Beyond Growth Rates: Macro Trends and Price Cycles*

Carlo A. Favero† Alessandro Melone‡ Sean Myers§
Andrea Tamoni¶

February 2, 2026

Abstract

This paper provides new evidence on the link between asset prices and macroeconomic fundamentals. We show that aggregate stock price and consumption levels share a common stochastic trend. Temporary deviations of stocks from this long-run relationship – which we label price-consumption cycles – are mean-reverting and contain valuable information about expected returns. Price cycles forecast future returns both in- and out-of-sample and yield economically meaningful improvements in investor welfare. The predictive content does not reflect forecasts of consumption growth and disappears when consumption is omitted from the long-run relation, indicating that the information originates from the consumption-based component of prices. Our results indicate that asset prices are anchored to consumption in the long run, and departures from this anchor signal time variation in expected returns.

Keywords: Consumption Levels, Cointegration, Time-Varying Equity Premium, Return Predictability.
JEL codes: C22, E32, E44, G12.

*We thank Ravi Bansal, Robert Engle, Stefano Giglio, Andrei Gonçalves, Ravi Jagannathan, Sydney Ludvigson, Konark Saxena, and Allan Timmermann for helpful discussions and suggestions.

†Bocconi University & CEPR. E-mail: carlo.favero@unibocconi.it.
‡Fisher College of Business, The Ohio State University. E-mail: melone.11@osu.edu.
§The Wharton School, University of Pennsylvania. E-mail: semyers@wharton.upenn.edu.
¶Mendoza College of Business, University of Notre Dame. E-mail: atamoni@nd.edu.

1 Introduction

A central challenge for consumption-based asset-pricing models is the weak empirical correlation between aggregate stock returns and consumption growth, a disconnect commonly referred to as the correlation puzzle (Cochrane and Hansen, 1992; Daniel and Marshall, 1997; Campbell and Cochrane, 1999). A large literature has sought to resolve this puzzle by modifying preferences, frictions, and information structures (e.g., Cochrane, 2005; Ludvigson, 2013; Campbell, 2017). Much less attention, however, has been devoted to a closely related question: what is the relationship between the level of consumption and the level of aggregate stock market?

This paper shows that consumption contains valuable information for asset prices once one studies *levels* rather than *growth rates*. We document that consumption levels and stock market prices share a long-run equilibrium relationship. This finding has two key implications for macro-finance.

First, if consumption and stock prices are cointegrated, the correlation between one-period consumption growth and one-period returns will be weak when temporary price shocks dominate consumption volatility. This mechanism offers a new perspective on the correlation puzzle: the puzzle arises naturally when the permanent component of stock prices comoves with consumption, but transitory price fluctuations obscure this link.

Second, we show that the existence of a correlation puzzle necessarily implies return predictability. When aggregate equity price and consumption levels share a common stochastic trend, any component of one-period returns not accounted for by consumption growth is by construction transitory and must eventually reverse. This implication is particularly relevant in light of the widespread view that time variation in expected returns explains a substantial fraction of stock price fluctuations (Shiller, 1981; LeRoy and Porter, 1981), and of recent evidence questioning both traditional valuation-based predictors (e.g., Goyal et al., 2024) and the stability of price–variance decompositions (e.g., Nagel, 2024).

Overall, our finding that deviations of aggregate stock prices from a broad macroeconomic trend forecast subsequent returns aligns with the trend–cycle, or permanent–transitory, decomposition of prices proposed by Fama and French (1988b). In this framework, log stock prices consist of a permanent component that follows a random walk with drift and a transitory component that mean-reverts over time, generating return predictability. We extend this idea by allowing the permanent component to be common across assets and tied to trends in the real economy, highlighting how deviations of prices

from their shared macroeconomic trend give rise to predictable return dynamics.

Our empirical analysis builds on this insight and documents a strong long-run relationship between asset prices and macroeconomic fundamentals, specifically aggregate consumption. Using standard cointegration methods, we find robust evidence that portfolio prices and consumption share a common stochastic trend. This result holds across a wide range of portfolios, including the market, value, size, and growth portfolios, for which we consistently reject the null of no long-run relationship.

Having established this long-run relationship, we turn to the short-run dynamics of return predictability, where our main methodological contribution arises. A long-standing idea in asset pricing is that, although stock prices may deviate from fundamentals in the short run, they are ultimately anchored to macroeconomic forces, and such deviations may contain information about future returns. Much of the empirical literature has pursued this idea through valuation ratios, such as price–dividend or price–earnings ratios, which implicitly assume a stable and known long-run relationship between prices and fundamentals. Under this approach, predictability arises because departures from a fixed mean valuation ratio are expected to reverse over time. While influential, this framework relies on strong restrictions that are difficult to reconcile with structural change and regime shifts, leading valuation ratios to be highly persistent and primarily informative about returns at very long horizons, with weaker and less stable predictive power in real-time.

A separate strand of the literature seeks to relax these restrictions by estimating long-run relationships among macroeconomic variables related to asset prices, such as consumption or aggregate wealth, and using deviations from these relationships to forecast returns. By excluding prices from the estimation of long-run trends, this approach avoids mechanically inducing predictability. Its effectiveness, however, depends on how closely these variables track stock prices: when they adjust only gradually to market valuations, deviations from such relationships naturally predict returns only at long horizons.

Our approach differs from both strands. Rather than assuming a fixed long-run relationship or inferring price dynamics indirectly through related quantities, we directly estimate the long-run relationship between the level of stock prices and the level of aggregate consumption. Crucially, we estimate this relationship in real-time, which prevents predictability from arising mechanically from in-sample fitting, while allowing the long-run relationship itself to evolve over time. Our framework therefore preserves the economic intuition that prices are tethered to fundamentals, without imposing rigid assumptions about the stability or exact form of that relationship.

Using this methodology, we document that the real-time market's long-run beta with respect to consumption hovers around one and displays transitory fluctuations during major recessions such as the global financial crisis and the COVID-19 downturn. Portfolios with higher expected returns, such as value and small-cap stocks, exhibit consistently higher long-run consumption betas than their growth and large-cap counterparts. Moreover, deviations of prices from consumption fundamentals co-move strongly across portfolios, and a single common factor explains most of their variation. This pervasive component suggests that the long-run relationship between financial markets and consumption is driven by an aggregate economic source, rather than portfolio-specific dynamics.

Price deviations from their consumption-implied long-run values – what we refer to as the price-consumption cycle – are moderately persistent, with estimated half-lives ranging from about 1.7 years for small and value portfolios to roughly three years for large and growth portfolios. These dynamics imply that prices tend to realign with consumption fundamentals over the business cycle. Consistent with this interpretation, price deviations are powerful predictors of future returns, both statistically and economically. When a portfolio's price rises above its consumption-implied long-run value, subsequent returns tend to be lower. This pattern holds for the market and for characteristic-sorted portfolios, with a one-unit deviation in log price predicting roughly a half percentage-point decline in next-year returns. The price-consumption cycle explains roughly 4 percent of the variation in quarterly returns—an explanatory power that is comparable to, or even exceeds, that of traditional valuation ratios. Predictability is strongest at short and medium horizons, gradually diminishing beyond five years, as expected price deviations are eventually reabsorbed. The results are robust to alternative return definitions and to substituting aggregate consumption with other measures of real activity such as GDP. Altogether, the evidence points to prices and returns being anchored to the real economy through consumption-based fundamentals, establishing a stable and economically meaningful linkage between asset markets and macroeconomic dynamics.

Out-of-sample tests confirm that this predictability is not a statistical artifact. Forecasts based solely on information available in real-time continue to outperform the historical mean benchmark at short horizons. Although predictive performance naturally declines relative to the in-sample regressions, the out-of-sample R^2 remains sizable for horizons up to one year. Some portfolios, such as value and growth, experience a deterioration in performance during the early 2000s and certain post-crisis periods. However, outside these episodes, the real-time forecasts track their in-sample counterparts closely, underscoring the economic relevance and

stability of the link between price-consumption cycles and future returns.

We next quantify the economic value of this predictability from the perspective of an investor. Timing strategies that condition on price deviations yield substantially higher risk-adjusted performance than simple buy-and-hold benchmarks. Sharpe ratios rise by about 45 percent for the market portfolio and nearly doubles for the growth portfolio, with smaller but still meaningful improvements for the size and value portfolios. When evaluated through the lens of a mean–variance investor, the information contained in price deviations translates into annualized utility gains between 2 and 3.5 percent of wealth—economically large figures comparable to the management fees investors would be willing to pay for such forecasting ability. These results demonstrate that our price cycle is not only statistically significant but also economically valuable for portfolio management and market timing.

Finally, we examine whether the price-cycles we identify merely captures consumption dynamics or known valuation effects. We find that deviations from the long-run consumption relationship do not predict future consumption growth at any horizon, ruling out the possibility that our results simply reflect the anticipation of macroeconomic conditions. Moreover, when we omit consumption from the estimation of the long-run relation, the resulting price cycle variable loses all predictive power for returns. The predictive information therefore originates specifically from the link between asset prices and consumption, rather than from a generic time trend. Compared with the consumption–wealth ratio of Lettau and Ludvigson (2001, 2005), our price cycle measure captures more pronounced medium-term fluctuations in prices that traditional wealth-based aggregates fail to track, particularly in the post-1980 period of accelerated equity growth.

Taken together, these findings reveal a stable relationship between asset prices and consumption that embeds information about expected returns at both short and long horizons. Price deviations from this consumption-implied equilibrium serve as a real-time indicator of variation in future returns, generating robust predictability across portfolios and horizons and delivering substantial economic value to investors. The results highlight a powerful and intuitive mechanism through which financial markets and the real economy are jointly determined: when prices rise faster than consumption, expected returns fall, and vice versa. Overall, the evidence establishes the return–consumption nexus as a central feature of asset price dynamics and provides a macroeconomic foundation for time-series variation in expected returns.

Related literature. It is important to immediately separate our paper from the large literature that has tried to establish a link between the macroeconomy and financial assets. Since the seminal contribution of Chen et al. (1986), and motivated by the theoretical work of Merton (1973), researchers have explored the possibility that state variables like inflation and economic growth are a source of systematic investment risk and can explain the cross-sectional dispersion of stock returns.[1] The standard approach is to extract innovations and to estimate risk compensation using a Fama-MacBeth procedure (or its variants). We instead propose a cointegration framework that exploits long-run comovement between macroeconomic trends and factor prices, and use it to predict the time-series variation in a given equity factor. In this respect, our work contributes to the literature that exploits cointegration between dividends and prices (Campbell and Shiller, 1988), or between consumption and wealth (Lettau and Ludvigson, 2001), to predict the aggregate market factor. We discuss our connection to these literatures in more detail in Section 2.1. Additionally, we show that factors beyond the aggregate market are predictable using macroeconomic variables and a similar cointegration logic. Hansen et al. (2008) and Bansal et al. (2009) exploit the cointegrating relationship between consumption and dividends of portfolios sorted on characteristics such as book-to-market to demonstrate how long-run risk contributes to cross-sectional variation in returns. Bansal and Kiku (2011) show that the cointegration between dividends and aggregate consumption has implications for the implementation of optimal portfolio rules, particularly at long investment horizons. We complement this line of research by showing that cointegration between asset prices and real economic trends is informative about the time-series predictability of market and characteristic-sorted portfolio returns.

Our paper is another step toward unifying cross-sectional and time-series predictability of returns. Despite the large literature documenting the variability of aggregate stock returns over time (e.g., Shiller, 1981; Fama and French, 1988a), factor models (e.g., Fama and French, 1993) often abstract from the predictability of factors and mainly focus on their ability to generate cross-sectional dispersion in asset risk premiums. Haddad et al. (2020) constitute a notable exception. These authors propose a new statistical approach to predict "anomaly" portfolios, by predicting their principal components (PCs) using their own book-to-market ratios. Rather than predicting PCs, we predict directly characteristics-based factor returns like those used in the Fama and French (2015) or Hou et al. (2015) models. Furthermore, we rely on macroeconomic fundamentals rather than financial ratios.

[1] See, e.g., Ferson and Harvey (1991); Vassalou (2003); Boons, Duarte, De Roon and Szymanowska (2020); Barroso, Boons and Karehnke (2021); Rapach and Zhou (2021).

