

Consumption Anchors Stock Prices^{*}

Carlo A. Favero¹

Alessandro Melone²

Sean Myers³

Andrea Tamoni⁴

August 28, 2026

Abstract

Aggregate stock prices and aggregate consumption share a common stochastic trend. We estimate this long-run relation in real time and recover a *price–consumption cycle* that captures transitory deviations of stock prices from their consumption-implied value. These deviations mean-revert over business-cycle horizons and predict future returns on the aggregate market and characteristics-sorted portfolios, both in- and out-of-sample, from one quarter to two years ahead. The cycle does not forecast consumption growth, but contains information about future dividend growth, and its return-predictive power disappears when consumption is excluded from the long-run relation. A simple model with permanent and transitory consumption shocks rationalizes these findings and the time variation in the estimated price–consumption loading. The evidence identifies consumption as a macroeconomic anchor for asset prices and departures from this anchor as a source of time-varying expected returns.

Keywords: Consumption Levels, Cointegration, Time-Varying Equity Premium, Return Predictability.

JEL codes: C22, E32, E44, G12.

^{*}We thank Stefano Giglio and Andrei Gonçalves for very helpful discussions of an earlier draft of this paper, titled “Macro Trends and Factor Timing.” We are also grateful to Ravi Bansal, Robert Engle, Ravi Jagannathan, Leonid Kogan, Dmitry Livdan, Sydney Ludvigson, Konark Saxena, Allan Timmermann, and Stijn Van Nieuwerburgh for comments and suggestions.

¹Bocconi University & CEPR. E-mail: carlo.favero@unibocconi.it.

²Fisher College of Business, The Ohio State University. E-mail: melone.11@osu.edu.

³The Wharton School, University of Pennsylvania. E-mail: semyers@wharton.upenn.edu.

⁴Mendoza College of Business, University of Notre Dame. E-mail: atamoni@nd.edu

1 Introduction

A large body of evidence shows that expected returns on aggregate stocks vary substantially over time, and that this variation accounts for a sizable share of stock price fluctuations (Shiller, 1981; LeRoy and Porter, 1981; Fama and French, 1988a). A central open question is whether these movements in expected returns are tied to the real economy and, if so, through which observable channel. The natural candidate – aggregate consumption – has proven elusive. One-period stock returns and one-period consumption growth are essentially uncorrelated, a disconnect known as the “correlation puzzle” (e.g., Cochrane and Hansen, 1992; Daniel and Marshall, 1997) that has motivated decades of work on preferences, risks, frictions, and information structures designed to reconcile asset prices with consumption-based pricing kernels (Cochrane, 2005; Ludvigson, 2013; Campbell, 2017).

This paper argues that an important part of the connection between asset prices and aggregate consumption becomes visible once attention shifts from growth rates to *levels*. Aggregate stock prices and consumption share a common stochastic trend. Prices fluctuate around a long-run value associated with consumption, and temporary departures from this relation mean-revert over business-cycle horizons. We label these departures the *price–consumption cycle* (PCC). When prices lie above their consumption-implied long-run value, subsequent returns tend to be low; when prices lie below it, subsequent returns tend to be high. Thus, weak contemporaneous comovement between returns and consumption growth need not imply a weak relationship between financial markets and the real economy. If prices exhibit large but temporary departures from their consumption-implied value, the relationship in *growth rates* can be weak even when the long-run relation in *levels* is strong.

To uncover this relationship, we estimate a general cointegrating relationship between log asset prices and log consumption rather than imposing the unit coefficient embedded in conventional price-to-fundamental ratios. Our specification allows the loading of prices on consumption to differ from one and includes a deterministic time trend, so that the estimated loading reflects comovement between prices and consumption rather than being mechanically determined by their average growth rates. This choice departs from existing approaches in two ways. Standard valuation ratios, such as the price-dividend or price–consumption ratio, impose a fixed unit long-run relationship and tie predictability to mean-reversion around a constant unconditional mean. Papers studying consumption in levels instead estimate long-run relationships between consumption and non-price variables, such as wealth or dividends (Lettau and Ludvigson, 2001; Hansen et al., 2008; Bansal et al., 2009).

Our approach estimates the relationship directly from asset prices and consumption, thereby using consumption to identify the long-run component around which prices fluctuate.

Directly estimating the relationship between prices and consumption creates an important econometric challenge. If the long-run relation is estimated over the full sample, deviations from it may predict subsequent price growth partly by construction, since future prices have already contributed to the estimated relation. We therefore estimate the price–consumption relationship recursively, using only information available at each point in time. The resulting price–consumption cycle is thus a genuinely real-time predictor. This procedure preserves the information contained in prices while avoiding look-ahead bias, and allows the estimated short-run relationship between prices and consumption to adapt as their joint dynamics evolve. We find that the resulting deviation $\hat{\delta}_t$ is substantially less persistent and mean-reverts more quickly than standard valuation ratios, which helps explain why it predicts returns at short and medium horizons rather than only at long horizons.

Four empirical findings support our interpretation. *First*, asset prices and consumption are cointegrated, both for the aggregate market and across value, growth, small, and big portfolios. At the same time, the recursively estimated loading of prices on consumption varies substantially over time. For the market, the loading is generally above one, remains close to 2 for much of the decade following the Great Financial Crisis (GFC), and declines sharply around the COVID-19 recession. These magnitudes are consistent with the consumption-leverage mechanism of, e.g., [Abel \(1999\)](#) and [Bansal and Yaron \(2004\)](#). The loadings are also systematically higher for value and small portfolios than for growth and big portfolios, mirroring well-known cross-sectional patterns in average returns.

Second, the price–consumption cycle predicts returns at horizons from one quarter to two years, both in sample and out of sample. A one percentage point positive deviation from the consumption-implied value forecasts a decline in returns of roughly half a percentage point over the following year. [Figure 1](#) summarizes these results. In sample, across portfolios, the price–consumption cycle explains about 4%, 18%, and 32% of the variation in quarterly, annual, and two-year returns, respectively. This predictive power substantially exceeds that of traditional valuation ratios and, despite the simplicity of our framework, is comparable to leading machine-learning forecasts of returns ([Gu et al., 2020](#)). The results are robust to alternative return definitions, including real returns, excess returns, and real capital gains.

Out-of-sample tests confirm that this predictability is not a statistical artifact. Although predictive performance naturally declines relative to the in-sample results, the out-of-sample R^2 remains sizable, averaging about 3.4%, 14.5%, and 21.1% across




Figure 1: IS vs OOS R^2

This figure reports the in-sample and out-of-sample R^2 from predictive regressions of future portfolio excess returns on real-time price–consumption cycles, across portfolios and horizons. Out-of-sample R^2 statistics are computed using an expanding-window exercise. The sample is 1997:Q1–2024:Q4.

portfolios at the quarterly, annual, and two-year horizons, respectively. With the exception of the growth portfolio, the out-of-sample R^2 lies close to the 45-degree line, underscoring the stability and economic relevance of the link between price–consumption cycles and future returns. Even for growth, where the one- and two-year out-of-sample R^2 is lower than in sample, the corresponding estimates remain economically large, at about 13.3% and 11.1%. These magnitudes are notable given that most in-sample return predictors have zero or negative out-of-sample R^2 (Goyal and Welch, 2008; Goyal et al., 2024). The economic value is also substantial: timing strategies based on the price–consumption cycle raise Sharpe ratios by roughly 50% for the market and nearly double them for the growth portfolio, generating economically meaningful utility gains for a mean-variance investor.

Third, the price–consumption cycle contains information beyond conventional return predictors. It mean-reverts substantially faster than traditional valuation ratios, consistent with its stronger predictive performance over horizons from one quarter to two years. Its forecasting ability remains robust after controlling for standard valuation ratios and predictors derived from alternative cointegrating relationships among consumption,

dividends, and wealth. Moreover, the price–consumption cycles across portfolios share a strong common component: the first principal component explains over 80% of the variation in $\widehat{\delta}_t$. This commonality points to an aggregate state variable linked to consumption fundamentals and connects our time-series evidence to the cross-sectional factor structure emphasized by [Haddad, Kozak and Santosh \(2020\)](#) and others.

Fourth, and crucial for interpretation, the evidence points to consumption as an anchor for asset prices. The price–consumption cycle does *not* forecast future consumption growth at either short or longer horizons. Moreover, when consumption is removed from the long-run estimation, leaving only a deterministic time trend, the resulting price deviation has no predictive power for returns. The information therefore originates from the link between prices and consumption rather than from generic mean reversion in stock prices. These findings help distinguish our interpretation from an alternative mechanism in which return predictability arises because investors sell assets when consumption is temporarily low to finance current expenditures. If that mechanism were central, PCC should contain information about subsequent consumption growth, and reversals should be concentrated among non-dividend-paying stocks, which investors must sell to raise cash. Neither implication holds. The evidence is instead consistent with consumption identifying the long-run macroeconomic component of asset values, while deviations of prices from that component forecast subsequent adjustments in financial-market quantities. Consistent with this interpretation, PCC also negatively predicts dividend growth at longer horizons, indicating that dividends tend to be above trend when PCC is high. This result provides additional economic intuition for why PCC contains return-predictive information beyond that captured by the price-dividend ratio.

To organize these empirical findings, we introduce a simple model in the main text. The model distinguishes between long-lived shocks to a macroeconomic fundamental and short-lived shocks to consumption. Asset cash flows load on the long-lived fundamental, while observed consumption contains both this permanent component and transitory consumption-specific disturbances. This framework rationalizes several features of the evidence simultaneously. In particular, it (i) generates time variation in the conditional projection of prices on consumption as the relative volatility of long- and short-lived consumption shocks changes, (ii) explains why the estimated loading can fall sharply when consumption is dominated by transitory disturbances, as during the COVID-19 episode, (iii) generates return predictability from PCC without consumption-growth predictability and (iv) clarifies why PCC can outperform the price-dividend ratio and contain

information about future dividend growth. The purpose of the model is to provide a parsimonious and coherent environment for interpreting the collection of empirical findings, rather than to offer a unique structural explanation. To highlight that this is not the only possible explanation, we also provide an alternative model in which price deviations are driven by mispricing due to biased expectations about future consumption growth.

