Consumption Fluctuations and Expected Returns

Victoria Atanasov, Stig Vinther Møller, and Richard Priestley*

Abstract

This paper introduces a novel consumption-based variable, cyclical consumption, and examines its predictive properties for stock returns. Future expected stock returns are high (low) when aggregate consumption falls (rises) relative to its trend and marginal utility from current consumption is high (low). We show that the empirical evidence ties consumption decisions of agents to time-variation in returns in a manner consistent with asset pricing models based on external habit formation. The predictive power of cyclical consumption is not confined to bad times and subsumes the predictability of many popular forecasting variables.

JEL Classification: G10; G12; G17

Keywords: cyclical consumption fluctuations; time-varying expected stock returns; predictability, habit formation.

*We would like to thank the editor, Stefan Nagel, an associate editor, and two anonymous referees for generous comments that improved the paper. We also thank John Y. Campbell, John H. Cochrane, Tim A. Kroencke, Ernst Maug, Erik Theissen, Jessica A. Wachter (WFA discussant), and participants at the AFFI meeting 2018, the SoFiE meeting 2018, the WFA meeting 2018, the EEA-ESEM congress 2018, and seminar participants at University of Mannheim and University of Oslo for helpful comments and suggestions. We also thank Amit Goyal, Martin Lettau, and Robert J. Shiller for making their data available. We express our gratitude to David E. Rapach for providing his code for the bootstrap analysis. Victoria Atanasov is from the Chair of Finance, University of Mannheim, Stig V. Møller is at CREATES, Aarhus University and affiliated with the Danish Finance Institute. Send correspondence to Richard Priestley, Department of Finance, BI Norwegian Business School, Nydalsveien 37, N-0444 Oslo, Norway; telephone: +47 46 410 515; e-mail: richard.priestley@bi.no. We have read the *Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

In this paper, we take a new approach to linking stock return predictability to both bad and good economic times. Consider an economy where investors exhibit external habit formation as in, for example, Campbell and Cochrane (1999), and therefore risk premia vary over time through variation in risk aversion. In good times, when consumption rises above its trend and hence the marginal utility of present consumption is low, investors are willing to give up current consumption and invest. This in turn forces stock prices to increase and future expected returns to decrease. Conversely, in bad times, when consumption falls below its trend and hence the marginal utility of current consumption is high, expected returns in the future need to be high in order to induce investors to postpone the valuable present consumption and to invest and consume in the future. It is our conjecture that cyclical fluctuations in aggregate consumption should be useful in picking out bad and good times in the economy as measured from a representative agents' point of view, and hence informative about future excess stock returns. If the argument holds true, there should exist an inverse relation between cyclical consumption and future expected returns in the data.

The empirical results that we present in this paper confirm the idea that future expected returns are high (low) when consumption is falling below (rising above) its trend and cyclical consumption is low (high). Cyclical fluctuations in consumption, which we intermittently refer to as cc, capture a significant fraction of the variation in future stock market returns. The results we document are important because they imply an intimate relation between expected returns and consumption suggesting that asset prices are driven by fundamental shocks reflecting changes in marginal utility.

An important and novel finding is that the predictive power of cyclical consumption is not confined to bad times alone. Cyclical consumption provides a consistent description of how positive and negative macroeconomic events, reflected through consumption decisions of investors, affect stock market returns. These results are notable because they stand in stark contrast to Rapach, Strauss, and Zhou (2010), Henkel, Martin, and Nardari (2011), Dangl and Halling (2012), and Golez and Koudijs (2018) who find that popular predictor variables can only forecast stock returns in bad times, whereas there is essentially no evidence of predictability in good times, that is, during business cycle expansions.

To extract the cyclical component of consumption, we employ a simple and robust linear projection method of Hamilton (2018). This procedure provides an alternative concept of what macroeconomists often refer to as the "cyclical component" of a time series and it is advantageous over other prominent detrending methods in two important respects. First, the procedure ensures that the identified cyclical component is stationary and consistently estimated for a wide range of nonstationary processes. Second, it produces a series which is accurately related to the underlying economic fluctuations as opposed to, for instance, the popular Hodrick and Prescott (1997) filter which can spuriously generate dynamic relations. This feature is particularly appealing because it implies that any predictive ability of cyclical consumption for stock returns is more likely to reflect actual predictability rather than arise as a result of a statistical artifact of the decomposition method (Hamilton (2018)). We explore a variety of alternative specifications and utilize other econometric procedures to isolate cyclical variation in consumption such as polynomial time trends and backward-looking moving averages, and find even stronger evidence of predictability. The choice of Hamilton's (2018) detrending procedure as a benchmark specification provides a conservative and robust view of return predictability.

Our findings are supportive of theoretical explanations of asset prices which generate time-varying expected returns such as models with time-varying risk aversion. In the external habit formation model of Campbell and Cochrane (1999), for example, habit acts like a trend for consumption. A decline in consumption relative to the trend, which can be thought of as bad times, leads to low stock prices and high expected returns. Conversely, an increase in consumption above trend, which can be thought of as good times, leads to high stock prices and low expected returns. Under relatively mild assumptions, there exists a tight relation between a finite-horizon version of the surplus consumption variable of Campbell and Cochrane (1999), which generates changes in equity prices in the model, and cyclical consumption.

To explore formally the link between cyclical consumption and habit models, we simulate data from the Campbell and Cochrane (1999) model and investigate the extent of the model-implied predictability, and examine its consistency with the time-series predictability that we observe in our actual data. The simulations show that the habit model produces an inverse relation between expected returns and cyclical consumption just as we find in the data. The degree of in-sample predictability implied by the model is qualitatively comparable to that in the data. The out-of-sample tests reinforce the results from in-sample regressions but typically indicate less predictable movements in expected returns. These findings open up a possibility to interpret our results as evidence of countercyclical variation in the market price of consumption risk.

We perform a battery of robustness checks of our empirical findings and address a number of econometric concerns surrounding predictive regressions with persistent predictors (Nelson and Kim (1993) and Stambaugh (1999)). Both the IVX testing approach of Kostakis, Magdalinos, and Stamatogiannis (2015) that robustifies the inference to the degree of regressor persistence, and an advanced bootstrap procedure that accounts for the regressor's time series properties indicate strong evidence of predictability at the one-quarter horizon which extends to horizons of about five years. The predictability does not vanish during the post-oil-crisis period in which standard popular business cycle indicators have proven dismal as predictive variables (Welch and Goyal (2008)).

We also show that the forecasting power of cyclical consumption fluctuations is not confined to the aggregate U.S. stock market. Robust patterns of predictability exist across industry portfolios. In addition, the strong predictive ability of cyclical consumption extends to international equity markets. A global measure of cyclical consumption computed as a simple average of respective developed market country-specific components captures a large part of time-variation in future expected returns on the world market portfolio as well as on the regional portfolios such as the European portfolio, the EAFE (Europe, Australia, and the Far East) portfolio or the G7 portfolio.

Explaining the dynamic behavior of asset returns using aggregate consumption data is a challenging task for financial economists. Very few studies find evidence in favor of returns being predictable from consumption. Perhaps the most prominent consumption-based predictive variable is Lettau and Ludvigson's (2001) consumption-wealth ratio, *cay*. We find that cyclical consumption contains predictive information which goes clearly over and above that of many well-recognized variables, including the consumption-wealth ratio of Lettau and Ludvigson (2001), the ratio of labor income to consumption of Santos and Veronesi (2006), and the conditional volatility of consumption of Bansal, Khatchatrian, and Yaron (2005). We consider nineteen alternative popular economic variables and find that very few of them have predictive power and none of them can systematically generate better out-of-sample forecasts than cyclical consumption.

While we have emphasized the connection between our empirical analysis and the external habit model of Campbell and Cochrane (1999), our result that stock returns are predictable by consumption fluctuations appears consistent with other classes of asset pricing models such as learning models which can generate countercyclical variation in risk premia (Collin-Dufresne, Johannes, and Lochstoer (2016) and Nagel and Xu (2018)). A series of positive fundamental shocks in a learning model makes the agent optimistic, asset prices high, and subsequent future returns, on average, low. For example, Nagel and Xu (2018) predict that the equity premium is negatively related to long-run weighted averages of past real per capita payout growth rates and they verify this empirically. Thus, in line with our empirical results, past growth rates of fundamentals generate slow-moving time-variation in expected returns. However, unlike the habit-based explanation of return predictability, the learning model of Nagel and Xu (2018) features constant relative risk aversion and return predictability which is induced by subjective belief dynamics rather than time-varying risk aversion.

Other models could also be congruous with our empirical findings that consumption fluctuations can predict future stock returns. For example, models with heterogeneous investors such as Constantinides and Duffie (1996) and Constantinides and Ghosh (2017) can generate a link between past fundamentals and expected market returns. In these models, countercyclical shocks to labor income risk imply a countercyclical variation in the equity premium and hence stock return behavior which could be reflected in consumption fluctuations. Relatedly, Chien, Cole, and Lustig (2016) show that in a model with agents with different asset trading technologies, a sequence of bad shocks can magnify cyclical fluctuations in the price of risk and drive up the Sharpe ratio.

Furthermore, recent models that include leverage offer a direct link where countercyclical variation in leverage generates predictability of the risk premium and affects aggregate consumption dynamics. For example, Gomes and Schmid (2017) develop a general equilibrium model with heterogeneous firms, where countercyclical leverage drives up risk premia on financial assets in downturns which is naturally reflected in credit spread changes. Because defaults tend to cluster in downturns, when the market price of risk is high, credit spreads spike up in recessions. These endogenous movements in credit prices amplify the effects of macroeconomic shocks and imply predictable patterns in expected stock returns over business cycle.

The paper proceeds as follows. Section I explains how cyclical consumption is constructed. Section II presents the empirical results. A number of robustness tests are summarized in Section III. Section IV compares the out-of-sample forecasting ability across alternative predictors. Section V lays out a simple economic framework based on the habit model of Campbell and Cochrane (1999) where cyclical consumption emerges as a relevant predictor variable for future stock returns. It also conducts a simulation analysis to compare the extent of predictability in the model and historical data. We conclude in Section VI.

I. Extracting cyclical consumption

As our primary measure of consumption, we use aggregate seasonally adjusted consumption expenditures on nondurables and services from the National Income and Product Accounts (NIPA) Table 7.1 constructed by the Bureau of Economic Analysis (BEA) in the Department of Commerce of the United States. The data are quarterly, in real per capita terms, measured in 2009 chain weighted dollars, and span the period from the first quarter of 1947 to the fourth quarter of 2017.

To extract the cyclical component of consumption, we employ a simple and robust linear projection method of Hamilton (2018) which provides an alternative means to identify what macroeconomists usually refer to as the "cyclical component" of a time series. We regress the log of real per capita consumption, c_t, on a constant and four lagged values of consumption as of date $t-k$:

$$c_t = b_0 + b_1 c_{t-k} + b_2 c_{t-k-1} + b_3 c_{t-k-2} + b_4 c_{t-k-3} + \omega_t, \tag{1}$$

where the regression error, ω_t, is our measure of cyclical consumption cc_t at time t:

$$cc_t = c_t - \widehat{b}_0 - \widehat{b}_1 c_{t-k} - \widehat{b}_2 c_{t-k-1} - \widehat{b}_3 c_{t-k-2} - \widehat{b}_4 c_{t-k-3}. \tag{2}$$

This procedure has several attractive features over other popular detrending methods. In particular, it offers a reasonable model-free way to construct a time series which is accurately related to the actual economic fluctuations as opposed to, for instance, the Hodrick and Prescott (1997) filter which can spuriously generate series with dynamics that have no relation to the underlying data-generating process. Under plausible assumptions, the Hamilton (2018) method ensures that the identified residual component is stationary and consistently estimated for a wide range of unknown and possibly nonstationary processes.[1]

[1] The detrending procedure of Hamilton (2018) allows us to remove the nonstationary component of c_t without modeling the nonstationarity, as the decomposition in Equation (1) will imply a stationary process ω_t, if either the kth difference of c_t or the deviation of c_t from a kth-order deterministic time polynomial is stationary for some k as the sample size becomes large, see Hamilton (2018) for a formal proof.

Furthermore, by virtue of the fact that cc is a one-sided filter, any finding that cc can predict future observations of some other variable should represent a true predictive ability rather than an artifact of a choice of a detrending method.[2]

An empirical implementation of Equation (1) requires a choice of k. Hamilton (2018) recommends using a two-year horizon as a standard benchmark for business cycle dynamics and values of around five years for capturing the effect of longer-term shocks which are "nevertheless still transient". We experimented with alternative specifications of k ranging from one to eleven years and generally found evidence of stock return predictability. The benchmark results we present in the paper are based on cc computed using a horizon of six years, i.e. $k = 24$ with quarterly data.[3]

[Figure 1 about here]

Figure 1 shows a time series plot of cc computed from Equation (2) for $k = 24$ along with recession dates as defined by the NBER. Cyclical consumption has an unconditional mean of zero by construction, a standard deviation of 3.74%, and a first order autocorrelation of 0.97 corresponding to a half-life of slightly over five years. This implies highly persistent expected returns in the return forecasting regressions as emphasized by Campbell and Cochrane (1999), Pastor and Stambaugh (2009), and van Binsbergen and Koijen (2010).[4] The figure illustrates that cc exhibits significant business cycle fluctuations in the post-war period in that it typically rises after recessions and reaches its highest values some time before the onset of recessions, and falls throughout economic contractions. Our contention is that these fluctuations in cyclical consumption constitute a more accurate description of

[2] In this respect, Hamilton (2018) argues that in contrast to the HP cyclical series which is readily forecastable from its own lagged values and likewise past values of other variables, the realizations of ω will by construction be difficult to predict.

[3] A choice of a six-year horizon turns out to be consistent with implications of the external habit model of Campbell and Cochrane (1999) as we show in Section V.

[4] For comparison, Lettau and Ludvigson (2013) identify a risk aversion shock with a half-life of over four years.

good and bad economic times than previously employed predictor variables. If so, cyclical consumption should contain predictive information about future expected stock returns. We test this hypothesis below.

II. Predictive regression analysis

We investigate the forecasting ability of cyclical consumption for two measures of aggregate stock market returns: the return on Standard and Poor's composite stock price index (S&P 500) and the return on the Center for Research in Security Prices (CRSP) value-weighted index of U.S. stocks listed on the NYSE, NASDAQ, and Amex. We compute excess returns by subtracting the return on the 30-day Treasury bill from the market return. We focus on excess returns but also examine nominal returns as well as real returns calculated by deflating nominal returns with the inflation rate of the aggregate U.S. Consumer Price Index (CPI). We download the data on returns from the Wharton Research Data Services (WRDS) database and the CPI inflation rate from the Bureau of Labor Statistics (BLS). Unless otherwise specified, we compute a measure of cyclical consumption from the most recently available figures for seasonally adjusted consumption of nondurables and services in real per capita terms and based on full-sample parameter estimates in Equation (2).

A. Return predictive regressions

We consider a standard predictive regression model for analyzing aggregate stock return predictability:

$$r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}, \tag{3}$$

where cc_t is one-quarter lagged cyclical consumption and $r_{t,t+h}$ is the h-quarter ahead log excess return on the stock market. We measure $r_{t,t+h}$ as the h-quarter continuously compounded log return on the market less the corresponding h-quarter continuously compounded log Treasury bill return.

To test the significance of β in Equation (3), we use the Newey and West (1987) heteroskedasticity- and autocorrelation-robust t-statistic (truncated at lag h; our results are robust towards other choices of truncation lags). In addition, we compute empirical p-values for the slope estimates from a wild bootstrap procedure that accounts for the persistence in regressors and correlations between equity stock return and predictor innovations, and allows for general forms of heteroskedasticity.[5] This simulation produces an empirical distribution that better approximates the finite sample distribution of the slope estimates in Equation (3). For more powerful tests, we follow the recommendation of Inoue and Kilian (2004) and calculate p-values for a one-sided alternative hypothesis.[6]

[Table I about here]

Panel A of Table I reports the OLS estimates of β, the corresponding t-statistics (in parentheses), and the adjusted R^2s, $\bar{R}^2$, (in square brackets) from predictive regressions in Equation (3). We find that the estimated coefficient on cc is negative and that there is an economically sizable predictive impact of cyclical consumption on future excess stock market returns. In particular, the point estimate of β in the quarterly regression on the S&P 500 index is -1.70 in annual terms (first row, second column in Table I). This implies that a fall in cc by one standard deviation below its mean leads to a rise in the expected return of about 6 percentage points at an annual rate. The estimate of the coefficient is strongly statistically significant and the associated $\bar{R}^2$ is 3.69%. Thus, expected returns are low when cyclical consumption is high in good times or economic upswings, and expected returns are high when cyclical consumption is low in bad times or economic downturns. This result is

[5] A general concern with predictability regressions is that their reliability can be undermined by the uncertainty regarding the order of integration of the predictor variable. Statistical inference can be unreliable when the predictor variable is persistent and its innovations are highly correlated with returns (Nelson and Kim (1993) and Stambaugh (1999)). Modelling the predictive variables as local-to-unity processes can lead to invalid inference if the regressor contains stationary or near-stationary components (Valkanov (2003), Lewellen (2004), Campbell and Yogo (2006), and Hjalmarsson (2011)).

[6] The bootstrap procedure we apply follows that of Rapach, Ringgenberg, and Zhou (2016).

consistent with investors responding rationally to countercyclical variation in the price of consumption risk over time: A fall in consumption relative to its past history indicates bad economic times where marginal utility of current consumption is high and future returns are expected to be high.

