of the state of the term structure.

A graphical look at the relation between stock returns and consumption growth will help summarize some of the empirical evidence. Figure 3 displays two scatter plots of er^s_{t+1} and Δc_{t+1}, based on the slope of the term structure at the end of quarter t. In the upper scatter plot, which corresponds to steeper slopes, mean consumption growth is higher than in the lower scatter plot, (compare the vertical solid lines), mean excess stock returns are higher (compare the horizontal solid lines), and both consumption growth and stock returns are less volatile (compare the dispersion of the clouds). In the upper scatter plot, there is no observable correlation, while in the lower plot the two variables are clearly positively correlated. For example, in the bottom panel, there are 11 quarters in which the stock market fell by at least 10 percent; contemporaneous consumption growth (again, measured by the change from quarter $t + 1$ to quarter $t + 2$) was below its sample mean in 10 of the 11 quarters.

4.2 Regression results

In this subsection we take a closer look at the joint behavior of consumption growth and asset returns. The evidence reported above concerned raw growth rates and excess returns. In this section I take a closer look at this joint behavior by focusing on (estimated) innovations in consumption growth and asset returns. In addition, I make some inroads in understanding *why* the slope is associated with both consumption volatilities and correlations between consumption and asset returns. As discussed above, the term structure slope is negatively associated with past consumption growth, and positively associated with future consumption growth. Thus we can ask whether the slope is simply proxying for either past or expected future consumption.

Following the two-step procedure used by Schwert (1989) and adopted by many others, I first construct innovations to the time series. Innovations in log consumption growth are produced by regressing growth in quarter $t+1$ on variables known at the end of quarter t. They are the slope of the term structure, the excess stock return in quarter t, inflation (the log change in the price level from quarter $t-1$ to quarter t), and two lags of both consumption growth and labor income growth.

$$\Delta c_{t+1} = b_0 + b_1 SL_t + b_2 \Delta c_{t-1} + b_3 \Delta c_{t-2} + b_4 \Delta li_{t-1} + b_5 \Delta li_{t-2} + b_6 infl_t + b_7 r_t^s + \tilde{c}_{t+1} \quad (12)$$

The real variables are twice-lagged to avoid overlapping with the dependent variable. I use the same regression to construct quarter-t forecasts of Δc_{t+1}, denoted $\widehat{\Delta c_{t+1}}$. In the second stage, we take a closer look at forecasts of consumption volatility. I use the quarter-t term structure slope and the quarter-t one-quarter-ahead forecast of consumption growth to predict the volatility of innovations in consumption growth. I use both squared residuals, and following Schwert (1989), absolute residuals in this stage. The use of absolute residuals is more robust to outliers.

$$|\tilde{c}_{t+1}| = c_0 + c_1 SL_t + c_2 \widehat{\Delta c_{t+1}} + \eta_{t+1} \quad (13)$$

$$(\tilde{c}_{t+1})^2 = c_0 + c_1 SL_t + c_2 \widehat{\Delta c_{t+1}} + \zeta_{t+1} \quad (14)$$

Missing from (13) and (14) is a measure of lagged consumption growth. In regressions not detailed here, I found that the measure of surplus consumption (9) had no explanatory power for the volatility of consumption growth, thus here I focus only on the slope and forecasted future consumption growth. The results of this second stage are reported in Table 5. Newey-West asymptotic t-statistics are adjusted for one lag of moving average residuals.

The message of Table 5 is that the relation between the slope and consumption volatility appears to proxy for a relation between forecasted consumption growth and consumption volatility. Over the full sample, high quarter-t forecasts of future consumption growth correspond to low consumption volatility in quarter $t+1$. The relation is strongly statistically significant for both absolute residuals and squared residuals. Because the slope tends to be steeper than usual in times of high forecasted consumption growth, the slope also forecasts volatility. However, the forecast power of the slope for volatility is subsumed by that of forecasted consumption growth. The table also reports subperiod results which give general support to this view. In both 1952–1977 and 1983–2000, forecasted consumption growth is more closely tied to subsequent consumption volatility than is the slope.[12]

The same basic approach is used to forecast covariances between innovations in consumption growth and asset returns. I construct innovations to excess stock returns and bond returns in quarter $t+1$ by regressing them on the the lagged slope of the term structure. The constructed innovations are denoted $\tilde{er}^s_{t+1}$ and $\tilde{er}^b_{t+1}$. I then estimate (15):

$$(\tilde{c}_{t+1})(\tilde{er}^k_{t+1}) = d_0 + d_1 SL_t + d_2 \widehat{\Delta c_{t+1}} + \eta_{k,t+1}, \;\; k \in \{s, b\} \tag{15}$$

Table 6 reports the results from (15). As in Table 5, three samples are examined: The full 1952:1 through 2000:4 period and the pre-experiment (1952:1–1977:4) and post-experiment (1983:1–2000:4) periods. Newey-West asymptotic t statistics are adjusted for one lag of moving average residuals.

The results are not surprising, given what we have seen in Tables 3 through 5. The quarter $t+1$ covariance between consumption growth and stock returns is closely associated with the quarter t term structure slope. In the full sample, we can reject the hypothesis of no relation at the 1% significance level. As in Table 5, this relation appears to proxy for an underlying relation between the covariance and expected consumption growth. When expected consumption growth is high, the covariance between consumption and stock returns is low. When both the slope and expected consumption growth are allowed to forecast this covariance, only the coefficient on expected consumption growth is statistically significant. Evidence from the smaller samples is less conclusive (in particular, in the latter sample the explanatory power of the slope slightly exceeds that of expected consumption growth), but the sample sizes are small.[13]

The evidence we saw in Table 5 is that consumption growth and stock returns are unrelated when the slope of the term structure is less than usual. The evidence in Table 6 does

[12] To produce the subperiod results, (12) was reestimated on the subperiod's data, then (13) and (14) were estimated on the subperiod-specific residuals.

[13] See footnote 12.

not directly address this result. We can, however, modify regression (15) by replacing the explanatory variables, including the constant term, with two dummy variables. The first dummy variable equals one when the slope is greater than its median, and the second equals one when the slope is less than its median. The results of this regression (which are not reported in any table) confirm the earlier evidence. When the slope exceeds its median, the covariance between consumption growth and stock returns is statistically indistinguishable from zero at any conventional significance level. When the slope is less than its median, the covariance differs from zero at the 0.0005 level. This result also holds when the slope dummies are replaced with dummies based on whether the expected growth in consumption is greater or less than its median.

The results for bond-return covariances are dramatically different. Over the period 1952 through 2000, neither the slope nor expected consumption growth are associated with the covariance between consumption growth and bond returns. The same (non) pattern holds in the pre-experiment subperiod. One puzzle is that in the post-experiment subperiod, the slope (but not expected consumption growth) is strongly *positively* associated with this covariance. Since this relation is not robust across the time periods, and there are relatively few observations in the post-experiment period, the true significance of this relation may be overstated by the statistical tests. Therefore I do not pursue this anomalous result further.