An influential literature examines how long-run betas explain the cross-section of stock returns. Parker and Julliard (2005) show that the covariance between stock returns and *future* three-year consumption growth explains a substantial portion of the variation in returns across market-to-book portfolios. Bansal et al. (2005) further demonstrate that covariances between cash flow growth and *past* consumption growth help explain the value premium. More recently, Bandi and Tamoni (2023) show that a business-cycle beta—computed from the covariance between consumption growth aggregated over a four-year horizon and four-year holding-period returns—provides a valuable pricing signal.

Given these results, evaluating consumption risk is naturally suited to approaches based on long-run factor loadings (e.g., Kamara et al., 2016) and spectral factor loadings (e.g., Bandi et al., 2021). Our analysis builds on this literature by employing a very long-run (beyond business-cycle) measure of betas, computed using cumulative stock returns as proxies for prices and cumulative consumption growth rates, i.e., consumption levels. Consistent with the literature on factor loadings derived from aggregation and frequency-specific decompositions, our long-run betas are higher for value and small portfolios and lower for growth and large portfolios.

Our contribution, however, differs in an important way. We show that the long-run correlation between asset prices and consumption also contains information about short-run price deviations—what we refer to as the price cycle. The time-series predictability associated with this short-run dislocation between financial markets and the real economy is the central contribution of our analysis.

2 The Time Series of Consumption and Asset Prices

Price growth, or alternatively returns, and consumption growth are commonly characterized in the literature as stationary time series. This, in turn, implies that log prices, p_t, and log consumption, c_t, are non-stationary unit-root processes. Given log consumption growth Δc_t and a generic asset with log price growth Δp_t, we have:

$$c_t = c_{t-1} + \Delta c_t \ , \tag{1}$$

$$p_t = p_{t-1} + \Delta p_t \ . \tag{2}$$

An interesting, and somewhat overlooked, question in the literature is whether the

stochastic trend in consumption and price growth is common, or whether the two series are driven by independent trends. While a vast literature links consumption growth to asset returns (see, e.g., Cochrane, 2005; Campbell, 2017), to the best of our knowledge, we are the first paper to estimate the relationship between the level of aggregate consumption and the level of stock prices,[2]

$$p_t = \alpha_0 + \alpha_1 t + \beta c_t + \delta_t \ . \tag{3}$$

As discussed in Section 2.1, existing studies either assume the parameters of the relationship by setting $\alpha_1 = 0$ and $\beta = 1$ and studying a fixed ratio of price to consumption or focus on a cointegrating relationship between consumption and non-price objects such as dividends. Section 2.2 also discusses why we should include a time trend in this relationship even if we believe that prices and consumption have the same average growth over time.

Inspired by Fama and French (1988b),[3] we consider a simple case in which the residuals δ_t follow an AR(1) process,

$$\delta_t = \rho \delta_{t-1} + v_t \ . \tag{4}$$

If the residuals δ_t are stationary, i.e., $|\rho| < 1$, then consumption and prices are cointegrated. In this case, they share a common stochastic trend, and δ_t captures the cyclical deviation of prices from the long-run value implied by the trend in consumption. We therefore refer to δ_t as the "price-consumption cycle" associated with the asset at time t.[4]

The presence of a common trend, together with the stationarity of the price-consumption cycle, has direct implications for the dynamics of one-period price growth. Differencing equation (3), we have

$$\Delta p_t = \alpha_1 + \beta \Delta c_t + (\rho - 1)\delta_{t-1} + v_t \ . \tag{5}$$

Comparing equations (3) and (5) yields three implications for traditional models relating

[2] Throughout the paper, both stock prices and consumption are in real terms, rather than nominal, to ensure there is no mechanical cointegrating relationship between stock prices and consumption purely through inflation.

[3] Fama and French (1988b), equations (1)–(4), propose a model in which prices are the sum of a random walk component, q_t, and an AR(1) cyclical component, z_t. This maps directly to our equations (3)–(4), with $z_t = \delta_t$ and $q_t = \alpha_1 t + \beta c_t$, under the assumption that consumption growth is i.i.d.

[4] In the terminology of cointegration studies (Engle and Granger, 1987; Hendry, 1995), the price-consumption cycle corresponds to the "error-correction term."

price growth, or similarly asset returns, to consumption growth.

First, if prices are cointegrated with consumption, one-period price growth may display only a weak correlation with one-period consumption growth. This occurs when the volatility of shocks to the price-consumption cycle v_t dominates the volatility of consumption growth. Second, cointegration between consumption and prices implies predictability in one-period price growth. The $(\rho-1)\delta_{t-1}$ term in equation (5) indicates that any short-term price growth not explained by consumption growth, $\beta\Delta c_t$, is temporary and will eventually be reversed. Predictability disappears only when the price-consumption cycle follows a unit root ($\rho = 1$), in which case prices and consumption do not share a common stochastic trend. Finally, while this framework of cointegration between consumption and prices is interpreted as informative about future price growth, it could alternatively indicate that deviations from the long-run relationship forecast future consumption growth (Engle and Granger, 1987; Johansen, 1995).

2.1 Connection to price ratios and other cointegrating relations

The central challenge in using equations (3)-(4) to predict future price growth is that if one estimates these equations in-sample, then the deviation δ_t will mechanically predict future price growth in-sample. To avoid this mechanical result, the literature has taken two approaches.

Approach 1 is to assume the coefficients of equation (3) so that these parameters do not need to be estimated in-sample. Specifically, many papers assume $\alpha_1 = 0$, $\beta = 1$, and $\alpha_0 = \bar{p} - \bar{c}$, where $\bar{p} - \bar{c}$ is the unconditional mean for the price-to-consumption ratio. In this case, the deviation is simply the current price-to-consumption ratio less its unconditional mean.[5] Because the parameter values are assumed rather than estimated, one can then use the deviation to predict future price growth or returns in-sample without fear of mechanical success.

The downside to approach 1 is that these assumptions may be too rigid. If the mean price-to-consumption ratio changes over time, either in a continuous way or via structural breaks, then the deviation may not be an accurate predictor of future price growth because it incorrectly predicts reversion to the unconditional mean. This might explain why these

[5]Other papers use a similar approach of assuming $\alpha_1 = 0$ and $\beta = 1$ using dividends or other measures of macroeconomic fundamentals rather than consumption. These predictors (e.g., the price-to-dividend ratio) carry similar benefits and potential downsides as what we discuss for the price-to-consumption ratio and Section 4.3 shows the empirical differences between our estimated price-consumption cycle δ_t and a variety of ratios of price to consumption, dividends, and earnings.

types of predictors often perform poorly out of sample (Goyal and Welch, 2007; Goyal et al., 2024). Similarly, the assumption of $\beta = 1$ may be inaccurate. Even small differences between the true value of β and 1 could meaningfully impact the assessment of whether prices are too high or too low relative to consumption. As discussed in Section 2.2, just because one believe prices and consumption have the same average growth rate (i.e., one believes the price-to-consumption ratio does not trend to +/-∞ over time), does not imply that $\beta = 1$.

Approach 2 is to estimate a long-term relationship that does not include price p_t. One can assume that stock prices are closely related to or even cointegrated with some other variable x_t, and then estimate the relationship

$$x_t = \alpha_0 + \alpha_1 t + \beta c_t + \varepsilon_t \tag{6}$$

in-sample. Because price data is not used when estimating α_0, α_1, and β, one can use ε_t to predict future price growth in-sample. For example, Bansal et al. (2001) use dividends as their measure of x_t and Lettau and Ludvigson (2001) use a combination of human capital and asset holdings as their measure of x_t.

The downside to approach 2 is that the parameters $(\alpha_0, \alpha_1, \beta)$ are trained on data about x_t rather than data on prices. This means that the ability to predict future price growth depends on how strongly x_t is related to prices. If x_t is only strongly related to prices at long horizons (i.e., there can be persistent deviations between x_t and p_t) or the relationship between x_t and p_t fluctuates over time, then ε_t may only predict future price growth at long horizons and may perform poorly out of sample.

Compared to these two approaches, we take a different path. We focus on estimating equation (3) in real-time and then predicting future price growth and returns. Like the two previous approaches, this addresses the issue that an in-sample estimate of equation (3) would mechanically predict future price growth. Relative to approach 1, our approach allows us to better capture changes in parameters over time and allows for $\beta \neq 1$.[6] Relative to approach 2, our approach directly uses data on prices to train the $(\alpha_0, \alpha_1, \beta)$ parameters. Rather than excluding all price data to avoid look-ahead bias (i.e., replacing p_t with x_t), we exclude future price data and estimate our parameters in real-time. Section 4.3 shows (i) that our estimated δ_t has a much lower persistence than typical approach 1 predictors (e.g. the price-to-consumption ratio, the price-to-dividend ratio), meaning that we find faster and more consistent reversion to the mean and (ii) our estimated δ_t reflects information that

[6] Section 3.2 shows that our parameter estimates change noticeably over the last 50 years.

is not captured in the Bansal et al. (2001), Hansen et al. (2008), and Bansal et al. (2009) dividend-consumption deviations nor the Lettau and Ludvigson (2001) *cay*. Further, Section 4.3 shows that, compared to both approach 1 and approach 2 predictors, δ_t provides stronger price growth predictability, particularly at short and medium horizons.

2.2 The role of a time trend

We know that both prices and consumption tend to increase over time, and we can define $\tilde{p}_t$ and $\tilde{c}_t$ by regressing prices and consumption onto time,[7]

$$c_t = \alpha_{0,c} + \alpha_{1,c} t + \tilde{c}_t \ , \tag{7}$$

$$p_t = \alpha_{0,p} + \alpha_{1,p} t + \tilde{p}_t \ . \tag{8}$$

Suppose that, instead of estimating equation (3), we estimated a restricted version without a time trend,

$$p_t = \alpha_0 + \beta c_t + \delta_t \ . \tag{9}$$

If $\tilde{p}_t$ and $\tilde{c}_t$ are both stationary, then the best fit β will be entirely determined by average growth rates, $\beta = \frac{\alpha_{1,p}}{\alpha_{1,c}}$, rather than any information about the comovement between prices and consumption. If we want β to reflect information about comovement rather than just information about average growth rates, we need to include a time trend in equation (3).

A more specific version of this point is the following. There is a reasonable belief that $\alpha_{1,c} = \alpha_{1,p}$, as $\alpha_{1,c} \neq \alpha_{1,p}$ would imply that the price-to-consumption ratio trends to +/- ∞ over time. If we are using equation (9), then $\alpha_{1,c} = \alpha_{1,p}$ implies β must be 1. This is how we get the common approach of assuming $\alpha_1 = 0$ and $\beta = 1$. However, once we consider equation (3) and allow $\alpha_1 \neq 0$, then $\alpha_{1,p} = \alpha_{1,c}$ places no restriction on β and we are free to estimate β from data on comovements instead of from an assumption on average growth rates. Rephrased, $\alpha_{1,p} = \alpha_{1,c}$ does not imply that we do not need time trend in equation (3). Having $\alpha_{1,p} = \alpha_{1,c}$ only implies that we do not need a time trend if we also believe that $\beta = 1$. If we want to consider $\beta \neq 1$, then $\alpha_{1,p} = \alpha_{1,c}$ directly implies that we need to include a time trend $\alpha_1 \neq 0$ in equation (3).

[7] Note that even if prices and consumption do not have constant time trends, over any finite sample we can always regress price and consumption onto time to get $\alpha_{1,c}$ and $\alpha_{1,p}$.

2.2.1 Simple model of asset prices

To aid in the understanding of Section 2.2, we provide a simple model in which we can have $\beta \neq 1$ even if prices and consumption have the same average growth rate. In Appendix B, we provide a more detailed model that also discusses the cointegrating relationship between cumultative returns and consumption, and show two specific examples of the model: one in which the price-consumption cycle is driven by time-varying risk and another in which the price-consumption cycle is driven by behavioral biases in consumption expectations.