Columns three to seven in Panel A of Table I show that predictability extends to longer horizons of one to five years. The extent of predictability increases with the horizon both in terms of the size of the estimated coefficient and $\bar{R}^2$ statistics, but at a decreasing rate. For example, at the four quarter horizon the estimated coefficient and $\bar{R}^2$ are almost four times as large as the ones recorded at the one quarter horizon. In contrast, the increase from the sixteenth to the twentieth quarter horizon for the coefficient size is around twenty five percent and for the $\bar{R}^2$ less than ten percent.[7]

The second row in Panel A of Table I reveals a similar pattern of predictability for the CRSP value-weighted returns. The table also shows that the predictive power of cyclical consumption applies to both real returns (Panel B) and actual returns (Panel C), although the evidence of predictability for actual returns is not quite as prominent.

[Table II about here]

Kostakis, Magdalinos, and Stamatogiannis (2015) develop a test that is robust to the regressor's degree of persistence (including unit root, local-to-unit root, near-stationary or stationary persistence classes) and has good size and power properties. This approach alleviates practical concerns about the quality of inference under possible misspecification of the (generally unobservable) time series properties of the regressor in long-horizon predictive regressions. Table II reports the results using their IVX estimator to test the significance of the estimate of β in Equation (3). We find that the null hypothesis of no predictability

[7] Following the advice of an anonymous referee, we compared the direct regression coefficients in Table I with coefficients implied from a first-order VAR model. Table AIII in the internet appendix shows that the indirect coefficients are very similar to the direct coefficients we obtain from the time-overlapping multi-horizon regressions.

can be rejected at the 5% level for excess returns and real returns at all horizons and for both the CRSP and S&P 500 indices. For actual returns, we obtain slightly lower IVX-Wald statistics but typically also reject the null hypothesis of no predictability.

In summary, we show that stock returns are predictable by cyclical consumption fluctuations at various horizons over the post-war period. Expected returns are predicted to be high when consumption falls relative to its trend and cyclical consumption is low and marginal utility is high. In bad times when the marginal utility of consumption is high, investors want to consume more and therefore require a higher expected return to give up valuable current consumption. In good times, marginal utility of consumption is low and investors are inclined to save through investing in stocks driving prices up and expected returns down. These findings constitute new evidence of time-varying risk premia which ties stock return predictability directly to fluctuations in consumption.[8]

B. Predicting stock returns in good and bad times

Some popular predictor variables are able to forecast returns in bad times as defined by recessions but not in good times, that is, during business cycle expansions (Rapach, Strauss, and Zhou (2010), Henkel, Martin, and Nardari (2011), Dangl and Halling (2012), and Golez and Koudijs (2018)). In light of this, Cujean and Hasler (2017) develop a theoretical mechanism with heterogeneous agents that causes recession-centric stock return predictability. Several other studies emphasize the usefulness of financial institutions and intermediation coupled with frictions and market segmentation since the 2007-2009 sub-prime financial crisis for rationalizing stock market behavior and capturing a propagation of a shock in bad times as opposed to normal and good times (see the discussion in Cochrane (2017)).

[8] As noted in Section I the benchmark results reported in the paper are based on a cyclical consumption measure computed from Equation (2) for $k = 24$. Table AIV in the internet appendix shows that the forecasting power of cyclical consumption is significant across various consumption horizons k ranging from one to eleven years ($k = 4, 8, ..., 44$). For any return holding period between one quarter and five years, the predictability is strongest at cycle lengths of five to six years (k values between 20 and 24).

The finding that returns are only predictable in bad times is a general concern for standard asset pricing models that emphasize the impact of time variation in risk premia as a common explanation of asset prices. To examine whether the relation between future returns and cyclical consumption is only present in bad economic times, we estimate a linear two-state predictive regression model in the spirit of Boyd, Hu, and Jagannathan (2005):

$$r_{t,t+h} = \alpha + \beta_{bad} I_{bad} cc_t + \beta_{good} \left(1 - I_{bad}\right) cc_t + \varepsilon_{t,t+h}, \tag{4}$$

where $r_{t,t+h}$ is the h-quarter ahead log excess return on the CRSP value-weighted index, I_{bad} is the state indicator that equals one during bad economic states and zero otherwise, and cc_t is one-quarter lagged cyclical consumption. Furthermore, β_{bad} and β_{good} denote the slope coefficients which measure the return predictability in bad and good states, respectively.

To evaluate the regression in Equation (4), we first follow Dangl and Halling (2012) and Henkel, Martin, and Nardari (2011) and employ the NBER-dated chronology of recessions for the identification of bad states. That is, the indicator variable I_{bad} takes on a value of unity during the NBER-dated recessions and zero otherwise. Panel A of Table III summarizes the results.

[Table III about here]

An important finding is that the predictive power of cyclical consumption is not confined to bad times alone. In particular, the results in Panel A of Table III indicate that cyclical consumption provides a consistent description of future stock returns both in good and bad economic states. In detail, at the one quarter horizon, the coefficient estimates in Panel A of Table III are -0.83 (with a t-statistic of -1.86) in bad times and -0.37 (with a t-statistic of -2.61) in good times, with bootstrap p-values indicating statistical significance at the 5% and 1% levels, respectively. To understand these units, note that a one-standard-deviation fall in cc in bad times leads to a rise in the expected excess return of approximately 3 percentage points at a quarterly horizon, roughly a 12-percentage-point increase at an annual rate. A corresponding change in annualized returns during good times amounts to slightly more than 5 percentage points. These estimates imply a total average reaction of future expected returns of close to 6.5 percentage points per annum. Differences in the level of statistical significance over bad and good times can be due to the fact that recessions are more infrequent than expansions (41 versus 243 data points in our seventy-year sample). These results are notable because they stand in marked contrast to several studies which document the presence of predictability in economic recessions and a lack of such in economic expansions. Cyclical consumption typically retains its significance at the various horizons that we consider. The $\bar{R}^2$ statistics in Panel A of Table III increase monotonically from 3.22% at a quarterly horizon to 35.01% at a horizon of five years.

To guard against the possibility that these results are due to the specific definition of a recession as identified by the NBER's Business Cycle Dating Committee, we next apply three alternative identifications of bad states. In particular, we follow Rapach, Strauss, and Zhou (2010) and measure bad states using the bottom third of sorted growth rates of real GDP in Panel B of Table III. We download the series of real seasonally adjusted GDP from the Federal Reserve Economic Data (FRED) of the Federal Reserve Bank of St. Louis. Panel C of Table III defines bad states as periods with the manufacturing purchasing managers index (PMI) issued by the Institute of Supply Management being below an optimal threshold value of 44.48 (Berge and Jordà (2011)). Finally, Panel D of Table III uses a further definition of bad states as periods when cyclical consumption is one standard deviation below its mean. During the full sample period, the four regime definitions classify 41, 94, 31 and 49 realizations as bad states, respectively.[9]

[9] To examine the robustness of the results, we experimented with a number of alternative definitions of bad states such as periods with cyclical consumption being 0.5, 1.5 or 2 standard deviations below its mean or periods with the lowest 5%, 10%, 15%, 20% or 25% of cyclical consumption realizations (in turn). We also defined bad regimes based on sorted values of real profit growth or real net cash flow growth as in Rapach, Strauss, and Zhou (2010), considered a measure of bad times from the Survey of Professional Forecasters (SPF) following Henkel, Martin, and Nardari (2011), and another one based on an unemployment recession

The results show that when we use the same measure of bad times as in the existing literature, there is consistent evidence of predictability both in bad and good times. The estimates of β_{bad} typically exceed the corresponding β_{good} counterparts (in absolute terms), but stock return predictability is not confined to relatively short recession periods alone. This result is in sharp contrast to the predictability pattern reported in, for example, Henkel, Martin, and Nardari (2011) and Dangl and Halling (2012) who find that return predictability is driven predominantly by rare recession periods.

These results indicate a strong predictive ability of cyclical consumption which is stable over time and across states of nature. This is a novel finding in the prevailing literature which documents that the forecasting power of many popular predictor variables is often concentrated in relatively short time spans of adverse macroeconomic changes.

C. Alternative detrending methods

Since there is no a priori theoretical guideline regarding the choice of an appropriate econometric procedure to isolate cyclical variation in consumption, it is instructive to compare the predictive ability of cc with other empirical measures of cyclical consumption. In the following, we consider five such definitions. First, we follow a voluminous literature in macroeconomics and finance and assume a secular linear upward trend in consumption:

$$c_t = d_0 + d_1 t + \omega_t, \tag{5}$$

where the residual measures cyclical consumption, cc. A second technique extends a linear trend formulation to allow for a breakpoint and hence makes it possible to account for a well-known fall in the macroeconomic risk, or the volatility of the aggregate economy, at the

gap of Stock and Watson (2010). Cyclical consumption generally emerges as a strong predictor of stock returns in both good and bad times.

beginning of the 1990s[10]:

$$c_t = \begin{cases} d_0 + d_1 t + \omega_t & \text{for } t \leq t_1 \\ d_0 + d_1 t + d_2\,(t - t_1) + \omega_t & \text{for } t > t_1, \end{cases} \tag{6}$$

where the breakpoint t_1 corresponds to the first quarter of 1992 (see also Lettau, Ludvigson, and Wachter (2008)). Essentially, Equation (6) presents a piecewise OLS regression which fits two separate lines to the disconnected data around the break date.

We also allow for higher order time polynomials such as a quadratic time trend model which conveniently accounts for slowly changing trends by establishing a quadratic exposure estimate d_2 that can intensify or diminish the linear time trend:

$$c_t = d_0 + d_1 t + d_2 t^2 + \omega_t, \tag{7}$$

and a corresponding cubic representation:

$$c_t = d_0 + d_1 t + d_2 t^2 + d_3 t^3 + \omega_t. \tag{8}$$

Finally, we follow Campbell (1991) and Hodrick (1992) and calculate a "stochastically detrended" consumption series as a backward-looking moving average based on a five-year window, where cc in quarter t is equal to the difference between the natural logarithm of consumption in quarter t and the average of the natural logarithm of consumption in quarters t-20 to t-1. The six measures of cyclical consumption that we identify display cross-correlations in the range of 0.34 to 0.91.

[Table IV about here]

Table IV reports estimation results for the predictive regression in Equation (3) based

[10] An extensive body of the macroeconomic literature finds evidence of a regime shift to lower volatility of real macroeconomic activity occurring in the last two decades of the 20th century (see, for example, McConnell and Perez-Quiros (2000) and Stock and Watson (2002)).

on alternative measures of cc. Cyclical consumption displays stable and robust predictive power regardless of how we detrend consumption. However, there are also some differences. Simple linear and quadratic trend specifications exhibit weaker long-run predictability, while the breaking and cubic detrending methods often yield stronger predictability than our benchmark results in Table I.

These results emphasize that our choice of the detrending procedure of Hamilton (2018) as a benchmark specification generally provides a conservative view of return predictability. Further, the question about which method should be employed to isolate cyclical variation in consumption appears largely irrelevant since all methods reveal substantial return predictability.

D. Temporal stability of estimates

Welch and Goyal (2008) highlight that many business cycle predictor variables have performed particularly poorly after the oil price crisis in the mid 1970s. To address this point, Table V reexamines the evidence of predictability over three subsamples: 1980-2017, 1990-2017, and 2000-2017. The results for the first two subsamples compare fairly with the full sample results in Table I. For the post-2000 sample period, we find systematically larger coefficient estimates, in absolute terms, and $\bar{R}^2$ statistics which are well beyond those reported in Table I. To provide an example, the β coefficients are -0.54 and -6.31 (t-statistics of -2.94 and -6.95) with $\bar{R}^2$ values of 7.35% and 55.64% for the S&P500 index at the one-quarter and twenty-quarter horizons, respectively, in the 2000-2017 sample, whereas the according estimates in the full sample are -0.43 and -5.33 (t-statistics of -3.28 and -4.28) with $\bar{R}^2$ statistics of 3.69% and 34.99%.

[Table V about here]

We obtain similar results in three other periods that we do not report in the table: in the post-1965 data (see also, Welch and Goyal (2008)); a period predating the global financial crisis; and a sample which omits the data in the aftermath of the run-up in prices in the early

2000s. We conclude that the predictive power of cyclical consumption fluctuations is not confined to any particular period and is not concentrated in sub-samples with severe crises, a pattern often found in the literature. This result is interesting in view of the fact that many traditional predictor variables tend to record a reduction in the extent of predictability in the data after the mid 1970s.

We also study the temporal stability of the β estimates in Equation (3) to structural breaks as prescribed by Elliott and Müller (2006). Their proposed $\widehat{qLL}$ test statistic for the hypothesis that $\beta_t = \beta$ for all t and any h is particularly useful in the context of predictive regressions because it is asymptotically efficient for a wide range of data-generating processes, has superior size properties in small samples than other popular statistics, and is simple to construct. Moreover, the simulation analysis in Paye and Timmermann (2006) shows that the test of Elliott and Müller (2006) possesses excellent finite sample size properties even in the presence of highly persistent lagged endogenous predictors. Table AV in the internet appendix documents that the $\widehat{qLL}$ statistics for our benchmark estimates in Table I are never significant at any horizon (we find similar results for subsamples). These results emphasize a stable relation between consumption fluctuations and future expected stock returns.

E. Out-of-sample analysis

Bossaerts and Hillion (1999) and Welch and Goyal (2008) point out that in-sample predictability of stock returns is not necessarily robust to out-of-sample validation and therefore in-sample predictability does not generally indicate that it is possible to obtain reliable out-of-sample forecasts. Reasons why out-of-sample results might differ from in-sample results include effects from loss of information when splitting up samples in out-of-sample tests, structural breaks, and parameter uncertainty (see, for example, Inoue and Kilian (2004), Paye and Timmermann (2006), Lettau and Van Nieuwerburgh (2008), and Cochrane (2008)). Furthermore, Nagel and Xu (2018) show that in models with learning, the presence of in-sample predictability does not necessarily imply that out-of-sample predictability will also be evident.

There are at least two possible interpretations of out-of-sample tests. One view of out-of-sample testing is that it is a means of helping to validate the in-sample relations. An alternative view of out-sample testing is that it is a way to assess whether a savvy investor could construct a real time trading strategy.[11]

On the one hand, we could assume that the econometrician has a limited information set relative to the economic agent. Unlike the econometrician, the economic agents are aware of the history of past consumption and its relation to the consumption trend, and hence they know how stock returns react to the real quantities. This perspective is termed the "economic agents knew" framework. On the other hand, we could adopt a "savvy investor" framework where the agent waits until consumption is reported, often with a lag, and from there calculates cyclical consumption based on real-time data to form a forecast of the next period return in order to subsequently trade.

The framework one chooses has implications on how the out-of-sample tests are conducted. Since our main aim is to examine the validity of the in-sample evidence on predictability of future returns by cyclical consumption, rather than to create a trading strategy, we follow the "economic agents knew" framework. This allows us to use current, that is, revised latest-available consumption data and a one period lag in the predictability regression.

We proceed as follows. First, using the revised consumption data, we recursively estimate cyclical consumption every quarter using data available at the time of the forecast. Then we employ these values of cc in recursive predictive regressions for stock returns to form out-of-sample forecasts. We use an expanding estimation window where the coefficients in the return forecasting regression are estimated recursively using only information available through time t for forecasting over the next h quarters. To ensure that our results are not sensitive to the choice of the evaluation period, we perform out-of-sample tests for three

[11] We thank an anonymous referee for making this distinction between the two perspectives for interpretation of out-of-sample predictability tests.

different out-of-sample forecasting periods: 1980Q1-2017Q4, 1990Q1-2017Q4, and 2000Q1-2017Q4.[12]

For nested forecast comparison tests, we specify a model of constant expected returns, that is, a benchmark model where a constant is the sole explanatory variable. The constant expected return model is a restricted nested version of an unrestricted model of time-varying expected returns, which includes both a constant and cc. To this end, we compare the forecasting error from a series of out-of-sample return forecasts obtained from a prediction equation that includes a constant and cc (the unrestricted model), to a prediction equation that includes a constant as the sole forecasting variable (the restricted model). For example, Welch and Goyal (2008) show that the historical average forecast is a very stringent out-of-sample benchmark.

E.1. Baseline out-of-sample results

In Table VI, we show results of out-of-sample predictions of the log excess return on the CRSP value-weighted index over various horizons ranging from one quarter to five years. We find that the unrestricted model typically generates significantly better forecasts than the restricted model. For instance, the ENC-NEW test of Clark and McCracken (2001) rejects the null hypothesis that the forecasts from the constant expected return model encompass the forecasts from the time-varying expected return model at the 1% level for all horizons and all forecasting periods that we consider. The MSE-F test of McCracken (2007) rejects the null hypothesis that the mean squared errors from the unrestricted model are bigger than or equal to those from the historical average return.

[Table VI about here]

[12] Starting the out-of-sample evaluation in 1980 provides a reasonably long initial in-sample period for reliably estimating the parameters used to generate the first predictive regression forecast. This issue is of particular relevance for us because a consistent estimation of the trend parameters in cc requires a large number of observations.

The out-of-sample R^2_{OOS} statistics in Table VI are all positive, meaning that cc systematically delivers a lower average forecasting error than the historical average forecast. For example, at the one-quarter horizon, the R^2_{OOS} is 0.64% (significant at the 10% level) when we forecast from 1980, 1.96% (significant at the 10% level) when we forecast from 1990, and 1.35% (albeit insignificant) when we forecast from 2000. It is instructive to compare these measures of fit with corresponding R^2 statistics from in-sample regressions in Table V of the order of 1.86% for the post-1980 period, 2.96% for the post-1990 period, and 5.56% for the post-2000 period. Consistent with Bossaerts and Hillion (1999) and Welch and Goyal (2008), we find a lower out-of-sample fit for every forecast evaluation period that we consider at a horizon of one quarter.