Recall from Table 5 that correlations between bond returns and consumption growth are small. Similarly, the sample mean of $(\tilde{c}_{t+1})(\tilde{er}^b_{t+1})$ is statistically indistinguishable from zero. (This is not reported in any table.) In the absence of the results from Tables 5 and 6, we might choose to chalk this result up to bad consumption data. Bad consumption data would also explain why the covariance between consumption growth and bond returns does not seem to vary with the slope of the term structure (aside from the short post-experiment period). But the stock-return results in these tables document that there is enough information in the consumption data to infer strong patterns in the covariance between consumption growth and stock returns. Thus the lack of a pattern in covariances involving bond returns likely reflects the true economic relation between aggregate consumption and the bond market.

4.3 Interpretations

The evidence reported in sections 2, 4.1 and 4.2 is difficult to incorporate into the standard representative-agent, consumption-based theories of asset pricing. This begs the question—what theories are consistent with this evidence? One possibility is to stay within the consumption-based framework but reject the hypothesis that a representative agent consumes per capita U.S. consumption. The representative agent may consume per capita

OECD consumption (i.e., markets are truly global), or there may be no representative agent. Either direction makes it easier to explain away the fact that the stock market earns high excess returns at times when its U.S. aggregate consumption beta is zero.

What is not so easy to explain away is the forecast power of the slope for expected excess returns. In the heterogeneous agent model of Constantinides and Duffie (1996), time-variation in the variance of household-idiosyncratic shocks creates time-variation in risk premia. But why should this variance increase at times when the slope is steep? One confounding piece of evidence documented above is that the variance of aggregate consumption is low at such times, which means that the variance of idiosyncratic shocks will have to move opposite the variance of aggregate shocks. In addition, it is hard to come up with a story linking the slope to the variance of idiosyncratic shocks that does not have some link to recessions. But then the story will have trouble explaining the results of Table 2. Surplus consumption, which we can think of as a recession variable, does not capture the forecast power of the slope for excess bond returns.

The same trouble befalls other models that attempt to provide a unified explanation of the slope's forecast power for stock and bond excess returns. Consider, for example, the approach in Brandt and Wang (1999). Essentially, they assert that investors' risk premia are affected by inflation; inflation exogenously raises the consumption habit (X_t) of a representative consumer, making investors more risk averse. Because the slope contains information about inflation, the slope will forecast excess returns. But again, we are then faced with the puzzle of why surplus consumption captures the forecast power of the slope for excess stock returns but not for bond returns. It is worth noting that almost any source of time-variation in bond risk premia implies that the slope forecasts excess bond returns. The reasoning follows Berk (1995). If a shock leads investors to require higher expected excess returns to bonds, long-maturity bond prices will fall (their yields will rise), and the term structure slope will rise accordingly.

The results discussed so far raise questions, not just about risk premia, but also about the relation between asset returns and business cycles. I investigate some of these questions next.

4.4 Predicting business cycles with asset returns

Asset returns forecast business cycles. Increases in stock prices correspond to contemporaneous and future increases in aggregate consumption, output, and income.[14] Returns to long-term bonds also forecast business cycles, although this forecasting power is typically

[14]See, e.g., Fischer and Merton (1984) or Barro (1989).

expressed in terms of changes in interest rates.[15] The evidence in Tables 4 and 6 suggests that the strength of these forecasts may depend on the slope of the term structure. We have seen that the relation between the excess stock return in quarter t and consumption growth from quarter t to quarter $t+1$ depends on the slope of the term structure. Is the same dependence to be found in more distant growth rates and with other macroeconomic variables? I explore this question here.

My approach is to slightly modify linear forecasts of macroeconomic growth rates by allowing the parameters to depend on the term structure slope. Define D_SLOPE_t as a dummy variable that equals one if the slope of the term structure is greater than its sample median at the end of quarter t. I estimate regressions of the following form:

$$\Delta x_{t+i} = b_0 + b_1 D\_SLOPE_t + (b_2 + b_3 D\_SLOPE_t) er^k_{t+1} + e_{k,t+i}, \; k = \{s, b\}, \qquad (16)$$

where Δx_{t+i} is the change in the log of aggregate consumption, GDP, or labor income from quarter $t+i$ to quarter $t+i+1$. (All variables are aggregate, real, per capita.) The results of estimating (16) from 1952:1 through 2000:3 are displayed in Table 7. The coefficient b_1 reflects the direct forecasting power of the quarter-t slope for future macroeconomic growth. The coefficient b_2 reflects the forecasting power of stock (er^s_{t+1}) or bond (er^b_{t+1}) return when the quarter-t slope is less steep than usual, while the sum b_2+b_3 reflects the forecasting power when the slope is steeper than usual. Newey-West asymptotic t-statistics are adjusted for one lag of moving average residuals.

There are three main points to take from Table 7. First, when the term structure slope at the end of quarter t is flatter than usual, the stock market excess return in quarter $t+1$ is a strong forecaster of the growth of consumption, GDP, and labor income from both $t+1$ to $t+2$ and $t+2$ to $t+3$. Each of these six estimates of b_2 is significantly greater than zero at the five percent level. The forecasting power dies off for growth from $t+3$ to $t+4$.

Second, the forecasting power of the stock market significantly differs when the term structure slope is steeper than usual. For five of the six regressions for which b_2 is significant, b_3 is negative and significantly different from zero. The sum $b_2 + b_3$ is statistically indistinguishable from zero in each of the regressions. (This is not reported in the table.) In other words, when the slope is steeper than usual, the forecasting ability of the stock market is nonexistent.

Third, the relation between bond returns and macroeconomic growth roughly corresponds to that between stock returns and macroeconomic growth. When the slope is flatter than usual, bond excess returns forecast future growth in consumption, output, and income; the

[15] See Estrella and Mishkin (1998) for some recent evidence.

growth is concentrated from quarter $t + 2$ to quarter $t + 3$. When the slope is steeper than usual, this forecasting power disappears.

These results are puzzling. Their main message is that the relation between asset returns and the business cycle is predictably strong or weak, depending on the slope of the term structure. More generally, these results suggest that the kinds of shocks that hit the economy when the term structure is steep are fundamentally different from the kinds of shocks that hit when the term structure is less steep. (When the dummy for the magnitude of the slope of the term structure is replaced with a dummy variable for the magnitude of next quarter's forecasted consumption growth, the results are qualitatively unchanged.) In this paper, I make no attempt to explore the reasons why these qualitative differences exist. Instead, I simply note that the results imply linear regressions (including vector autoregressions) that link asset returns with business cycles are misspecified. The usual joke about the stock market is that it has forecasted $x + y$ of the past x recessions, $y > 0$.[16] The results here suggest that the y recessions that did not materialize may have been forecasted when the slope of the term structure was steep.

4.5 What drives stock and bond prices?

Stock prices should respond to news about expected future cash flows and discount rates, while bond prices should respond to news about expected future interest rates. Realizations of cash flows, discount rates, and future interest rates are noisy measures of these expectations. Thus a standard test of models of asset-price determination is to regress asset returns on future realizations of relevant variables.[17] I follow the same approach here in examining whether the explanatory power of future variables for current asset returns depends on the slope of the yield curve.