Taken together, our results identify aggregate consumption as a macroeconomic anchor for asset prices and characterize departures from this anchor as a parsimonious and economically interpretable source of return predictability. Consumption identifies a slow-moving macroeconomic trend shared by asset prices, while financial prices can move substantially around that trend. These deviations are transitory and contain information about subsequent returns. The central contribution of the paper is therefore not simply to introduce another return predictor, nor only to document a long-run relation between consumption and asset prices. It is to show that a consumption-based macroeconomic anchor for asset prices can be estimated directly and in real time, that departures from this anchor robustly forecast returns, and that the resulting empirical patterns can be rationalized by familiar economic forces involving long- and short-lived macroeconomic shocks, time-varying risk, and beliefs.

Related literature. Our finding that the price–consumption cycle negatively predicts subsequent returns is related to the trend-cycle, or permanent-transitory, decomposition of stock prices in [Fama and French \(1988b\)](#). In that 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 shared across assets and tied to macroeconomic trends, thereby showing how deviations of prices from their common real-side trend give rise to predictable return dynamics.

More broadly, our paper differs from the large literature linking the macroeconomy to asset prices. Since the seminal contribution of [Chen, Roll and Ross \(1986\)](#), building on the theoretical insights of [Merton \(1973\)](#), researchers have studied whether state variables such as inflation and economic growth proxy for systematic investment risk and help explain the cross-section of stock returns.¹ The standard approach extracts innovations in macroeconomic variables and estimates their risk prices using Fama-MacBeth procedures

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

or related methods. More recently, [Bryzgalova, Huang and Julliard \(2026\)](#) use financial returns to identify the stochastic process of consumption growth using Bayesian filtering techniques. Instead, we exploit the long-run relation between macroeconomic trends and asset prices, and use it to predict time-series variation in a given equity factor. In this respect, our paper also contributes to the literature that uses cointegration between dividends and prices ([Campbell and Shiller, 1988](#)), or between consumption and wealth ([Lettau and Ludvigson, 2001](#)), to forecast the aggregate market factor. Whereas [Lettau and Ludvigson \(2001\)](#) start from the household budget constraint but do not impose the absence of arbitrage, [Lustig, Van Nieuwerburgh and Verdelhan \(2013\)](#) estimate total wealth in a way that is consistent with both the intertemporal budget constraint and no-arbitrage restrictions. They do so using a no-arbitrage affine model to price bonds and stocks.² We discuss these connections in more detail in Section 2.1.

[Hansen, Heaton and Li \(2008\)](#) and [Bansal, Dittmar and Kiku \(2009\)](#) 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 consumption is informative about the time-series predictability of market and characteristics-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, Kozak and Santosh \(2020\)](#) 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 directly predict characteristics-based factor returns like those used in the [Fama and French \(2015\)](#) or [Hou, Xue and Zhang \(2015\)](#) models. Furthermore, we rely on the macroeconomic fundamentals, as pinned down

²[Lettau and Ludvigson \(2001\)](#) estimate *cay*, a measure of the inverse wealth-consumption ratio, and document substantial in-sample predictability for risk premia. In contrast, [Lustig et al. \(2013\)](#) find that the wealth-consumption ratio is considerably less variable than the price-dividend ratio and that it primarily predicts future variation in interest rates, rather than variation in risk premia.

by the consumption cycle, i.e. deviation of asset prices from their long-run stochastic trend shared with consumption, rather than on financial ratios.

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, Korajczyk, Lou and Sadka, 2016](#)) and spectral factor loadings (e.g., [Bandi, Chaudhuri, Lo and Tamoni, 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 big portfolios.

Our contribution, however, differs in an important way. We show that the long-run correlation between asset prices and consumption also helps to predict asset returns – what we refer to as the price–consumption 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

The time-series relationship between asset prices and consumption is crucial for consumption-based asset pricing, as consumption-based preferences are relevant for pricing stocks only to the extent that stocks prices comove with consumption. The first row of Table [1](#) tests for a direct relationship between price growth and consumption growth and shows no significant correlation for the aggregate stock market nor for any of the portfolios of Small, Big, Value and Growth stocks. Log consumption growth Δc_t is based on real personal consumption expenditures, and log price growth Δp_t is based on the real capital gains (i.e., the real

Table 1: Motivating Evidence

This table reports descriptive statistics. $AC(\cdot)$ is the autocorrelation of a given series. ***, **, and * indicates respectively 1%, 5%, and 10% level of significance using Newey–West standard errors. The sample period is 1947:Q1 to 2024:Q4. The last row reports results for $\hat{\delta}$ estimated on an expanding window with the first 30 years used as a burn-in sample. Quarterly observations.

	Mkt	Small	Big	Value	Growth
$\text{Corr}(\Delta p_t, \Delta c_t)$	0.079	0.090	0.097	0.125*	0.058
$\text{Corr}(r_t, \Delta c_t)$	0.072	0.088	0.094	0.122*	0.057
1 - $AC(pc_t)$	0.026*	0.009	0.020*	0.009	0.033*
1 - $AC(\delta_t)$	0.044***	0.106***	0.043***	0.086***	0.073***
1 - $AC(\hat{\delta}_t)$	0.084***	0.186***	0.080***	0.110***	0.144***

return excluding dividends)³ for the aggregate market, as well as the individual legs of the SMB (Small and Big) and HML (Value and Growth) factors provided by Ken French’s data library.⁴ The sample uses non-overlapping log quarterly growth for 1947:Q1-2024:Q4. The second row of Table 1 shows that the results are virtually unchanged if we replace log price growth with exact total returns. This is known as the “correlation puzzle,” as documented in Cochrane and Hansen (1992), Daniel and Marshall (1997), and Campbell and Cochrane (1999).

While these estimates of correlation allow for a general linear relationship between price growth and consumption growth, to the best of our knowledge, we are the first to estimate a general linear relationship between the log level of consumption, $c_t = c_{t-1} + \Delta c_t$, and the log level of stock prices, $p_t = p_{t-1} + \Delta p_t$.⁵ We specify the long-run relationship as:

$$p_t = \alpha_0 + \alpha_1 t + \beta c_t + \delta_t . \quad (1)$$

Section 2.2 details why a time trend is needed in this relationship even if prices and consumption are believed to share the same average growth rate over time. Inspired by Fama and French (1988b),⁶ we consider a simple case in which the deviations δ_t follow an

³Throughout the paper, both stock prices and consumption are in real terms, rather than nominal, to ensure there is no mechanical relationship between stock prices and consumption purely through inflation.

⁴The 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.

⁵Lowercase variables denote natural logarithms. For brevity, throughout the rest of the paper we shorten “log consumption” and “log prices” to simply consumption and prices.

⁶Fama 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 (1)–(2), with $z_t = \delta_t$ and $q_t = \alpha_1 t + \beta c_t$, under the assumption that consumption growth is i.i.d.

AR(1) process,

$$\delta_t = \rho\delta_{t-1} + v_t . \quad (2)$$

If the deviations are stationary, i.e., $1 - \rho > 0$, then consumption and prices are related in levels and δ_t captures the cyclical deviation of prices from the long-run value implied by consumption. We therefore refer to δ_t as the “price–consumption cycle” associated with the asset at time t .⁷

Many papers focus on the special case of a one-to-one relationship in levels, which corresponds to $\alpha_1 = 0$ and $\beta = 1$. In this case, the deviation δ_t is simply the demeaned price–consumption ratio. However, as shown in the third row of Table 1, the price–consumption ratio is highly persistent and provides little evidence of a stable one-to-one relationship between prices and consumption in levels. Across the five assets, we find only limited evidence of mean-reversion.⁸

When we consider a more general relationship, the evidence for mean-reversion becomes stronger. The fourth row of Table 1 shows the value for $1 - \rho$ from a general estimation of equations (1)-(2). Specifically, these equations can be combined into

$$\Delta p_t = \alpha_1 + \beta\Delta c_t + (\rho - 1)\delta_{t-1} + v_t . \quad (3)$$

$$= \tilde{\alpha}_0 + \tilde{\alpha}_1 t + (\rho - 1)p_{t-1} - (\rho - 1)\beta c_{t-1} + \beta\Delta c_t + v_t . \quad (4)$$

where $\tilde{\alpha}_0 = (1 - \rho)\alpha_0 + \rho\alpha_1$ and $\tilde{\alpha}_1 = (1 - \rho)\alpha_1$. Thus, we report the coefficient $1 - \rho$ from a regression of current price growth onto negative lagged prices, controlling for current consumption growth, lagged consumption, and the time trend.

While the standard errors for this coefficient account for the fact that β is an estimated object, there is an obvious concern that this mean-reversion may still arise from fitting prices to consumption in-sample. Thus, we treat this row solely as motivating evidence and the rest of the paper will focus on a real-time estimation of the relationship. As a preview, the final row of Table 1 shows the results for mean-reversion in our real-time estimate for the

⁷In the terminology of cointegration studies (Engle and Granger, 1987; Hendry, 1995), the price–consumption cycle corresponds to the “error-correction term.”

⁸Furthermore, imposing $\beta = 1$ restricts the relation between prices and consumption and may therefore be poorly suited to settings with structural change or regime shifts. This concern is related to broader evidence that predictive relationships are often unstable. For example, Paye and Timmermann (2006) document breaks in the OLS coefficient in regressions of returns on the lagged dividend-price ratio; Pastor and Stambaugh (2001) estimate breaks in the equity premium; and Lettau and Van Nieuwerburgh (2008) show that shifts in the mean of financial ratios can make forecasting relationships unstable when ignored.

deviation, $\widehat{\delta}_t$. In both of these two final rows, we find highly significant evidence for a level relationship, as the deviations are clearly mean-reverting across all five assets.