We assume that consumption growth is iid,

$$\Delta c_{t+1} = \alpha_{1,c} + u_{t+1} \ . \tag{10}$$

Consider an asset that pays $\exp(\alpha_1 t + \beta c_t)$ each period. We assume that the expected return on the asset is stationary. Using the Campbell-Shiller log-linear approximation, we have that the log price of this asset is

$$p_t = \alpha_0 + \alpha_1 t + \beta c_t - \sum_{j=1}^{\infty} \phi^{j-1} \left(E_t \left[r_{t+j} \right] - \bar{r} \right) \ . \tag{11}$$

where $\bar{r}$ is the average return on the asset and (α_0, ϕ) are log-linearization constants. The average growth for prices is $\alpha_{1,p} = \alpha_1 + \beta \alpha_{1,c}$. Even if $\alpha_{1,p} = \alpha_{1,c}$, we will not have $\alpha_1 = 0$ if $\beta \neq 1$. If expected returns follow an AR(1) process,

$$E_{t+1} \left[r_{t+2} \right] = \bar{r} + \rho \left(E_t \left[r_{t+1} \right] - \bar{r} \right) + u^r_{t+1} \tag{12}$$

then prices follow our exact specification in equations (3)-(4). Specifically, the price-consumption cycle will be

$$\delta_t = \rho \delta_{t-1} - \frac{1}{1 - \rho \phi} u^r_t. \tag{13}$$

2.3 Real-Time Estimation

For clarity, we use j to distinguish the different assets, meaning that our empirical specification is

$$p_{j,t} = \alpha_{0,j} + \alpha_{1,j} t + \beta_j c_t + \delta_{j,t} \ , \tag{14}$$

$$\delta_{j,t} = \rho_j \delta_{j,t-1} + v_{j,t} \ . \tag{15}$$

As mentioned previously, estimating equations (14)-(15) over the entire sample will mechanically generate a relationship between the current price-consumption cycle $\delta_{j,t}$ and future price growth. To avoid this mechanical relationship and look-ahead bias more generally, we estimate equations (14)-(15) for each asset j in real-time.

A naive approach is to simply estimate equation (14) in real-time by regressing price onto a time trend and consumption using the observations available up to time t. One could then use the residuals of this regression as an estimate $\widehat{\delta}_{j,t}$ for the price-consumption cycle. However, this places no restrictions on the estimated $\widehat{\delta}_{j,t}$ and risks overfitting. The regression residuals will simply be whatever values perfectly fit the observed sample, even if those values imply an erratic time series $\widehat{\delta}_{j,t}$.

To reduce this risk of overfitting, we incorporate the information from equation (15) into our estimation, which restricts the price-consumption cycle to be AR(1). Imposing this AR(1) restriction onto equation (14) gives a linear relationship that conveniently isolates the iid shock $v_{j,t}$,

$$p_{j,t} \ = \ \tilde{\alpha}_{0,j} + \tilde{\alpha}_{1,j} t + \rho_j p_{j,t-1} + \beta_j c_t - \rho_j \beta_j c_{t-1} + v_{j,t}, \tag{16}$$

where $\tilde{\alpha}_{0,j} \equiv \alpha_{0,j} (1 - \rho_j) + \rho_j \alpha_{1,j}$ and $\tilde{\alpha}_{1,j} \equiv (1 - \rho_j) \alpha_{1,j}$. The price-consumption cycle $\delta_{j,t}$ is then a combination of its lag and the $v_{j,t}$ shock following equation (15).[8]

To implement equations (15) and (16) in real-time, we use the following procedure. We enter into period t with an estimate $\widehat{\delta}_{j,t-1}$ for the previous value of the price-consumption cycle. We then estimate a regression of $p_{j,\tau}$ on a time trend, its lag, contemporaneous log

[8] The estimation of our (16) builds on a well-established tradition in time-series econometrics and is known as the COMFAC (Common Factor) approach. This method was originally introduced by John Denis Sargan in the 1970s and later popularized by Hendry and Mizon (1978).

consumption, and lagged log consumption,

$$p_{j,\tau} = a_{0,j} + a_{1,j}\tau + b_{j,1}p_{j,\tau-1} + b_{j,2}c_\tau + b_{j,3}c_{\tau-1} + v_{j,\tau}, \tag{17}$$

using all observations up to time t (i.e., all observations $\tau \leq t$). Given the estimated $\hat{b}_{j,1}$ and $\hat{v}_{j,t}$ from this regression, we then update the estimated price-consumption cycle $\widehat{pcc}_{j,t}$,

$$\widehat{\delta}_{j,t} = \hat{b}_{j,1}\widehat{\delta}_{j,t-1} + \hat{v}_{j,t}. \tag{18}$$

We then repeat this process for $t+1$ and subsequent periods to construct our price-consumption cycle predictor $\widehat{\delta}_{j,t+1}$ that is available in real-time. For all assets j, we initialize $\widehat{\delta}_{j,0}$ to 0 for the first period.

3 Empirical Price-Consumption Cycle

3.1 Data

We bring our framework to the data by examining the relationship between consumption and the Fama and French (1993) factors. Specifically, we use quarterly observations from 1947:Q1 to 2024:Q4 for the returns on the market, as well as on the individual legs of the SMB (Small and Big) and HML (Value and Growth) factors available on Ken French's website.[9] We complement these data with real per-capita personal consumption expenditures. The market portfolio is proxied by the value-weighted return of all CRSP firms incorporated in the United States and listed on the NYSE, AMEX, or NASDAQ. Unless otherwise noted, we focus on non-overlapping three-month holding-period returns.

From the research portfolio returns, we construct prices as the cumulative capital gains, and we also consider log consumption c_t. Specifically, we define:

$$c_t = c_{t-1} + \Delta c_t, \tag{19}$$

$$p_{j,t} = p_{j,t-1} + \Delta p_{j,t}, \tag{20}$$

$$j \in \{\text{Mkt}, \text{Small}, \text{Big}, \text{Value}, \text{Growth}\}.$$

We first examine the long-run relationships between prices and consumption, and the

[9] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

related presence of cointegration:

$$p_{j,t} = \alpha_{0,j} + \alpha_{1,j}t + \beta_j c_t + \delta_{j,t}. \tag{21}$$

Having established cointegration, we then investigate predictability within the error-correction representation:

$$\Delta p_{j,t+1} = \alpha_{1,j} + \beta_j \Delta c_{t+1} + \delta_j \delta_{j,t} + v_{j,t+1}. \tag{22}$$

Here, short-run dynamics in returns are captured by consumption growth, while the error-correction term $u_{j,t}$ reflects long-run deviations from the cointegrating equilibrium, providing a direct channel through which misalignments between consumption and factor returns translate into return predictability.

3.2 Long-Run Relationships

In Panel A of Table 1 we implement the cointegration tests proposed by Johansen (1991). The Johansen procedure presumes a p-dimensional vector autoregressive model with k lags, and provides two tests for cointegration. The Trace test evaluates the null hypothesis that there are exactly r cointegrating relationships against the alternative of p such relations. The L-max statistic instead tests the null of r cointegrating relations to the alternative of $r + 1$. In our bivariate system, the null hypothesis corresponds to no cointegration ($r = 0$), while the alternative implies a single cointegrating vector. Both the Trace and L-max statistics indicate strong evidence of cointegration between portfolios prices and consumption.

In Panel B, we estimate an error correction model:

$$\Delta p_{j,t+1} = \alpha_{1,j} + \beta_j \Delta c_{t+1} + (\rho_j - 1)\delta_{j,t} + \ v_{j,t+1}$$

Whereas Panel A focuses on the long-run relationship between prices and consumption levels, the error correction model captures both short-run dynamics and long-run adjustment. Evidence of a long-run relationship requires rejecting not only $\beta_j = 0$ but also $\rho_j - 1 = 0$. To test this joint hypothesis, we use the bounds testing procedure of Pesaran, Shin, and Smith (2001) (PSS). Rejection of the null under the PSS test implies $(\rho_j - 1 < 0)$ and $\beta_j \neq 0$, confirming the existence of a levels relationship. The PSS tests reject the null of no levels relationship for all portfolios.

Table 1: Tests for Long-Run Relationships

This table reports results for the Johansen cointegration test (Panel A) and the Pesaran, Shin, and Smith (2001) (PSS) test (Panel B). We use real per-capita personal consumption expenditures. The number of lags for the Johansen test is chosen according to the Bayesian information criterion, and we include a linear time-trend. In Panel B, we report the F-statistics for the Wald bounds-test for no cointegration in PSS (level of significance are also computed as in PSS); the lag length p and q are set to 1, so to test the specification (16). The sample period is 1947:Q1 to 2024:Q4. Quarterly observations.

Panel A: Johansen Cointegration Test

$$p_{j,t} = \alpha_0 + \alpha_{1,j}t + \beta c_t + \delta_{j,t}$$

	Mkt	Small	Big	Value	Growth
r = 0 (Trace) |	28.47**	29.35**	28.59**	30.32**	28.42**
r = 0 (MaxEigen) |	24.96***	25.66***	24.90***	25.66***	24.72***

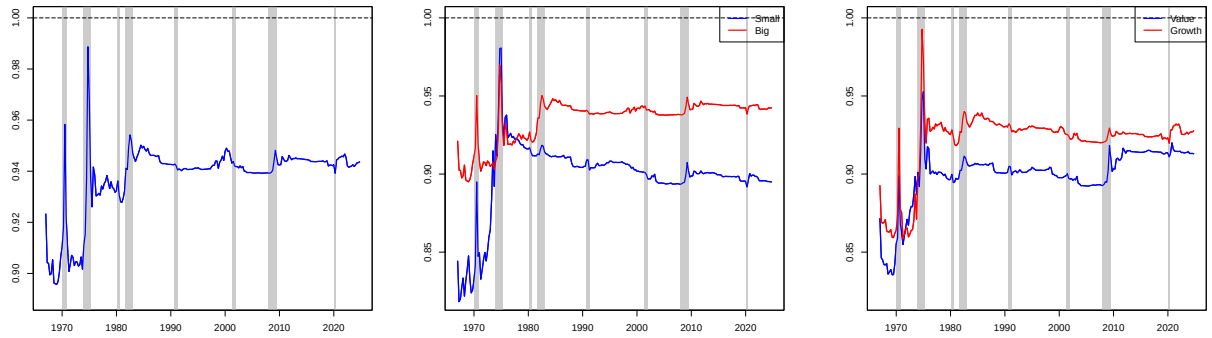
Panel B: PSS Test $\beta_j \cap (\rho_j - 1) \neq 0$

$$r_{j,t+1} = \alpha_{1,j} + \beta_j \Delta c_{t+1} + (\rho_j - 1)\delta_{j,t} + \ v_{j,t+1}$$

	Mkt	Small	Big	Value	Growth
PSS	11.640***	11.380***	10.343***	8.603**	12.342***

Having established the presence of a long-run relationship between asset prices and consumption, we now examine the properties of price deviations from consumption. Figure 1 reports the real-time estimates of the persistence parameter, ρ, computed over an expanding sample. Consistent with the PSS test, all estimates satisfy $\rho < 1$, confirming mean reversion in price deviations. The estimates are relatively stable over time, averaging around 0.9 for small and value portfolios—corresponding to a half-life of $\ln(1/2)/\ln(\rho) = 6.6$ quarters (or about 1.7 years). For the market, big, and growth portfolios, persistence is somewhat higher, with $\rho \approx 0.95$, implying a half-life of roughly 13.5 quarters (about three years). Across all cases, price deviations exhibit clear business-cycle fluctuations but are notably less persistent than those observed for valuation ratios, like dividend–to-price.

Figure 2 presents the real-time estimates of β_j (left panel) and the price deviations from long-run consumption levels, δ_j (right panel), obtained from equation (16) for the market portfolio. The left panel shows that the market's long-run beta with respect to consumption is close to one, exhibiting relatively mild fluctuations over time, with the exception of pronounced and rapid movements during the 2008 financial crisis and the COVID-19 recession. The right panel illustrates that price deviations display a clear business-cycle pattern, with prices falling below their long-run level during recessions.



(a) Market **(b)** Size **(c)** Value

Figure 1: Real-Time ρ Estimates

This figure shows the real-time ρ estimates for different portfolios (see Section 2.3 for details on the estimation procedure). The sample period starts in 1947:Q1 and we use the first 20 years as a burn-in sample. Estimates span the sample 1967:Q1 to 2024:Q4. Quarterly observations.

These negative deviations forecast higher expected returns (since $\rho < 1$), implying that discount rates are countercyclical.