At horizons greater than one quarter, the R^2_{OOS} statistics are all statistically significant. They are often close to, but remain systematically below, their in-sample counterparts both in the early 1980-2017 and in the late 2000-2017 evaluation periods. In the post-1990 sample, the out-of-sample R^2 estimates are less than the corresponding in-sample measures of fit for time horizons of up to two years while the reverse holds true for longer-term returns at horizons of between three and five years.

E.2. Additional out-of-sample results

In a robustness test, we follow Lettau and Ludvigson (2001) and consider a scenario where the predictive regression is estimated recursively each time a forecast is made but the parameters in cc are fixed at their values estimated over the full sample. This technique might be advantageous because it does not induce a sampling error in the estimation of parameters in cc, especially in the early estimation recursions. Table AVII in the internet appendix shows that using full sample estimates to measure cc often leads to stronger out-of-sample predictive power (exceptions include the results for longer-horizon returns in the post-1980 sample). Overall, this suggests that the reestimation of the parameters in cc induces sampling error into the parameter estimates, which may lead to less accurate forecasts.[13]

To summarize, our results show that cyclical fluctuations in consumption that we identify display statistically significant out-of-sample predictive power for aggregate stock market returns. This is the case when the out-of-sample forecasting starts in 1980, 1990, or 2000. These results are in contrast to Welch and Goyal (2008) who accentuate that a long list of popular business cycle predictor variables have been unsuccessful out-of-sample in the last few decades, an issue we return to in Section IV.

III. Further robustness tests and extensions

In this section, we investigate the predictive ability of cyclical consumption for stock returns sorted into industry portfolios, explore the robustness of our results to alternative ways of defining consumption, and examine international evidence.

A. Industry portfolios

In the preceding analysis, we have assessed the predictability of stock returns by means of two commonly used stock market indices that give a broad view of the behavior of the aggregate equity premium. In this section, we investigate how well cyclical consumption can forecast returns on portfolios of stocks sorted on industry SIC codes.[14]

[Table VII about here]

Table VII reports the estimation results from univariate predictive regressions for each of the ten industry portfolios. In line with our results for the total market portfolios in Table I, cyclical consumption emerges as a powerful predictor of a cross section of industry returns. The inverse relation between cyclical consumption and future expected returns is

[13] We also investigate the out-of-sample predictive power of *cc* using real-time data instead of revised data. The results are in Table AVIII in the internet appendix and they are largely consistent with our benchmark findings in Table VI.

[14] The portfolio data are from Ken French's online data library.

visible for all the industry portfolios. The results are strongly significant (usually at the 1% significance level) across all industries, apart from the energy category. In addition, we find that the regression slopes and $\bar{R}^2$ statistics vary across industries, illustrating cross-sectional differences in the sensitivities. In particular, returns on durable goods and hi-tech business equipment have the highest level of predictability. Overall, the results in Table VII reinforce our conclusion that cyclical consumption fluctuations predict stock returns and emphasize that time-varying expected rates of return contain a common macroeconomic component.

B. Alternative consumption measures

Thus far, our main empirical analysis has focused on real per capita NIPA expenditure on nondurable goods and services as a proxy of aggregate consumption. In this section, we consider the predictive ability of cyclical consumption extracted from various subcategories of personal consumption expenditure (PCE), including i) nondurable goods (NON); ii) services (SERV); iii) durable goods (DUR); iv) the stock of durable goods (SDUR) constructed from the year-end estimates of the chained quantity index for the net stock of consumer durable goods published by the Bureau of Economic Analysis (BEA) following Yogo (2006); v) nondurable and durable goods (GOODS); and vi) total PCE.

[Table VIII about here]

Table VIII shows results from the benchmark regression (3) applied to the log excess return on the value-weighted CRSP index. The predictive power of cyclical consumption is generally qualitatively similar in terms of coefficient magnitudes, statistical significance, and $\bar{R}^2$ measures across the six different expenditure aggregates that we consider. According to the $\bar{R}^2$ statistics, nondurable goods emerge as the strongest predictor of stock returns with $\bar{R}^2$ values of 3.18% and 45.42% for quarterly and five-year returns, respectively. However, the extent of predictability is very similar across consumption categories except those that involve durables where the extent of predictability is weaker, in particular at the one quarter horizon. It is interesting to note that the predictive ability of the aggregate consumption proxy in Table I compares fairly to that of nondurables and services, measured separately. Note also that at horizons of three years and above, we often find stronger results based on alternative PCE categories in Table VIII than in Table I. This evidence reinforces our main findings and further highlights the conservative nature of our benchmark results.

C. International evidence

To mitigate concerns regarding over-fitting or "data snooping" (Lo and MacKinley (1990) and Bossaerts and Hillion (1999)), we investigate the predictability of stock returns in international equity markets. We follow Ang and Bekaert (2007), Hjalmarsson (2010), and Rapach, Strauss, and Zhou (2013) and collect international total return indices in national currency from Morgan Stanley Capital International (MSCI) recorded since the beginning of 1970. We consider seven major developed market regions around the world, including the MSCI World, the MSCI World ex USA, the MSCI EAFE, the MSCI Europe, the MSCI Pacific, the MSCI Far East, and the MSCI G7 indices.[15] We focus our attention on actual returns because appropriate proxies for regional market risk-free rates and inflation rates are not available. Similar results are obtained for returns denominated in U.S. dollars, excess returns computed by subtracting the U.S. Treasury bill rate as a proxy for the world risk-free rate, and real returns computed by subtracting the U.S. CPI inflation rate as a proxy for the global inflation rate.

In what follows, we examine whether a global measure of cyclical variation in consumption reveals significant predictive power for future stock returns around the world. This approach

[15] The MSCI World equity index consists of 23 developed market countries including Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Hong Kong, Ireland, Israel, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, the United Kingdom, and the United States. The World index covers approximately 85% of the free float-adjusted market capitalization in each country. The MSCI EAFE index represents 21 developed market countries, not including the United States and Canada. The MSCI Europe consists of 15 major developed European countries. The MSCI Pacific index consists of 5 developed market countries, including Australia, Hong Kong, Japan, New Zealand, and Singapore, and the MSCI Far East index includes Hong Kong, Japan, and Singapore.

is motivated by the fact that world-wide rather than local fluctuations in the business cycle have gained on importance over recent decades (Lumsdaine and Prasad (2003) and Kose, Otrok, and Whiteman (2003)). To the extent that a global cyclical consumption component can adequately capture common business cycle related risks, our analysis contributes to the debate about the level of integration in financial markets (Pukthuanthong and Roll (2009), and Rangvid, Santa-Clara, and Schmeling (2016)).

To this end, we compute a global measure of cyclical consumption as a simple arithmetic average of country-specific cyclical consumption components. The latter are obtained by fitting the regression in Equation (1) to the logarithm of real seasonally-adjusted consumption expenditures in 20 developed market countries from the MSCI World index for which consumption data is available from the OECD database over the full sample period (not including Hong Kong, Israel, and Singapore).[16]

[Table IX about here]

The results for international predictability are reported in Table IX. We find a stable negative relation between cyclical consumption and future stock returns. This relation is always economically and statistically significant. In economic magnitudes, the international estimates imply an even stronger impact of cyclical consumption on expected returns than our benchmark findings for the United States. For instance, we find that a fall in the global cyclical consumption by one standard deviation below its mean would lead to a rise in the expected return on the MSCI World index of the order of about 7.5 percentage points per annum. The corresponding Newey and West (1987) t-statistic of -3.59 and the bootstrap p-value indicate significance at the 1% level. Variation in cyclical consumption accounts for 5.32% of the variation in the quarterly world market return. Cyclical consumption retains

[16] We have considered a number of alternative global measures of cyclical consumption such as a GDP-weighted average or the first principal component of the national cyclical consumption series. We also experimented with consumption data for the G7 countries only and found generally similar conclusions.

its predictive power at any return horizon that we consider with associated $\bar{R}^2$ statistics climbing to 46.75%, 53.97%, and 51.86% for $h = 12$, 16, and 20 quarters, respectively.

The results are broadly similar across the different regions with predictability, perhaps not surprisingly due to the more recent sample period, being strongest at the G7 level and weakest for the Pacific region. The consistency of the estimated sign, its size, and the statistical significance provides evidence that cyclical consumption is useful in tracking future movements in global equity returns. These results are in line with our benchmark findings and they suggest that our main results are specific to the U.S. stock market.

IV. Alternative predictor variables

How does the predictive information contained in cyclical consumption compare to other well known predictor variables that have been rationalized by their ability to track business cycle conditions? To address this question, we consider a set of out-of-sample tests with alternative business cycle variables that have been used in the extant literature. The forecasting variables that we consider include the fifteen predictors studied by Welch and Goyal (2008),[17] the share of labor income to consumption (s^w) of Santos and Veronesi (2006), the consumption-wealth ratio (cay) of Lettau and Ludvigson (2001), consumption volatility (σ_c) of Bansal, Khatchatrian, and Yaron (2005), and the output gap (gap) of Cooper and Priestley (2009).

We use revised macroeconomic data to compute s^w, cay, σ_c, gap, and cc. We compute the share of labor income to consumption following Santos and Veronesi (2006) using the definition of labor income in Lettau and Ludvigson (2001). The data for total personal consumption expenditures, labor income and asset wealth that are used to compute the consumption-wealth ratio are downloaded from the website of Martin Lettau. We calculate consumption volatility as $\sigma_{c,t-1,J} \equiv \log\left(\sum_{j=1}^{J} \left|\eta_{c,t-j}\right|\right)$, where $\eta_{c,t}$ is the residual from an

[17] The source of these data is the online library of Amit Goyal.

AR(1) process of log growth rate in real per capita nondurables and services and $J = 4$ quarters following Bansal, Khatchatrian, and Yaron (2005). The output gap is constructed from the industrial production data available at the Federal Reserve Bank of St. Louis following Cooper and Priestley (2009).

This gives us a total of nineteen alternative predictor variables:

1. Log dividend-price ratio (dp): log of a 12-month moving sum of dividends paid on the S&P 500 index minus the log of prices on the S&P 500 index.

2. Log dividend yield (dy): log of a 12-month moving sum of dividends paid on the S&P 500 index minus the log of lagged prices on the S&P 500 index.

3. Log earnings-price ratio (e/p): log of a 12-month moving sum of earnings on the S&P 500 index minus the log of prices on the S&P 500 index.

4. Log dividend-payout ratio (d/e): log of a 12-month moving sum of dividends minus the log of a 12-month moving sum of earnings on the S&P 500 index.

5. Stock variance ($svar$): sum of squared daily returns on the S&P 500 index.

6. Book-to-market ratio (b/m): ratio of book value to market value for the Dow Jones Industrial Average.

7. Net equity expansion ($ntis$): ratio of a 12-month moving sum of net equity issues by NYSE-listed stocks to the total end-of-year market capitalization of NYSE stocks.

8. Treasury bill rate (tbl): interest rate on a three-month Treasury bill (secondary market).

9. Long-term yield (lty): long-term government bond yield.

10. Long-term return (ltr): return on long-term government bonds.

11. Term spread (tms): long-term yield on government bonds minus the Treasury bill rate.

12. Default yield spread (dfy): difference between the BAA- and AAA-rated corporate bond yields.

13. Default return spread (dfr): long-term corporate bond return minus the long-term

government bond return.

14. Inflation (*infl)*: calculated from the Consumer Price Index (CPI) for all urban consumers.

15. Investment-to-capital ratio (i/k): log ratio of aggregate private nonresidential fixed investment to aggregate capital for the whole economy (Cochrane (1991)).

16. Share of labor income to consumption (s^w): ratio of the compensation of employees to the consumption of nondurables plus services (Santos and Veronesi (2006)).

17. Consumption-wealth ratio (cay): residual from a cointegrating relation between log consumption, log asset (nonhuman) wealth, and log labor income (Lettau and Ludvigson (2001)).

18. Consumption volatility (σ_c): log of a backward-looking moving average of the absolute innovations in consumption growth based on a four-quarter window (Bansal, Khatchatrian, and Yaron (2005)).

19. Output gap (gap): residual from a regression of log of industrial production on a time trend which contains linear and quadratic components (Cooper and Priestley (2009)).

We employ a recursive out-of-sample methodology as in Section II.E. to calculate equity premium forecasts for each predictor. We use the 1953Q4-1979Q4 sample as the initial estimation period and expand it by one quarter in each recursion. The forecasting ability is evaluated by means of the out-of-sample R^2 statistic (R^2_{OOS}).[18]

[Table X about here]

Table X presents results of out-of-sample horse races pitting the forecasts for each predic-

[18] Table AI in the internet appendix provides descriptive statistics of the predictor varaibles, and Table AII shows that the in-sample predictive power of cc compares favorably with that of standard business cycle predictor variables. Out of the 19 alternative economic predictors that we consider, only four variables including the term spread (tms), the investment-to-capital ratio (i/k) of Cochrane (1991), the consumption-wealth ratio (cay) of Lettau and Ludvigson (2001), and the output gap (gap) of Cooper and Priestley (2009) exhibit significant and strong predictive ability for stock returns. The remaining variables are typically insignificant at the 5% level and/or generate low R^2s.

tor variable against the historical average return benchmark forecast. The overall picture is that the traditional predictive variables have rather weak out-of-sample predictive power. At a horizon of one quarter, eighteen out of nineteen alternative predictors generate a negative R^2_{OOS} statistic and thus fail to outperform the historical average forecast. This result echoes the message of Welch and Goyal (2008) that many economic variables deliver a very erratic out-of-sample performance in the period after the oil price shocks of the 1970s. In marked contrast to this observation, the predictive power of cc clearly stands out. The quarterly R^2_{OOS} statistic for cc is positive at 0.64% (significant at the 10% level), meaning that unlike many popular predictors, cc outperforms the prevailing mean benchmark and clears the out-of-sample hurdle.

A similar picture prevails at longer horizons. For example, at horizons of one, two, and three years, we register negative R^2_{OOS} statistics for seventeen out of nineteen alternative predictor variables (the two exceptions with positive R^2_{OOS}'s are tms and i/k), whereas cc generates positive R^2_{OOS} statistics of 4.14%, 10.55%, and 17.63% (significant at a 1% level), respectively. Overall, for return holding periods of up to three years, cc emerges as the most powerful predictor in our sample. At horizons of four and five years, i/k is the only variable that yields R^2_{OOS} statistics which slightly exceed those produced by cc.

To summarize, the results in Table X reinforce a stable and strong predictive ability of cyclical consumption relative to numerous popular business cycle predictors. We find that none of the nineteen alternative traditional predictor variables usually considered in the literature can systematically generate better out-of-sample forecasts of the equity premium than cyclical consumption. These results attempt to address a concern of Welch and Goyal (2008) who demonstrate that it is very difficult to identify individual economic variables capable of generating reliable out-of-sample forecasts. Against this backdrop, we show that cyclical consumption outperforms the historical average by meaningful margins and generates better forecasts than popular forecasting variables.

V. The external habit model

The predictive regression analysis in Section II documents an inverse relation between cyclical consumption and future expected stock returns: A fall (rise) in consumption below (above) trend indicates bad (good) times when marginal utility of current consumption and future expected returns are high (low). A natural question is how this empirical evidence relates to consumption-based asset pricing theory which aims to explain the dynamic behavior of asset returns using aggregate consumption data.[19]

Campbell and Cochrane (1999), for example, assume that investors evaluate current consumption relative to a habit level of consumption that can be thought of as a weighted moving average of past consumption expenditures.[20] In their model, habit acts as a trend for consumption: a decline in consumption relative to the trend in a recession leads to high expected returns and low asset prices. This begs a question about how our detrended consumption variable relates to consumption habit, and what restrictions such a relation may impose on the consistency of our choice of the cycle parameter k in the Hamilton (2018) filter with respect to return predictability. To address these issues, we study the implications of the habit model of Campbell and Cochrane (1999) for time-varying expected returns. Specifically, Section V.A. sets down a simple economic framework wherein cyclical consumption emerges as a relevant predictor variable for future stock returns. Section V.B.

[19]Countercyclical variation in risk premia has been incorporated in prominent equilibrium models which can generate time-varying expected returns, including models with time-varying risk aversion (Campbell and Cochrane (1999)), time-varying aggregate consumption risk (Bansal and Yaron (2004)), and time-varying disaster risk (Farhi and Gabaix (2016) and Wachter (2013)).

[20]The empirical analysis in Bansal, Kiku, and Yaron (2012) and Beeler and Campbell (2012) points to an important distinction between the habit model of Campbell and Cochrane (1999) and the long-run risks model of Bansal and Yaron (2004). Specifically, the long-run risks model implies that past or current consumption cannot explain future dividend-price ratios or returns, while the habit model in contrast suggests that asset prices are backward-looking and that past consumption growth forecasts future price-dividend ratios and returns.

discusses details on the calibration and Section V.C. examines the compatibility of our results with the predictability generated by the model.

A. Consumption habit, cyclical consumption, and time-varying returns

Campbell and Cochrane (1999) augment the standard power utility function with a time-varying subsistence level X_t, which represents the agent's "external habit" and is defined indirectly through the surplus consumption ratio $S_t \equiv \frac{C_t - X_t}{C_t}$. To ensure stationarity and prevent habit from falling below consumption, Campbell and Cochrane (1999) assume that the log surplus consumption ratio, $s_t \equiv \log(S_t)$, follows a mean-reverting heteroscedastic first-order autoregressive process:

$$s_{t+1} = (1-\phi)\,\overline{s} + \phi s_t + \lambda(s_t)\,\upsilon_{t+1}, \tag{9}$$

where $\overline{s}$ is the steady state value of s_t, ϕ is the habit persistence parameter, and $\lambda(s_t)$ is a nonlinear monotonically decreasing sensitivity function that determines how innovations in consumption growth υ_{t+1} influence s_{t+1}. Surplus consumption is the only state variable in the model, and it controls the price of risk and generates time-variation in expected returns.