How can a relationship between prices and consumption in levels be consistent with little to no correlation between prices and consumption in growth? Conveniently, equation (3) summarizes how the relationship in levels of equations (1)-(2) connects to the relationship between current price growth Δp_t and current consumption growth Δc_t . When the volatility of shocks to the price–consumption cycle v_t dominates the volatility of consumption growth, then price growth and consumption growth may only be weakly correlated even if they have a strong relationship in levels. In fact, a strong relationship in levels plus a weak relationship in growth implies predictability of price growth. The $(\rho - 1)$ term in equation (3) indicates that any short-term price growth not explained by consumption growth, $\beta\Delta c_t$, is temporary and will eventually be reversed.⁹ Predictability disappears only when the price–consumption cycle is not mean-reverting ($\rho = 1$), in which case prices and consumption do not have a level relationship.

Before presenting the empirical results on the estimated level relationship and price-growth predictability in Sections 3 and 4, Sections 2.1–2.3 discuss how our approach relates to prior work on level relationships between prices and consumption, the role of including a time trend, and the details of the real-time estimation.

2.1 Connection to price ratios and other cointegrating relations

The central challenge in using equations (1)–(2) to predict future price growth is that, when estimated in sample and $|\rho| < 1$, the deviation δ_t predicts future price growth by construction. To avoid this “near mechanical” predictability, the literature has taken two approaches.

The first approach assumes the coefficients in equation (1), thereby avoiding in-sample estimation. Specifically, one sets $\alpha_1 = 0$, $\beta = 1$, and $\alpha_0 = \bar{p} - \bar{c}$, where $\bar{p} - \bar{c}$ denotes the unconditional mean of the price-to-consumption ratio. As discussed earlier, under these assumptions, the deviation reduces to the current price-to-consumption ratio minus its unconditional mean.¹⁰ Because the parameter values are assumed rather than estimated,

⁹This 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).

¹⁰A vast literature follows this approach, i.e., by assuming $\alpha_1 = 0$ and $\beta = 1$ and using dividends, or other measures of macroeconomic fundamentals rather than consumption, such as output (e.g., Rangvid, 2006). These predictors, such as the price-to-dividend ratio, share similar benefits and potential drawbacks to those we discuss for the price-to-consumption ratio. Section 4.4 shows the empirical differences between

one can then use the deviation to predict future price growth or returns in-sample without fear of mechanical success.

The downside of this approach is that the imposed restrictions may be overly rigid. If the mean price-to-consumption ratio changes over time, either gradually 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 types of predictors often perform poorly out of sample (Goyal and Welch, 2008; 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 believes that prices and consumption have the same average growth rate (i.e., that the price-to-consumption ratio does not trend to $\pm\infty$ over time) does not imply that $\beta = 1$.

The second approach estimates a long-run 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 in-sample relationship

$$x_t = \alpha_0 + \alpha_1 t + \beta c_t + \varepsilon_t . \quad (5)$$

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 limitation of this approach 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 related to prices only at long horizons, so that persistent deviations between x_t and p_t can arise, or the relationship between x_t and p_t fluctuates over time, then ε_t may only predict future price growth in the distant future and may perform poorly out of sample.

Compared to these two approaches, we take a different path by estimating equation (1) in real time and then using the resulting deviation to predict future price growth and returns. As in the previous approaches, this avoids the near mechanical predictability that arises when equation (1) is estimated in sample. Relative to the first approach, our method better

our estimated price-consumption cycle δ_t and a variety of ratios of prices to consumption, dividends, and earnings.

captures changes in parameters over time and allows for $\beta \neq 1$.¹¹ Relative to the second approach, we directly use 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 4.4 shows (i) that our estimated δ_t has a much lower persistence than typical predictors from the first approach (e.g. 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 is not captured in the Bansal et al. (2001), Hansen et al. (2008), and Bansal et al. (2009) dividend-to-consumption deviations nor the Lettau and Ludvigson (2001) consumption-to-wealth (cay_t). Further, Section 4.4 shows that, compared to predictors from both approaches, 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,¹²

$$c_t = \alpha_{0,c} + \alpha_{1,c}t + \tilde{c}_t , \quad (6)$$

$$p_t = \alpha_{0,p} + \alpha_{1,p}t + \tilde{p}_t . \quad (7)$$

Suppose that, instead of estimating equation (1), we estimated a restricted version without a time trend,

$$p_t = \alpha_0 + \beta c_t + \delta_t . \quad (8)$$

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 (1).

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 $\pm\infty$ over time. If we are using equation (8), then $\alpha_{1,c} = \alpha_{1,p}$ implies β must be 1. This explains

¹¹Section 3 shows that we generally estimate β much higher than 1.

¹²Note 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 $\alpha_{1,p}$.

the common approach of imposing $\alpha_1 = 0$ and $\beta = 1$ in (1). However, once we consider equation (1) 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 a time trend in equation (1). 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 (1).

2.2.1 Simple model of asset prices

To aid in the understanding of Section 2.2, we provide a simple model in which $\beta \neq 1$ even though prices and consumption have the same average growth rate. Section 6 and Appendix C provide more detailed models based on time-varying risk and mispricing, respectively.

We assume that consumption growth is i.i.d.,

$$\Delta c_{t+1} = \alpha_{1,c} + u_{t+1} . \quad (9)$$

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

$$p_t = \alpha_0 + \alpha_1 t + \beta c_t - \sum_{j=1}^{\infty} \phi^{j-1} (E_t[r_{t+j}] - \bar{r}) . \quad (10)$$

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,

$$E_{t+1}[r_{t+2}] = \bar{r} + \rho (E_t[r_{t+1}] - \bar{r}) + u_{t+1}^r \quad (11)$$

then prices follow our exact specification in equations (1)-(2). Specifically, the price–consumption cycle will be

$$\delta_t = \rho \delta_{t-1} - \frac{1}{1 - \rho \phi} u_t^r . \quad (12)$$

2.3 Real-time estimation

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

$$p_{j,t} = \alpha_{j,0} + \alpha_{j,1}t + \beta_j c_t + \delta_{j,t} , \quad (13)$$

$$\delta_{j,t} = \rho_j \delta_{j,t-1} + v_{j,t} . \quad (14)$$

As mentioned previously, estimating equations (13)-(14) 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 (13)-(14) for each asset j in real time.

A naive approach is to estimate equation (13) in real time by regressing price onto a time trend and consumption using only observations available up to time t . One could then use the residuals from this regression as an estimate of the price-consumption cycle, $\hat{\delta}_{j,t}$. This approach, however, places no restrictions on the estimated $\hat{\delta}_{j,t}$ and therefore risks overfitting: the regression residuals simply absorb whatever variation is needed to fit the observed sample, even if doing so implies an erratic time series for the $\hat{\delta}_{j,t}$.

To reduce the risk of overfitting, we incorporate the information from equation (14) into our estimation, which restricts the price-consumption cycle to be AR(1). Imposing this AR(1) restriction onto equation (13) gives a linear relationship that conveniently isolates the iid shock $v_{j,t}$,

$$p_{j,t} = \tilde{\alpha}_{j,0} + \tilde{\alpha}_{j,1}t + \rho_j p_{j,t-1} + \beta_j c_t - \rho_j \beta_j c_{t-1} + v_{j,t}, \quad (15)$$

where $\tilde{\alpha}_{j,0} \equiv \alpha_{j,0}(1 - \rho_j) + \rho_j \alpha_{j,1}$ and $\tilde{\alpha}_{j,1} \equiv (1 - \rho_j) \alpha_{j,1}$. The price-consumption cycle $\delta_{j,t}$ is then a combination of its lag and the $v_{j,t}$ shock following equation (14).¹³

To implement equations (14) and (15) in real time, we use the following procedure. We enter into period t with an estimate $\hat{\delta}_{j,t-1}$ for the previous value of the price-consumption cycle. We then estimate a regression of $p_{j,t}$ on a time trend, its lag, contemporaneous log

¹³The estimation of our Eq. (15) 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 \(1978\)](#). Also, although (4) and (15) are algebraically related, they differ in their left-hand-side variables. Equation (4) is written in terms of price growth, Δp_t , whereas (15) is written in terms of the price level, p_t . Thus, (4) can be interpreted as the short-run representation of the long-run price-level relation in (15).

Table 2: Tests for Long-Run Relationships

This table reports results for the [Engle and Granger \(1987, EG\)](#) cointegration test (Panel A) and the [Pesaran et al. \(2001, PSS\)](#) test (Panel B). In Panel A, we report t-statistics for Dickey–Fuller tests for the cointegrating residuals and use critical values from [MacKinnon \(2010\)](#). In Panel B, we report the F-statistics for Wald bounds-tests for no cointegration and use critical values from [Kripfganz and Schneider \(2020\)](#); the lag lengths p and q are set to 1 to test specification (15). ***, **, and * indicates respectively 1%, 5%, and 10% level of significance. Annual data span 1929–2024.