I start with stock returns and follow Kothari and Shanken (1992) by using dividend growth, investment growth, and future stock returns as explanatory variables. Because of the persistence of dividends, news about today's dividends corresponds to changes in the current stock price. News today about future investment opportunities also shows up in today's stock price.[18] Kothari and Shanken include future stock returns in the regression to help correct for the errors-in-variables problem associated with using the future realization of investment as a proxy for the current news about future investment.

A bit of data-mining revealed that the vast majority of information in quarterly dividend growth for stock returns is the contemporaneous dividend shock. I construct this shock by

[16] This formulation of the joke seems to take much of the humor out of it...

[17] See Fama (1990) for an implementation of this approach and references to earlier research.

[18] The causality could be reversed if managers use q-theory to make investment decisions.

fitting the quarterly log change in dividends to an AR(4).[19] The shocks, denoted $\tilde{d}_t$, are in-sample residuals from the regression. Quarterly investment growth, denoted ΔI_t, is defined as the log change, from quarter t to quarter $t+1$, in real gross domestic private investment. The timing is chosen to be consistent with the convention used here for other growth rates of real variables. The excess aggregate stock return during quarters t through $t+4$ is denoted R^s_{t+4}. The regression is

$$er^s_t = b_0 + b_1\tilde{d}_t + b_2\Delta I_t + b_3\Delta I_{t+1} + b_4\Delta I_{t+2} + b_5\Delta I_{t+3} + b_6\Delta I_{t+4} + b_7 R^s_{t+4} + e_{s,t}. \quad (17)$$

I estimate (17) on the full 1952–2000 sample, and on two samples formed by splitting the data based on the slope of the term structure at the end of quarter $t-1$. For each sample, three regressions are estimated: One uses only dividend growth, another uses only future investment growth and future stock returns, and the third includes all explanatory variables. Newey-West asymptotic t-statistics are adjusted for one lag of moving average residuals. The results are displayed in Table 8.

The table reports that in the full sample, both innovations in dividends and future investment growth help explain (in an adjusted-R^2 sense) excess stock returns. In the full sample, dividend innovations alone produce an adjusted R^2 of 25 percent. However, when the slope of the term strucure in quarter $t-1$ is steeper than usual, the explanatory power is substantially lower (adjusted R^2 is 12 percent) than when the slope is flatter than usual (adjusted R^2 is 40 percent).[20] This pattern carries over to the explanatory power of future investment growth. In the full sample, five quarters of future changes in investment (along with future stock returns to correct for the errors-in-variables problem) produce an adjusted R^2 of 19 percent. But when the slope is steeper than usual, the adjusted R^2 is only three percent, versus 28 percent when the slope is flatter than usual. When all explanatory variables are used, they capture only 17 percent of the variability in excess stock returns conditioned on a steep slope of the term structure. When the slope is flatter, the variables explain over half of the variability in excess stock returns.

Before attempting to interpret this evidence, I note another result in the table. Notwithstanding the differences in R^2, the residual standard errors are not substantially different

[19] I use quarter-to-quarter changes in log quarterly dividends, not quarter-to-quarter changes in a trailing four-quarter average of log quarterly dividends. Thus there are seasonals in the dividend series that I use. Including seasonal dummies in the forecasting regression did not materially affect the results.

[20] These results are not sensitive to the construction of dividend innovations or the exclusion of future dividends from the regression. If lags -4 through 4 of dividend growth are used as explanatory variables instead of contemporary dividend innovations, the adjusted R^2 is 16 percent for the steep-slope sample and 41 percent for the low-slope sample.

across the slope-sorted samples. (In fact, when the dividend innovation is excluded from the regression, the residual standard error is *higher* for the flat-slope regression.) Recall from Table 3 that excess stock returns are more volatile when the slope is flatter. It appears that during such times, the economy is subject to a type of shock that jointly affects stock returns, dividends, and future investment opportunities. When the slope is steeper than usual, this source of uncertainty disappears. As with the macro forecasting regression (16), the results are qualitatively unchanged if the sample is split based on the magnitude of the one-quarter-ahead forecast of consumption growth.[21]

I now turn to bond returns. Bond returns can be decomposed into news about future short-term interest rates and news about future excess bond returns. For example, bond prices can fall today because short-term interest rates jumped up and are expected to remain high, because investors revise upward their expectations of future changes in short-term interest rates, or because investors revise upward their expectations of future excess returns to bonds. Here I ask whether the kinds of news that moves bond prices depends on the slope of the term structure. I regress the quarter-t excess bond return on contemporaneous and future changes in short-term interest rates y_t and on the excess bond return realized during quarters $t+1$ through $t+4$, denoted R^b_{t+4}. (The future excess bond return also helps correct for the errors-in-variables problem discussed in stock returns.)

$$er^b_t = b_0 + b_1\Delta y_t + b_2\Delta y_{t+1} + b_3\Delta y_{t+2} + b_4\Delta y_{t+3} + b_5\Delta y_{t+4} + b_6 R^b_{t+4} + e_{b,t} \qquad (18)$$

The change Δy_t is defined as the three-month bill yield at the end of quarter t less the three-month bill yield at the end of quarter $t-1$. As with (17), the number of leads included in (18) was determined by data-mining. I estimate (18) on the full 1952–2000 sample and on two subsamples splitting the data based on the slope of the term structure at the end of quarter $t-1$. The results are displayed in Table 9.

The results show that in the full sample, the news in excess bond returns is primarily news about contemporaneous changes in short-term interest rates. Neither changes in short-term interest rates over the next year nor excess bond returns over the next year help explain current returns. This result carries over to periods when the slope is flatter than usual. However, when the slope is steeper than usual, the determinants of current bond returns appear to change. The contemporaneous change in short-term yields remains important. However, bond returns also appear to respond to news about expected changes in short-term yields over the next year. In addition, future excess returns have explanatory power, although it is not clear whether this explanatory power is just picking up the errors-in-variables problem

[21] Using all explanatory variables, the adjusted R^2 in the high-expected-growth sample is 24 percent, versus 53 percent in the low-expected-growth sample.

or whether bond prices today are reacting to changes in expected future excess bond returns. (The sign of the estimated coefficient is consistent with both explanations.)

The evidence in Tables 8 and 9 does not do much to advance our understanding of what drives stock and bond returns. If anything, they deepen our confusion. In particular, when the term structure is steep, stock returns do not appear to be closely related to dividends or future investment. In results that are not detailed here, I have run many additional regressions, attempting to find a variable–such as real interest rates, inflation, or GDP–that has more explanatory power. The results of these regressions can be inferred from their absence in this paper.

5 Concluding comments

This paper attempts to explain, in the context of consumption-based, representative-agent asset-pricing models, why the slope of the term structure forecasts excess returns to stocks and bonds. In one sense, the exercise is successful because it finds that the joint dynamics of consumption and returns are closely linked to the slope. But this evidence goes the wrong way. The evidence, combined with the models, imply that a higher slope corresponds to *lower* expected excess returns for stocks and unchanged expected excess returns for bonds. This implication is avoided in heterogeneous-agent models, but the results in this paper pose problems for these models as well.