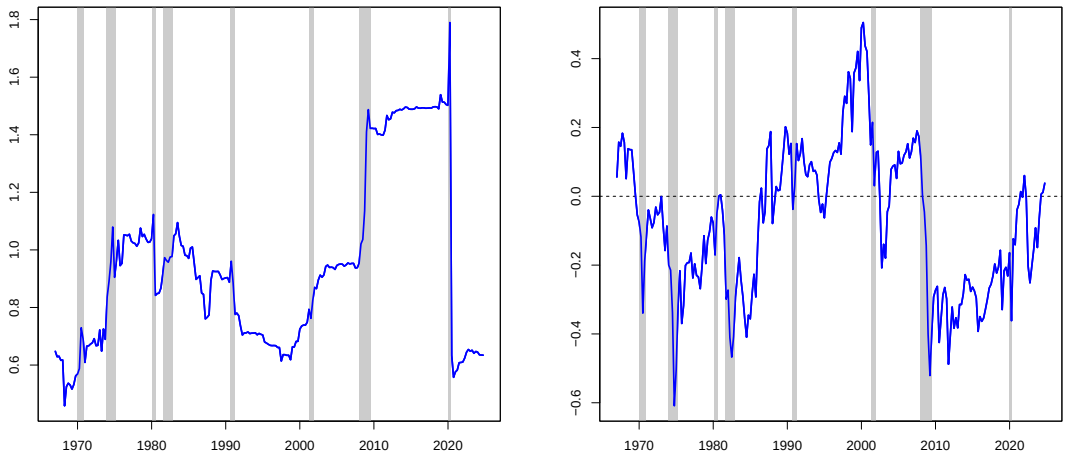
Figure 3 reports the corresponding betas for the anomaly portfolios. Portfolios associated with higher expected returns—such as value and small-cap portfolios—exhibit higher long-run betas with consumption compared to growth and large-cap portfolios.

Figure 4 plots the price deviations for the individual legs of the value and size portfolios. As with the market, price deviations exhibit clear procyclical behavior. We also observe a strong common component across portfolios. This pattern is corroborated in Table 2, Panel B, which reports high correlations in price deviations across portfolios, and in Panel C, which shows that the first principal component explains approximately 80% of the total variance.

4 Return Predictability

4.1 In-Sample

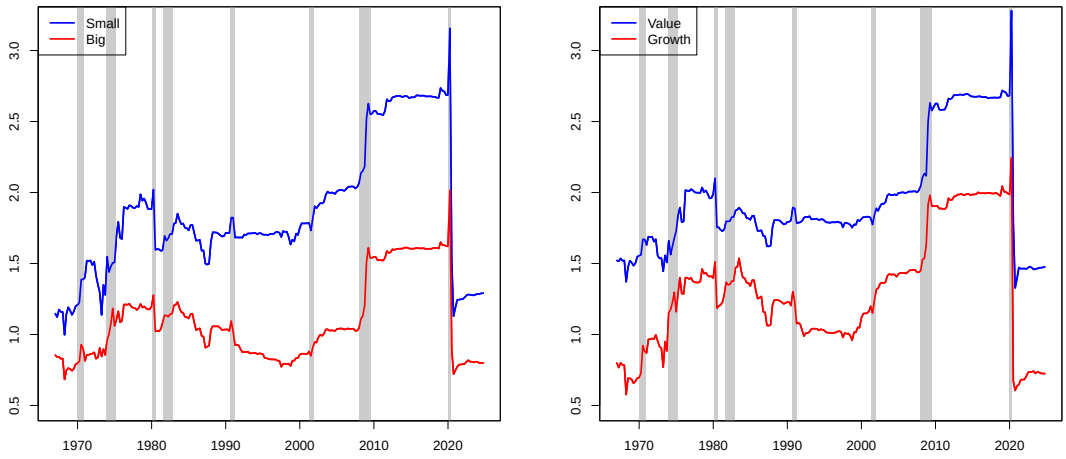
Given the evidence of cointegration established in the previous section, we next evaluate the predictive ability of price deviations for the market and for characteristics-based factor returns. Table 3 reports the results of the in-sample one-quarter-ahead predictive regressions. Although the regression coefficients are estimated using in-sample data, it is important to



(a) Time-Varying Beta **(b)** Price Deviations

Figure 2: Market Beta and Price Deviations

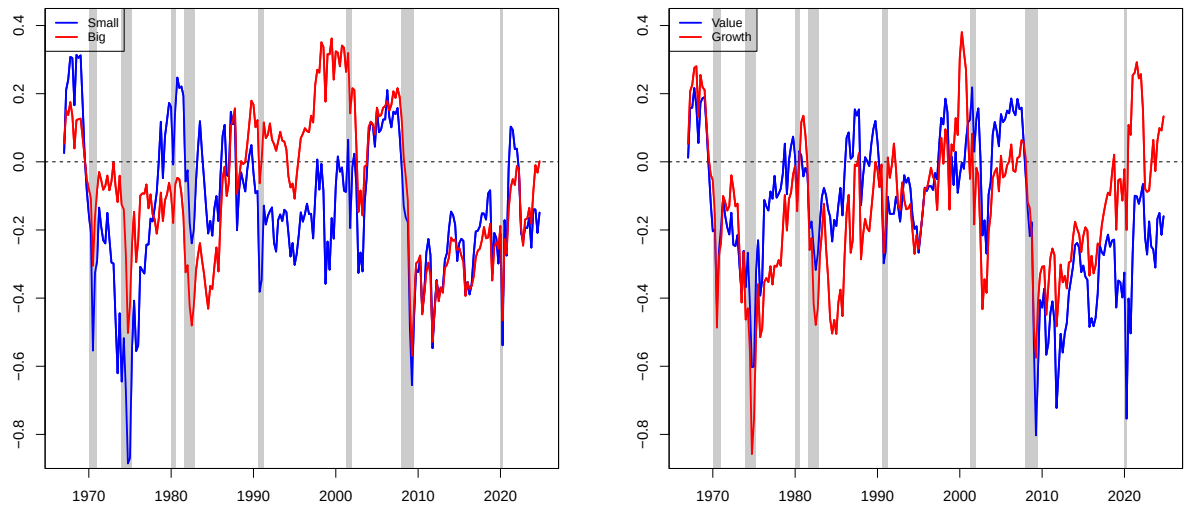
This figure shows the real-time factor loadings of market price on consumption and its price deviations (see Section 2.3 for details on the estimation procedure). The sample period starts in 1947:Q1 and we use the first 20 years as a burn-in sample. Estimates span the sample 1967:Q1 to 2024:Q4. Quarterly observations.



(a) Size **(b)** Value

Figure 3: Time-Varying Betas: Size and Value

This figure shows the real-time factor loadings of prices on consumption for different portfolios. The sample period starts in 1947:Q1 and we use the first 20 years as a burn-in sample. Estimates span the sample 1967:Q1 to 2024:Q4. Quarterly observations.



(a) Size **(b)** Value

Figure 4: Price Deviations: Size and Value

This figure shows the real-time portfolio deviations for different portfolios. The sample period starts in 1947:Q1 and we use the first 20 years as a burn-in sample. Estimates span the sample 1967:Q1 to 2024:Q4. Quarterly observations.

note that the price deviations themselves are computed fully out of sample. In other words, we do not use a fixed cointegrating vector based on full-sample estimates; instead, we re-estimate the cointegrating parameters recursively at each point in time (see Figure 2 for the market and Figure 3 for the characteristic-sorted portfolios).

The loading coefficient on the price deviation is economically and statistically significant, with a negative sign: a positive deviation of (log) portfolio prices from their long-run relationship with consumption predicts lower expected returns in the subsequent period. For the market, as well as for the individual legs of the value and size portfolios, the magnitude of this effect is approximately –0.1 per unit of deviation. Further supporting evidence on the role of price deviations in explaining the time-series dynamics of factor returns is provided by the R^2 values reported in the last row of Table 3, which show that these deviations account for about 4% of the total variance in quarterly factor returns.

Figure 5 illustrates how return predictability varies across horizons. Panel (a) reports the coefficients on price deviations from the one-quarter-ahead to the five-year-ahead predictive regressions, while Panel (b) displays the corresponding R^2 values. The

Table 2: Portfolio Deviations: Summary Statistics

This table reports summary statistics for the real-time portfolio deviations estimated as in Eq. (17)–(18). The sample period starts in 1947:Q1 and we use the first 20 years as a burn-in sample. Estimates span the sample 1967:Q1 to 2024:Q4. Quarterly observations.

Panel A: Descriptive Statistics

	Mkt	Small	Big	Value	Growth
Mean	-0.07	-0.14	-0.08	-0.14	-0.14
Std	0.21	0.21	0.21	0.21	0.21
AR1	0.92	0.85	0.92	0.88	0.87

Panel B: Correlation Matrix

	Mkt	Small	Big	Value	Growth
Mkt	1.00				
Small	0.58	1.00			
Big	0.98	0.56	1.00		
Value	0.78	0.84	0.83	1.00	
Growth	0.79	0.74	0.72	0.63	1.00

Panel C: PCA Loadings and Variance Explained

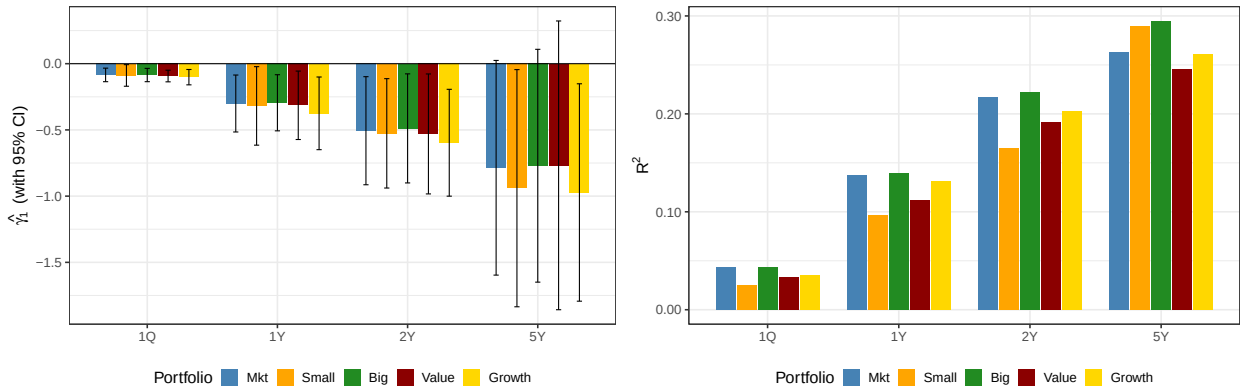
	PC1	PC2	PC3	PC4	PC5	VarianceExplained	CumulativeVariance
P1	-0.47	-0.45	0.04	0.60	-0.47	0.80	0.80
P2	-0.41	0.73	-0.02	0.46	0.30	0.12	0.91
P3	-0.46	-0.47	-0.18	-0.14	0.72	0.08	0.99
P4	-0.46	0.20	-0.59	-0.49	-0.41	0.01	1.00
P5	-0.43	0.08	0.79	-0.42	-0.11	0.00	1.00

Table 3: Predictive Regressions

This table reports results for the predictive regression $r_{j,t+1} = \gamma_{0,j} + \gamma_{1,j}\widehat{\delta}_{j,t} + \varepsilon_{j,t+1}$ for different portfolios, where $\widehat{\delta}_{j,t}$ are the real-time price cycle for each portfolio j computed as in Eq. (17)–(18). Values in parentheses are Newey-West standard errors. The sample period is 1967:Q1 to 2024:Q4. Quarterly observations.

	Mkt	Small	Big	Value	Growth
$\widehat{\delta}_{j,t}$	-0.085^{***}	-0.089^{**}	-0.086^{***}	-0.093^{***}	-0.102^{***}
	(0.026)	(0.042)	(0.026)	(0.022)	(0.030)
Observations	231	231	231	231	231
R^2	0.043	0.025	0.044	0.033	0.036

coefficients are negative across all horizons but gradually lose statistical significance beyond approximately five years, consistent with the idea that price deviations are fully reabsorbed over longer horizons. The R^2 values in Panel (b) increase with the horizon, and indicate



(a) Predictive Coefficients **(b)** R^2

Figure 5: Predictive Regressions Across Portfolios and Different Time Horizons

This figure shows predictive coefficients and R^2 for the predictive regression $r_{j,t+i} = \gamma_{0,j} + \gamma_{1,j}\widehat{\delta}_{j,t} + \varepsilon_{j,t+i}$ for different portfolios and time horizons, where $\widehat{\delta}_{j,t}$ are the real-time price deviations for each portfolio j computed as in Eq. (17)–(18). We report 95% confidence intervals computed using Newey-West standard errors for $i = 1$ or Hodrick standard errors for $i > 1$. The sample period is 1967:Q1 to 2024:Q4. Quarterly observations.

that most of the predictability is concentrated in the market, as well as in the big and growth legs, with somewhat weaker effects for the small and value portfolios.