Panel A: EG Test. $H_0: \gamma_j = 0$ $p_{j,t} = \alpha_{j,0} + \alpha_{j,1}t + \beta_j c_t + \delta_{j,t}$ $\Delta\delta_{j,t} = \gamma_j \delta_{j,t-1} + \varepsilon_{j,t}$					
	Mkt	Small	Big	Value	Growth
Annual	−3.818**	−4.998***	−3.953**	−4.887***	−4.405***
Panel B: PSS Test. $H_0: \beta_j \cap (\rho_j - 1) = 0$ $\Delta p_{j,t+1} = \alpha_{j,1} + \beta_j \Delta c_{t+1} + (\rho_j - 1)\delta_{j,t} + v_{j,t+1}$					
	Mkt	Small	Big	Value	Growth
Annual	11.690***	21.149***	12.970***	17.028***	16.932***

consumption, and lagged log consumption,

$$p_{j,\tau} = a_{j,0} + a_{j,1}\tau + b_{j,1}p_{j,\tau-1} + b_{j,2}c_\tau + b_{j,3}c_{\tau-1} + v_{j,\tau}, \quad (16)$$

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 $\hat{\delta}_{j,t}$,

$$\hat{\delta}_{j,t} = \hat{b}_{j,1}\hat{\delta}_{j,t-1} + \hat{v}_{j,t}. \quad (17)$$

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

3 Long-Run Relation Between Price and Consumption

We begin by testing for cointegration between consumption and factor prices using standard procedures. Because cointegration tests are known to have low power, we use a long annual sample (1929–2024) to maximize statistical power.

Panel A of Table 2 reports results from the [Engle and Granger \(1987\)](#) two-step

cointegration test. This approach assesses whether the price–consumption cycle $\delta_{j,t}$ exhibits mean-reversion. Specifically, the null hypothesis is that $\delta_{j,t}$ follows a random walk, which implies $\gamma_j = 0$ in the regression $\Delta\delta_{j,t} = \gamma_j\delta_{j,t-1} + \varepsilon_{j,t}$. Across all portfolios, we reject the null at conventional significance levels. The negative test statistics indicate that $\gamma_j = \rho_j - 1 < 0$, consistent with $\delta_{j,t}$ being stationary and mean-reverting.

Panel B reports estimates of the error-correction model in equation (3). Whereas Panel A focuses solely on the long-run relationship between prices and consumption levels – captured by the price–consumption cycle – the error-correction model incorporates both short-run return 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 et al. (2001, PSS). Although rejection of the null does not guarantee that both conditions are rejected, we later verify through predictability tests and additional evidence on cointegration that both conditions are supported by the data. The PSS tests reject the null of no levels relationship for all portfolios.

Having established a long-run relationship between asset prices and consumption using in-sample cointegration tests, we next examine the real-time properties of price deviations from consumption following the framework in Section 2.3. The sample begins in 1947:Q1, and we estimate the price–consumption cycle recursively using an expanding window based on equations (16) and (17). We set the initial estimation window to 30 years, so that the first observation of the price–consumption cycle is in 1977:Q1.¹⁴

Figure 2 reports expanding-window estimates of the persistence parameter ρ_j . Consistent with the EG and PSS results, all estimates satisfy $\rho_j < 1$, confirming mean-reversion in price deviations. For the market, big, and growth portfolios, the average persistence is approximately $\rho \approx 0.95$,¹⁵ implying a half-life of $\ln(1/2)/\ln(\rho) \approx 13.5$ quarters (about three years), consistent with business-cycle fluctuations. Although the estimates are relatively stable over time, we observe a gradual decline in persistence, particularly for the small portfolio, indicating faster mean-reversion in the price–consumption cycle.

Figure 3 presents real-time estimates of β_j (Panel (a)) and the price–consumption cycle, $\hat{\delta}_j$ (Panel (b)), from equation (15) for the market portfolio.¹⁶ Panel (a) shows that the long-

¹⁴We also consider shorter initial windows. While most results are qualitatively similar for windows of at least 20 years, the parameter estimates – particularly the persistence of the price–consumption cycle – become erratic. The estimates stabilize once the initial window reaches approximately 30 years.

¹⁵See Appendix Table A.1, Panel A.

¹⁶The price–consumption cycle is a mean-zero object. In practice, however, the out-of-sample estimate is not centered around zero when plotted after excluding the burn-in period. To facilitate visual interpretation,




Figure 2: 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 30 years as a burn-in sample. Estimates span the sample 1977:Q1 to 2024:Q4. Quarterly observations.

run consumption beta for the market varies substantially over time, with a mean estimate of 1.66 and values ranging from 0.95 to 2.34. The beta declines gradually from the late 1970s through the 1990s, increases around the GFC, and remains close to 2 during the 2010s. During the COVID-19 recession, the point estimate falls sharply and slightly below 1, consistent with the exceptional disconnect between consumption dynamics and asset prices during this period. The 95% confidence bands indicate that the beta is statistically significant throughout the sample, with the exception of the post-COVID period. Overall, the estimates point to economically meaningful consumption leverage in market cash flows, consistent with theoretical models such as [Bansal and Yaron \(2004\)](#), in which stock cash flows are a levered version of consumption, and [Abel \(1999\)](#), who argues that risk premia on stocks are related to their consumption leverage, which he estimates to be around 2.

Panel (b) of Figure 3 shows that price deviations display a clear business-cycle pattern, with prices falling below their long-run level during recessions. These negative price deviations forecast higher expected returns (since $\rho < 1$), implying that discount rates are countercyclical.

Figure 4 reports the corresponding long-run betas for the characteristics-sorted factor portfolios. As with the market, these betas vary over time. Despite this time variation, a

we adjust only the series shown in the figure by adding the average full-sample estimate over the burn-in period. Appendix Figure A.1 reports the full-sample estimate of the price-consumption cycle together with its real-time counterpart.



(a) Time-Varying Beta

(b) Price Deviations


Figure 3: Market Beta and Price Deviations

Panel (a) shows the real-time factor loadings of market prices on consumption, together with their respective 95% confidence intervals; panel (b) shows the market price–consumption cycle. See Section 2.3 for details on the estimation procedure. The sample period starts in 1947:Q1 and we use the first 30 years as a burn-in sample. Estimates span the sample 1977:Q1 to 2024:Q4. Quarterly observations.

clear pattern emerges: portfolios with higher average returns – such as value and small-cap portfolios – exhibit larger consumption betas than growth and large-cap portfolios.

Overall, the evidence in Figures 2–4 supports our hypothesis that both long-run betas and the persistence of the price–consumption cycle – and thus the coefficient governing short-run return dynamics in the error-correction framework – vary over time. Imposing a stable relationship may therefore distort the analysis of return dynamics. These results also highlight the limitations of relying on full-sample cointegration tests over relatively short samples.

Appendix Figure A.2 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 Appendix Table A.1, Panel B, which reports high correlations in price–consumption cycles across portfolios, and in Panel C, which shows that the first principal component explains approximately 80% of the total variance.



(a) Size

(b) Value

Figure 4: 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 30 years as a burn-in sample. Estimates span the sample 1977:Q1 to 2024:Q4. Quarterly observations.

4 Price–Consumption Cycle and Return Predictability

This section examines the return-predictive content of the price–consumption cycle, i.e. the deviation of asset prices from their long-run stochastic trends, from several complementary perspectives. Section 4.1 begins with the in-sample evidence, assessing whether real-time price–consumption cycles forecast subsequent returns across the aggregate market and characteristics-sorted portfolios and over different horizons. Section 4.2 then evaluates the stability of this relationship out-of-sample using an expanding-window forecasting exercise, thereby assessing whether the predictive gains could have been achieved in real time. Section 4.3 turns from statistical to economic significance and measures the value of the forecasts for an investor implementing market- and factor-timing strategies. Finally, Section 4.4 compares the predictive performance of the price–consumption cycle with traditional valuation ratios and with predictors derived from alternative cointegrating relationships, in order to assess whether the signal contains information beyond that captured by existing return predictors.

Table 3: Predictive Regressions

This table reports estimates from 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 $r_{j,t+1}$ denotes the excess return on portfolio j and $\widehat{\delta}_{j,t}$ is its real-time price–consumption cycle, computed as described in Eqs. (16)–(17). Values in parentheses are Newey–West standard errors. The sample period is 1977:Q1 to 2024:Q4. Quarterly observations.

	Mkt	Small	Big	Value	Growth
$\widehat{\delta}_{j,t}$	−0.098*** (0.030)	−0.127*** (0.038)	−0.082*** (0.027)	−0.089*** (0.033)	−0.110*** (0.036)
Observations	191	191	191	191	191
R ²	0.057	0.041	0.044	0.034	0.040

4.1 In-sample

Given the evidence of cointegration established in the previous section, we next evaluate the predictive ability of price–consumption cycles for the market and for characteristics-based factor returns. Although statistically significant in-sample predictability does not guarantee out-of-sample performance, in-sample tests typically have higher power (e.g., [Inoue and Kilian, 2005](#)). For this reason, we start with the in-sample evidence.

Table 3 presents results from one-quarter-ahead predictive regressions.¹⁷ While the regression coefficients are estimated in sample, the price–consumption cycles are constructed fully out-of-sample (see Panel B of Figure 3 for the market and Figure A.2 for the characteristics-sorted portfolios). Specifically, we do not rely on a fixed cointegrating vector estimated over the full sample; instead, we recursively re-estimate the cointegrating parameters at each point in time.

The loading coefficient on the price–consumption cycle 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

¹⁷We study the predictability of excess returns to make our results comparable to previously proposed return predictors. All results in this section and the next, both in sample and out of sample, are robust to forecasting real returns and real capital gains.




Figure 5: Predictive Regressions Across Portfolios and Different Time Horizons

This figure shows predictive coefficients and R^2 for the predictive regression $r_{j,t+h} = \gamma_{0,j} + \gamma_{1,j}\hat{\delta}_{j,t} + \varepsilon_{j,t+h}$ for different portfolios j and time horizons h , where $\hat{\delta}_{j,t}$ are the real-time price deviations for each portfolio j computed as in Eq. (16)–(17). We report 95% confidence intervals computed using Newey–West standard errors for $h = 1$ or Hodrick standard errors for $h > 1$. The sample period is 1977:Q1 to 2024:Q4. Quarterly observations.