Appendix C in the working paper of Wachter (2006) formally demonstrates that a first-order approximation around $s_t = \overline{s}$ implies that surplus consumption adjusts gradually to the history of current and past consumption with coefficient ϕ:

$$s_t \approx \kappa + \lambda(\overline{s}) \sum_{j=0}^{\infty} \phi^j \Delta c_{t-j}, \tag{10}$$

where κ is a constant depending on model parameters. The model requires a high, but less than unity, value of ϕ to match stock market data. While surplus consumption is, in the theory, influenced by past consumption going back to infinity, a "cut-off" horizon can be used to obtain an empirical proxy for s_t. If we omit the constant κ and the proportionality parameter $\lambda(\overline{s})$, and assume a close to unity value of the persistence parameter ($\phi \approx 1$), it follows that there exists a close link between a finite-horizon proxy of surplus consumption and cyclical consumption:

$$\hat{s}_t \approx c_t - c_{t-k} \approx cc_t, \tag{11}$$

where k determines how long habit reacts to past consumption. The second approximation in (11) follows from the fact that under the random walk hypothesis for consumption, Hamilton's (2018) detrending procedure reduces to a difference filter because, for large samples, the OLS estimates in Equation (1) converge to $b_1 = 1$ and all other $b_j = 0$. The resulting cyclical component is then simply given by the difference over a k-quarter horizon, or, equivalently, the sum of the observed changes over k periods.

If excess returns on the stock market and consumption growth are jointly conditionally lognormally distributed, the Campbell and Cochrane (1999) model implies:

$$E_t\left(r_{t+1}\right) + \frac{1}{2}\sigma_t^2 = \gamma_t cov_t\left(r_{t+1}, \Delta c_{t+1}\right), \tag{12}$$

where $E_t\left(r_{t+1}\right)$ is the expected log excess stock return, γ_t is the state-dependent price of consumption risk defined as $\gamma_t = \gamma\left(1 + \lambda\left(s_t\right)\right)$, $cov_t\left(r_{t+1}, \Delta c_{t+1}\right)$ is the amount of risk, and $\frac{1}{2}\sigma_t^2$ is a Jensen's inequality term. Since $\lambda\left(s_t\right)$ is inversely related to s_t, and cc_t and s_t are tightly linked as they both depend on past consumption growth, it follows that low levels of cyclical consumption increase γ_t and forecast high expected returns. This prediction turns out consistent with the empirical evidence that we presented in Section II. The inverse relation between s_t, and therefore cc_t, and risk premia operates also via the conditional covariance term in Equation (12) because a fall in consumption toward the habit in bad times is associated with a rise in $cov_t\left(r_{t+1}, \Delta c_{t+1}\right)$ in the model. Overall, the dependence of expected returns on surplus consumption is close to linear as illustrated in Figure 4 in Campbell and Cochrane (1999), except for very low values of the surplus consumption ratio.[21]

[21] Note also that the effect of the Jensen's inequality term is quite small empirically. For instance, we obtain a coefficient of -0.41 (t-statistic of -3.07) and an $\bar{R}^2$ measure of 3.11% in the benchmark predictive

B. Parameter calibration

We proceed by conducting a simple simulation study with two purposes. First, we would like to evaluate the extent to which the habit model of Campbell and Cochrane (1999) can match the time-series predictability of stock returns in the data. Secondly, we aim to analyze the frequency at which consumption innovations are related to expected returns. We employ Campbell and Cochrane's (1999) parameter values to simulate a sample path of 1,000,000 quarters of artificial data for returns on stocks and consumption growth from the model.[22] We then calculate population values for a variety of statistics. Table XI compares simulated means and standard deviations implied by the model to corresponding statistics in our empirical sample over the 1947Q1-2017Q4 period. Following Campbell and Cochrane (1999), we set the average log consumption growth at 1.89% and its standard deviation at 1.50%, which should be measured against the values of 1.89% and 1.00% in our empirical sample (all terms per annum). As for the further parameter choices, the utility curvature parameter is set at $\gamma = 2.00$, the persistence parameter of the log surplus consumption ratio at $\phi = 0.87$, and the subjective discount factor at $\delta = 0.89$. As Table XI shows, we can match quite closely the means and standard deviations of consumption growth and excess returns as well as their Sharpe ratios, using either the consumption claim or dividend claim to model the market return.[23]

[Table XI about here]

regression for quarterly excess CRSP returns in Equation (3) compared to the estimate of -0.42 (t-statistic of -3.06) and an $\bar{R}^2$ statistic of 3.16% for log excess CRSP returns reported in Table I.

[22] Campbell and Cochrane (1999) pick parameters by calibrating the model to match certain moments in the post-war U.S. data over the 1947-1995 period. Our 1947-2017 sample displays similar statistics as in Table 2 in Campbell and Cochrane (1999). For instance, the mean of the log consumption growth is 1.89 in both samples; the standard deviation of the log consumption growth is 1.50 and 1.00; the Sharpe ratio for log stock returns is 0.43 and 0.40; the Sharpe ratio for simple returns is 0.50 and 0.49; the average of the log excess stock returns is 6.64 and 6.52; and the standard deviation of log stock returns is 15.20 and 16.38 in percent per annum in the two samples, accordingly.

[23] See the Internet Appendix for further details about the simulations.

C. Implications for stock return predictability

In this section, we use our simulations to investigate the model's ability to reproduce the results in Section II, that is, an inverse relation between cyclical consumption and future expected stock returns. In addition, the simulations allow us to assess the impact of cycle length on asset prices and study the frequency at which consumption innovations are incorporated into risk prices.

[Table XII about here]

To begin with, we analyze the population properties of the model for in-sample predictability of stock returns. Based on a sample path of 1,000,000 quarterly simulations, we compute a long series of artificial realizations of log excess returns and our cyclical consumption variable as defined in (11) for $k = 24$.[24] We then examine the extent of model-implied predictability by estimating the standard predictive regression (3). The upper row of Table XII, Panel A displays the estimates in simulated data implied by the consumption claim; the middle row presents simulation results for the dividend claim; and the entries in the bottom row replicate the findings in the historical data (see also, Table I). We show OLS estimates of slope coefficients and adjusted $\bar{R}^2$ statistics in percent in square brackets. We do not report t-statistics of the simulation results because the large sample size makes them meaningless.

As one can see, cyclical consumption can predict returns at various horizons ranging from one quarter to five years. The slope coefficients have the right (negative) sign in the model as in the empirical regressions. The model's predictions for the consumption claim are somewhat weaker compared to the actual data but the general patterns are similar. In the model, the slope coefficients increase (in absolute terms) from -0.22 at a horizon of one quarter to -3.04 at a horizon of five years, for the consumption claim, compared with values

[24] We also examined predictability in the historical data for a measure of cyclical consumption computed from a difference filter in (11) and obtained results which were similar to our benchmark findings.

of -0.42 and -5.01 in the actual data.[25] For the dividend claim, the predictive coefficients on cyclical consumption are similar to the consumption claim, but the adjusted $\bar{R}^2$ statistics are reduced. For instance, the population $\bar{R}^2$'s for the consumption claim rise from 1.75% at a one-quarter horizon to 23.80% at the five-year horizon, whereas with values of 1.46% and 18.13%, respectively, the measures of regression fit are lower using the dividend claim (see also, Campbell and Cochrane (1999)). For comparison, the corresponding $\bar{R}^2$'s in the actual postwar sample are of around 3% for quarterly returns and close to 35% for the five-year returns.

We turn next to an assessment of the out-of-sample predictability in the model by generating 2,500 artificial samples of size 284, which matches the number of observations in our postwar historical sample. For each artificial sample, we then compute 112-(h-1) out-of-sample forecasts, which matches the number of forecasts in the 1990-2017 evaluation window. The upper row of Table XII, Panel B shows average out-of-sample R^2 values across the 2,500 artificial samples for the consumption claim; the middle row gives corresponding statistics for the dividend claim; and the bottom row replicates the results in the historical data (see also, Table VI). For the consumption claim, we find that cyclical consumption consistently outperforms the historical mean in forecasting returns out-of-sample for each forecast horizon h, with out-of-sample R^2 measures rising from 1.00% for quarterly returns to 9.51% for five-year returns. The out-of-sample R^2 values with the dividend claim are genereally lower across horizons but still positive. Overall, as expected, we see more predictability in the historical sample compared to the model out-of-sample.

Finally, Table AX in the internet appendix examines the impact of cycle length on the model fit. For consistency with Table AIV, we study a total of eleven specifications with consumption cycles k varying from 4 to 44 quarters. Similar to the patterns in the historical

[25] Table AIX in the internet appendix shows that cyclical consumption displays a comparable degree of volatility and autocorrelation in actual and simulated data. For example, the standard deviation and first-order autocorrelation are 3.74% and 0.97 in actual data compared to 3.66% and 0.96 in simulated data.

sample, the simulations indicate that the predictive power of cc is increasing in k up to cycle lengths of around five or six years, almost stagnating for values of k between six and eight years, and slightly deteriorating thereafter. The predictive coefficients on cyclical consumption are similar for both claims, but the $\bar{R}^2$ statistics tend to be lower when using the dividend claim compared to the consumption claim. The findings generally indicate that returns adjust to changing economic conditions at frequencies of around five to six years, giving a macroeconomic, consumption-related foundation for the existence of risk determinants in asset prices which are due to low frequency dynamics. This insight appears consistent with our choice of a six-year cycle in the benchmark application of Hamilton's (2018) filter in Section II.[26]

In summary, we make two main points. Our first point is that simulated data from the Campbell and Cochrane (1999) model produce an inverse relation between cyclical consumption and future expected stock returns as we find in the empirical data. The model generates qualitatively similar patterns but yields lower predictability, especially for longer horizon returns. Our second point is that values of around five to six years are optimal for the parameter k in the Hamilton (2018) filter in terms of capturing predictable variation in expected returns.

VI. Conclusion

The predictability of stock returns has been rationalized as evidence suggesting the existence of time-variation in expected risk premia. Common predictor variables have, however, for the most part been unsuccessful in establishing a sound relation to fundamentals (Cochrane (2005), Welch and Goyal (2008), and Henkel, Martin, and Nardari (2011)).

[26]While cc captures a slow-moving business cycle related risk premium component, cay is less persistent and works particularly well at business cycle frequencies of relatively short horizons as shown by Lettau and Ludvigson (2001). In our sample, the half-life of cc is about 5 years, while that of cay is about 2 years, illustrating that cc works at a lower frequency than cay.

In this paper, we propose a novel consumption-based predictive variable, called cyclical consumption, and show that it captures a significant fraction of variation in expected stock returns. To identify a cyclical component of consumption, which measures deviations of aggregate consumption from its trend, we employ the robust linear projection method of Hamilton (2018). We document a robust inverse relation between cyclical consumption and future expected returns: When economic conditions deteriorate, often referred to as bad times, consumption drops below its trend, leading to a rise in the marginal utility of current consumption. As a consequence, prices fall and future expected returns rise. Conversely, in good times when consumption rises above trend and marginal utility from consumption is low, prices rise and future expected returns fall. The empirical evidence we find ties consumption decisions of agents to time-variation in expected returns in a manner consistent with rational asset pricing and suggests that stock return predictability arises as a rational response to changing business conditions.

Our findings are supportive of theoretical explanations of asset prices which emphasize the role of habit formation in consumption such as Campbell and Cochrane (1999). Using simulations, we show that the habit model produces a similar inverse relation between expected returns and cyclical consumption. Our analysis emphasizes that low frequency fluctuations in consumption capture slow-moving countercyclical variation in expected returns.

We conduct a battery of robustness checks and conclude that the predictive power of cyclical consumption is higher than that of many well-recognized forecasting variables, is stable over time, not confined to bad times or times of crises, is evident in industry portfolio and international data, and is robust to a variety of alternative specifications and methods applied to isolate cyclical variation in consumption. Taken together, our evidence lends support to asset pricing models based on external habit formation where the dynamics of expected returns are driven by changes in the level of current consumption relative to its past history.

References

Ang, Andrew, and Geert Bekaert, 2007, Stock return predictability: Is it there? *Review of Financial Studies* 20, 651-707.

Bansal, Ravi, Varoujan Khatchatrian, and Amir Yaron, 2005, Interpretable asset markets? *European Economic Review* 49, 531-560.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An empirical evaluation of the long-run risks model for asset prices, *Critical Finance Review* 1, 183-221.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.

Beeler, Jason, and John Y. Campbell, 2012, The long-run risks model and aggregate asset prices: An empirical assessment, *Critical Finance Review* 1, 141-182.

Berge, Travis J., and Òskar Jordà, 2011, Evaluating the classification of economic activity into recessions and expansions, *American Economic Journal: Macroeconomics* 3, 246-277.

Van Binsbergen, Jules H., and Ralph S.J. Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439-1471.

Bossaerts, Peter, and Pierre Hillion, 1999, Implementing statistical criteria to select return forecasting models: What do we learn? *Review of Financial Studies* 12, 405-428.

Boyd, John H., Jian Hu, and Ravi Jagannathan, 2005, The stock market's reaction to unemployment news: Why bad news is usually good for stocks, *Journal of Finance* 60, 649-672.

Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157-179.

Campbell, John Y., and James H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205-251.

Campbell, John Y., and Motohiro Yogo, 2006, Efficient tests of stock return predictability, *Journal of Financial Economics* 81, 27-60.

Chien, YiLi, Harold Cole, and Hanno Lustig, 2016, Implications of heterogeneity in preferences, beliefs, and asset trading technologies in an endowment economy, *Review of Economic Dynamics* 20, 215-239.

Clark, Todd E., and Michael W. McCracken, 2001, Tests of equal forecasting accuracy and encompassing for nested models, *Journal of Econometrics* 105, 85-110.

Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291-311.

Cochrane, John H., 1991, Production-based asset pricing and the link between stock returns and economic fluctuations, *Journal of Finance* 46, 207-234.

Cochrane, John H., 2005, Financial markets and the real economy, in Richard Roll, ed.: *The International Library of Critical Writings in Financial Economics* (Edward Elgar Series).

Cochrane, John H., 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533-1575.

Cochrane, John H., 2017, Macro-finance, *Review of Finance* 21, 945-985.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2016, Parameter learning in general equilibrium: The asset pricing implications, *American Economic Review* 106, 664-698.

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

Constantinides, George M., and Anisha Ghosh, 2017, Asset pricing with countercyclical household consumption risk, *Journal of Finance* 122, 415-459.

Cooper, Ilan, and Richard Priestley, 2009, Time-varying risk premiums and the output gap, *Review of Financial Studies* 22, 2801-2833.

Cujean, Julien, and Michael Hasler, 2017, Why does return predictability concentrate in bad times? *Journal of Finance* 72, 2717-2758.

Dangl, Thomas, and Michael Halling, 2012, Predictive regressions with time-varying coefficients, *Journal of Financial Economics* 106, 157-181.

Elliott, Graham, and Ulrich K. Müller, 2006, Efficient tests for general persistent time variation in regression coefficients, *Review of Economic Studies* 73, 907-940.

Fahri, Emmanual, and Xavier Gabaix, 2016, Rare disasters and exchange rates, *Quarterly Journal of Economics* 131, 1-52.

Golez, Benjamin, and Peter Koudijs, 2018, Four centuries of return predictability, *Journal of Financial Economics* 127, 248-263.

Gomes, João F., and Lukas Schmid, 2017, Equilibrium asset pricing with leverage and default, *Journal of Finance*, forthcoming.

Hamilton, James D., 2018, Why you should never use the Hodrick-Prescott filter, *Review of Economics and Statistics* 100, 831-843.

Henkel, Sam James, J. Spencer Martin, and Federico Nardari, 2011, Time-varying short-horizon predictability, *Journal of Financial Economics* 99, 560-580.

Hjalmarsson, Erik, 2010, Predicting global stock returns, *Journal of Financial and Quantitative Analysis* 45, 49-80.

Hjalmarsson, Erik, 2011, New methods for inference in long-horizon regressions, *Journal of Financial and Quantitative Analysis* 46, 815-839.

Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357-386.

Hodrick, Robert J., and Edward C. Prescott, 1997, Postwar U.S. business cycles: An empirical investigation, *Journal of Money, Credit and Banking* 29, 1-16.

Inoue, Atsushi, and Lutz Kilian, 2004, In-sample or out-of-sample tests of predictability: Which one should we use? *Econometric Reviews* 23, 371-402.

Kose, M. Ayhan, Christopher Otrok, and Charles H. Whiteman, 2003, International business cycles: World, region, and country-specific factors, *American Economic Review* 93, 1216-1239.

Kostakis, Alexandros, Tassos Magdalinos, and Michalis P. Stamatogiannis, 2015, Robust econometric inference for stock return predictability, *Review of Financial Studies* 28, 1506-1553.

Lettau, Martin, and Sydney C. Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 66, 815-849.

Lettau, Martin, and Sydney C. Ludvigson, 2013, Shocks and crashes, in Jonathan Parker and Michael Woodford, ed.: *National Bureau of Economic Research Macroeconomics Annual* (MIT Press, Cambridge MA).

Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter, 2008, The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies* 21, 1653-1687.

Lettau, Martin, and Stijn Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607-1652.

Lewellen, Jonathan, 2004, Predicting returns with financial ratios, *Journal of Financial Economics* 74, 209-235.

Lo, Andrew W., and A. Craig MacKinley, 1990, Data-snooping biases in tests of financial asset pricing models, *Review of Financial Studies* 3, 431-467.