The evidence documented here also tells us that the relation between asset returns and the macroeconomy is closely linked to the expected growth of aggregate consumption. Since the term-structure slope forecasts future consumption, this relation is also linked to the slope. When the slope is flatter than usual (expected growth is low), the economy appears to be subject to a kind of shock that affects the current stock market and future short-run (one or two quarters out) economic activity. This shock drives the observed positive correlation between the stock market and future output. But when the slope is steeper than usual (expected growth is high), whatever drives the shock is missing. At those times, stock returns are largely unrelated to future output. What is the source of this shock? When the slope is steep, what determines stock returns? These questions are left for future work.

Appendix

This appendix describes the sources of the data used in this paper.

1. Excess stock returns

   Monthly raw aggregate stock returns are measured by the continuously-compounded return to the Center for Research in Security Prices (CRSP) NYSE/Amex/Nasdaq value-weighted index. Quarterly raw returns are produced by summing monthly returns. Quarter t excess returns are produced by subtracting the continuously-compounded yield to a three-month Treasury bill observed at the end of quarter $t-1$ (also from CRSP).

2. Excess bond returns

   Monthly raw returns to long-maturity Treasury bonds are measured by the continuously-compounded return to a CRSP-constructed portfolio of Treasury bonds with maturities between five and ten years. Construction of excess quarterly returns follows the procedure for stock returns.

3. Term structure slope

   The slope of the term structure is measured by the difference between a ten-year Treasury bond yield and a three-month Treasury bill yield. Both are from CRSP and are observed at quarter-end. The yields are expressed in percent/year. Alternative slope measures that are referred to in the paper replace the ten-year yield with either a thirty-year Treasury bond yield, a five-year zero-coupon Treasury yield (implied by coupon bond prices), or an average yield on long-term, Aaa-rated corporate bonds. All but the corporate bond yields are from CRSP and measured at quarter-end. The corporate bond yields are from Moody's by way of the Federal Reserve's database, and are the average yields during the last month of the quarter.

4. NIPA data

   National Income and Product Accounts (NIPA) data are from the Bureau of Economic Analysis (BEA) web site (publication date April 2001). Gross domestic product (GDP) is divided by the GDP deflator and by mid-quarter U.S. population (also from the BEA web site). Aggregate consumption is measured by the sum of real expenditures on nondurables (divided by the nondurables deflator) and services (divided by the services deflator), divided by the U.S. population. The definition of labor income follows Lettau and Ludvigson (2001). It is divided by the deflator for personal consumption expenditures (PCE) and U.S. population. After-tax corporate profits are measured by the NIPA measure of profits after tax, adjusted for inventory valuation and capital consumption. The result is divided by the PCE deflator. The NIPA data range from 1947:1 through 2000:4.

5. Dividends

   Monthly D/P ratios are constructed by subtracting the return, excluding dividends, to the CRSP value-weighted index from the return, including dividends, to the same index. An artificial series of monthly nominal dividends is constructed by multiplying these D/P ratios by the cumulated gross return, excluding dividends, to the CRSP value-weighted index. Quarterly nominal dividends are the sum of the monthly dividends. Real dividends are produced by dividing by the PCE deflator.

Quarters Ahead (i)	Stock Return Regressions Constant	Slope	R^2	Bond Return Regressions Constant	Slope	R^2
1	1.506	1.149	0.021	0.228	0.483	0.025
	(2.60)	(2.08)		(1.07)	(1.43)	
2	1.501	0.975	0.014	0.223	0.568	0.036
	(2.58)	(1.53)		(1.04)	(2.37)	
3	1.496	1.012	0.015	0.219	0.645	0.048
	(2.54)	(1.68)		(1.03)	(2.97)	
4	1.508	0.343	-0.003	0.220	0.499	0.026
	(2.54)	(0.66)		(1.01)	(2.08)	

Table 1: Forecasting Excess Returns With the Term Structure Slope, 1952:1–2000:4

Quarter t excess log returns to stock and Treasury bond portfolios are regressed on the demeaned slope of the yield curve as of the end of quarter $t - i$. The stock portfolio is the CRSP value-weighted index. The bond portfolio is a portfolio of Treasury bonds with maturities between five and ten years. Excess returns are produced by subtracting the yield on a three-month Treasury bill. The slope is measured by the difference between ten-year and three-month Treasury yields. All variables are expressed in percent. Asymptotic t-statistics, adjusted for generalized heteroskedasticity and one lag of moving-average residuals, are in parentheses. The reported R^2's are adjusted for degrees of freedom.

Quarters	Stock Return Regressions			Bond Return Regressions		
Ahead (i)	Slope	Cons Growth	R^2	Slope	Cons Growth	R^2
1	0.730	-12.260	0.037	0.441	-1.515	0.021
	(1.31)	(-2.15)		(1.20)	(-0.61)	
2	0.590	-11.302	0.026	0.537	-1.211	0.032
	(0.94)	(-1.96)		(1.98)	(-0.48)	
3	0.589	-11.717	0.028	0.611	-1.484	0.045
	(0.97)	(-1.86)		(2.65)	(-0.66)	
4	-0.143	-13.839	0.021	0.485	-1.359	0.025
	(-0.29)	(-2.20)		(1.94)	(-0.60)	

Table 2: Forecasting Excess Returns With the Slope and Historical Consumption Growth, 1957:2–2000:4

Quarter t excess log returns to stock and Treasury bond portfolios are regressed on the quarter $t-i$ yield curve slope and on the quarter $t-i$ "consumption surplus," defined as a weighted average of the past 40 quarters of aggregate consumption growth (equation (8) in the text). The stock portfolio is the CRSP value-weighted index. The bond portfolio is a portfolio of Treasury bonds with maturities between five and ten years. Excess returns are produced by subtracting the yield on a three-month Treasury bill. The slope is measured by the difference between ten-year and three-month Treasury yields. All variables are expressed in percent. Asymptotic t-statistics, adjusted for generalized heteroskedasticity and one lag of moving-average residuals, are in parentheses. The reported R^2's are adjusted for degrees of freedom.

Sample period	Variable	Slope > median	Slope < median	F-test
1952:1-2000:3	Δc_{t+1}	0.377	0.546	0.000
	Δc_{t+2}	0.377	0.541	0.000
	Δc_{t+3}	0.367	0.551	0.000
	er^s_{t+1}	7.112	8.679	0.052
	er^s_{t+2}	7.504	8.578	0.191
	er^s_{t+3}	7.179	8.916	0.036
	er^b_{t+1}	3.157	3.205	0.884
	er^b_{t+2}	2.952	3.445	0.132
	er^b_{t+3}	3.124	3.278	0.640
1952:1-1977:4	Δc_{t+1}	0.485	0.537	0.472
	Δc_{t+2}	0.403	0.579	0.011
	Δc_{t+3}	0.456	0.545	0.204
	er^s_{t+1}	7.318	8.628	0.243
	er^s_{t+2}	7.692	8.641	0.409
	er^s_{t+3}	7.081	9.114	0.074
	er^b_{t+1}	2.330	1.908	0.157
	er^b_{t+2}	2.298	2.068	0.453
	er^b_{t+3}	1.935	2.407	0.123
1983:1-2000:3	Δc_{t+1}	0.331	0.400	0.284
	Δc_{t+2}	0.307	0.414	0.089
	Δc_{t+3}	0.288	0.428	0.028
	er^s_{t+1}	6.998	8.493	0.264
	er^s_{t+2}	7.520	8.246	0.594
	er^s_{t+3}	7.412	8.274	0.531
	er^b_{t+1}	3.723	2.529	0.027
	er^b_{t+2}	3.765	2.456	0.015
	er^b_{t+3}	3.647	2.700	0.089