Appendix Figures A.2 and A.3 replicate Figure 5, replacing cumulative excess returns with cumulative real returns and cumulative capital gains, respectively. Appendix Figures A.4 and A.5 replicate the analysis in Figure 5 under two alternative specifications. In the first, we replace real per-capita PCE with real aggregate PCE; in the second, we substitute real per-capita PCE with real per-capita GDP, following Rangvid (2006). Across all cases, the results confirm that the patterns are robust to alternative measures of returns and to different proxies for real economic activity.

In summary, our results indicate that the prices of the aggregate market and prominent cross-sectional factors are anchored to macroeconomic fundamentals, thereby establishing a novel and robust link between financial markets and the real economy.

4.2 Out-of-Sample

Figure 6 provides out-of-sample evidence supporting a link between financial markets and the real economy. The figure plots the in-sample R^2 on the x-axis and the out-of-sample R^2 statistic proposed by Campbell and Thompson (2008) on the y-axis. The latter is obtained

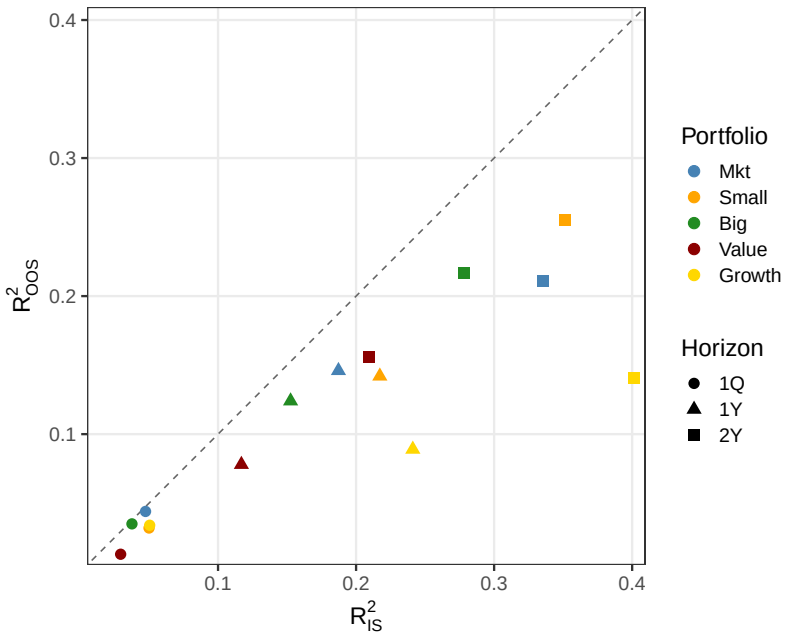


Figure 6: IS vs OOS R^2

This figure shows the IS vs OOS R^2 from the predictive regressions $r_{j,t+i} = \gamma_{0,j} + \gamma_{1,j}\widehat{\delta}_{j,t} + \varepsilon_{j,t+i}$ for different portfolios and time horizons, where $\widehat{\delta}_{j,t}$ are the real-time price deviations for each portfolio j computed as in Eq. (17)–(18). We use the first half of the sample as training sample (1967:Q1–1995:Q4) and an expanding window until the end of the sample (2024:Q4). The IS R^2 is computed over the sample 1996:Q1–2024:Q4

by re-estimating the predictive regression using only information available at time t to forecast returns at $t+1$, benchmarking these forecasts against the historical average based on an expanding window. As expected, all data points—corresponding to different portfolios and forecast horizons—lie below the 45-degree line, reflecting the typical deterioration in predictive performance when moving from in-sample to out-of-sample evaluation. What is noteworthy, however, is that the out-of-sample R^2 values remain relatively large and only modestly lower for the one-quarter and one-year horizons, with the only exceptions being the value portfolio at the one-quarter horizon and the growth portfolio at the one-year horizon.

Figure 7 focuses on the one-quarter-ahead predictive relationship by showing the in-sample and out-of-sample performance of the corresponding predictive regressions over time. Specifically, the figure plots the cumulative squared prediction errors of the null model—based on historical average returns—minus those of our model, which relies on the long-run relationship between prices and consumption. An increase in a line indicates

periods when our model outperforms the historical mean benchmark, while a decrease in a line indicates better performance of the prevailing mean return.

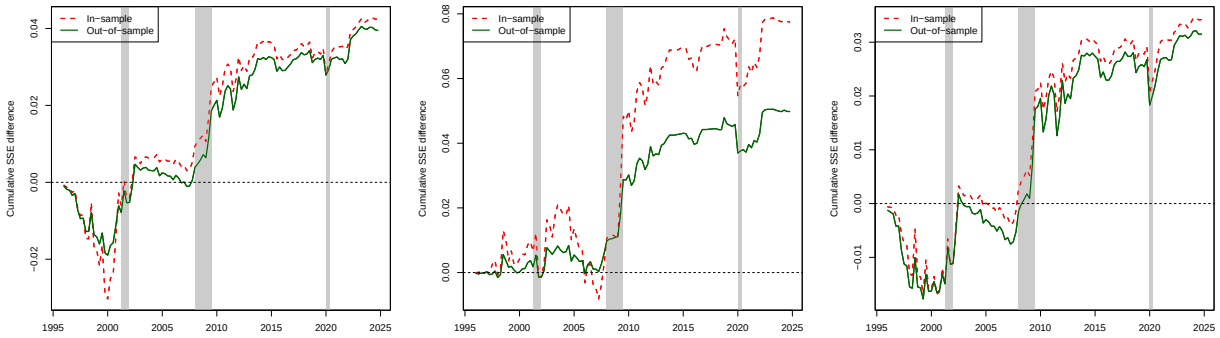
Panels (a) and (c) indicate that our model generally outperforms the historical mean for the market and big portfolios, with the sole exception of the period around 2000. Moreover, there is little difference between the in-sample and out-of-sample predictive regressions for these portfolios. Panel (d) shows that the value portfolio is the most difficult to predict. The lines decline in several periods, including non-recessionary episodes such as the post-2000 years and the 2010–2020 decade, particularly when the model is evaluated out of sample. Panels (b) and (e) show that the weaker out-of-sample performance for the small and growth portfolios is largely driven by two recession periods—the 2007–08 financial crisis for small portfolios and the 2001 recession for growth portfolios. Outside these episodes, the out-of-sample performance of the model closely tracks its in-sample counterpart.

4.3 Comparison to Other Return Predictors

How does the predictive performance of our price-consumption cycle compare to more traditional price ratios or predictors based on other cointegrating relationships? Tables 4-5 provide a direct side-by-side comparison. Given that these previous predictors were designed to predict returns rather than price growth, we focus on return predictability for all of the tests.

Starting with the results for price ratios, the first column of Table 4 shows the results of a univariate regression predicting the next quarter market excess return using our price-consumption cycle (PCC). For comparison, column (2) shows the results using the price-to-consumption ratio, which gives a marginally significant, smaller magnitude coefficient and an adjusted R^2 less than 1%. Importantly, column (7) shows that in a joint regression, the predictability of returns comes entirely from PCC, with the adjusted R^2 of column (7) actually being lower than the adjusted R^2 of column (1). For robustness, we also consider a range of additional price-to-fundamental ratios using output, dividends, M&A-adjusted dividends, and ten-year cyclically adjusted earnings as measures of fundamentals. As shown in columns (3)-(6) and (8)-(11), the results are similar to the price-to-consumption ratio, with little to no significant return predictability in univariate regressions, and no significant coefficients on the price ratios in joint regressions with PCC.

Why do we see such a large difference in performance between PCC and traditional price ratios? The key distinction is the persistence of the predictors. As discussed in Section



(a) Market **(b)** Small **(c)** Big



(d) Value **(e)** Growth

Figure 7: IS and OOS R^2 over Time

This figure shows the performance of IS and OOS predictive regressions. For the IS regressions, the graphs show the cumulative squared returns minus the cumulative squared residuals from the predictive regression using the real-time price-consumption cycles. For the OOS regressions, the graphs show the cumulative squared prediction errors using the historical mean minus the cumulative squared prediction error of the predictive regression using the price cycle. We use the first half of the sample as training sample (1967:Q1–1995:Q4) and an expanding window until the end of the sample. Quarterly observations.

2.1, predictability of price growth or returns using traditional price ratios relies on the idea that these price ratios will revert back to their unconditional means. However, if the mean changes over time and/or β differs from 1, then these price ratios can stay above or below their unconditional means for many years. Columns (1)-(6) show the persistence of each predictor minus 1 neae the bottom of the table. For PCC, this value is -0.08 and highly significant. For the traditional price ratios, this value is closer to 0 and less significant. In fact, for many price ratios, we cannot reject that these ratios have a unit root. Thus, we find faster and more consistent mean reversion for PCC, which unsurprisingly leads to larger magnitude and more significant return predictability at short horizons.

Table 4: Multivariate Predictive Regressions: PCC and Fixed Long-Run Relationships

This table presents results on aggregate market return predictability. We regress the one-quarter-ahead log market excess return on several aggregate predictors. The variables P/C, P/Y, and P/D denote the (log) price-to-consumption, price-to-output and price-to-dividends ratios, respectively. P/D M&A-adj is P/D with dividends adjusted for M&A activity, and CAPE is the cyclically adjusted price-to-earnings ratio (divided by 10). Values in parentheses are Newey–West standard errors. The sample period is 1968:Q1–2024:Q4. Quarterly observations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
PCC	−0.090***						−0.096***	−0.096***	−0.102***	−0.103***	−0.111***
	(0.027)						(0.029)	(0.029)	(0.026)	(0.027)	(0.029)
P/C		−0.029*					0.006				
		(0.016)					(0.016)				
P/Y			−0.025*					0.007			
			(0.013)					(0.012)			
P/D				−0.016					0.013		
				(0.016)					(0.015)		
P/D M&A-adj					−0.012					0.014	
					(0.016)					(0.014)	
CAPE						−0.004					0.010
						(0.007)					(0.007)
Persistence - 1	-0.08***	-0.03**	-0.03**	-0.02	-0.03	-0.01					
Observations	231	231	231	231	231	231	231	231	231	231	231
Adjusted R^2	0.048	0.008	0.006	0.001	−0.001	−0.003	0.044	0.045	0.047	0.048	0.052

Table 5 compares our PCC to predictors based on other cointegrating relationships. As shown in equation (6), one can estimate a cointegrating relationship between consumption and another variable x_t and then use the deviations from this long-term relationship to predict future returns. Menzly et al. (2004) use the cointegrating residuals from the relationship between dividends and consumption assuming $\alpha_1 = 0$ and $\beta = 1$, Hansen et al. (2008) use the residuals assuming $\alpha_1 = 0$, and Bansal et al. (2001) use the residuals without restricting α_1 or β; and Bansal et al. (2009) incorporate a VAR structure to measure the return predictor implied by the dividend-consumption relationship. Additionally, Lettau and Ludvigson (2001) use a combination of human capital and asset holdings as their measure of x_t and utilize the deviation *cay* to predict returns.