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 the price–consumption cycle from one-quarter-ahead, one- to two-year-ahead, predictive regressions, while Panel (b) displays the corresponding R^2 values. The coefficients are negative across all horizons but gradually lose statistical significance as the horizon lengthens, 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 that predictability is particularly strong for the market and the big portfolio.

Appendix Figure A.3 assesses robustness using Philadelphia Fed vintage first-release consumption data, lagged by one quarter to reflect realistic data-release timing.¹⁸ This sample starts in 1969:Q3, when vintage data first became available. The results are similar to those in Figure 5 and, if anything, show somewhat stronger predictive power.

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

¹⁸The observation for 1995:Q3 is missing in the first release, so we use the second-release vintage for that quarter.




Figure 6: IS vs OOS R^2

This figure shows the IS vs OOS R^2 from the predictive regressions $r_{j,t+i} = \gamma_{0,j} + \gamma_{1,j}\hat{\delta}_{j,t} + \varepsilon_{j,t+i}$ for different portfolios and time horizons, where $r_{j,t+1}$ denotes the excess return on portfolio j and $\hat{\delta}_{j,t}$ is its real-time price-consumption cycle, computed as described in Eqs. (16)–(17). OOS R^2 statistics are computed using an expanding window that begins 20 years after the start of the in-sample period (1997:Q1–2024:Q4). The IS R^2 is computed over the sample 1997:Q1–2024:Q4.

4.2 Out-of-sample

We next turn to the out-of-sample analysis. Recall that the sample begins in 1947:Q1, and we use the first 30 years to obtain the initial estimate of the price-consumption cycle using equations (16) and (17). As a result, the first observation of the price-consumption cycle is in 1977:Q1. Following, e.g., [Goyal et al. \(2024\)](#), we employ an expanding window with a 20-year burn-in training sample to generate the first forecast. The out-of-sample evaluation period therefore begins in 1997:Q1.

Figure 6 provides out-of-sample evidence consistent with 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 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.

Notably, the out-of-sample R^2 values remain relatively large and are only modestly lower than their in-sample counterparts at the one-quarter and one-year horizons, with the growth portfolio at the one-year horizon being the only exception. More pronounced deterioration appears at the two-year horizon, yet even in this case the out-of-sample R^2 values remain positive and economically sizable: approximately 10% for growth, 20% for value and small-cap portfolios, and just below 30% for the market and the big portfolio.

Appendices A.3-A.5 show that our results are robust to various measures of the price–consumption cycle. Appendix Figure A.4 repeats our analysis using Philadelphia Fed first-release consumption data, lagged by one quarter to reflect realistic data-release timing. Appendix Figure A.5 use real nondurable and services consumption in place of real aggregate PCE. Appendix Figures A.6 and A.7 use cumulative excess returns and cumulative real returns, respectively, instead of cumulative capital gains.

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 price–consumption cycle. 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 (e) show that our model generally outperforms the historical mean for the market and growth portfolios, with the main exception occurring around 2000. For these portfolios, the gap between in-sample and out-of-sample performance is small. Panels (b) and (c) indicate more erratic performance for the small and big portfolios prior to 2007–08, with results largely comparable to, or only slightly better than, the historical benchmark. After that period, however, the out-of-sample performance improves markedly and closely tracks its in-sample counterpart. Panel (d) reveals that the value portfolio is the most difficult to predict. Performance declines during several episodes, including non-recessionary periods such as the post-2000 years and the 2010–2020 decade, particularly in out-of-sample evaluation.

Overall, across portfolios, the price–consumption cycle performs well during the 2007–08 financial crisis but poorly during the COVID-19 recession.




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. OOS statistics are computed using an expanding window that begins 20 years after the start of the in-sample period (1997:Q1–2024:Q4). Quarterly observations.

4.3 Economic value

We also assess the economic value of market 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. The timing strategy relies on predictive regressions estimated both in-sample and out-of-sample. In all cases, the cointegrating relationship is estimated in real time, ensuring that the signal is available to the investor at the time of portfolio choice.

First, Sharpe ratios based on out-of-sample forecasts are close to those obtained using in-sample predictions, consistent with Section 4.2, which shows that out-of-sample R^2 values are

Table 4: 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. Quarterly observations. The sample period is 1977:Q1–2024:Q4. OOS statistics are computed using an expanding window that begins 20 years after the start of the in-sample period (1997:Q1–2024:Q4).

	Mkt	Small	Big	Value	Growth
B&H SR	0.40	0.32	0.41	0.34	0.32
IS SR	0.65	0.49	0.56	0.40	0.65
OOS SR	0.61	0.45	0.51	0.34	0.66
ΔU	4.26	2.31	2.42	0.06	3.17

only modestly lower than their in-sample counterparts. Second, timing generally improves performance relative to buy-and-hold. The Sharpe ratio increases by about 52% for the market (from 0.40 to 0.61) and doubles for the growth portfolio (from 0.32 to 0.66). Gains are more modest for the small and big portfolios, at approximately 40% and 24%, respectively.

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

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

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. Following, e.g., [Campbell and Thompson \(2008\)](#), we assume $\gamma = 3$. Also, we impose no short selling and limit leverage to at most 100%, implying the constraint $0 \leq w_{t+1|t} \leq 2$.

Finally, the average utility realized by the investor is given by

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

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 $\Delta = \bar{U}_1 - \bar{U}_0$. When annualized (by multiplying by four), this value can be interpreted as the portfolio

Table 5: 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). AC(.) denotes the first-order autocorrelation of each predictor. Values in parentheses are Newey–West standard errors. The sample period is 1977:Q1–2024:Q4. Quarterly observations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
PCC	−0.098*** (0.030)						−0.096*** (0.026)	−0.097*** (0.031)	−0.094*** (0.025)	−0.097*** (0.027)	−0.096*** (0.027)
P/C		−0.021 (0.019)					−0.018 (0.018)				
P/Y			−0.018 (0.014)					−0.017 (0.013)			
P/D				−0.019 (0.016)					−0.009 (0.014)		
P/D M&A-adj					−0.008 (0.016)					−0.003 (0.014)	
CAPE						−0.006 (0.007)					−0.001 (0.006)
1 - AC(.)	0.084***	0.027	0.022	0.024	0.027	0.014					
Observations	191	191	191	191	191	191	191	191	191	191	191
Adjusted R ²	0.052	0.002	0.000	0.003	−0.003	−0.001	0.052	0.052	0.049	0.047	0.047

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 4 reports these utility gains for a mean–variance investor exploiting the information conveyed by the price–consumption cycle. The forecasts based on deviations of factor prices from their long-run relation with consumption generate substantial utility gains. The annualized gains are 4.3% for the market, 3.2% for growth, and 2.4% and 2.3% for the big and small portfolios — economically significant magnitudes. Value is the exception, with a gain of essentially zero, mirroring its weaker out-of-sample performance in Figure 7.

4.4 Comparison to other return predictors

How does the predictive performance of the price–consumption cycle compare with more traditional price ratios or predictors derived from alternative cointegrating relationships? Tables 5–6 provide a direct side-by-side comparison. Because these alternative predictors were primarily developed to forecast returns – particularly aggregate market returns – we

focus on market return predictability in all comparison tests.

We begin with traditional price ratios. Column (1) of Table 5 reports a univariate regression of next-quarter market excess returns on our price–consumption cycle (PCC). Column (2) instead uses the price-to-consumption ratio, imposing a $(1, -1)$ cointegrating relationship between prices and consumption. In this case, the coefficient is statistically insignificant.

We next consider alternative price-to-fundamental ratios, using output (column 3), dividends (column 4), M&A-adjusted dividends (column 5), and ten-year cyclically adjusted earnings (column 6) as measures of fundamentals. As shown in columns (3)–(6), none of these ratios exhibit significant return predictability in univariate regressions. Moreover, columns (7)–(11) show that including them jointly with PCC leaves both the magnitude and the statistical significance of PCC coefficient largely unchanged.

Why does PCC outperform traditional price ratios? One key difference is its persistence. As discussed in Section 2.1, return predictability based on standard price ratios relies on the assumption that these ratios revert to a constant unconditional mean. However, if the mean evolves over time and/or $\beta \neq 1$, then price ratios can remain above or below their historical averages for extended periods.

Near the bottom of the table, columns (1)–(6) report $1-AC(\cdot)$, where $AC(\cdot)$ is the first-order autocorrelation of each predictor. For PCC, the estimate is 0.084 and is highly significant, indicating meaningful mean-reversion. For the traditional price ratios, this value is much closer to zero and generally statistically insignificant. In fact, for many price ratios, we cannot reject that these ratios have a unit root. Thus, the faster and more stable mean-reversion of PCC naturally translates into stronger short-horizon return predictability. While this evidence provides a statistical interpretation for the superior predictive performance of PCC relative to standard valuation ratios, Section 5.2 provides an economic interpretation through the Campbell–Shiller identity.

Table 6 compares our PCC to predictors based on other cointegrating relationships. As shown in equation (5), 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 $\alpha_1 = 0$, 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

Table 6: 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 cointegrating 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 1977:Q1–2024:Q4. Quarterly observations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
PCC	−0.098*** (0.030)						−0.118*** (0.033)	−0.108*** (0.029)	−0.096*** (0.028)	−0.093*** (0.029)	−0.096*** (0.031)
<i>dc</i>		0.032 (0.030)					−0.036 (0.032)				
<i>dc no trend</i>			0.035 (0.036)					−0.025 (0.031)			
<i>dc trend</i>				−0.060 (0.049)					−0.052 (0.045)		
EC-VAR					0.848 (0.591)					0.644 (0.446)	
<i>cay</i>						−0.185 (0.139)					−0.034 (0.165)
Observations	191	191	191	191	191	191	191	191	191	191	191
Adjusted R ²	0.052	0.000	0.000	0.002	0.010	0.001	0.052	0.049	0.052	0.055	0.047

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 5, column (1) of Table 6 reports a univariate regression of next-quarter excess returns on the price–consumption cycle. Columns (2)–(6) instead use predictors derived from alternative cointegrating relationships. Across these specifications, the R^2 values remain modest – between 0% and 1% – well below the 5% achieved by PCC.