Table 3: The volatilities of asset returns and aggregate consumption growth, conditioned on the slope of the yield curve

The log change in aggregate per capita real consumption on nondurables and services from quarter $t+i$ to quarter $t+i+1$ is denoted Δc_{t+i}. The excess log return to the aggregate stock market in quarter $t+i$ is denoted er^s_{t+i}. The excess log return to a portfolio of long-term Treasury bonds in quarter $t+i$ is denoted er^b_{t+i}. The quarter-end term-structure slope is measured by the difference between ten-year and three-month Treasury yields. All variables are expressed in percent. The reported dates are the maximum ranges for quarter t. Because the last observation for consumption growth is 2000:3 and the last observation for asset returns is 2000:4, the range is occasionally shorter than labeled. The column labeled "F-test" reports the p-value of the F-statistic that the population variances are equal.

Sample period	Variable	Slope > median				Slope < median			
		Δc_{t+1}	Δli_{t+1}	er^s_{t+1}	er^b_{t+1}	Δc_{t+1}	Δli_{t+1}	er^s_{t+1}	er^b_{t+1}
1952:1-2000:3	Δc_{t+1}	1.00				1.00			
	Δli_{t+1}	0.40	1.00			0.66	1.00		
	er^s_{t+1}	0.01	-0.08	1.00		0.37	0.27	1.00	
	er^b_{t+1}	0.08	-0.03	0.21	1.00	0.08	0.05	0.18	1.00
1952:1-1977:4	Δc_{t+1}	1.00				1.00			
	Δli_{t+1}	0.52	1.00			0.67	1.00		
	er^s_{t+1}	0.09	0.03	1.00		0.43	0.38	1.00	
	er^b_{t+1}	0.05	-0.07	0.21	1.00	-0.14	-0.14	0.02	1.00
1983:1-2000:3	Δc_{t+1}	1.00				1.00			
	Δli_{t+1}	0.34	1.00			0.49	1.00		
	er^s_{t+1}	-0.21	-0.09	1.00		0.40	0.08	1.00	
	er^b_{t+1}	0.21	0.04	0.30	1.00	-0.20	0.18	-0.08	1.00

Table 4: Conditional correlations among asset returns and real variables

Log-differenced aggregate consumption, Δc_t, is defined in Table 3. Log-differenced labor income, Δli_t, is defined similarly. The excess log return to the aggregate stock market in quarter t+1 is denoted er^s_{t+1}. The excess log return to a portfolio of long-term Treasury bonds in quarter $t+1$ is denoted er^b_{t+1}. All variables are expressed in percent. The correlations are conditioned on whether the slope of the yield curve at the end of quarter t is steeper or flatter than its median. Details on the ranges of data used are in Table 3.

Sample Period	Dependent Variable	Coef on: Slope	Coef on: Fitted Cons Growth
1952-2000	$\|\tilde{c}_{t+1}\|$	-0.032	
		(-1.82)	
1952-2000	$\|\tilde{c}_{t+1}\|$		-0.295
			(-2.65)
1952-2000	$\|\tilde{c}_{t+1}\|$	-0.007	-0.271
		(-0.36)	(-2.27)
1952-2000	$(\tilde{c}_{t+1})^2$	-0.043	
		(-2.00)	
1952-2000	$(\tilde{c}_{t+1})^2$		-0.317
			(-2.35)
1952-2000	$(\tilde{c}_{t+1})^2$	-0.021	-0.244
		(-1.00)	(-1.92)
1952-1977	$\|\tilde{c}_{t+1}\|$	0.007	-0.187
		(0.11)	(-0.62)
1952-1977	$(\tilde{c}_{t+1})^2$	0.004	-0.270
		(0.06)	(-0.82)
1983-2000	$\|\tilde{c}_{t+1}\|$	-0.012	-0.169
		(-0.62)	(-1.29)
1983-2000	$(\tilde{c}_{t+1})^2$	-0.013	-0.197
		(-0.83)	(-1.61)

Table 5: Forecasting the volatility of innovations in consumption growth

The absolute and squared innovations in log consumption growth $\tilde{c}_{t+1}$ are regressed on the previous slope of the term structure and the previous quarter's forecast of consumption growth. Innovations in consumption, as well as the one-quarter-ahead forecasts, are produced with an OLS regression of consumption growth on lagged consumption growth, lagged labor income growth, the slope of the term structure, inflation, and stock returns. All variables are in percent. Asymptotic t statistics, adjusted for generalized heteroskedasticity and one lag of moving average residuals, are in parentheses.

Sample Period	Regressions of stock return covariances: Slope	Regressions of stock return covariances: Fitted Cons Growth	Regressions of bond return covariances: Slope	Regressions of bond return covariances: Fitted Cons Growth
1952-2000	−0.665		−0.204	
	(−3.05)		(−0.84)	
		−5.324		−1.096
		(−3.43)		(−0.89)
	−0.255	−4.419	−0.153	−0.553
	(−1.42)	(−2.98)	(−0.70)	(−0.71)
1952-1977	−0.984		0.189	
	(−1.87)		(1.39)	
		−5.329		0.895
		(−1.96)		(1.37)
	−0.371	−4.016	0.113	0.496
	(−0.67)	(−1.30)	(0.64)	(0.59)
1983-2000	−0.715		0.243	
	(−1.91)		(3.34)	
		−2.872		0.258
		(−1.17)		(0.67)
	−0.701	−2.784	0.242	0.228
	(−1.91)	(−1.23)	(3.30)	(0.60)

Table 6: Forecasting covariances between consumption growth and asset returns

A two-step procedure is used to forecast one-quarter-ahead covariances between consumption growth and asset returns. In step one, innovations in log consumption growth, excess stock returns, and excess bond returns are produced by regressing the raw variables on explanatory variables. In step two, the product of innovations in consumption growth and innovations in asset returns are regressed on the lagged slope of the term structure and forecasted consumption growth. The consumption forecasts are produced with the same regression used to construct one-quarter-ahead innovations. All variables are in percent. Asymptotic t statistics, adjusted for generalized heteroskedasticity and one lag of moving average residuals, are in parentheses.