As in Table 4, the first column of Table 5 shows the results of a univariate regression of next quarter excess returns on the current price-consumption cycle. Columns (2)-(6) show the results using the alternative predictors based on other cointegrating relationships. We find that several of these approaches yield meaningful predictability, particularly those in Bansal et al. (2009). Columns (7)–(11) show that our PCC remains a significant predictor

Table 5: Multivariate Predictive Regressions: PCC and Other Cointegration-Based Variables

This table presents results on aggregate market return predictability. We regress the one-quarter-ahead log market excess return on several aggregate predictors. dc is the difference between (log) aggregate dividends and consumption; the variables dc no trend and dc trend are computed from cointergating regressions without and with a time trend as in Hansen et al. (2008) and Bansal et al. (2001), respectively; EC-VAR is the return predictor implied by the EC-VAR in Bansal et al. (2009). Values in parentheses are Newey–West standard errors. The sample period is 1968:Q1–2024:Q4. Quarterly observations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
PCC	-0.090^{***}						-0.090^{***}	-0.089^{***}	-0.083^{***}	-0.084^{***}	-0.092^{***}
	(0.027)						(0.027)	(0.027)	(0.029)	(0.031)	(0.027)
dc		0.010					-0.002				
		(0.027)					(0.024)				
dc no trend			0.019					0.006			
			(0.024)					(0.021)			
dc trend				-0.078^{**}					-0.038		
				(0.038)					(0.039)		
EC-VAR					0.883^{**}					0.209	
					(0.421)					(0.405)	
cay						-0.010					0.066
						(0.205)					(0.213)
Observations	231	231	231	231	231	231	231	231	231	231	231
Adjusted R^2	0.048	-0.004	-0.002	0.012	0.015	-0.004	0.044	0.044	0.048	0.045	0.045

even after controlling for these alternative measures, suggesting that it captures a distinct dimension of return predictability. Rephrased, by using direct data on prices to train the parameters of our cointegrating relationship, our PCC incorporates information that is not embedded in these previous relationships.

5 The Role of Consumption

The evidence in Table 6 shows that our price cycle does not predict consumption growth, either over the next quarter (Panel A) or at longer horizons (Panel B). The results are robust to the choice of factor leg used to extract the price cycle—whether from the market or any of the SMB and HML components. Furthermore, there is no evidence of predictability either in-sample or out-of-sample, as confirmed by the very small R^2 and OOS R^2.

Table 7 shows that when we estimate equation (16) including only the time trend—while excluding consumption—the resulting price cycle fails to predict returns.

Table 6: Predicting Consumption

This table reports results for the predictive regressions of real PCE log growth rate on price cycles for different horizons. The sample period starts in 1947:Q1 and we use the first 20 years as a burn-in sample. The real-time price cycles span the sample 1967:Q1 to 2024:Q4. The OOS R^2 are computed using the first half of the sample as training sample (1967:Q1–1995:Q4) and an expanding window until the end of the sample. Values in parenthesis are Newey-West standard errors in Panel A and Hodrick standard errors in Panels B. Quarterly observations.

Panel A: 1-quarter-ahead

	Mkt	Small	Big	Value	Growth
$\widehat{\delta}_{j,t}$	0.0004	−0.001	−0.001	−0.001	0.001
	(0.003)	(0.004)	(0.003)	(0.004)	(0.003)
Observations	231	231	231	231	231
R^2	0.0001	0.0001	0.0003	0.0001	0.0002
OOS R^2	−0.010	−0.007	−0.013	−0.019	−0.005

Panel B: 5-years-ahead

	Mkt	Small	Big	Value	Growth
$\widehat{\delta}_{j,t}$	−0.009	−0.016	−0.024	−0.020	0.002
	(0.054)	(0.037)	(0.059)	(0.059)	(0.023)
Observations	212	212	212	212	212
R^2	0.002	0.006	0.013	0.008	0.000
OOS R^2	−0.083	−0.037	−0.153	−0.129	−0.008

Panel (a) of Figure 8 shows that our price cycle differs from the consumption–wealth ratio proxy—*cay*—proposed by Lettau and Ludvigson (2005). What drives this difference? The consumption series we use is identical to that in $cay_t = c_t - (\omega a_t + (1-\omega)\, h_t)$. Hence, the discrepancy must arise from the cumulative return series r_t^{cum} differing from the aggregate wealth component, $(\omega a_t + (1-\omega)h_t)$, estimated by Lettau and Ludvigson (2005).

Indeed, Panel (b) shows that the aggregate wealth component in Lettau and Ludvigson (2005) is much smoother: while it captures the overall trend in cumulative ex-dividend prices, it fails to track several sizable medium-term fluctuations. Furthermore, the wealth trend fit is generally good until the 1980s but deteriorates thereafter, as the stock market growth accelerates.

Table 7: Removing Consumption

This table reports results for the predictive regression $r_{j,t+1} = \gamma_{0,j} + \gamma_{1,j}\hat{\delta}_{j,t} + \varepsilon_{j,t+1}$ for different portfolios, where $\hat{\delta}_{j,t}$ denotes the real-time deviation of each portfolio's cumulative return from its time trend (i.e., we use Eqs. (17)–(18), but include only the time trend component, excluding consumption levels). The sample period starts in 1947:Q1 and we use the first 20 years as a burn-in sample. The real-time cycle spans the sample 1967:Q1 to 2024:Q4. The OOS R^2 are computed using the first half of the sample as training sample (1967:Q1–1995:Q4) and an expanding window until the end of the sample. Values in parenthesis are Newey-West standard errors in Panel A and Hodrick standard errors in Panels B. Quarterly observations.

Panel A: 1-quarter-ahead

	Mkt	Small	Big	Value	Growth
$\hat{\delta}_{j,t}$	−0.011	−0.063	−0.022	−0.073***	−0.015
	(0.024)	(0.041)	(0.022)	(0.024)	(0.025)
Observations	231	231	231	231	231
R^2	0.001	0.018	0.005	0.026	0.002
OOS R^2	-0.019	0.020	-0.013	0.023	-0.007

Panel B: 5-years-ahead

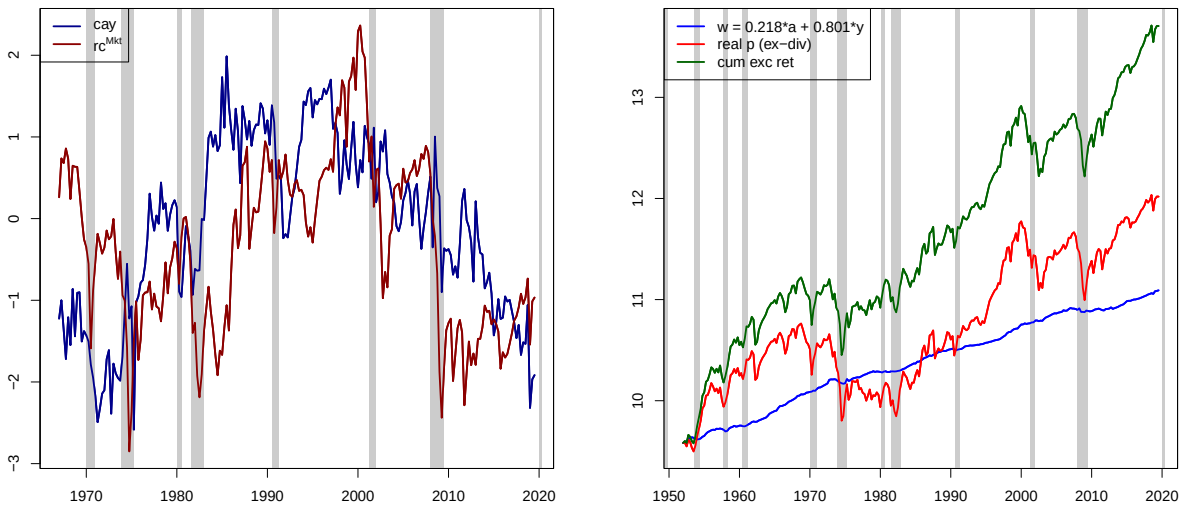
	Mkt	Small	Big	Value	Growth
$\hat{\delta}_{j,t}$	−0.055	−0.783*	−0.156	−0.664	−0.079
	(0.358)	(0.446)	(0.382)	(0.453)	(0.338)
Observations	212	212	212	212	212
R^2	0.002	0.277	0.018	0.235	0.004
OOS R^2	−0.496	0.190	−0.325	0.310	−0.311

6 Economic Value

In our final tests, we assess the economic value of market and factor timing to an investor.

We begin by comparing the Sharpe ratio of a simple buy-and-hold strategy with that of a timing strategy applied to each portfolio. For the timing strategy, we consider both in-sample predictive regressions and out-of-sample predictive relations. In all cases, the cointegrating relation is estimated in real-time, ensuring that the predictive signal is available to our investor. We find that timing generally enhances the Sharpe ratio relative to buy-and-hold, increasing it by about 45% for the market (from 0.42 to 0.61) and up to 93% for the growth portfolio (from 0.33 to 0.64). The gains are more modest for the small and big portfolios, at approximately 36% and 20%, respectively.

Next, we consider a mean–variance investor who, at the end of quarter t, solves the



(a) **(b)**

Figure 8: Market price cycle vs Consumption-Wealth Ratio

Panel (a) shows the consumption-wealth ratio measure, *cay*, proposed by Lettau and Ludvigson (2005) along with the market's price cycle. Panel (b) shows the wealth trend in Lettau and Ludvigson (2005) (the blue line) together with cumulative excess returns (the dark green line) and real prices (the red line) for the market portfolio.

following optimization problem:

$$\arg\max_{w_{t+1|t}} E_t\left[R_{p,t+1}\right] - 0.5\gamma \mathrm{Var}\left(R_{p,t+1}\right)$$

where $R_{p,t+1} = w_{t+1|t} f_{t+1}$, $w_{t+1|t}$ denotes the portfolio weights on the factors for period $t+1$ given the return forecasts at time t, and γ represents the coefficient of relative risk aversion. The average utility realized by the investor is given by

$$\overline{U}_j = \overline{R}_{p,t+1} - 0.5\gamma\widehat{\mathrm{Var}}\left(R_{p,t+1}\right) \text{ , for } j = 0, 1,$$

where $j = 0$ or 1 indicates whether the investor relies on the prevailing mean return or on the competing forecast of f_{t+1} implied by our long-run model linking financial markets and economic activity. The average utility gain (or increase in certainty-equivalent return) from using our predictive model instead of the historical mean benchmark is given by $\Delta = \overline{U}_1 - \overline{U}_0$. When annualized (by multiplying by four), this value can be interpreted as the portfolio

Table 8: Economic Value of Predictive Regressions

This table reports the Sharpe ratio of a buy-and-hold strategy (first row), the Sharpe ratio of a timing strategy based on full-sample regressions of future returns (of the asset in the column) on the price cycle (second row), and the Sharpe ratio of a timing strategy based on out-of-sample regressions of future returns (of the asset in the column) on the price cycle (third row). The fourth row reports the certainty equivalent difference between the buy-and-hold and the out-of-sample timing strategies. For the certainty equivalent calculation, we assume a coefficient of risk aversion $\gamma = 3$. Quarterly observations. The sample period is 1967:Q1–2024:Q4. For the out-of-sample, we use the first half of the sample as training sample (1967:Q1–1995:Q4) and an expanding window until the end of the sample.

	Mkt	Small	Big	Value	Growth
B&H SR	0.42	0.33	0.43	0.35	0.33
IS SR	0.62	0.51	0.54	0.42	0.63
OOS SR	0.61	0.45	0.52	0.39	0.64
ΔU	3.54	2.31	2.24	-0.15	2.75

management fee—expressed as a percentage of wealth—that an investor would be willing to pay for access to the information contained in our forecasts. The last row of Table 8 reports these utility gains for a mean–variance investor exploiting the information conveyed by price deviations from consumption.

The forecasts based on deviations of factor prices from their long-run relation with consumption generate substantial utility gains. The annualized gains reach 2.3% for the timing of small and big portfolios, 2.8% for growth, and 3.5% for the market—figures that are economically significant.

7 Conclusion

This paper revisits the relationship between asset prices and macroeconomic fundamentals by shifting attention from consumption growth to consumption levels. While an influential literature emphasizes the difficulty of linking one-period returns to one-period consumption growth, our analysis shows that studying levels reveals an economically meaningful connection. We document that cumulative stock returns and aggregate consumption levels share a common stochastic trend, implying a stable long-run relationship between financial markets and the real economy.

Building on this insight, we characterize temporary deviations of prices from their consumption-implied long-run values—the price cycle—and show that these deviations are

mean-reverting. Their dynamics provide a decomposition of returns into permanent and temporary components: the permanent component comoves with consumption, while the temporary component reflects transitory price fluctuations. This structure offers a simple organizing principle for understanding asset-price variation. When prices move away from their consumption-implied benchmark, the resulting deviations carry information about future returns as markets gradually realign with fundamentals.