Lumsdaine, Robin L., and Eswar S. Prasad, 2003, Identifying the common component in international economic fluctuations: A new approach, *Economic Journal* 113, 101-127.

McConnell, Margaret M., and Gabriel Perez-Quiros, 2000, Output fluctuations in the United States: What has changed since the early 1980s? *American Economic Review* 90, 1464-1476.

McCracken, Michael W., 2007, Asymptotics for out of sample tests of Granger causality, *Journal of Econometrics* 140, 719-752.

Nagel, Stefan, and Zhengyang Xu, 2018, Asset pricing with fading memory, Working paper.

Nelson, Charles R., and Myung J. Kim, 1993, Predictable stock returns, *Journal of Finance* 48, 641-661.

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

Pastor, Lubos, and Robert F. Stambaugh, 2009, Predictive systems: Living with imperfect predictors, *Journal of Finance* 4, 1583-1628.

Paye, Bradley S., and Allan Timmermann, 2006, Instability of return predictor models, *Journal of Empirical Finance* 13, 274-315.

Pukthuanthong, Kuntara, and Richard Roll, 2009, Global market integration: An alternative measure and its application, *Journal of Financial Economics* 94, 214-232.

Rangvid, Jesper, Pedro Santa-Clara, and Maik Schmeling, 2016, Capital market integration and consumption risk sharing over the long run, *Journal of International Economics* 103, 27-43.

Rapach, David E., Matthew C. Ringgenberg, and Guofu Zhou, 2016, Short interest and aggregate stock returns, *Journal of Financial Economics* 121, 46-65.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *Review of Financial Studies* 23, 821-862.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2013, International stock return predictability: What is the role of the United States? *Journal of Finance* 68, 1633-1662.

Santos, Tano, and Pietro Veronesi, 2006, Labor income and predictable stock returns, *Review of Financial Studies* 19, 1-44.

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

Stock, James H., and Mark W. Watson, 2002, Has the business cycle changed and why? in Mark Gertler and Kenneth Rogoff ed.: *NBER Macroeconomics Annual* (MIT Press, Cambridge MA).

Stock, James H., and Mark W. Watson, 2010, Modeling inflation after the crisis, NBER Working paper 16488.

Yogo, Motohiro, 2006, A consumption-based explanation of expected stock returns, *Journal of Finance* 61, 539-580.

Valkanov, Rossen, 2003, Long-horizon regressions: Theoretical results and applications, *Journal of Financial Economics* 68, 201-232.

Wachter, Jessica A., 2006, A consumption-based model of the term structure of interest rates, *Journal of Financial Economics* 79, 365-399.

Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance* 68, 987-1035.

Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455-1508.

Table I

Benchmark Predictive Regressions

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log stock market return, and cc_t is one-quarter lagged cyclical consumption. The table shows results for log excess market returns (Panel A), log real market returns (Panel B), and log market returns (Panel C) for the S&P 500 index and the CRSP value-weighted index. For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
			Panel A: Excess Market Returns			
SP500	-0.43	-1.60	-2.72	-3.36	-4.34	-5.33
	(-3.28)***	(-3.83)***	(-4.21)***	(-4.65)***	(-4.69)***	(-4.28)***
	[3.69]	[12.93]	[20.79]	[24.70]	[32.36]	[34.99]
CRSP	-0.42	-1.55	-2.57	-3.13	-4.08	-5.01
	(-3.06)***	(-3.59)***	(-3.96)***	(-4.47)***	(-4.57)***	(-4.17)***
	[3.16]	[11.43]	[18.68]	[22.46]	[31.53]	[34.46]
			Panel B: Real Market Returns			
SP500	-0.41	-1.56	-2.67	-3.33	-4.38	-5.54
	(-3.11)***	(-3.55)***	(-3.84)***	(-4.12)***	(-4.17)***	(-3.97)***
	[3.34]	[11.71]	[18.47]	[21.51]	[28.19]	[31.45]
CRSP	-0.40	-1.51	-2.52	-3.10	-4.12	-5.22
	(-2.91)***	(-3.34)***	(-3.62)***	(-3.95)***	(-3.99)***	(-3.80)***
	[2.87]	[10.45]	[16.83]	[19.93]	[28.04]	[31.66]
			Panel C: Market Returns			
SP500	-0.35	-1.28	-2.11	-2.51	-3.33	-4.26
	(-2.65)***	(-3.00)***	(-3.05)***	(-3.05)***	(-3.16)***	(-3.15)***
	[2.36]	[8.51]	[12.86]	[13.94]	[19.06]	[22.16]
CRSP	-0.34	-1.23	-1.96	-2.28	-3.07	-3.94
	(-2.47)**	(-2.80)***	(-2.84)***	(-2.87)***	(-3.02)***	(-3.05)***
	[1.98]	[7.38]	[11.18]	[12.06]	[17.87]	[21.11]

Table II

IVX-Wald Statistics

This table shows IVX-Wald statistics of Kostakis, Magdalinos, and Stamatogiannis (2015) for predictive OLS regressions summarized in Table I. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
Panel A: Excess Market Returns						
SP500	9.72***	9.38***	7.72***	6.38**	7.11***	7.83***
CRSP	8.43***	7.89***	6.19**	4.95**	5.63**	6.23**
Panel B: Real Market Returns						
SP500	8.54***	8.38***	7.00***	5.87**	6.80***	7.89***
CRSP	7.45***	7.09***	5.63**	4.57**	5.42**	6.35**
Panel C: Market Returns						
SP500	5.99**	5.60**	4.29**	3.25*	3.81**	4.51**
CRSP	5.15**	4.63**	3.61*	2.37	2.87*	3.45*

Table III

Predictability in Good and Bad Times

The table presents results of two-state predictive regressions of the form $r_{t,t+h} = \alpha + \beta_{bad} I_{bad} cc_t + \beta_{good} (1 - I_{bad})\, cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess return on the CRSP value-weighted index, cc_t is one-quarter lagged cyclical consumption, and I_{bad} is the state indicator that equals one during bad economic states and zero otherwise. Panel A employs the NBER-dated chronology of recessions to define bad states following Rapach, Strauss, and Zhou (2010) and Henkel, Martin, and Nardari (2011); Panel B defines bad states using the bottom third of sorted growth rates of real GDP following Rapach, Strauss, and Zhou (2010); Panel C defines bad states as periods with the Purchasing Managers Index below a threshold value of 44.48 specified as in Berge and Jordà (2011); Panel D defines bad states as periods with cyclical consumption realizations below its mean by more than one standard deviation. For each regression, the table reports slope estimates, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
			Panel A: NBER Business Cycle Dates			
β_{bad}	-0.83	-3.85	-3.53	-3.01	-5.86	-7.01
	(-1.86)**	(-2.91)***	(-2.09)***	(-2.39)***	(-4.91)***	(-3.74)***
β_{good}	-0.37	-1.26	-2.45	-3.15	-3.84	-4.71
	(-2.61)***	(-2.93)***	(-3.76)***	(-4.12)***	(-4.17)***	(-3.97)***
	[3.22]	[14.42]	[18.69]	[22.14]	[32.07]	[35.01]
			Panel B: Real GDP Growth			
β_{bad}	-0.76	-2.61	-3.08	-3.24	-4.54	-5.95
	(-3.56)***	(-4.41)***	(-3.84)***	(-4.32)***	(-5.73)***	(-5.01)***
β_{good}	-0.24	-1.00	-2.30	-3.07	-3.84	-4.50
	(-1.37)*	(-2.35)**	(-3.35)***	(-3.88)***	(-3.65)***	(-3.30)***
	[3.99]	[13.95]	[18.74]	[22.15]	[31.46]	[34.84]
			Panel C: Purchasing Managers Index (PMI)			
β_{bad}	-0.90	-5.02	-5.71	-5.26	-7.24	-7.66
	(-1.61)*	(-6.07)***	(-3.91)***	(-3.74)***	(-3.80)***	(-3.12)**
β_{good}	-0.38	-1.29	-2.33	-2.97	-3.82	-4.77
	(-2.73)***	(-3.05)***	(-3.70)***	(-4.18)***	(-4.30)***	(-3.93)***
	[3.14]	[15.56]	[20.50]	[22.94]	[32.76]	[35.03]
			Panel D: Cyclical Consumption			
β_{bad}	-0.46	-1.17	-1.58	-2.61	-2.75	-3.92
	(-2.15)**	(-2.34)**	(-2.18)**	(-3.22)***	(-2.88)**	(-3.12)***
β_{good}	-0.38	-1.88	-3.44	-3.60	-5.21	-5.81
	(-1.79)**	(-2.86)***	(-3.34)***	(-2.94)***	(-3.35)***	(-3.09)***
	[2.80]	[11.48]	[19.97]	[22.50]	[33.08]	[34.97]

Table IV

Alternative Detrending Methods

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess return on the CRSP value-weighted index, and cc_t is one-quarter lagged cyclical consumption. We compute cc by fitting a linear, linear with a break, quadratic or cubic time trend specification as indicated in the first column. The stochastic method computes cyclical consumption as a five-year backward-looking moving average. For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
Linear	-0.22	-0.83	-1.43	-1.87	-2.40	-2.77
	(-2.63)***	(-2.89)***	(-2.90)***	(-3.21)***	(-3.28)***	(-2.68)**
	[1.49]	[5.76]	[9.29]	[11.70]	[14.69]	[13.44]
Break	-0.62	-2.35	-3.99	-5.05	-6.00	-6.46
	(-3.37)***	(-3.30)***	(-3.63)***	(-4.78)***	(-5.50)***	(-4.52)***
	[3.54]	[13.15]	[22.36]	[28.67]	[33.87]	[29.92]
Quadratic	-0.44	-1.64	-2.69	-3.21	-3.66	-3.62
	(-2.50)**	(-2.60)***	(-2.61)***	(-2.87)***	(-3.25)***	(-2.85)***
	[1.79]	[7.03]	[11.07]	[12.44]	[13.28]	[9.61]
Cubic	-0.83	-3.24	-5.55	-6.84	-8.00	-8.59
	(-3.51)***	(-3.78)***	(-4.29)***	(-5.90)***	(-6.28)***	(-4.84)***
	[3.87]	[15.27]	[26.44]	[32.54]	[37.66]	[33.20]
Stochastic	-1.09	-4.60	-8.57	-11.62	-15.26	-18.88
	(-1.92)*	(-3.08)***	(-3.67)***	(-5.18)***	(-6.42)***	(-6.15)***
	[1.35]	[7.11]	[14.87]	[22.11]	[31.90]	[36.29]

Table V

Predictive Regressions for Subsamples

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess return on the S&P 500 index or the CRSP value-weighted index, and cc_t is one-quarter lagged cyclical consumption. The sample of returns covers the period from 1980Q1 to 2017Q4 (Panel A), from 1990Q1 to 2017Q4 (Panel B), and from 2000Q1 to 2017Q4 (Panel C). For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
Panel A: Post-1980 Period						
SP500	-0.35	-1.34	-2.53	-3.43	-4.55	-5.55
	(-2.24)**	(-2.96)***	(-3.10)***	(-3.25)***	(-3.52)***	(-3.70)***
	[2.22]	[9.21]	[17.92]	[23.26]	[30.42]	[33.27]
CRSP	-0.35	-1.30	-2.31	-2.98	-3.91	-4.61
	(-2.11)**	(-2.80)***	(-2.89)***	(-3.03)***	(-3.45)***	(-3.76)***
	[1.86]	[7.90]	[15.48]	[19.44]	[26.35]	[27.54]
Panel B: Post-1990 Period						
SP500	-0.44	-1.67	-3.00	-4.13	-5.61	-7.09
	(-2.52)***	(-3.07)***	(-2.82)***	(-3.20)***	(-3.74)***	(-4.58)***
	[3.83]	[14.35]	[21.34]	[26.99]	[36.29]	[40.49]
CRSP	-0.42	-1.53	-2.59	-3.43	-4.69	-5.85
	(-2.25)***	(-2.64)***	(-2.34)***	(-2.62)***	(-3.29)***	(-4.43)***
	[2.96]	[11.16]	[16.44]	[20.53]	[29.89]	[33.20]
Panel C: Post-2000 Period						
SP500	-0.54	-2.06	-3.28	-4.06	-4.91	-6.31
	(-2.94)***	(-3.40)***	(-2.93)***	(-3.35)***	(-4.62)***	(-6.95)***
	[7.35]	[23.72]	[30.40]	[37.38]	[47.44]	[55.64]
CRSP	-0.52	-1.92	-2.83	-3.33	-3.99	-5.14
	(-2.58)***	(-2.86)***	(-2.33)**	(-2.57)***	(-3.58)**	(-5.41)***
	[5.56]	[18.30]	[21.66]	[25.70]	[34.79]	[42.77]

Table VI

Benchmark Out-of-sample Tests

The table presents results of out-of-sample forecasts of h-quarter-ahead log excess returns on the CRSP value-weighted index where a time-varying expected returns model with cyclical consumption as regressor is compared against a constant expected returns model. The parameters used to calculate cyclical consumption are estimated recursively from the current latest-available consumption data. R^2_{OOS} is the out-of-sample R^2 in percent. ENC-NEW is the encompassing test statistic of Clark and McCracken (2001) and MSE-F is the F-statistic of McCracken (2007). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values in case of the ENC-NEW and MSE-F statistics, and according to the Clark and West (2007) test in case of the R^2_{OOS} statistics. The first observation in the out-of-sample period is 1980Q1, 1990Q1, or 2000Q1, and the predictive model is estimated recursively until 2017Q4.

	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
			Panel A: Forecasting from 1980			
ENC-NEW	2.85***	12.73***	25.26***	35.79***	54.53***	60.56***
MSE-F	0.98***	6.43***	17.09***	30.17***	42.93***	41.07***
R^2_{OOS}	0.64*	4.14***	10.55***	17.63***	23.86***	23.59***
			Panel B: Forecasting from 1990			
ENC-NEW	3.51***	13.77***	22.47***	29.46***	46.57***	53.95***
MSE-F	2.24***	9.76***	17.63***	26.92***	44.09***	52.80***
R^2_{OOS}	1.96*	8.22***	14.37***	21.05***	31.25***	36.21***
			Panel C: Forecasting from 2000			
ENC-NEW	2.40***	7.89***	11.76***	17.01***	29.89***	39.64***
MSE-F	0.99***	4.15***	5.85***	10.97***	20.33***	33.08***
R^2_{OOS}	1.35	5.67***	8.26***	15.24***	26.29***	38.43***

Table VII

Portfolios Sorted on Industry

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess portfolio return, and cc_t is one-quarter lagged cyclical consumption. The table shows results for industry categories including Nondurable Goods (NON), Durable Goods (DUR), Manufacturing (MAN), Energy (ENG), HiTech Business Equipment (HT), Telephone and Television Transmission (TEL), Wholesale and Retail (SHOPS), Healthcare and Medical Equipment (HLTH), Utilities (UTILS) and Other industry categories (OTHER). For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
NON	-0.31	-1.12	-1.94	-2.54	-3.47	-4.51
	(-2.25)**	(-2.42)***	(-2.70)***	(-2.72)***	(-2.82)***	(-2.81)**
	[1.63]	[6.40]	[10.71]	[14.24]	[19.95]	[23.97]
DUR	-0.66	-2.36	-3.71	-4.47	-5.69	-6.99
	(-3.47)***	(-4.31)***	(-5.17)***	(-6.18)***	(-6.62)***	(-6.36)***
	[4.07]	[13.38]	[19.44]	[26.59]	[35.94]	[39.88]
MAN	-0.39	-1.40	-2.15	-2.42	-3.12	-3.84
	(-2.61)***	(-3.07)***	(-3.34)***	(-3.55)***	(-3.27)***	(-2.91)***
	[2.11]	[7.97]	[12.55]	[14.57]	[20.19]	[21.21]
ENG	-0.18	-0.68	-0.93	-0.56	-0.69	-1.26
	(-1.22)	(-1.36)	(-1.05)	(-0.50)	(-0.56)	(-0.88)
	[0.18]	[1.49]	[1.68]	[0.26]	[0.56]	[2.22]
HT	-0.62	-2.35	-4.04	-5.23	-6.85	-8.29
	(-2.93)***	(-3.42)***	(-3.72)***	(-4.68)***	(-5.18)***	(-5.03)***
	[3.42]	[12.04]	[18.81]	[23.12]	[32.31]	[34.21]
TEL	-0.46	-1.74	-3.24	-4.29	-5.19	-5.89
	(-3.11)***	(-3.19)***	(-3.22)***	(-3.23)***	(-3.19)***	(-2.87)***
	[3.99]	[11.74]	[18.48]	[21.10]	[23.34]	[22.74]
SHOPS	-0.42	-1.49	-2.32	-2.91	-3.97	-5.05
	(-2.76)***	(-3.19)***	(-3.43)***	(-3.46)***	(-3.52)***	(-3.37)***
	[2.28]	[8.83]	[12.16]	[15.81]	[23.39]	[26.96]
HLTH	-0.37	-1.38	-2.48	-3.45	-4.92	-6.48
	(-2.48)***	(-2.85)***	(-2.76)***	(-2.87)***	(-3.28)***	(-3.61)***
	[1.92]	[8.01]	[12.67]	[16.02]	[23.69]	[30.01]
UTILS	-0.31	-1.24	-2.17	-2.56	-3.11	-3.54
	(-2.81)***	(-3.55)***	(-3.92)***	(-3.41)***	(-3.04)***	(-2.46)**
	[2.20]	[9.49]	[16.84]	[18.15]	[20.41]	[20.26]
OTHER	-0.43	-1.64	-2.75	-3.53	-4.74	-5.95
	(-2.66)***	(-3.36)***	(-4.31)***	(-4.90)***	(-4.42)***	(-3.90)***
	[2.23]	[8.87]	[14.66]	[18.89]	[26.28]	[29.99]