Dependent Variable	Stock return regressions b_1	b_2	b_3	R^2	Bond return regressions b_1	b_2	b_3	R^2
Δc_{t+1}	0.144	0.024	-0.024	0.105	0.137	0.015	-0.005	0.016
	(2.01)	(4.18)	(-3.11)		(1.88)	(0.72)	(-0.23)	
Δc_{t+2}	0.195	0.012	-0.015	0.049	0.177	0.030	-0.027	0.045
	(2.67)	(2.78)	(-2.20)		(2.39)	(2.28)	(-1.38)	
Δc_{t+3}	0.161	0.003	-0.005	0.013	0.155	0.019	-0.030	0.022
	(2.22)	(0.50)	(-0.71)		(2.21)	(0.91)	(-1.29)	
Δgdp_{t+1}	0.515	0.037	-0.032	0.114	0.581	-0.018	-0.045	0.093
	(3.64)	(3.46)	(-2.16)		(4.12)	(-0.42)	(-0.96)	
Δgdp_{t+2}	0.604	0.049	-0.040	0.191	0.634	0.086	-0.111	0.137
	(4.38)	(5.30)	(-2.72)		(4.14)	(4.01)	(-3.55)	
Δgdp_{t+3}	0.430	0.021	-0.012	0.059	0.441	0.057	-0.054	0.058
	(2.89)	(1.62)	(-0.75)		(2.90)	(1.54)	(-1.25)	
Δli_{t+1}	0.210	0.029	-0.037	0.038	0.187	0.014	-0.022	-0.003
	(1.69)	(2.24)	(-2.09)		(1.37)	(0.45)	(-0.57)	
Δli_{t+2}	0.308	0.018	-0.011	0.037	0.329	0.018	-0.020	0.022
	(2.34)	(1.81)	(-0.84)		(2.51)	(1.20)	(-0.62)	
Δli_{t+3}	0.054	-0.014	0.035	0.013	0.131	0.005	-0.019	-0.009
	(0.40)	(-0.87)	(1.97)		(0.98)	(0.16)	(-0.46)	

Table 7: Forecasting real activity with asset returns and the term structure slope, 1952:1 through 2000:3

Log-differenced quarterly real per capita consumption, Δc_t, is defined in Table 1. Log-differenced quarterly real per capita GDP, Δgdp_t, and labor income, Δli_t, are defined similarly. The variable D_SLOPE_t is a dummy variable that equals one if the slope of the Treasury term structure at the end of quarter t exceeds its median. Excess log returns to stock and Treasury bond markets, er^s_{t+1} and er^b_{t+1}, are defined in Table 1. All variables are expressed in percent. This table reports results from the regression

$$\Delta x_{t+i} = b_0 + b_1 D\_SLOPE_t + (b_2 + b_3 D\_SLOPE_t) er^k_{t+1} + e_{k,t+i},\ x = \{c, gdp, li\},\ k = \{s, b\}.$$

Heteroskedasticity-consistent t-statistics, adjusted for one lag of moving average residuals, are in parentheses. The reported R^2's are adjusted for degrees of freedom.

Sample	$\tilde{d}_t$	ΔI_t	ΔI_{t+1}	ΔI_{t+2}	ΔI_{t+3}	ΔI_{t+4}	R^s_{t+4}	SEE	R^2
Full sample	0.652							6.966	0.254
	(5.96)								
Slope > median	0.413							6.714	0.116
	(2.91)								
Slope < median	0.855							6.691	0.403
	(6.89)								
Full sample		0.182	0.535	0.366	-0.050	0.261	-0.059	7.260	0.187
		(1.51)	(3.39)	(2.78)	(-0.61)	(2.64)	(-1.62)		
Slope > median		0.210	0.127	0.408	0.182	0.122	0.005	7.047	0.026
		(1.17)	(0.61)	(1.97)	(1.12)	(0.73)	(0.10)		
Slope < median		0.211	0.732	0.388	-0.098	0.416	-0.120	7.293	0.284
		(1.50)	(4.08)	(2.37)	(-0.92)	(3.11)	(-2.14)		
Full sample	0.615	0.131	0.512	0.293	-0.005	0.247	-0.061	6.145	0.417
	(6.34)	(1.41)	(4.59)	(3.17)	(-0.07)	(2.56)	(-2.00)		
Slope > median	0.451	0.187	0.254	0.381	0.179	0.209	-0.011	6.510	0.169
	(3.45)	(1.13)	(1.42)	(2.09)	(1.14)	(1.28)	(-0.26)		
Slope < median	0.722	0.095	0.567	0.248	-0.051	0.276	-0.090	5.688	0.565
	(6.72)	(0.79)	(4.53)	(2.24)	(-0.49)	(2.31)	(-2.20)		

Table 8: Explaining stock returns with current and future real variables, 1952:1 through 2000:4

Excess log stock returns in quarter t are regressed on the contemporaneous innovation in log quarterly dividends ($\tilde{d}_t$, the innovations are residuals from an AR(4) regression), current and future changes in log real private fixed domestic investment (ΔI_{t+i}), and on the excess stock return over the next four quarters (R^s_{t+4}). The data are split into two samples, based on whether the slope of the Treasury term structure in quarter $t-1$ is greater or less than its median. All variables are expressed in percent. The column labeled SEE reports the standard error of the estimated residual. Heteroskedasticity-consistent t-statistics, adjusted for one lag of moving average residuals, are in parentheses. The reported R^2's are adjusted for degrees of freedom.

Sample	Δy_t	Δy_{t+1}	Δy_{t+2}	Δy_{t+3}	Δy_{t+4}	R^b_{t+4}	SEE	R^2
Full sample	-2.167	-0.112	-0.197	0.196	0.080	0.003	2.369	0.456
	(-13.82)	(-0.37)	(-0.99)	(0.91)	(0.32)	(0.09)		
Slope > median	-2.588	-1.160	-0.432	-0.900	-0.231	-0.114	2.209	0.505
	(-8.44)	(-2.76)	(-1.18)	(-2.05)	(-0.64)	(-2.19)		
Slope < median	-2.046	-0.024	-0.399	0.122	-0.143	-0.088	2.089	0.581
	(-9.20)	(-0.06)	(-1.11)	(0.44)	(-0.58)	(-1.20)		

Table 9: Explaining bond returns with current and future short-term interest rates and future bond returns, 1952:1 through 2000:4

Excess log bond returns in quarter t are regressed on current and future changes in three-month Treasury bill yields y_t and on the excess bond return over the next four quarters (R^b_{t+4}). The data are split into two samples, based on whether the slope of the Treasury term structure in quarter $t-1$ is greater or less than its median. All variables are expressed in percent. The column labeled SEE reports the standard error of the estimated residual. Heteroskedasticity-consistent t-statistics, adjusted for one lag of moving average residuals, are in parentheses. The reported R^2's are adjusted for degrees of freedom.

References

Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* 80, 38-42.

Backus, David K., Allan W. Gregory, and Stanley E. Zin, 1989, Risk premiums in the term structure: Evidence from artificial economies, *Journal of Monetary Economics* 24, 371-399.

Barro, Robert J., 1989, The stock market and the macroeconomy: Implications of the October 1987 crash, in R. W. Kamphuis, R. C. Kormendi, and J. W. H. Watson (eds.), *Black Monday and the Future of Financial Markets,* Dow Jones Irwin, Homewood, Ill.

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, On biases in tests of the expectations hypothesis of the term structure of interest rates, *Journal of Financial Economics* 44, 309-348.

Berk, Jonathan B., 1995, A critique of size-related anomalies, *Review of Financial Studies* 8, 275-286.

Boudoukh, Jacob, 1993, An equilibrium model of nominal bond prices with inflation-output correlation and stochastic volatility, *Journal of Money, Credit, and Banking* 25, 636-665.

Brandt, Michael W., and Kevin Q. Wang, 1999, Time-varying risk aversion, unexpected inflation, and the dynamics of the term structure, Working paper, The Wharton School and the Rotman School of Management.

Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373-399.

Campbell, John Y., 1999, Asset prices, consumption, and the business cycle, Chapter 19 in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics, Volume 1*, North-Holland: Amsterdam, 1231-1303.

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., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* Princeton University Press, Princeton, NJ.

Chen, Nai-Fu, Financial investment opportunities and the macroeconomy, 1991, *Journal of Finance* 46, 529-554.

Constantinides, George M., 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519-543.

Constantinides, George M., and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219-240.

Epstein, Lawrence, and Zin, Stanley, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937-969.

Estrella, Arturo, and Gikas A. Hardouvelis, 1991, The term structure as a predictor of real economic activity, *Journal of Finance* 46, 555-576.

Estrella, Arturo, and Frederic S. Mishkin, 1998, Predicting U.S. recessions: Financial variables as leading indicators, *Review of Economics and Statistics* 80, 45-61.

Fama, Eugene F., 1990, Stock returns, expected returns, and real activity, *Journal of Finance* 45, 1089-1108.

Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23-49.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3-56.

Fischer, Stanley, and Robert C. Merton, 1984, Macroeconomics and finance: The role of the stock market, *Carnegie-Rochester Conference Series on Public Policy* 21, 57-108.

Harvey, Campbell R., 1989, Forecasts of economic growth from the bond and stock markets, *Financial Analysts Journal* 45, 38-45.

Kothari, S. P., and Jay Shanken, 1992, Stock return variation and expected dividends, *Journal of Financial Economics* 31, 177-210.

Labadie, Pamela, 1994, The term structure of interest rates over the business cycle, *Journal of Economic Dynamics and Control* 18, 671-697.

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance*, forthcoming.

Lucas, Robert E., Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429-1446.

Mehra, Rajnish, and Prescott, Edward, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145-161.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703-708.

Pennacchi, George G., 1991, Identifying the dynamics of real interest rates and inflation: Evidence from survey data, *Review of Financial Studies* 4, 53-86.

Roll, Richard R., A critique of the asset pricing theory's tests, part I: On past and potential testability of the theory, *Journal of Financial Economics* 4, 129-176.

Rubinstein, Mark, 1976, The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics* 7, 407-425.

Schwert, G. William, 1989, Why does stock market volatility change over time, *Journal of Finance* 44, 1115-1153.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375-421.

Sundaresan, Suresh, 1989, Intertemporally dependent preferences and the volatility of consumption and wealth, *Review of Financial Studies* 2, 73-88.

Wachter, Jessica A., 2001, Habit formation and returns on bonds and stocks, Working paper, Stern School of Business, NYU.

Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401-421.

Whitelaw, Robert F., 2000, Stock market risk and return: An equilibrium approach, *Review of Financial Studies* 13, 521-547.

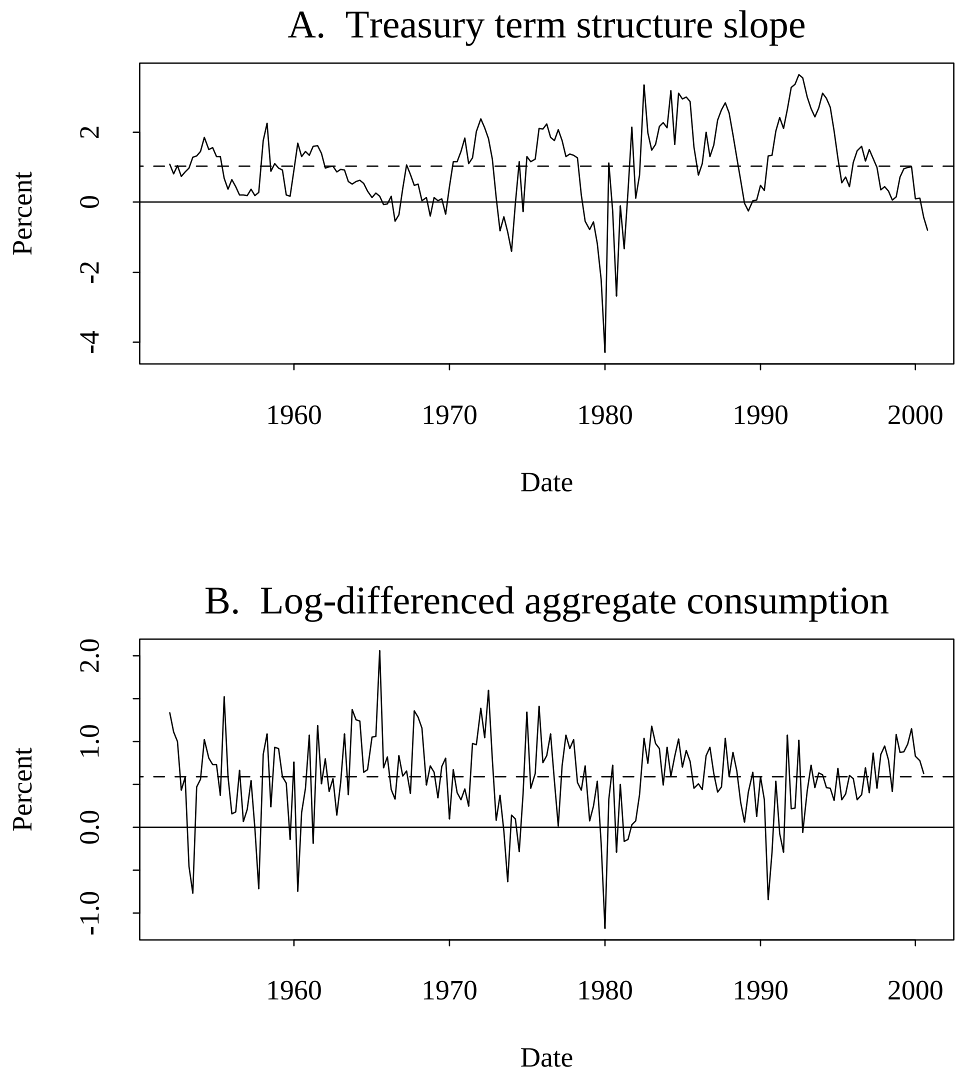

Figure 1: The slope of the term structure and consumption growth

The term structure slope at the end of quarter t is the difference between ten-year and three-month Treasury yields. Consumption growth in quarter t is measured by 100 times the log change, from quarter t to quarter $t+1$, of aggregate per capita real consumption on nondurables and services. Dashed lines are drawn at median values.