Our empirical analysis demonstrates that the price cycle forecasts returns across portfolios and horizons, remains robust in real-time, and generates economically meaningful utility gains for a mean–variance investor. The predictive content does not reflect expectations of future consumption growth and disappears when consumption is omitted from the long-run relation, indicating that the information originates specifically from the consumption-based component of prices. These findings complement and extend existing work that seeks to identify macroeconomic origins of return predictability and the sources of time-varying expected returns.

References

Bandi, Federico M, Shomesh E Chaudhuri, Andrew W Lo, and Andrea Tamoni (2021) "Spectral factor models," *Journal of Financial Economics*, 142 (1), 214–238.

Bandi, Federico M. and Andrea Tamoni (2023) "Business-cycle consumption risk and asset prices," *Journal of Econometrics*, 237 (2), None.

Bansal, Ravi, Robert F Dittmar, and Christian T Lundblad (2001) "Consumption, Dividends, and the Cross-Section of Equity Returns," *Working Paper, Duke University*.

Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad (2005) "Consumption, Dividends, and the Cross Section of Equity Returns," *Journal of Finance*, 60 (4), 1639–1672.

Bansal, Ravi, Robert Dittmar, and Dana Kiku (2009) "Cointegration and consumption risks in asset returns," *Review of Financial Studies*, 22 (3), 1343–1375.

Bansal, Ravi and Dana Kiku (2011) "Cointegration and long-run asset allocation," *Journal of Business & Economic Statistics*, 29 (1), 161–173.

Barroso, Pedro, Martijn Boons, and Paul Karehnke (2021) "Time-varying state variable risk premia in the ICAPM," *Journal of Financial Economics*, 139 (2), 428–451.

Boons, Martijn, Fernando Duarte, Frans De Roon, and Marta Szymanowska (2020) "Time-varying inflation risk and stock returns," *Journal of Financial Economics*, 136 (2), 444–470.

Campbell, John Y (2017) *Financial decisions and markets: a course in asset pricing*: Princeton University Press.

Campbell, John Y and John H Cochrane (1999) "By force of habit: A consumption-based explanation of aggregate stock market behavior," *Journal of Political Economy*, 107 (2), 205–251.

Campbell, John Y and Robert J Shiller (1988) "The dividend-price ratio and expectations of future dividends and discount factors," *The Review of Financial Studies*, 1 (3), 195–228.

Campbell, John Y and Samuel B Thompson (2008) "Predicting excess stock returns out of sample: Can anything beat the historical average?" *The Review of Financial Studies*, 21 (4), 1509–1531.

Chen, Nai-Fu, Richard Roll, and Stephen A Ross (1986) "Economic forces and the stock market," *Journal of Business*, 383–403.

Cochrane, John H (2005) *Asset pricing: Revised edition*: Princeton university press.

Cochrane, John H and Lars Peter Hansen (1992) "Asset pricing explorations for macroeconomics," *NBER Macroeconomics Annual*, 7, 115–165.

Daniel, Kent and David Marshall (1997) "Equity-premium and risk-free-rate puzzles at long horizons," *Macroeconomic Dynamics*, 1 (2), 452–484.

Engle, Robert F and Clive WJ Granger (1987) "Co-integration and error correction: representation, estimation, and testing," *Econometrica*, 251–276.

Fama, Eugene F and Kenneth R French (1988a) "Dividend yields and expected stock returns," *Journal of financial economics*, 22 (1), 3–25.

——— (1988b) "Permanent and temporary components of stock prices," *Journal of Political Economy*, 96 (2), 246–273.

——— (1993) "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*.

——— (2015) "A five-factor asset pricing model," *Journal of Financial Economics*, 116 (1), 1–22.

Ferson, Wayne E and Campbell R Harvey (1991) "The variation of economic risk premiums," *Journal of political economy*, 99 (2), 385–415.

Goyal, Amit and Ivo Welch (2007) "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction," *The Review of Financial Studies*, 21 (4), 1455–1508.

Goyal, Amit, Ivo Welch, and Athanasse Zafirov (2024) "A comprehensive 2022 look at the empirical performance of equity premium prediction," *The Review of Financial Studies*, 37 (11), 3490–3557.

Haddad, Valentin, Serhiy Kozak, and Shrihari Santosh (2020) "Factor Timing," *The Review of Financial Studies*, – (7), –2649.

Hansen, Lars Peter, John C Heaton, and Nan Li (2008) "Consumption strikes back? Measuring long-run risk," *Journal of Political Economy*, 116 (2), 260–302.

Hendry, David F (1995) *Dynamic econometrics*: Oxford university press.

Hendry, David F and Grayham E Mizon (1978) "Serial correlation as a convenient simplification, not a nuisance: A comment on a study of the demand for money by the Bank of England," *The Economic Journal*, 88 (351), 549–563.

Hou, Kewei, Chen Xue, and Lu Zhang (2015) "Digesting anomalies: An investment approach," *The Review of Financial Studies*, 28 (3), 650–705.

Johansen, Søren (1991) "Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models," *Econometrica*, 1551–1580.

——— (1995) *Likelihood-based inference in cointegrated vector autoregressive models*: Oxford University Press on Demand.

Kamara, Avraham, Robert A Korajczyk, Xiaoxia Lou, and Ronnie Sadka (2016) "Horizon pricing," *Journal of Financial and Quantitative Analysis*, 51 (6), 1769–1793.

LeRoy, Stephen F and Richard D Porter (1981) "The present-value relation: Tests based on implied variance bounds," *Econometrica*, 555–574.

Lettau, Martin and Sydney Ludvigson (2001) "Consumption, aggregate wealth, and expected stock returns," *Journal of Finance*, 56 (3), 815–849.

Lettau, Martin and Sydney C Ludvigson (2005) "Expected returns and expected dividend growth," *Journal of Financial Economics*, 76 (3), 583–626.

Ludvigson, Sydney C (2013) "Advances in consumption-based asset pricing: Empirical tests," *Handbook of the Economics of Finance*, 2, 799–906.

Menzly, Lior, Tano Santos, and Pietro Veronesi (2004) "Understanding predictability," *Journal of Political Economy*, 112 (1), 1–47.

Merton, Robert C (1973) "An intertemporal capital asset pricing model," *Econometrica*, 867–887.

Nagel, Stefan (2024) "Stock Market Valuation: Explanations, Non-Explanations, and Some Open Questions," in *Comments given at NBER SI Asset Pricing Meeting*.

Parker, Jonathan A and Christian Julliard (2005) "Consumption risk and the cross section of expected returns," *Journal of Political Economy*, 113 (1), 185–222.

Rangvid, Jesper (2006) "Output and expected returns," *Journal of Financial Economics*, 81 (3), 595–624.

Rapach, David and Guofu Zhou (2021) "Sparse Macro Factors," Working Paper.

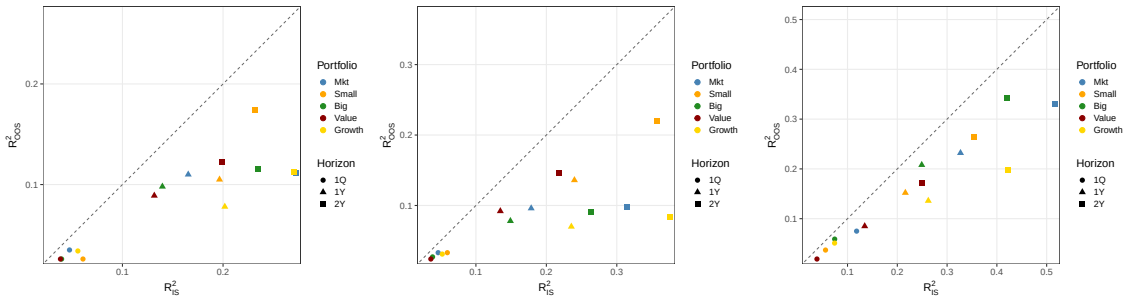
Shiller, Robert J (1981) "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *The American Economic Review*, 71 (3), 421–436.

Vassalou, Maria (2003) "News related to future GDP growth as a risk factor in equity returns," *Journal of Financial Economics*, 68 (1), 47–73.

Appendix

A Additional Results

A.1 Sub-Samples

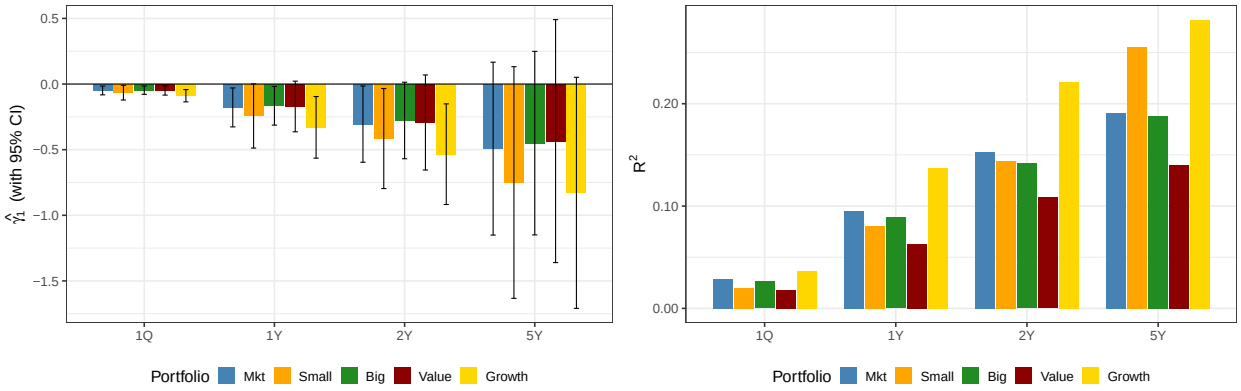


(a) From 1980 **(b)** From 1990 **(c)** From 2000

Figure A.1: Predictive Regressions Across Portfolios and Different Time Horizons

This figure replicates

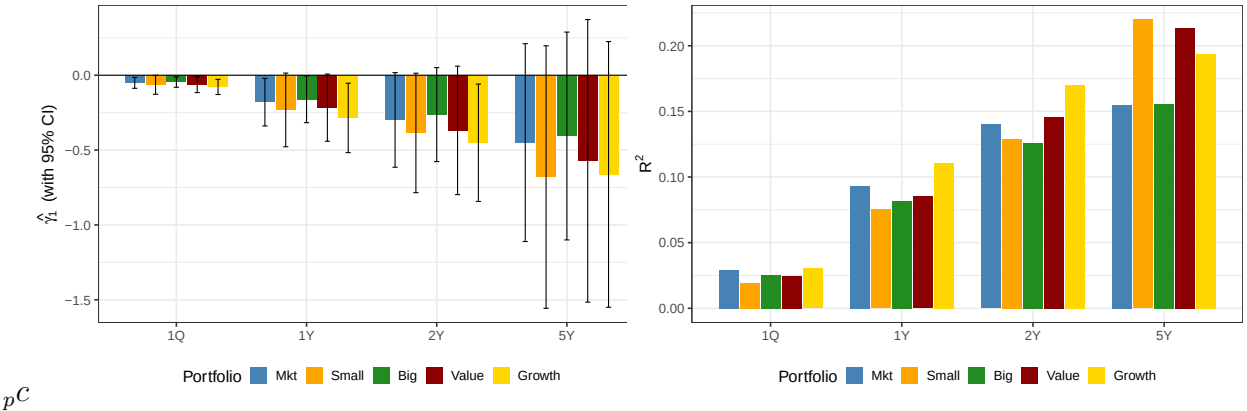
A.2 Alternative Return Measures



(a) Predictive Coefficients **(b)** R^2

Figure A.2: Predictive Regressions Across Portfolios and Different Time Horizons

This figure replicates Figure 5 using cumulative real returns to compute price deviations.