Table VIII

Alternative Consumption Measures

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess return on the CRSP value-weighted index, and cc_t is one-quarter lagged cyclical consumption. We compute cc by applying the robust linear projection method of Hamilton (2018) to the logarithm of real per capita consumption expenditure for nondurable goods (NON), services (SERV), durable goods (DUR), the stock of durable goods (SDUR), nondurable and durable goods (GOODS), or aggregate personal consumption expenditure (PCE) as indicated in the first column. For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
NON	-0.38	-1.44	-2.50	-3.22	-4.28	-5.14
	(-3.04)***	(-3.73)***	(-3.95)***	(-4.65)***	(-5.93)***	(-6.58)***
	[3.18]	[11.83]	[21.14]	[28.51]	[42.41]	[45.42]
SERV	-0.38	-1.39	-2.33	-2.80	-3.57	-4.33
	(-2.84)***	(-3.33)***	(-3.72)***	(-4.16)***	(-4.05)***	(-3.51)***
	[2.66]	[9.42]	[15.90]	[18.58]	[24.81]	[26.51]
DUR	-0.07	-0.29	-0.58	-0.82	-1.06	-1.28
	(-1.79)**	(-2.49)***	(-3.24)***	(-3.79)***	(-4.24)***	(-4.59)***
	[0.96]	[5.85]	[14.31]	[23.25]	[33.33]	[37.79]
SDUR	-0.10	-0.39	-0.79	-1.10	-1.39	-1.61
	(-1.99)**	(-2.38)***	(-3.05)***	(-3.47)***	(-3.81)***	(-3.98)***
	[1.16]	[5.86]	[14.44]	[22.82]	[30.62]	[31.66]
GOODS	-0.17	-0.69	-1.30	-1.76	-2.31	-2.79
	(-2.42)***	(-3.02)***	(-3.49)***	(-4.02)***	(-4.70)***	(-5.16)***
	[1.93]	[8.57]	[18.05]	[27.12]	[39.67]	[43.78]
PCE	-0.29	-1.08	-1.91	-2.48	-3.20	-3.91
	(-2.85)***	(-3.20)***	(-3.80)***	(-4.36)***	(-4.58)***	(-4.45)***
	[2.72]	[9.93]	[18.58]	[25.41]	[35.40]	[39.37]

Table IX

International Evidence

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log return on world or regional MSCI total equity indices and cc_t is one-quarter lagged global cyclical consumption. We compute global cyclical consumption as a cross-country average for 20 developed countries, including Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, the United Kingdom and the United States. For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. We consider the longest possible sample period for each set of test asset returns. The sample of returns on the aggregate G7 index covers the period from 1977Q1 to 2017Q4 and from 1970Q1 to 2017Q4 otherwise.

	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
World	-0.48	-1.95	-3.53	-4.73	-5.78	-6.32
	(-3.59)***	(-4.04)***	(-4.63)***	(-5.33)***	(-6.51)***	(-7.35)***
	[5.32]	[22.47]	[37.46]	[46.75]	[53.97]	[51.86]
World ex USA	-0.48	-1.96	-3.55	-4.76	-5.79	-6.23
	(-3.56)***	(-3.90)***	(-4.18)***	(-4.52)***	(-5.20)***	(-5.33)***
	[4.87]	[19.73]	[33.05]	[41.82]	[48.06]	[44.70]
EAFE	-0.51	-2.01	-3.62	-4.90	-6.07	-6.74
	(-3.65)***	(-3.82)***	(-4.00)***	(-4.21)***	(-4.74)***	(-5.18)***
	[5.24]	[19.70]	[32.09]	[40.50]	[47.14]	[46.73]
Europe	-0.51	-2.08	-3.78	-5.07	-6.16	-6.67
	(-3.75)***	(-4.11)***	(-4.35)***	(-4.64)***	(-5.22)***	(-5.31)***
	[5.54]	[21.68]	[35.69]	[44.11]	[49.99]	[47.14]
Far East	-0.51	-2.11	-3.85	-5.06	-6.03	-6.33
	(-2.97)***	(-3.65)***	(-3.63)***	(-3.53)***	(-3.54)***	(-3.14)***
	[3.38]	[13.88]	[22.58]	[27.17]	[30.66]	[26.76]
Pacific	-0.49	-2.03	-3.65	-4.77	-5.66	-5.94
	(-3.07)***	(-3.68)***	(-3.60)***	(-3.41)***	(-3.36)***	(-2.98)***
	[3.63]	[14.78]	[23.85]	[28.54]	[31.21]	[27.06]
G7	-0.44	-1.86	-3.51	-4.83	-5.91	-6.50
	(-3.13)***	(-3.66)***	(-4.18)***	(-4.85)***	(-6.11)***	(-7.40)***
	[5.14]	[22.80]	[40.93]	[50.31]	[56.95]	[55.14]

Table X

Out-of-sample Results with Alternative Predictor Variables

The table presents the out-of-sample R^2 statistics in percent from h-quarter ahead forecasts of log excess returns on the CRSP value-weighted index where the time-varying expected returns model with one of the predictive variables listed in the second column as a regressor is compared against a constant expected returns model. Section IV contains definitions of the forecasting variables. The parameters used to calculate variables 17-20 are estimated recursively from the current latest-available data. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to the Clark and West (2007) test statistics. The first observation in the out-of-sample period is 1980Q1; the predictive model is estimated recursively until 2017Q4.

#	var	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
1	dp	-5.29	-23.40	-49.33	-40.09	-30.04	-40.84
2	dy	-6.45	-21.81	-41.22	-33.53	-25.92	-39.13
3	e/p	-2.63	-10.62	-22.59	-21.79	-27.60	-37.16
4	d/e	-3.30	-5.77	-8.57	-15.48	-10.88	-8.36
5	$svar$	-17.32	-17.82	-18.72	-25.54	-18.04	-24.82
6	b/m	-1.72	-6.12	-15.93	-20.37	-24.04	-28.74
7	$ntis$	-2.51	-13.75	-9.92	-9.37	-15.42	-31.11
8	tbl	-4.03	-11.83	-5.78	-25.16	-64.17	-127.11
9	lty	-2.45	-9.93	-13.35	-38.98	-75.43	-133.55
10	ltr	-1.55	-2.12	-1.58	-3.86	-7.58	-9.14
11	tms	-2.78	-0.93	10.31***	12.77***	13.88***	5.32***
12	dfy	-2.70	-4.90	-3.29	-15.30	-21.02	-16.52
13	dfr	-0.49	-2.48	-3.94	-6.80	-9.24	-12.59
14	$infl$	-1.86	-0.84	1.20**	-6.36	-10.65	-22.00
15	i/k	0.23	2.21**	8.79***	16.76***	27.39***	24.88***
16	s^w	-3.17	-9.69	-17.69	-30.03	-39.36	-45.37
17	cay	-2.18	-13.33	-20.03	-15.01	-10.90	-24.23
18	σ_c	-0.44	0.06	-3.69	-13.05	-10.67	-10.49
19	gap	-2.02	-5.51	-4.45	-1.47	7.50***	9.69***
20	cc	0.64*	4.14***	10.55***	17.63***	23.86***	23.59***

Table XI

Summary Statistics of Simulated and Historical Data

The table presents summary statistics for log consumption growth rates (Δc) and the log aggregate stock market returns (r) expressed in annualized percentages. It shows the time-series averages (E), standard deviations (σ), and Sharpe ratios computed as the mean excess return ($r - r_f$) divided by the standard deviation. In the model, $r - r_f$ is the log return on the consumption or dividend claim minus the log risk-free rate. In the data, $r - r_f$ is the log return on the value-weighted CRSP index minus the log Treasury bill return. Data statistics reported in columns "Consumption claim" and "Dividend claim" present the moments in the simulated data. We generate 1,000,000 quarterly observations based on the calibrated parameter values of Campbell and Cochrane (1999). Column "Actual data" summarizes the moments in our empirical sample covering the period from 1947Q1 to 2017Q4.

Statistic	Consumption claim	Dividend claim	Actual data
$E(\Delta c)$	1.89	1.89	1.89
$\sigma(\Delta c)$	1.50	1.50	1.00
$E(r - r_f) / \sigma(r - r_f)$	0.43	0.38	0.40
$E(r - r_f)$	5.22	5.01	6.52
$\sigma\left(r - r^f\right)$	12.02	13.31	16.38

Table XII

Model-implied Predictability

Panel A of the table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess market return, and cc_t is one-quarter lagged cyclical consumption computed as specified in Equations (1) and (11) in the actual and simulated data for $k = 24$, respectively. For each regression, the table reports the slope estimate and the adjusted R^2 statistic in percent in square brackets. The rows "Consumption claim" and "Dividend claim" show results from 1,000,000 quarterly simulated observations. The row "Actual data" displays results in the historical data. Panel B of the table presents results of out-of-sample forecasts of h-quarter ahead log excess stock returns where the time-varying expected returns model with cyclical consumption as regressor is compared against a constant expected returns model. The rows "Consumption claim" and "Dividend claim" show average out-of-sample R^2 statistics in percent in simulated data from 2,500 artificial samples of size 284, which matches the number of observations in our post-war sample. For each artificial sample, we compute 112-(h-1) out-of-sample forecasts for consistency with the number of forecasts in the post-1990 evaluation window. The row "Actual data" shows corresponding results in the historical data.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
Panel A: In-sample Predictability						
Consumption claim	-0.22	-0.82	-1.51	-2.10	-2.60	-3.04
	[1.75]	[6.53]	[12.02]	[16.59]	[20.48]	[23.80]
Dividend claim	-0.22	-0.82	-1.52	-2.11	-2.61	-3.05
	[1.46]	[5.36]	[9.68]	[13.13]	[15.90]	[18.13]
Actual data	-0.42	-1.55	-2.57	-3.13	-4.08	-5.01
	[3.16]	[11.43]	[18.68]	[22.46]	[31.53]	[34.46]
Panel B: Out-of-sample Predictability						
Consumption claim	1.00	3.47	5.99	7.56	8.54	9.51
Dividend claim	0.68	2.15	3.18	3.26	2.69	1.71
Actual data	1.96	8.22	14.37	21.05	31.25	36.21

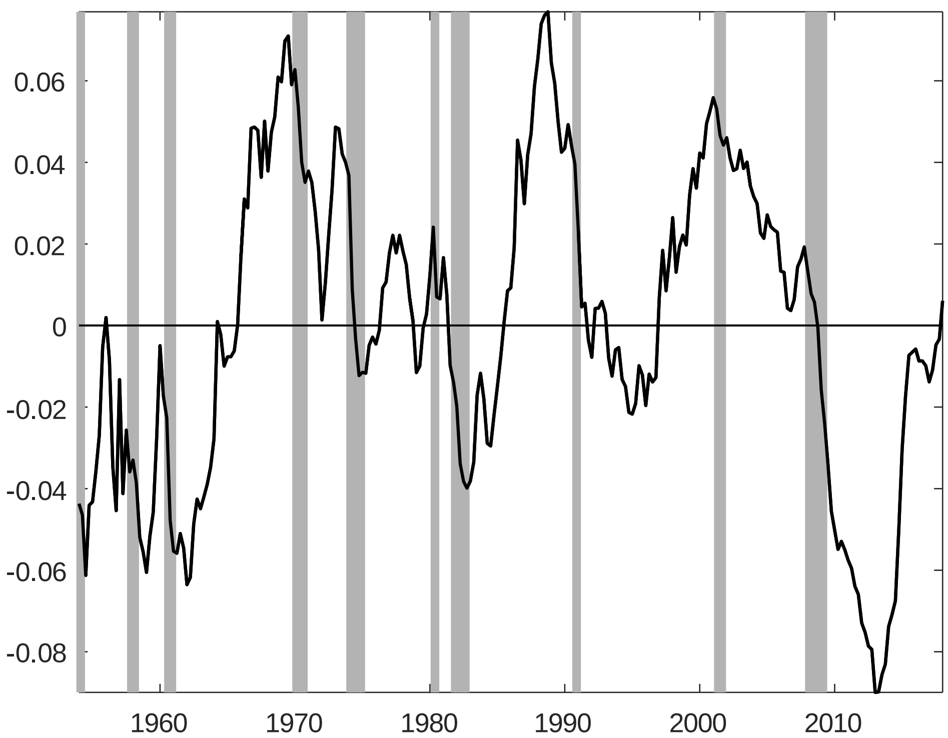

Figure 1. Cyclical consumption. The figure plots the series of cyclical consumption along with NBER recessions represented by shaded bars over the period from 1953Q4 to 2017Q4.

Internet Appendix for

"Consumption Fluctuations and Expected Returns"

Internet Appendix A: The out-of-sample procedure and tests

The out-of-sample recursive scheme we employ is similar to e.g. Cooper and Priestley (2009). Below we briefly summarize how we proceed for the case of an out-of-sample window covering the period from 1980Q1 to 2017Q4. We start by using consumption data since 1947Q1 through 1979Q4 to estimate

$$c_t = b_{0,\tau c} + b_{1,\tau c} c_{t-k} + b_{2,\tau c} c_{t-k-1} + b_{3,\tau c} c_{t-k-2} + b_{4,\tau c} c_{t-k-3} + \omega_t, \tag{A1}$$

where $\tau c = 1979\text{Q}4$, $k = 24$, and the residual ω_t is the measure of cc from 1953Q4 to 1979Q4. The time subscript τc for the parameters in Equation (A1) indicates that they are updated recursively in each quarter. Next, we expand the sample period of consumption by one quarter and estimate over the period from 1947Q1 to 1980Q1:

$$c_t = b_{0,\tau c+1} + b_{1,\tau c+1} c_{t-k} + b_{2,\tau c+1} c_{t-k-1} + b_{3,\tau c+1} c_{t-k-2} + b_{4,\tau c+1} c_{t-k-3} + \omega_t, \tag{A2}$$

where $\tau c + 1 = 1980\text{Q}1$. We add on the last estimate of cc for 1980Q1 from Equation (A2) to the time series of cc over the period from 1953Q4 to 1979Q4 generated above. We repeat this procedure, quarter-by-quarter, recursively estimating the parameters and saving the values of cc until the end of the sample.

In a next step, we employ these values of cc in recursive predictive regressions for stock returns to form out-of-sample forecasts. We use an expanding estimation window where the coefficients in the return forecasting regression are estimated recursively using only information available through time t for forecasting over the next h quarters.

The assessment of out-of-sample predictability involves three metrics. The first statistic we report is the powerful ENC-NEW statistic of Clark and McCracken (2001) which extends the encompassing test of Harvey, Leybourne, and Newbold (1998) by deriving a nonstandard asymptotic distribution of this test statistic under the null of nested forecasts. The ENC-NEW statistic tests the null hypothesis that the restricted forecasting model encompasses the unrestricted forecasting model; the alternative is that the time-varying expected return model contains information that could be used to significantly improve the forecast of the constant expected return model. The second is the MSE-F statistic of McCracken (2007) which tests the null hypothesis that the restricted forecasting model has a mean squared error (MSE) that is less than or equal to that of the unrestricted forecasting model; the alternative is that the unrestricted model has a smaller MSE. The third test is the out-of-sample R^2_{OOS} statistic which measures the proportional reduction (or increase) in the MSE of the unrestricted model relative to the MSE of the prevailing mean benchmark forecast. The R^2_{OOS} statistic is measured in units that are comparable to the in-sample R^2. The R^2_{OOS} takes positive (negative) values when the predictive regression model predicts better (worse) than the historical mean.[1] The critical values for the ENC-NEW and MSE-F statistics are obtained from a bootstrap procedure. To assess the statistical significance of the R^2_{OOS}'s, we employ the Clark and West (2007) test statistic, which tests the null hypothesis that the historical average MSE is less than or equal to the predictive regression MSE against the alternative hypothesis that the historical average MSE is greater than the predictive regression MSE. This statistic is a correction of the Diebold and Mariano (1995) statistic and is demonstrated to be more suitable for nested models.

[1] The results for predictive regression forecasts that implement the Campbell and Thompson (2008) economically motivated sign restrictions in order to reduce parameter estimation uncertainty and help stabilize predictive regression forecasts are qualitatively similar. These restrictions entail setting the slope coefficient of cc in the bivariate predictive regression forecast to zero if the estimated slope coefficient is positive, and setting the forecast of the equity premium to zero if the bivariate predictive regression forecast of the equity premium is negative.

Internet Appendix B: The model of Campbell and Cochrane (1999)

Campbell and Cochrane (1999) augment the standard power utility function with a time-varying subsistence level X_t that adapts nonlinearly to current and past average consumption in the economy:

$$E_t \sum_{t=0}^{\infty} \delta^t \frac{(C_t - X_t)^{1-\gamma} - 1}{1-\gamma}, \tag{B1}$$

where C_t is aggregate consumption at time t, δ is the subjective time discount factor, and γ is the utility curvature parameter. The reference point X_t represents the agent's external habit level, which is defined indirectly through the surplus consumption ratio $S_t \equiv \frac{C_t - X_t}{C_t}$. The local coefficient of relative risk aversion is γ/S_t.

Campbell and Cochrane (1999) assume a mean-reverting heteroskedastic first-order autoregressive process for the log surplus consumption ratio, $s_t \equiv \log(S_t)$:

$$s_{t+1} = (1-\phi)\,\overline{s} + \phi s_t + \lambda(s_t)\,\upsilon_{t+1}, \tag{B2}$$

where $\overline{s}$ is the steady state value of s_t, ϕ is the habit persistence parameter, and $\lambda(s_t)$ is a nonlinear monotonically decreasing sensitivity function that determines how innovations in consumption growth υ_{t+1} influence s_{t+1}.