Columns (7)–(11) include PCC jointly with each alternative predictor. PCC remains statistically significant in all cases, indicating that it captures a distinct dimension of return predictability. Put differently, by using direct data on prices to train the parameters of our cointegrating relationship, PCC incorporates information that is not embedded in these alternative specifications.

Finally, Appendix A.7 shows that our price–consumption cycle captures variation in expected equity excess returns beyond that contained in standard yield-based predictors.

5 Consumption as an Anchor

So far, we have documented a robust long-run relationship between asset prices and consumption and shown that deviations from this relationship, PCC, predict future returns. These findings, however, do not speak to the economic mechanism underlying this predictability. One possibility is that PCC reflects temporary fluctuations in consumption: when consumption is low relative to asset prices, households sell assets to finance current expenditures, generating subsequent return predictability as consumption recovers. An alternative interpretation is that consumption provides a slow-moving macroeconomic anchor for asset prices. Under this interpretation, consumption does not adjust to deviations in PCC; instead, prices and cash flows subsequently move toward the consumption-implied fundamental.

We distinguish between these interpretations in two steps. Section 5.1 examines whether PCC predicts future consumption growth and assesses the role of consumption in capturing the trend in asset prices relative to a deterministic-trend benchmark. Section 5.2 turns to cash flows, asking whether PCC predicts subsequent dividend growth and whether its return predictability differs between dividend-paying and non-dividend-paying firms.

5.1 The role of consumption

The evidence in Figure 8 shows that the price–consumption cycle does not predict consumption growth, either over the next quarter or at longer horizons, with an R^2 that is essentially zero. Furthermore, Figure 9 shows that when we estimate equation (16) using only a deterministic time trend, i.e., excluding consumption, the resulting detrended price series has little predictive power for returns. Together, these results underscore that consumption provides the long-run macroeconomic anchor toward which stock prices mean-revert, rather than simply proxying for generic low-frequency variation in prices. In contrast, prior work suggests that consumption may be less important for other ratios, such as the consumption–wealth ratio (Brennan and Xia, 2005).

Figure 10 further examines the differences between our price–consumption cycle and the consumption–wealth ratio proxy, *cay*, proposed by Lettau and Ludvigson (2001).

Panel (a) plots the two series. Because *cay* is expected to predict returns positively, while PCC forecasts return reversals, we compare *cay* with the negative of PCC. The two series comove during some periods but diverge substantially during others, resulting in a negative




Figure 8: Predicting Consumption Growth

This figure shows predictive coefficients and R^2 for the predictive regression $\Delta c_{t+h} = b_{0,i} + b_{1,i}PCC_t + \varepsilon_{j,t+h}$, where Δc_{t+h} is real PCE log growth rate, and PCC_t is the real-time market price deviation computed as in Eq. (16)–(17). We report 95% confidence intervals computed using Newey–West standard errors for $h = 1$ or Hodrick standard errors for $h > 1$. The sample period is 1977:Q1 to 2024:Q4. Quarterly observations.

correlation over the full sample.

To understand these differences, recall that *cay* is constructed as $c_t - (\omega a_t + (1 - \omega)h_t)$. Because both approaches use real PCE as the consumption measure, c_t , the discrepancy must arise from the asset-price component: our price series, p_t , versus the aggregate wealth component of Lettau and Ludvigson (2001), $\omega a_t + (1 - \omega)h_t$.

Panel (b) confirms this interpretation. The aggregate wealth component in Lettau and Ludvigson (2001) is considerably smoother than our price series. While it captures the long-run trend in cumulative ex-dividend prices, it misses several pronounced medium-term fluctuations. The two series track each other relatively closely through the 1980s but diverge thereafter, particularly as stock prices grow more rapidly.




Figure 9: Return Predictability Without Consumption

This figure shows predictive coefficients and R^2 for the predictive regression $r_{t+h} = \gamma_0 + \gamma_1 \tilde{\delta}_{time,t} + \varepsilon_{t+h}$, where r_{t+h} are market portfolio excess returns, and $\tilde{\delta}_{time,t}$ denotes the real-time deviation of the market's price from its deterministic time trend (i.e., we use Eqs. (16)–(17), but include only the time trend component, excluding consumption levels). We report 95% confidence intervals computed using Newey–West standard errors for $h = 1$ or Hodrick standard errors for $h > 1$. The sample period is 1977:Q1 to 2024:Q4. Quarterly observations.

5.2 Dividends

In Section 4, we show that the ability of the price–consumption cycle to predict market returns holds even when controlling for common return predictors such as the price–dividend ratio. The Campbell–Shiller identity imposes additional restrictions on this return predictability: a predictor beyond the price–dividend ratio must either alter the term structure of expected returns or be associated with offsetting variation in expected dividend growth. Appendix A.8 examines these restrictions through a variance decomposition of the price–dividend ratio for the market. The price–consumption cycle has little effect on the variance attribution of the price–dividend ratio, but it contains information about future dividend growth.¹⁹

¹⁹The limited effect of the price–consumption cycle on the variance attribution of the price–dividend ratio, together with its lack of long-horizon return predictability, is consistent with short- to medium-term price movements that subsequently mean-revert.




Figure 10: Market Price–Consumption Cycle vs Consumption-Wealth Ratio

Panel (a) shows the consumption-wealth ratio measure proposed by Lettau and Ludvigson (2001) (*cay*) along with the (negative of the) market’s price–consumption cycle; both variables are normalized to have zero mean and unit standard deviation. Panel (b) shows the wealth trend (*w*) in *cay* (Lettau and Ludvigson, 2001) together with real prices (*p*) for the market portfolio; *p* is normalized to start at the same level as *w*.

We examine this cash-flow predictability directly by repeating our return-predictability analysis with dividend growth as the dependent variable. Figure 11 shows the results. At the one-quarter horizon, there is little evidence that the price–consumption cycle predicts dividend growth. Predictability becomes more pronounced at the one-year horizon, with a negative coefficient and an in-sample R^2 of 3.1%, although the estimate remains imprecisely measured. At the two-year horizon, the relationship is substantially stronger: the coefficient is -0.180 , with an in-sample R^2 of 11.1% and a positive out-of-sample R^2 of 8.1%.

These results are consistent with our interpretation of the fitted consumption trend as a long-run fundamental, or cash-flow, anchor. When prices are high relative to this consumption-implied anchor, future consumption does not adjust upward. Instead, prices subsequently decline, and Figure 11 shows that dividend growth also falls. The price-dividend ratio does not fully capture this deviation because dividends are themselves high when the price–consumption cycle is high, attenuating the increase in the price-dividend ratio. Thus, the subsequent adjustment occurs in prices and dividends




Figure 11: Predicting Dividend Growth

This figure shows predictive coefficients and R^2 for the predictive regression $\Delta d_{t+h} = b_{0,i} + b_{1,i}PCC_t + \varepsilon_{j,t+h}$, where Δd_{t+h} is real dividend log growth rate, and PCC_t is the real-time market price deviation computed as in Eq. (16)–(17). We report 95% confidence intervals computed using Newey–West standard errors for $h = 1$ or Hodrick standard errors for $h > 1$. The sample period is 1977:Q1 to 2024:Q4. Quarterly observations.

rather than in consumption: both move back toward the fundamental anchor implied by consumption.

We further assess this interpretation by separating stocks according to whether they pay dividends. This split provides a way to distinguish the macro-anchor interpretation from an alternative mechanism in which households sell stocks to finance consumption when consumption is low relative to stock prices. Under this liquidation mechanism, return predictability should be stronger among non-dividend-paying firms because dividend-paying stocks already provide cash distributions that can finance consumption without requiring asset sales. In contrast, if consumption proxies for aggregate fundamentals, the price–consumption cycle should predict returns similarly across the two groups.²⁰

²⁰We split the market into dividend-paying, non-dividend-paying, and Other firms using the dividend classification of Fama and French (2001). The Other category consists of unclassified firms, defined as firms with no month in year t for which cum-dividend returns are greater than ex-dividend returns ($\text{mthret} > \text{mthretx}$) and fewer than seven months with valid returns in year t . Our baseline specification uses the same price–consumption cycle for each portfolio rather than recomputing the predictor from the dividend-paying




Figure 12: Dissecting Return Predictability: Dividend-Paying vs Non-Dividend-Paying Firms

This figure shows predictive coefficients and R^2 for the predictive regression $r_{i,t+h} = \gamma_{0,i} + \gamma_{1,i}PCC_t + \varepsilon_{j,t+h}$, where $r_{i,t+h}$ are the excess returns for dividend-paying (Div) and non-dividend-paying (NoDiv) firms, and PCC_t is the real-time market price deviation computed as in Eq. (16)–(17). We report 95% confidence intervals computed using Newey–West standard errors for $h = 1$ or Hodrick standard errors for $h > 1$. The sample period is 1977:Q1 to 2024:Q4. Quarterly observations.

Figure 12 shows little evidence that return predictability differs systematically between dividend-paying and non-dividend-paying firms. Although statistical significance and R^2 are generally stronger for dividend-paying firms, the predictive coefficients are similar across the two groups. At each horizon, we cannot reject equality of the dividend-paying and non-dividend-paying coefficients, nor can we reject that either equals the corresponding market coefficient. Indeed, the point estimates for all three portfolios lie comfortably within one another’s 95% confidence intervals.²¹ Overall, the evidence provides little support for the hypothesis that return predictability arises from selling pressure to finance consumption. That mechanism predicts stronger price declines among non-dividend-paying firms, whereas the price–consumption cycle forecasts similar declines for dividend-paying and non-dividend-paying stocks. The evidence is instead more consistent with consumption serving as a macroeconomic anchor that captures fundamentals common across firms. This interpretation is not definitive, but it is consistent with the dividend-growth evidence in Figure 11: when prices are high relative to

and non-dividend-paying return series. Results are nearly identical when we instead construct separate price–consumption cycles for the two groups using our baseline methodology.