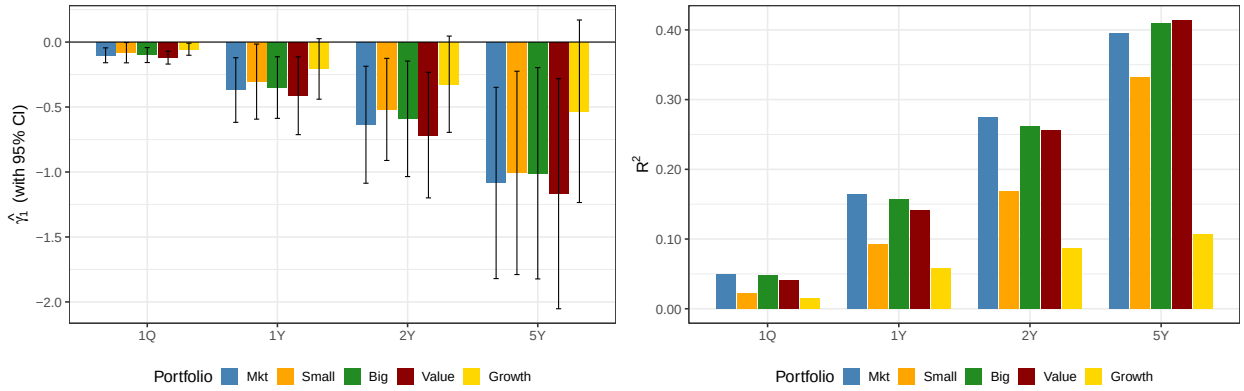
p^C

(a) Predictive Coefficients **(b)** R^2

Figure A.3: Predictive Regressions Across Portfolios and Different Time Horizons

This figure replicates Figure 5 using real capital gains to compute price deviations.

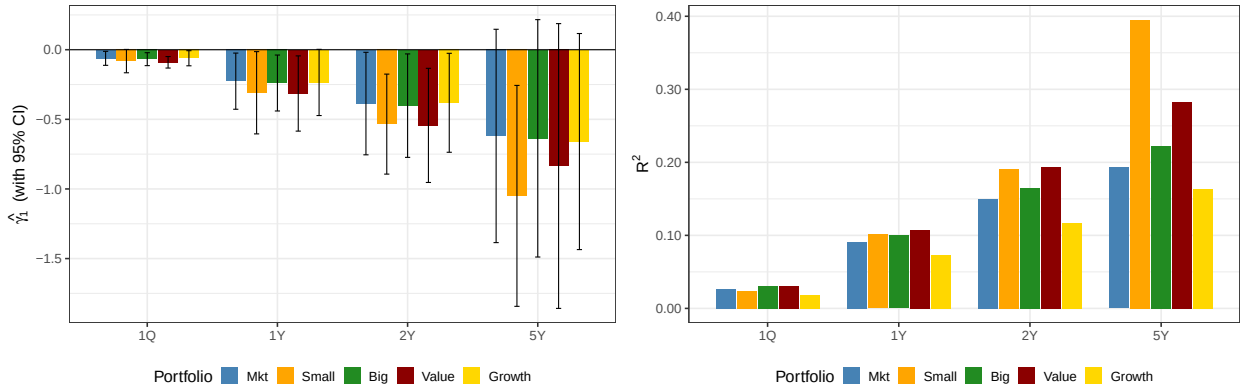
A.3 Alternative Macro-Trend Measures



(a) Predictive Coefficients **(b)** R^2

Figure A.4: Predictive Regressions Across Portfolios and Different Time Horizons

This figure replicates Figure 5 using real aggregate PCE to compute price deviations.



(a) Predictive Coefficients **(b)** R^2

Figure A.5: Predictive Regressions Across Portfolios and Different Time Horizons

This figure replicates Figure 5 using real per-capita GDP to compute price deviations.

B Time-varying Discount Rates and Mispricing

B.1 General form

We focus on what is arguably the most straightforward asset to price in a consumption-based framework: a consumption strip. Consider an asset that pays a single cash flow $\exp(\beta c_T)$

at some time T. For example, $\beta = 0$ represents a risk-free zero-coupon bond that pays 1 with certainty. We assume expected consumption growth is constant ($\bar{g}$) over time. In the terminal period, the return on this asset is

$$r_T \equiv \beta c_T - p_{T-1} \tag{B.1}$$

where p_{T-1} is the price of the claim and all variables are in logs. In all other periods, the return is

$$r_t \equiv p_t - p_{t-1}. \tag{B.2}$$

By iterating this equation for returns and taking expectations, we have the standard result that the price is simply the expected cash flow less the expected returns,

$$\begin{aligned} p_t &= \beta c_T - \sum_{j=1}^{T-t} r_{t+j} \\ &= E_t\left[\beta c_T\right] - E_t\left[\sum_{j=1}^{T-t} r_{t+j}\right] \\ &= \beta c_t - E_t\left[\sum_{j=1}^{T-t}\left(r_{t+j} - \beta \Delta c_{t+j}\right)\right] \end{aligned} \tag{B.3}$$

and that the one-period return is

$$r_t = \beta\left(\Delta c_t - \bar{g}\right) + E_{t-1}\left[r_t\right] - \left(E_t - E_{t-1}\right)\left[\sum_{j=1}^{T-t} r_{t+j}\right]. \tag{B.4}$$

For finite horizons T, returns will not be stationary over time, as the number of terms in the final summation deterministically declines to 0. So, we consider the limit as $T \to \infty$ (i.e., a strip that pays arbitrarily far in the future), which then gives that the one-period return can be written as

$$r_t = \bar{r} + \beta\left(\Delta c_t - \bar{g}\right) - E_t\left[\sum_{j=1}^{\infty}\left(r_{t+j} - \bar{r}\right)\right] + E_{t-1}\left[\sum_{j=1}^{\infty}\left(r_{t-1+j} - \bar{r}\right)\right] \tag{B.5}$$

where $\bar{r}$ is the unconditional average return. The cumulative return $r_t^{cum} = \sum_{\tau=0}^{t} r_\tau$ is then

$$r_t^{cum} \equiv \alpha_0 + \alpha_1 t + \beta c_t + \delta_t \tag{B.6}$$

where $\delta_t \equiv -E_t\left[\sum_{j=1}^{\infty}\left(r_{t+j} - \bar{r}\right)\right]$, $\alpha_1 \equiv \bar{r} - \beta\bar{g}$, and $\alpha_0 \equiv -\beta c_0 - \delta_0$. Thus, we have an instance of the cointegrating relationship (3) introduced in the main text for cumulative returns.

Our final assumption is that expected returns are AR(1). Specifically, $E_t[r_{t+1}] = \bar{r} + \mu_t$ where

$$\mu_t = \rho\mu_{t-1} + \varepsilon_t. \tag{B.7}$$

AR(1) expected returns are common in models with AR(1) underlying state variables. Appendices B.2 and B.3 show two concrete cases. With AR(1) expected returns, we have that the price-consumption cycle δ_t is also AR(1),

$$\delta_t = \rho\delta_{t-1} + v_t, \tag{B.8}$$

with shocks $v_t = -\frac{1}{1-\rho}\varepsilon_t$. This gives us the equivalent of equation (4) in the main text.

B.2 Time-varying volatility

To give a concrete example of the general model shown above, suppose consumption growth has time-varying volatility,

$$\Delta c_{t+1} = \bar{g} + \sigma_t\varepsilon^c_{t+1} \tag{B.9}$$

$$\sigma^2_{t+1} = \bar{\sigma}^2 + \rho\left(\sigma^2_t - \bar{\sigma}^2\right) + \sigma_t\sigma_\sigma\varepsilon^\sigma_{t+1} \tag{B.10}$$

where $Cov\left(\varepsilon^c_{t+1}, \varepsilon^\sigma_{t+1}\right) = -\phi < 0$. The shocks to volatility are also scaled by σ_t to ensure that volatility remains positive. All assets are priced by a representative agent with CRRA utility over consumption $\hat{\beta}^t\frac{C_t^{1-\gamma}-1}{1-\gamma}$. The negative covariance between the consumption shock and the volatility shock ensures that positive volatility shocks tend to occur in bad consumption growth states.

We consider a strip that pays $\exp(\beta c_T)$ at time T and take the limit as $T \to \infty$. We guess that the expected return on the strip is affine in σ^2_t,

$$E_t[r_{t+1}] = \bar{r} + \theta\left(\sigma^2_t - \bar{\sigma}^2\right). \tag{B.11}$$

This means that the one-period return is

$$r_{t+1} = \bar{r} - \beta\bar{g} + \beta\Delta c_{t+1} + \theta\left(\sigma^2_t - \bar{\sigma}^2\right) - \frac{1}{1-\rho}\theta\sigma_t\sigma_\sigma\varepsilon^\sigma_{t+1}. \tag{B.12}$$

We then solve for $\bar{r}$ and θ using the pricing equation for the strip's return, $E_t\left[\hat{\beta}\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\exp(r_{t+1})\right] = 1$. This gives

$$\theta = \left(\frac{1-\rho}{\sigma_\sigma}\right)^2\left[\lambda\phi - 1 + \sqrt{(\lambda\phi-1)^2 - \lambda^2}\right] \tag{B.13}$$

$$\bar{r} = -\ln\hat{\beta} + \gamma\bar{g} + \theta\bar{\sigma}^2 \tag{B.14}$$

where $\lambda \equiv \frac{\sigma_\sigma}{1-\rho}(\gamma - \beta)$. Translating to the general model, we have

$$r_t^{cum} \equiv \alpha_0 + \alpha_1 t + \beta c_t + r c_t \tag{B.15}$$

where

$$\delta_t = -\frac{1}{1-\rho}\theta\left(\sigma_t^2 - \bar{\sigma}\right). \tag{B.16}$$

B.3 Biased consumption expectations

Again, we focus on a representative agent with CRRA utility over consumption $\hat{\beta}^t \frac{C_t^{1-\gamma}-1}{1-\gamma}$. Suppose the agent believes that she lives in a long-run risk model. Specifically, she believes that consumption growth follows

$$\Delta c_{t+1} = \bar{g} + z_t + \sigma\varepsilon_{t+1}^c \tag{B.17}$$

$$z_{t+1} = \rho z_t + \sigma_z \varepsilon_{t+1}^z \tag{B.18}$$

where z_t is an observable variable and $Cov\left(u_{t+1}^c, u_{t+1}^z\right) = \frac{\sigma_{c,z}}{\sigma\sigma_z}$. We assume that consumption growth is objectively iid,

$$\Delta c_{t+1} = \bar{g} + \sigma u_{t+1}^c \tag{B.19}$$

meaning that z_t is irrelevant for growth. The variable z_t therefore captures waves of biased optimism or pessimism about future consumption growth. For example, if $\sigma_{c,z} > 0$, then the agent is overly optimistic about future consumption growth after periods of objectively good consumption shocks.

We now consider a strip that pays $\exp(\beta c_T)$ at time T, where we take the limit as $T \to \infty$. We guess that the agent's subjective expected return on the strip is $E_t^*[r_{t+1}] = \bar{r} + \theta^* z_t$. This means that the agent believes the realized return will follow

$$\begin{aligned} r_{t+1} &= E_t^*[r_{t+1}] + \beta\left(\Delta c_{t+1} - E_t^*[\Delta c_{t+1}]\right) + \left(E_{t+1}^* - E_t^*\right)\left[\sum_{j=2}^{\infty}\left(\beta\Delta c_{t+j} - r_{t+j}\right)\right] \\ &= \bar{r} + \theta^* z_t + \beta\left(\Delta c_{t+1} - \bar{g} - z_t\right) + \frac{\beta - \theta^*}{1-\rho}\sigma_z\varepsilon_{t+1}^z. \end{aligned} \tag{B.20}$$

We then solve for θ^* using the agent's pricing equation. Specifically, we need that $E_t^*\left[\hat{\beta}\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\exp(r_{t+1})\right] = 1$ where the expectation is under the agent's subjective

beliefs. This gives

$$\begin{aligned}\theta^* &= \gamma \\ \bar{r} &= -\ln\hat{\beta} + \gamma\bar{g} - \frac{1}{2}(\beta-\gamma)^2\left[\sigma^2 + \left(\frac{\sigma_z}{1-\rho}\right)^2 + 2\frac{\sigma_{c,z}}{1-\rho}\right].\end{aligned}$$

Now we deal with objective expectations. The realized one-period return follows equation (B.20). In other words, the agent is correct that realized returns follow equation (B.20), but she has the wrong expectations about what future Δc_{t+1} will be. This means that the objective expected return is

$$E_t[r_{t+1}] = \bar{r} - (\beta-\gamma)z_t. \tag{B.21}$$

In the language of the general model, we have

$$r_t^{cum} \equiv \alpha_0 + \alpha_1 t + \beta c_t + r c_t \tag{B.22}$$

where

$$\delta_t = \frac{1}{1-\rho}(\beta-\gamma)z_t. \tag{B.23}$$