The log consumption growth $\Delta c_{t+1} = \log(C_{t+1}/C_t)$ is given by:

$$\Delta c_{t+1} = g + \upsilon_{t+1}, \quad \upsilon_{t+1} \sim niid\left(0, \sigma_c^2\right), \tag{B3}$$

where g and σ_c^2 are constant parameters governing the mean and volatility. The sensitivity function is specified as

$$\lambda(s_t) = \begin{cases} \frac{1}{\overline{S}}\sqrt{1 - 2(s_t - \overline{s})} - 1 & \text{if} \quad s_t \leq s_{\max} \\ 0 & \text{if} \quad s_t \geq s_{\max}, \end{cases} \tag{B4}$$

where

$$\overline{S} = \sigma_c \sqrt{\frac{\gamma}{1-\phi}} \tag{B5}$$

is the steady state level of S_t and

$$s_{\max} \equiv \overline{s} + \frac{1}{2}\left(1 - \overline{S}^2\right) \tag{B6}$$

is the value of s_t at which the expression in Equation (B4) becomes zero. Specifying $\lambda(s_t)$ in this way implies that the real risk-free rate is constant over time. In particular, from the Euler equation

$$1 = E_t\left[R_{i,t+1} M_{t+1}\right], \tag{B7}$$

where $R_{i,t+1}$ is the real gross return on any traded asset i and M_{t+1} is the stochastic discount factor:

$$M_{t+1} = \delta \left(\frac{S_{t+1}}{S_t}\frac{C_{t+1}}{C_t}\right)^{-\gamma} = \delta e^{-\gamma\{g+(\phi-1)(s_t-\overline{s})+[1+\lambda(s_t)]v_{t+1}\}}, \tag{B8}$$

it follows that the log real risk-free rate is:

$$r_{f,t+1} = -\log(\delta) + \gamma g - \frac{\gamma}{2}(1-\phi). \tag{B9}$$

The aggregate market is represented as a claim to the future consumption stream and the price-consumption ratio for a consumption claim satisfies:

$$\frac{P_t}{C_t}(s_t) = E_t\left[M_{t+1}\frac{C_{t+1}}{C_t}\left[\frac{P_{t+1}}{C_{t+1}}(s_{t+1}) + 1\right]\right]. \tag{B10}$$

The surplus consumption ratio is the only state variable in the model and the price-consumption ratio can therefore be written as a function of s_t only. Furthermore, since consumption is the dividend paid by the market, the price-consumption ratio is analogous to the price-dividend ratio, and the return on the aggregate market is given as:

$$R_{t+1} = \frac{(P_{t+1}/C_{t+1}) + 1}{P_t/C_t} \frac{C_{t+1}}{C_t}. \tag{B11}$$

Following Campbell and Cochrane (1999), we solve the model by substituting for M_{t+1} from (B8) and consumption growth from (B3) and then use numerical integration over the normally distributed shock υ_{t+1}.[2] We also consider returns to the dividend claim (see the appendix to Campbell and Cochrane, 1999). We set the correlation between consumption and dividends to 0.2 when using the dividend claim such that dividends and consumption are imperfectly correlated. Furthermore, we set the dividend volatility to 6.1% to match the standard deviation of dividend growth in the post-war sample.[3]

[2] We solve the model using a Matlab program similar to the Gauss program available on John Y. Cochrane's website. See the appendix of Costa and Vasconcelos (2009) for a conversion of Campbell and Cochrane's (1999) Gauss code to Matlab. We calculate the integral using the Gauss-Legendre 40-point quadrature and bound the integral by -8 and $+8$ standard deviations. To solve the model, we use the GaussLegendre.m function available on Pavel Holoborodko's website.

[3] We follow Chen (2009) and use non-reinvested dividends when computing the dividend volatility.

Internet Appendix C: Additional results

Table AI

Descriptive Statistics

The table provides summary statistics for alternative predictive variables listed in the second column. Panel A reports for each variable the time-series average (mean), median, standard deviation (std), skewness (skew), kurtosis (kurt), minimum (min), maximum (max), and first-order autocorrelation ($\rho(1)$). Panel B displays the Pearson correlation coefficients. Section IV contains definitions of the forecasting variables. The sample covers the period from 1953Q4 to 2017Q4.

Panel A: Summary Statistics									
#	var	mean	median	std	skew	kurt	min	max	$\rho(1)$
1	dp	-3.56	-3.50	0.39	-0.31	2.37	-4.49	-2.78	0.98
2	dy	-3.54	-3.49	0.40	-0.32	2.45	-4.50	-2.77	0.98
3	e/p	-2.82	-2.86	0.42	-0.85	6.56	-4.81	-1.90	0.94
4	d/e	-0.74	-0.73	0.31	2.83	19.66	-1.24	1.38	0.89
5	$svar$	0.01	0.00	0.01	7.42	71.89	0.00	0.11	0.42
6	b/m	0.51	0.49	0.25	0.73	2.92	0.13	1.20	0.98
7	$ntis$	0.01	0.02	0.02	-0.88	3.63	-0.05	0.05	0.93
8	tbl	0.04	0.04	0.03	0.83	4.09	0.00	0.15	0.95
9	lty	0.06	0.06	0.03	0.82	3.27	0.02	0.15	0.98
10	ltr	0.02	0.01	0.05	0.90	5.63	-0.15	0.24	-0.04
11	tms	0.02	0.02	0.01	-0.25	3.04	-0.04	0.05	0.84
12	dfy	0.01	0.01	0.00	1.84	8.05	0.00	0.03	0.87
13	dfr	0.00	0.00	0.02	0.31	14.89	-0.12	0.16	-0.05
14	$infl$	0.01	0.01	0.01	0.19	6.51	-0.04	0.04	0.45
15	i/k	0.04	0.04	0.00	0.31	2.48	0.03	0.04	0.97
16	s^w	0.89	0.88	0.03	0.26	1.92	0.84	0.94	0.97
17	cay	-0.00	-0.00	0.02	-0.26	2.41	-0.05	0.04	0.91
18	σ_c	-4.52	-4.48	0.54	-0.43	3.14	-6.37	-3.41	0.81
19	gap	-0.00	-0.01	0.07	0.04	1.95	-0.14	0.12	0.96
20	cc	0.00	0.00	0.04	-0.22	2.40	-0.09	0.08	0.97

Panel B: Correlation Coefficients

#	var	dy	e/p	d/e	$svar$	b/m	$ntis$	tbl	lty	ltr	tms	dfy	dfr	$infl$	i/k	s^w	cay	σ_c	gap	cc
1	dp	0.98	0.72	0.29	-0.06	0.90	0.20	0.52	0.47	0.01	-0.24	0.28	0.01	0.34	-0.12	-0.12	0.09	0.38	-0.34	-0.13
2	dy	1.00	0.71	0.28	-0.15	0.88	0.19	0.50	0.46	0.01	-0.22	0.27	0.06	0.31	-0.15	-0.14	0.09	0.37	-0.38	-0.15
3	e/p		1.00	-0.45	-0.27	0.78	0.16	0.58	0.48	0.03	-0.34	0.08	-0.14	0.45	0.12	-0.03	0.05	0.06	-0.04	-0.10
4	d/e			1.00	0.29	0.08	0.04	-0.12	-0.06	-0.03	0.16	0.25	0.20	-0.19	-0.32	-0.11	0.05	0.40	-0.38	-0.03
5	$svar$				1.00	-0.09	-0.22	-0.08	-0.02	0.28	0.15	0.45	-0.10	-0.23	-0.03	0.10	0.15	0.06	0.02	0.08
6	b/m					1.00	0.23	0.60	0.53	0.00	-0.28	0.37	0.01	0.49	0.04	0.12	-0.14	0.35	-0.13	-0.09
7	$ntis$						1.00	0.10	0.04	-0.12	-0.13	-0.34	0.07	0.09	0.05	-0.16	-0.10	0.21	0.08	0.02
8	tbl							1.00	0.89	-0.03	-0.48	0.30	-0.07	0.57	0.51	-0.15	0.11	0.12	0.07	0.39
9	lty								1.00	0.04	-0.03	0.46	-0.01	0.50	0.32	-0.09	0.24	0.11	-0.16	0.32
10	ltr									1.00	0.13	0.27	-0.41	-0.24	-0.02	-0.02	0.09	-0.05	-0.12	-0.02
11	tms										1.00	0.23	0.15	-0.30	-0.51	0.15	0.22	-0.05	-0.46	-0.23
12	dfy											1.00	0.03	0.11	-0.14	0.31	-0.04	0.12	-0.30	-0.03
13	dfr												1.00	-0.03	-0.17	0.08	-0.09	0.03	-0.16	-0.06
14	$infl$													1.00	0.31	0.22	-0.13	0.07	0.20	0.25
15	i/k														1.00	-0.02	-0.03	-0.09	0.74	0.59
16	s^w															1.00	-0.67	-0.10	0.28	0.02
17	cay																1.00	-0.03	-0.32	-0.02
18	σ_c																	1.00	-0.14	-0.19
19	gap																		1.00	0.51

Table AII

In-sample Regressions with Alternative Predictive Variables

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta x_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess return on the CRSP value-weighted index, and x_t is one-quarter lagged predictive variable. Section IV contains definitions of the forecasting variables. For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

#	var	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
1	*dp*	0.02	0.09	0.14	0.16	0.18	0.23
		(1.59)**	(1.87)**	(1.69)*	(1.55)	(1.66)*	(2.48)**
		[0.73]	[3.91]	[5.83]	[5.89]	[6.01]	[8.40]
2	*dy*	0.03	0.09	0.13	0.15	0.16	0.21
		(1.78)**	(1.83)**	(1.54)*	(1.47)	(1.56)	(2.31)**
		[1.08]	[3.74]	[4.85]	[5.16]	[5.16]	[7.21]
3	e/p	0.01	0.04	0.05	0.06	0.04	0.04
		(0.52)	(0.81)	(0.64)	(0.64)	(0.41)	(0.32)
		[-0.16]	[0.62]	[0.39]	[0.56]	[0.06]	[-0.10]
4	d/e	0.02	0.07	0.14	0.15	0.20	0.29
		(0.93)	(1.70)*	(1.72)*	(1.40)	(1.71)*	(2.39)**
		[0.10]	[1.21]	[3.43]	[2.95]	[4.66]	[8.33]
5	*svar*	0.18	1.98	3.66	2.68	3.68	6.00
		(0.23)	(2.19)**	(2.87)***	(1.48)	(1.72)*	(2.22)**
		[-0.35]	[0.96]	[2.32]	[0.77]	[1.49]	[3.57]
6	b/m	0.02	0.08	0.08	0.05	0.04	0.10
		(0.74)	(1.01)	(0.66)	(0.32)	(0.24)	(0.56)
		[-0.11]	[0.87]	[0.41]	[-0.16]	[-0.26]	[0.29]
7	*ntis*	-0.12	-0.32	-0.41	-0.46	-1.29	-2.50
		(-0.34)	(-0.25)	(-0.25)	(-0.27)	(-0.57)	(-0.87)
		[-0.31]	[-0.27]	[-0.29]	[-0.30]	[0.33]	[1.79]
8	*tbl*	-0.33	-0.95	-1.29	-1.67	-1.96	-1.96
		(-1.85)**	(-1.71)**	(-2.05)**	(-1.82)**	(-1.46)*	(-1.07)
		[1.11]	[2.61]	[2.81]	[3.76]	[4.32]	[3.19]

#	var	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
9	lty	-0.24	-0.52	-0.48	-0.36	-0.26	0.03
		(-1.28)	(-0.84)	(-0.60)	(-0.33)	(-0.18)	(0.01)
		[0.25]	[0.29]	[-0.07]	[-0.26]	[-0.36]	[-0.43]
10	ltr	0.15	0.39	0.40	0.51	0.43	0.77
		(1.39)*	(2.34)**	(2.25)**	(2.69)**	(2.33)**	(3.01)**
		[0.56]	[1.14]	[0.49]	[0.80]	[0.30]	[1.37]
11	tms	0.67	2.62	4.30	6.28	7.72	8.54
		(1.63)**	(2.26)**	(3.41)***	(4.54)***	(3.81)***	(2.73)**
		[0.92]	[4.42]	[7.24]	[12.70]	[16.38]	[15.47]
12	dfy	0.89	4.24	5.29	5.55	9.38	16.85
		(0.55)	(1.16)	(1.12)	(1.01)	(1.34)	(2.01)*
		[-0.17]	[0.83]	[0.71]	[0.58]	[1.98]	[5.68]
13	dfr	0.51	0.39	0.02	0.31	0.50	0.39
		(1.84)**	(0.77)	(0.05)	(0.57)	(0.98)	(0.63)
		[1.67]	[-0.11]	[-0.40]	[-0.33]	[-0.23]	[-0.34]
14	$infl$	-0.91	-3.50	-3.88	-3.99	-5.15	-5.80
		(-1.21)	(-2.05)**	(-2.68)***	(-3.25)***	(-2.63)**	(-1.96)**
		[0.61]	[3.17]	[2.15]	[1.73]	[2.55]	[2.51]
15	i/k	-4.18	-13.73	-21.89	-32.21	-43.08	-50.46
		(-2.81)***	(-2.76)***	(-2.78)***	(-4.38)***	(-6.70)***	(-6.51)***
		[2.87]	[8.09]	[12.09]	[21.33]	[32.71]	[35.07]
16	s^w	-0.25	-0.77	-1.28	-1.63	-2.02	-2.26
		(-1.17)	(-1.11)	(-1.06)	(-0.92)	(-0.87)	(-0.83)
		[0.21]	[1.01]	[1.82]	[2.37]	[3.15]	[3.04]
17	cay	0.63	2.59	5.06	6.76	7.78	8.35
		(2.48)***	(2.93)***	(3.89)***	(4.78)***	(4.92)***	(4.55)***
		[1.75]	[8.18]	[18.16]	[25.00]	[27.55]	[24.79]
18	σ_c	-0.01	-0.04	-0.03	0.04	0.03	0.08
		(-0.73)	(-1.45)*	(-0.73)	(0.65)	(0.46)	(1.24)
		[-0.18]	[1.01]	[0.15]	[0.42]	[-0.03]	[1.54]
19	gap	-0.28	-0.95	-1.46	-2.02	-2.56	-2.99
		(-3.53)***	(-3.34)***	(-3.46)***	(-5.10)***	(-6.17)***	(-6.12)***
		[4.47]	[13.01]	[18.31]	[28.29]	[39.10]	[42.17]
20	cc	-0.42	-1.55	-2.57	-3.13	-4.08	-5.01
		(-3.06)***	(-3.59)***	(-3.96)***	(-4.47)***	(-4.57)***	(-4.17)***
		[3.16]	[11.43]	[18.68]	[22.46]	[31.53]	[34.46]

Table AIII

Implied Coefficients from a Restricted VAR

The table compares direct slope estimates obtained from long-run return regressions summarized in Table I against long-horizon coefficients implied by a restricted first-order VAR:

$$
\begin{aligned}
r_{t+1} &= \alpha + \beta cc_t + \varepsilon_{t+1}^r \\
cc_{t+1} &= \delta + \phi cc_t + \varepsilon_{t+1}^{cc},
\end{aligned}
$$

where r_{t+1} is the quarterly log stock market return and cc_t is one-quarter lagged cyclical consumption. We compute the implied long-horizon coefficients following Cochrane (2008) as $\beta^{(h)} = \beta\left(1-\phi^h\right)/(1-\phi)$, where h is the forecasting horizon. The table shows results for log excess market returns (Panel A), log real market returns (Panel B), and log market returns (Panel C) for the CRSP value-weighted index. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
			Panel A: Excess Market Returns			
Direct	-0.42***	-1.55***	-2.57***	-3.13***	-4.08***	-5.01***
Implied	-0.42***	-1.59***	-2.98***	-4.20***	-5.26***	-6.19***
			Panel B: Real Market Returns			
Direct	-0.40***	-1.51***	-2.52***	-3.10***	-4.12***	-5.22***
Implied	-0.40***	-1.54***	-2.88***	-4.05***	-5.08***	-5.97***
			Panel C: Market Returns			
Direct	-0.34***	-1.23***	-1.96***	-2.28***	-3.07***	-3.94***
Implied	-0.34**	-1.29**	-2.41**	-3.40**	-4.26**	-5.01**

Table AIV

In-sample Regressions with Alternative Values of Parameter k

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess return on the CRSP value-weighted index, and cc_t is one-quarter lagged cyclical consumption. Cyclical consumption is computed as specified in Equation (1) for different values of parameter k as indicated in the first column. For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1958Q4 to 2017Q4.