²¹Differences in the number of dividend-paying and non-dividend-paying firms affect the precision, and hence the statistical significance, of the estimates but should not systematically affect the estimated coefficients. We therefore emphasize differences in coefficient magnitudes rather than statistical significance across the two groups.

the consumption-implied anchor, aggregate dividend growth subsequently declines. Taken together, the evidence suggests that the predictive content of the price–consumption cycle reflects adjustment toward a common macroeconomic fundamental rather than differential selling pressure on stocks that do not provide current cash distributions.

6 An Organizing Framework

Across Sections 3-5, we have documented a number of results regarding estimates of β , return predictability, consumption growth predictability, and the connection between pcc_t , the price-dividend ratio, and dividend growth predictability. To help understand these results, we present a simple model that reconciles these results and shows how they can be related to familiar forces of short-lived and long-lived shocks to consumption.

Importantly, we are not claiming that this is the only model that can explain these results. Instead, the goal is to show that these results can be understood in a straightforward, coherent environment. Appendix C shows an alternative model based on biased expectations and mispricing. Combined, these two models underscore that the cointegrating relationship can emerge in fundamentally different economic environments, including through variation in discount rates or departures from rational expectations.

We assume that there is some macroeconomic fundamental x_t that follows

$$x_t = x_{t-1} + \sigma_{\ell,t-1}\varepsilon_t \quad (18)$$

$$\sigma_{\ell,t}^2 = \bar{\sigma}_{\ell}^2 + \rho(\sigma_{\ell,t-1}^2 - \bar{\sigma}_{\ell}^2) + \sigma_{w,\ell}w_{\ell,t} . \quad (19)$$

We assume that consumption depends on this fundamental plus a short-lived shock,

$$c_t = x_t + \sigma_{s,t-1}\eta_t \quad (20)$$

$$\sigma_{s,t}^2 = \bar{\sigma}_s^2 + \rho_s(\sigma_{s,t-1}^2 - \bar{\sigma}_s^2) + \sigma_{w,s}w_{s,t} . \quad (21)$$

The variables $\sigma_{\ell,t-1}$ and $\sigma_{s,t-1}$ control the volatility of the long-lived shocks, which impact consumption permanently through x_t , and the volatility of the short-lived shocks, which only affect consumption for one period. We focus on an asset that pays dividends

$$d_t = \phi x_t + \sigma_d \eta_{d,t}$$

which are a levered version of the macroeconomic fundamental, ϕx_t , plus noise.²² For simplicity, we assume the shocks are all independent white noise, $(\varepsilon_t, w_{\ell,t}, \eta_t, w_{s,t}, \eta_{d,t}) \sim N(0, I)$.

This setup is similar in spirit to [Bansal and Yaron \(2004\)](#) in which there are long-lived shocks which impact consumption and the risky asset dividends through a long-lived latent variable x_t , as well as short-lived shocks to consumption. It is also similar to [Schorfheide, Song and Yaron \(2018\)](#), who have separate volatility processes for the short-lived and long-lived shocks to consumption. The key difference relative to both models is that we apply these mechanisms to the levels of consumption and dividends rather than their growth rates, preserving their economic intuition while generating cointegration between prices and consumption.

We assume that our representative agent has a log SDF

$$m_{t+1} = -r^f - \frac{\gamma^2}{2} \sigma_{c,t}^2 - \gamma \sigma_{c,t} \varepsilon_{c,t+1} \quad (22)$$

where $\sigma_{c,t-1} \varepsilon_{c,t} \equiv \sigma_{\ell,t-1} \varepsilon_t + \sigma_{s,t-1} \eta_t$ denotes the total shock to consumption and $\sigma_{c,t-1}^2 \equiv \sigma_{\ell,t-1}^2 + \sigma_{s,t-1}^2$. We can then easily verify that the risk-free rate is always constant $R^f = 1/E_t[\exp(m_{t+1})] = \exp(r^f)$. This ensures that return predictability and movements in prices due to discount rates are due to movements in expected excess returns and eliminates the possibility that results can be due to mechanical movements in risk-free discount rates.

With the process for consumption, dividends, and the SDF specified, we can now derive the log price p_t of the risky asset that solves the agent's Euler equation,

$$E_t[\exp(m_{t+1} + r_{t+1})] = 1$$

where r_{t+1} is the one-period return. For tractability, we use the Campbell-Shiller log-linearized return identity,

$$r_{t+1} = \kappa + \Delta d_{t+1} + \varphi(p_{t+1} - d_{t+1}) - p_t + d_t. \quad (23)$$

We guess and verify that

$$p_t = \theta - \lambda(\sigma_{\ell,t}^2 - \bar{\sigma}_{\ell}^2) + \phi x_t \quad (24)$$

²²We can easily extend the model to also include time variation in σ_d . To make the model as concise as possible, we simply use constant σ_d .

solves the agent's Euler equation for

$$\lambda = \frac{\phi\gamma - \frac{1}{2}\phi^2}{1 - \varphi\rho} \quad (25)$$

$$\theta = \frac{\kappa - r^f - (1 - \varphi\rho) \lambda \bar{\sigma}_\ell^2 + \frac{1}{2} (1 - \varphi)^2 \sigma_d^2 + \frac{1}{2} \varphi^2 \lambda^2 \sigma_{w,\ell}^2}{1 - \varphi} . \quad (26)$$

While we have a consumption-based SDF and prices do depend on the volatility $\sigma_{\ell,t}^2$ of shocks to the long-lived fundamental, note that the volatility of shocks to the short-term component $\sigma_{s,t}^2$ does not affect prices. This means that there is consumption volatility that is not priced, in line with [Bryzgalova, Huang and Julliard \(2026\)](#).

In the subsections below, we discuss how this model explains our findings that (i) the estimated $\hat{\beta}$ fluctuates over time with a large drop during COVID-19, (ii) PCC strongly negatively predicts future returns, (iii) PCC does not predict future consumption growth, (iv) PCC is a stronger predictor of future returns than the price-dividend ratio, and (v) even after controlling for the price-dividend ratio, PCC negatively predicts future dividend growth. Section 6.4 also discusses including time trends in the model variables and how this relates to the discussion in Section 2.2.²³

Before turning to these subsections, we make one remark. The model uses ingredients that are familiar from the existing literature, including separate short-lived and long-lived consumption shocks. Through the no-arbitrage condition and the Campbell–Shiller log-linearization, it generates an equilibrium price that is tied to consumption through the common persistent fundamental x_t . We view the model primarily as a parsimonious rationalization of the empirical patterns documented in the paper, rather than as a unique structural explanation. In particular, we do not specify the economic forces underlying x_t ; the model simply assumes that consumption and dividends are both driven by this common persistent fundamental.

6.1 Time-varying parameter estimates

Define $\delta_t \equiv -\lambda (\sigma_{\ell,t}^2 - \bar{\sigma}_\ell^2)$ as our AR(1) deviation. Then, we have that

$$\begin{aligned} p_t &= \theta + \phi x_t + \delta_t \\ &= \theta + \phi c_t + \delta_t - \phi \sigma_{s,t-1} \eta_t . \end{aligned}$$

²³For simplicity and to focus on intuition, we remove the time trends from our baseline model.

If we do a conditional projection, we have

$$p_t = \alpha_{0,t-1} + \beta_{t-1}c_t + \delta_t + \sigma_{u,t-1}u_t \quad (27)$$

where

$$\alpha_{0,t-1} = \theta + (\phi - \beta_{t-1})x_{t-1} \quad (28)$$

$$\beta_{t-1} = \frac{\phi}{1 + \sigma_{s,t-1}^2/\sigma_{\ell,t-1}^2} \quad (29)$$

$$\sigma_{u,t-1}^2 = \phi^2 \frac{\sigma_{\ell,t-1}^2 \sigma_{s,t-1}^2}{\sigma_{\ell,t-1}^2 + \sigma_{s,t-1}^2} \quad (30)$$

and $u_t \sim N(0, 1)$ is iid white noise. Thus, we end up with a final equation relating prices and consumption where prices depend on time-varying parameters $(\alpha_{0,t-1}, \beta_{t-1})$, an AR(1) deviation δ_t , and a residual $\sigma_{u,t-1}u_t$. So, δ_t would be our price–consumption cycle.²⁴

Importantly, note that $\beta_t = \frac{\phi}{1 + \sigma_{s,t}^2/\sigma_{\ell,t}^2}$ moves over time depending on $\sigma_{s,t}^2/\sigma_{\ell,t}^2$. During periods where the volatility of shocks to the short-lived component of consumption dominate the volatility of shocks to the long-lived component, β_t is lower. Thus, through the lens of this model, the notable decline in β_t shown in Figures 3 and 4 suggests that, during the COVID-19 period, consumption volatility was largely driven by shocks to the short-term component. In fact, if we push this interpretation further, the model would imply that there was actually an increase in the relative volatility of shocks to the long-lived component of consumption at the very start of COVID (which generates the uptick in estimated β at the start of 2020) but that this quickly changed and consumption volatility became dominated by shocks to the short-term component shortly thereafter. Similarly, the model implies that the Great Recession of 2008-2009 corresponds to an increase in the relative volatility of shocks to the long-lived component. Ex post, these explanations seem plausible, as the long-term impact of COVID was initially quite uncertain but then was dominated by the huge short-term shocks to consumption caused by the imposition and later removal of lockdowns, and the Great Recession was followed by a very prolonged, slow recovery.