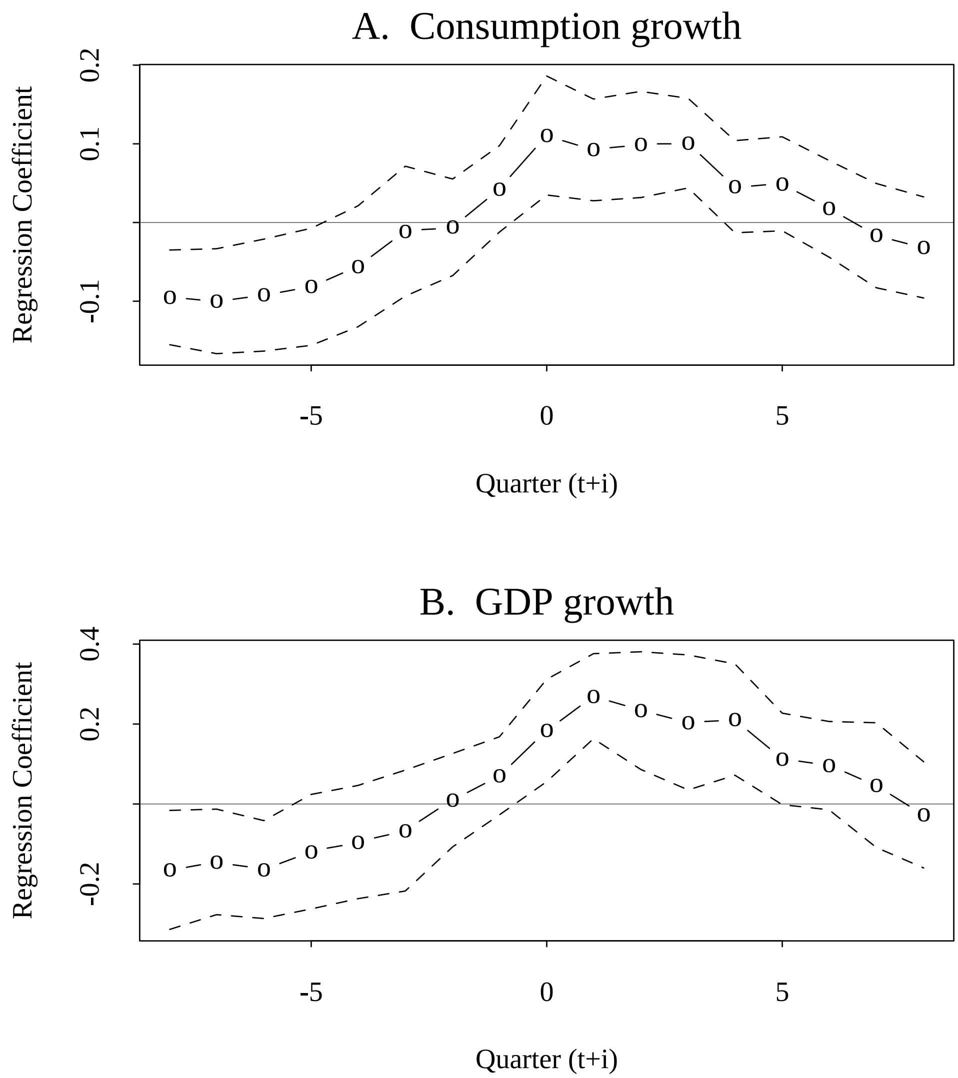

Figure 2: Regressions of lags and leads of consumption and GDP growth on the slope of the term structure, 1952:1 through 2000:4

Quarterly changes in log consumption and GDP Δc_t and Δgdp_t are defined in the notes to Table 3. This figure reports estimated coefficients from regressions of Δc_t and Δgdp_t on the slope of the Treasury term structure at the end of quarter $t - i$. For each real variable, 17 regressions are estimated ($i = -8, \ldots, 8$). The dashed lines are +/- two heteroskedasticity-consistent standard errors, adjusted for one lag of moving average residuals.

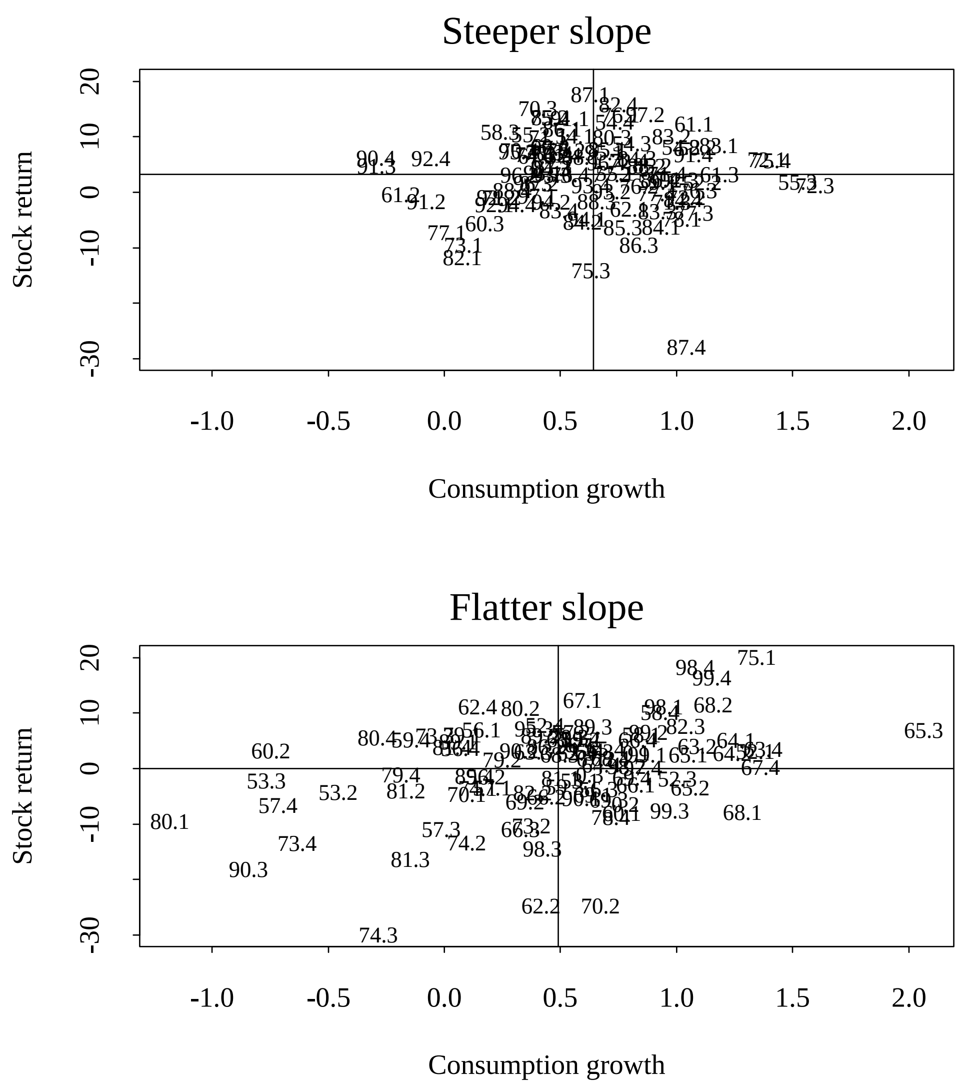

Figure 3: Scatter plots of excess stock market returns and contemporaneous consumption growth, conditioned on the slope of the yield curve

Quarterly changes in log consumption and quarterly excess stock returns are defined in the notes to Table 3. Quarter $t+1$ realizations from 1952 through 2000 are sorted into two groups, based on whether the slope of the term structure at the end of quarter t was greater or less than its median. These panels display scatter plots of the two groups, where each observation is labeled with the quarter $t+1$ date (YY:Q). The solid lines are the group-specific means of the variables.