	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
$k=4$	-0.59	-2.37	-5.38	-5.96	-7.55	-10.57
	(-0.92)	(-1.60)*	(-2.40)**	(-3.01)***	(-3.24)***	(-4.10)***
	[0.22]	[2.22]	[7.53]	[7.17]	[10.00]	[15.72]
$k=8$	-0.52	-2.10	-3.43	-4.26	-6.18	-7.97
	(-1.60)*	(-2.31)**	(-2.58)***	(-3.02)***	(-3.96)***	(-4.68)***
	[1.06]	[5.70]	[9.06]	[11.02]	[20.05]	[26.09]
$k=12$	-0.33	-1.38	-2.58	-3.87	-5.46	-6.30
	(-1.37)*	(-2.14)**	(-2.57)***	(-3.54)***	(-4.80)***	(-5.39)***
	[0.60]	[4.04]	[8.58]	[15.51]	[26.34]	[27.38]
$k=16$	-0.28	-1.23	-2.65	-3.86	-4.85	-5.28
	(-1.53)*	(-2.40)**	(-3.37)***	(-4.52)***	(-5.40)***	(-4.77)***
	[0.61]	[4.47]	[12.85]	[21.55]	[28.92]	[27.09]
$k=20$	-0.36	-1.45	-2.79	-3.60	-4.29	-5.22
	(-2.13)**	(-3.43)***	(-4.01)***	(-5.04)***	(-5.24)***	(-4.66)***
	[1.63]	[8.24]	[18.15]	[23.72]	[28.68]	[31.98]
$k=24$	-0.36	-1.39	-2.39	-2.99	-4.01	-4.85
	(-2.55)***	(-3.23)***	(-3.67)***	(-4.12)***	(-4.28)***	(-3.84)***
	[2.19]	[9.13]	[16.11]	[19.84]	[29.33]	[31.05]
$k=28$	-0.30	-1.12	-1.95	-2.76	-3.72	-4.38
	(-2.38)***	(-2.94)***	(-3.18)***	(-3.61)***	(-3.91)***	(-3.54)***
	[1.65]	[6.77]	[12.12]	[18.95]	[27.32]	[26.94]
$k=32$	-0.26	-0.95	-1.92	-2.64	-3.40	-3.87
	(-2.36)**	(-2.58)***	(-2.94)***	(-3.34)***	(-3.62)***	(-3.02)***
	[1.27]	[5.49]	[13.20]	[18.83]	[24.55]	[22.93]
$k=36$	-0.28	-1.09	-1.97	-2.54	-3.20	-3.42
	(-2.44)**	(-2.80)***	(-2.98)***	(-3.38)***	(-3.50)***	(-2.65)***
	[1.82]	[7.85]	[14.70]	[18.04]	[22.62]	[19.00]
$k=40$	-0.29	-1.04	-1.79	-2.31	-2.74	-2.82
	(-2.71)***	(-2.72)***	(-2.85)***	(-3.26)***	(-2.97)***	(-2.36)***
	[1.93]	[7.21]	[12.09]	[15.00]	[16.91]	[13.23]
$k=44$	-0.24	-0.90	-1.62	-1.94	-2.26	-2.33
	(-2.15)**	(-2.34)***	(-2.66)***	(-2.61)***	(-2.53)***	(-2.08)***
	[1.28]	[5.44]	[10.17]	[10.90]	[11.96]	[9.48]

Table AV

qLL-statistics

The table shows qLL-statistics of Elliott and Müller (2006) for the benchmark predictive regressions summarized in Table I. The qLL-statistic tests the hypothesis H_0: $\beta_t = \beta$ for all t in $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log stock market return, and cc_t is one-quarter lagged cyclical consumption. The 10%, 5%, and 1% critical values are -12.80, -14.32, and -17.57. The table shows results for log excess market returns (Panel A), log real market returns (Panel B), and log market returns (Panel C) for the S&P 500 index and the CRSP value-weighted index.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
Panel A: Excess Market Returns						
SP500	-2.38	-3.91	-3.54	-3.60	-3.95	-4.04
CRSP	-2.09	-3.54	-3.35	-3.26	-3.66	-3.92
Panel B: Real Market Returns						
SP500	-2.89	-4.45	-4.12	-3.87	-3.96	-3.85
CRSP	-2.59	-4.17	-4.22	-4.13	-4.21	-4.06
Panel C: Market Returns						
SP500	-2.67	-4.34	-4.22	-4.24	-4.36	-4.07
CRSP	-2.41	-4.05	-4.21	-4.23	-4.29	-3.94

Table AVI

Two-state Predictive Regressions with Alternative Detrending Methods

The table presents results of two-state predictive regressions of the form $r_{t,t+h} = \alpha + \beta_{bad} I_{bad} cc_t + \beta_{good} (1 - I_{bad}) cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess return on the CRSP value-weighted index, cc_t is one-quarter lagged cyclical consumption, and I_{bad} is the state indicator that equals one during bad economic states and zero otherwise. Bad states are defined as periods with cyclical consumption realizations below its mean by more than one standard deviation. We compute cc by fitting a linear (Panel A), linear with a break (Panel B), quadratic (Panel C) or cubic (Panel D) time trend specification; or compute cyclical consumption as a five-year backward-looking moving average (Panel E). For each regression, the table reports the slope estimate, Newey-West corrected t-statistics in parentheses (h lags), and adjusted R^2 statistics in percent in square brackets. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values. The sample covers the period from 1953Q4 to 2017Q4.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
			Panel A: Linear			
β_{bad}	-0.14	-0.46	-0.81	-1.66	-2.40	-3.03
	(-1.25)	(-1.25)	(-1.47)*	(-2.51)***	(-3.32)***	(-2.75)***
β_{good}	-0.31	-1.24	-2.03	-2.04	-2.40	-2.63
	(-1.85)**	(-2.31)***	(-2.59)***	(-2.47)**	(-2.41)**	(-1.92)*
	[1.26]	[6.16]	[9.96]	[11.41]	[14.33]	[13.11]
			Panel B: Break			
β_{bad}	-0.47	-1.32	-1.36	-3.07	-5.55	-7.06
	(-1.49)*	(-1.48)*	(-1.53)*	(-2.44)**	(-2.63)***	(-2.26)**
β_{good}	-0.71	-2.97	-5.62	-6.31	-6.28	-6.11
	(-3.01)***	(-3.00)***	(-3.78)***	(-4.90)***	(-5.24)***	(-3.45)***
	[3.25]	[13.94]	[26.34]	[30.32]	[33.67]	[29.72]
			Panel C: Quadratic			
β_{bad}	-0.38	-1.24	-0.74	-1.67	-2.27	-2.00
	(-1.46)	(-1.68)*	(-0.90)	(-1.57)*	(-1.52)*	(-1.11)
β_{good}	-0.49	-2.04	-4.58	-4.65	-4.87	-4.91
	(-1.69)*	(-1.93)**	(-2.79)***	(-2.46)***	(-2.62)***	(-2.91)***
	[1.43]	[6.98]	[15.03]	[14.15]	[14.21]	[10.44]
			Panel D: Cubic			
β_{bad}	-0.52	-2.10	-2.78	-5.03	-6.44	-5.50
	(-1.20)	(-1.90)**	(-2.09)**	(-2.49)**	(-2.87)***	(-2.76)***
β_{good}	-1.04	-4.03	-7.47	-8.10	-9.08	-10.61
	(-3.20)***	(-3.28)***	(-4.34)***	(-5.99)***	(-5.28)***	(-4.70)***
	[3.76]	[15.84]	[29.30]	[33.35]	[38.06]	[34.83]
			Panel E: Stochastic			
β_{bad}	-6.48	-9.19	-15.00	-20.86	-20.35	-15.08
	(-1.89)*	(-1.38)	(-1.94)*	(-3.25)***	(-3.79)***	(-2.48)**
β_{good}	-1.29	-4.77	-8.82	-11.98	-15.40	-18.80
	(-2.33)**	(-3.20)***	(-3.59)***	(-5.10)***	(-6.38)***	(-6.12)***
	[2.53]	[7.01]	[14.85]	[22.32]	[31.73]	[36.07]

Table AVII

Out-of-sample Tests with Fixed Parameter Estimates

The table presents results of out-of-sample forecasts of h-quarter-ahead log excess returns on the CRSP value-weighted index where a time-varying expected returns model with cyclical consumption as regressor is compared against a constant expected returns model. The parameters used to calculate cyclical consumption are estimated over the full sample from the current latest-available consumption data. R^2_{OOS} is the out-of-sample R^2 in percent. ENC-NEW is the encompassing test statistic of Clark and McCracken (2001) and MSE-F is the F-statistic of McCracken (2007). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values in case of the ENC-NEW and MSE-F statistics, and according to the Clark and West (2007) test in case of the R^2_{OOS} statistics. The first observation in the out-of-sample period is 1980Q1, 1990Q1, or 2000Q1, and the predictive model is estimated recursively until 2017Q4.

	$h=1$	$h=4$	$h=8$	$h=12$	$h=16$	$h=20$
			Panel A: Forecasting from 1980			
ENC-NEW	4.35***	16.37***	30.65***	32.39***	41.51***	36.29***
MSE-F	2.06***	7.87**	16.15***	14.65***	15.19***	-0.93
R^2_{OOS}	1.33*	5.02***	10.02***	9.41***	9.98***	-0.70***
			Panel B: Forecasting from 1990			
ENC-NEW	4.21***	14.87***	20.98***	22.88***	32.58***	31.52***
MSE-F	3.71***	12.98***	17.33***	20.48***	30.49***	26.46***
R^2_{OOS}	3.21**	10.64***	14.17***	16.86***	23.92***	22.15***
			Panel C: Forecasting from 2000			
ENC-NEW	3.86***	12.22***	16.60***	21.32***	36.56***	46.12***
MSE-F	3.92***	12.70***	15.71***	21.75***	37.20***	50.11***
R^2_{OOS}	5.16**	15.55***	19.47***	26.29***	39.49***	48.60***

Table AVIII

Out-of-sample Tests with Vintage Data

The table presents results of out-of-sample forecasts of h-quarter-ahead log excess returns on the CRSP value-weighted index where a time-varying expected returns model with cyclical consumption as regressor is compared against a constant expected returns model. The parameters used to calculate cyclical consumption are estimated recursively from historical vintages of consumption data from the ALFRED database at St. Louis Fed, and two-quarter lag is used in the out-of-sample forecasting regressions to take account for delays in macroeconomic releases. R^2_{OOS} is the out-of-sample R^2 in percent. ENC-NEW is the encompassing test statistic of Clark and McCracken (2001) and MSE-F is the F-statistic of McCracken (2007). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, according to one-sided wild bootstrap p-values in case of the ENC-NEW and MSE-F statistics, and according to the Clark and West (2007) test in case of the R^2_{OOS} statistics. The first observation in the out-of-sample period is 1980Q1, 1990Q1 or 2000Q1, and the predictive model is estimated recursively until 2017Q4.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
			Panel A: Forecasting from 1980			
ENC-NEW	4.68***	20.27***	38.50***	54.90***	93.22***	95.43***
MSE-F	2.36***	13.24***	30.65***	47.58***	76.56***	72.31***
R^2_{OOS}	1.53**	8.16***	17.45***	25.23***	35.85***	35.22***
			Panel B: Forecasting from 1990			
ENC-NEW	4.25***	17.16***	30.61***	46.50***	77.78***	89.86***
MSE-F	2.47***	11.53***	23.35***	40.68***	67.67***	80.39***
R^2_{OOS}	2.16*	9.56***	18.19***	28.71***	41.09***	46.36***
			Panel C: Forecasting from 2000			
ENC-NEW	3.72***	11.65***	17.51***	26.44***	44.68***	61.13***
MSE-F	2.45***	8.23***	11.29***	16.95***	26.48***	45.67***
R^2_{OOS}	3.29*	10.66***	14.80***	21.74***	31.72***	46.28***

Table AIX

Dynamics of Simulated and Historical Data

The table presents descriptive statistics for cyclical consumption. It shows the time-series averages (Mean), standard deviations (Std), and the first five autocorrelation coefficients. Cyclical consumption is computed as specified in Equations (1) and (11) in the actual and simulated data for $k = 24$, respectively. We generate 1,000,000 quarterly observations based on the calibrated parameter values of Campbell and Cochrane (1999). The empirical sample covers the period from 1953Q4 to 2017Q4.

	Mean	Std	AC1	AC2	AC3	AC4	AC5
Simulated data	0.00	0.04	0.96	0.92	0.87	0.83	0.79
Actual data	0.00	0.04	0.97	0.93	0.88	0.83	0.77

Table AX

Model-implied Predictability Across Various Cycle Lengths

The table presents results of predictive regressions of the form $r_{t,t+h} = \alpha + \beta cc_t + \varepsilon_{t,t+h}$, where h denotes the horizon in quarters, $r_{t,t+h}$ is the h-quarter ahead log excess market return, and cc_t is one-quarter lagged cyclical consumption computed as specified in Equation (11) for different values of parameter k as indicated in the first column. For each regression, the table reports OLS estimates of the regressor and adjusted R^2 statistics in percent in square brackets based on 1,000,000 quarterly simulated observations.

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
Panel A: Consumption Claim						
$k = 4$	-0.32	-1.19	-2.18	-3.03	-3.75	-4.35
	[0.64]	[2.34]	[4.23]	[5.81]	[7.13]	[8.21]
$k = 8$	-0.29	-1.09	-2.01	-2.80	-3.46	-4.02
	[1.07]	[3.94]	[7.18]	[9.89]	[12.13]	[13.97]
$k = 12$	-0.27	-1.01	-1.87	-2.59	-3.21	-3.73
	[1.36]	[5.05]	[9.24]	[12.72]	[15.60]	[18.05]
$k = 16$	-0.25	-0.94	-1.73	-2.41	-2.98	-3.47
	[1.56]	[5.81]	[10.62]	[14.61]	[17.99]	[20.80]
$k = 20$	-0.23	-0.88	-1.61	-2.24	-2.78	-3.24
	[1.69]	[6.28]	[11.47]	[15.86]	[19.51]	[22.63]
$k = 24$	-0.22	-0.82	-1.51	-2.10	-2.60	-3.04
	[1.75]	[6.53]	[12.02]	[16.59]	[20.48]	[23.80]
$k = 28$	-0.20	-0.76	-1.41	-1.97	-2.44	-2.85
	[1.79]	[6.71]	[12.30]	[17.02]	[21.05]	[24.41]
$k = 32$	-0.19	-0.72	-1.33	-1.85	-2.29	-2.67
	[1.80]	[6.72]	[12.39]	[17.19]	[21.21]	[24.59]
$k = 36$	-0.18	-0.67	-1.25	-1.74	-2.16	-2.51
	[1.78]	[6.72]	[12.39]	[17.13]	[21.12]	[24.51]
$k = 40$	-0.17	-0.64	-1.18	-1.64	-2.03	-2.37
	[1.78]	[6.65]	[12.22]	[16.89]	[20.84]	[24.17]
$k = 44$	-0.16	-0.60	-1.11	-1.55	-1.92	-2.23
	[1.73]	[6.49]	[11.94]	[16.54]	[20.40]	[23.67]

	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$
			Panel B: Dividend Claim			
$k = 4$	-0.32	-1.20	-2.20	-3.06	-3.78	-4.40
	[0.53]	[1.91]	[3.42]	[4.62]	[5.57]	[6.34]
$k = 8$	-0.30	-1.10	-2.03	-2.82	-3.50	-4.06
	[0.88]	[3.23]	[5.81]	[7.86]	[9.52]	[10.79]
$k = 12$	-0.27	-1.02	-1.88	-2.62	-3.24	-3.77
	[1.12]	[4.15]	[7.47]	[10.13]	[12.23]	[13.89]
$k = 16$	-0.25	-0.95	-1.75	-2.43	-3.01	-3.50
	[1.29]	[4.76]	[8.59]	[11.63]	[14.06]	[15.96]
$k = 20$	-0.24	-0.88	-1.63	-2.26	-2.80	-3.26
	[1.41]	[5.16]	[9.29]	[12.60]	[15.22]	[17.31]
$k = 24$	-0.22	-0.82	-1.52	-2.11	-2.61	-3.05
	[1.46]	[5.36]	[9.68]	[13.13]	[15.90]	[18.13]
$k = 28$	-0.21	-0.77	-1.42	-1.97	-2.45	-2.85
	[1.48]	[5.47]	[9.68]	[13.41]	[16.27]	[18.53]
$k = 32$	-0.19	-0.72	-1.33	-1.85	-2.30	-2.67
	[1.48]	[5.47]	[9.91]	[13.50]	[16.36]	[18.63]
$k = 36$	-0.18	-0.68	-1.25	-1.74	-2.16	-2.52
	[1.47]	[5.44]	[9.88]	[13.42]	[16.27]	[18.55]
$k = 40$	-0.17	-0.64	-1.18	-1.64	-2.03	-2.37
	[1.46]	[5.39]	[9.73]	[13.23]	[16.06]	[18.30]
$k = 44$	-0.16	-0.60	-1.11	-1.55	-1.92	-2.24
	[1.42]	[5.25]	[9.51]	[12.96]	[15.72]	[17.92]

References

Campbell, John Y., and James H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205-251.

Campbell, John Y., and Samuel B. Thompson, 2008, Predicting excess stock returns out-of-sample: Can anything beat the historical average? *Review of Financial Studies* 21, 1509-1531.

Chen, Long, 2009, On the reversal of return and dividend predictability, *Journal of Financial Economics* 92, 128-151.

Clark, Todd E., and Michael W. McCracken, 2001, Tests of equal forecasting accuracy and encompassing for nested models, *Journal of Econometrics* 105, 85-110.

Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291-311.

Cochrane, John H., 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533-1575.

Cooper, Ilan, and Richard Priestley, 2009, Time-varying risk premiums and the output gap, *Review of Financial Studies* 22, 2801-2833.

da Costa, Carlos Eugênio, and Jivago B. Ximenes de Vasconcelos, 2009, Can a habit formation model really explain the forward premium anomaly? Working paper.

Diebold, Francis X., and Roberto S. Mariano, 1995, Comparing predictive accuracy, *Journal of Business & Economic Statistics* 13, 134-144.

Elliott, Graham, and Ulrich K. Müller, 2006, Efficient tests for general persistent time variation in regression coefficients, *Review of Economic Studies* 73, 907-940.

Harvey, David I., Stephen J. Leybourne, and Paul Newbold, 1998, Tests for forecast encompassing, *Journal of Business and Economic Statistics* 16, 254-259.

McCracken, Michael W., 2007, Asymptotics for out of sample tests of Granger causality, *Journal of Econometrics* 140, 719-752.

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

Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455-1508.