²⁴Appendix B clarifies the distinction between short- and long-run price–consumption relationships: the time-varying β_t is a conditional short-run projection coefficient, whereas prices and consumption remain cointegrated with a constant long-run coefficient ϕ . We also show how a time-varying intercept reconciles these two relationships.

6.2 Price growth, return and consumption growth predictability

The expected price growth is

$$E_t [\Delta p_{t+1}] = -(1 - \rho) \delta_t$$

and the expected return is

$$E_t [r_{t+1}] = \bar{r} - (1 - \varphi \rho) \delta_t \quad (31)$$

where $\bar{r} = \kappa - \theta(1 - \varphi)$. This means that the deviation δ_t predicts future price growth and future returns at horizon h with geometrically decaying coefficients $-(1 - \rho)\rho^{h-1}$ and $-(1 - \varphi\rho)\rho^{h-1}$. This matches our empirical findings, where the predictability coefficients for cumulative returns generally follow a geometric sum.

In contrast, expected consumption growth is

$$E_t [\Delta c_{t+1}] = -\sigma_{s,t-1} \eta_t$$

which is orthogonal to the deviation δ_t . So the deviation does not predict consumption growth at any horizon.

6.3 Comparison to the price-dividend ratio and the predictability of dividend growth

The price-dividend ratio is

$$p_t - d_t = \theta + \delta_t - \sigma_d \eta_{d,t} .$$

From equation (31), δ_t is the best predictor of future returns in this model. So, δ_t will negatively predict returns even after controlling for the price-dividend ratio. Additionally, δ_t will also negatively predict future dividend growth after controlling for the price-dividend ratio, which we see from

$$E_t [\Delta d_{t+1}] = -\theta - \delta_t + (p_t - d_t) .$$

We can also compare the price-dividend ratio to an alternative measure of the deviation

that includes the noisy shock, $\tilde{\delta}_t \equiv p_t - \alpha_{0,t-1} - \beta_{t-1}c_t = \delta_t + \sigma_{u,t-1}u_t$. In this case, $\tilde{\delta}_t$ and $p_t - d_t = \theta + \delta_t - \sigma_d\eta_{d,t}$ are both noisy proxies for the true return predictor δ_t . This means that if σ_d^2 is large relative to the average value of $\sigma_{u,t-1}^2$, then $\tilde{\delta}_t$ will be a better predictor of returns than the price-dividend ratio. This also means that $\tilde{\delta}_t$ will predict future dividend growth after controlling for the price-dividend ratio.

Why would large σ_d^2 be plausible? Our interpretation of large σ_d^2 is not that dividends are literally highly volatile, but that dividends can have large deviations from the macroeconomic fundamentals of the firms ϕx_t , i.e., $\sigma_d^2 = \text{Var}(d_t - \phi x_t)$. In fact, in a more complex model where dividends are allowed to have persistent deviations from the fundamentals ϕx_t , an incentive to smooth dividends can be a microfoundation for high $\text{Var}(d_t - \phi x_t)$, as firms would choose to gradually adjust dividends in response to changes in ϕx_t . We focus on a simpler model with only one-period deviations from fundamentals as this allows straightforward closed-form solutions for prices, and we use high σ_d^2 as a short-hand for the concept that dividends may be a poor proxy of the macroeconomic fundamentals. Broadly speaking, the core idea of the return predictability result is that if $\beta_{t-1}c_t$ is a better proxy than d_t for the true macrofundamentals ϕx_t , then $\tilde{\delta}_t = p_t - \alpha_{0,t-1} - \beta_{t-1}c_t$ will be a better way to gauge when prices are “high” or “low” than the price-dividend ratio.

6.4 Including a time trend

For simplicity, our baseline model has no trends in any variables. We can generate $p_t = \alpha_{0,t-1} + \alpha_{1,t-1}t + \beta_{t-1}c_t + \delta_t + \sigma_{u,t-1}u_t$ with a non-zero time trend $\alpha_{1,t-1}t$ by making two modifications: (i) consumption is changed to $c_t = \alpha_{1,c}t + x_t + \sigma_{s,t-1}\eta_t$ and (ii) the dividend is changed to $\alpha_{1,p}t + \phi x_t + \sigma_d\eta_{d,t}$. Then, we will have

$$p_t = \alpha_{0,t-1} + \alpha_{1,t-1}t + \beta_{t-1}c_t + \delta_t + \sigma_{u,t-1}u_t \quad (32)$$

where $\alpha_{1,t-1} = \alpha_{1,p} - \beta_{t-1}\alpha_{1,c}$. This nicely fits with the discussion in Section 2.2 that even when $\alpha_{1,p} = \alpha_{1,c}$ (i.e., the price-consumption ratio does not trend to $\pm\infty$) we still need to include a time trend in equation (32) if we are allowing $\beta_{t-1} \neq 1$. Estimates of the time trend $\alpha_{1,t-1}$ will move over time as β_{t-1} moves over time.

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–consumption cycle – and show that these deviations are mean-reverting. Their dynamics provide a decomposition of prices 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–consumption 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

- Abel, Andrew B. (1999) “Risk premia and term premia in general equilibrium,” *Journal of Monetary Economics*, 43 (1), 3–33.
- 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), 105447.
- 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.
- Bansal, Ravi and Amir Yaron (2004) “Risks for the long run: A potential resolution of asset pricing puzzles,” *Journal of Finance*, 59 (4), 1481–1509.
- 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.
- Brennan, Michael J. and Yihong Xia (2005) “tay’s as good as cay,” *Finance Research Letters*, 2 (1), 1–14.
- Bryzgalova, Svetlana, Jiantao Huang, and Christian Julliard (2026) “Consumption in Asset Returns,” *The Journal of Finance*, 81 (4), 2271–2330.
- 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*, 55 (2), 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*.
- Fama, Eugene F. and Kenneth R. French (2001) “Disappearing dividends: changing firm characteristics or lower propensity to pay?” *Journal of Financial Economics*, 60 (1), 3–43.
- Fama, Eugene F and Kenneth R French (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 (2008) “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.
- Gu, Shihao, Bryan Kelly, and Dacheng Xiu (2020) “Empirical Asset Pricing via Machine Learning,” *The Review of Financial Studies*, 33 (5), 2223–2273.

- Haddad, Valentin, Serhiy Kozak, and Shrihari Santosh (2020) “Factor timing,” *Review of Financial Studies*, 33 (5), 1980–2018.
- 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.
- Inoue, Atsushi and Lutz Kilian (2005) “In-sample or out-of-sample tests of predictability: Which one should we use?” *Econometric Reviews*, 23 (4), 371–402.
- Kamara, Avraham, Robert A Korajczyk, Xiaoxia Lou, and Ronnie Sadka (2016) “Horizon pricing,” *Journal of Financial and Quantitative Analysis*, 51 (6), 1769–1793.
- Kripfganz, Sebastian and Daniel C Schneider (2020) “Response surface regressions for critical value bounds and approximate p-values in equilibrium correction models 1,” *Oxford Bulletin of Economics and Statistics*, 82 (6), 1456–1481.
- 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 C Ludvigson (2001) “Consumption, aggregate wealth, and expected stock returns,” *Journal of Finance*, 56 (3), 815–849.
- Lettau, Martin and Stijn Van Nieuwerburgh (2008) “Reconciling the Return Predictability Evidence,” *The Review of Financial Studies*, 21 (4), 1607–1652.
- Ludvigson, Sydney C (2013) “Advances in consumption-based asset pricing: Empirical tests,” *Handbook of the Economics of Finance*, 2, 799–906.
- Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan (2013) “The wealth-consumption ratio,” *Review of Asset Pricing Studies*, 3 (1), 38–94.
- MacKinnon, James G (2010) “Critical values for cointegration tests,” *QED Working Paper No. 1227*.
- 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.

- Parker, Jonathan A and Christian Julliard (2005) “Consumption risk and the cross section of expected returns,” *Journal of Political Economy*, 113 (1), 185–222.
- Pastor, Lubos and Robert F. Stambaugh (2001) “The Equity Premium and Structural Breaks,” *Journal of Finance*, 56 (4), 1207–1239.
- Paye, Bradley S. and Allan Timmermann (2006) “Instability of return prediction models,” *Journal of Empirical Finance*, 13 (3), 274–315.
- Pesaran, M. Hashem, Yongcheol Shin, and Richard J. Smith (2001) “Bounds testing approaches to the analysis of level relationships,” *Journal of Applied Econometrics*, 16 (3), 289–326.
- 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.
- Schorfheide, Frank, Dongho Song, and Amir Yaron (2018) “Identifying long-run risks: A Bayesian mixed-frequency approach,” *Econometrica*, 86 (2), 617–654.
- 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 Descriptive Statistics

Table A.1: Price–Consumption Cycles: Summary Statistics

This table reports summary statistics for the real-time price–consumption cycles estimated as in Eq. (16)–(17). The sample period starts in 1947:Q1 and we use the first 30 years as a burn-in sample. Estimates span the sample 1977:Q1 to 2024:Q4. Quarterly observations.

Panel A: Descriptive Statistics					
	Mkt	Small	Big	Value	Growth
Mean	-0.27	-0.09	-0.28	-0.12	-0.15
Std	0.20	0.18	0.21	0.21	0.19
AR1	0.92	0.81	0.92	0.89	0.86
Panel B: Correlation Matrix					
	Mkt	Small	Big	Value	Growth
Mkt	1.00				
Small	0.79	1.00			
Big	0.96	0.81	1.00		
Value	0.83	0.89	0.92	1.00	
Growth	0.69	0.66	0.72	0.62	1.00
Panel C: PCA Loadings and Variance Explained					
	PC1	PC2	PC3	PC4	PC5
Mkt	-0.46	0.08	-0.56	-0.48	0.48
Small	-0.45	0.20	0.67	-0.52	-0.20
Big	-0.47	0.12	-0.41	0.21	-0.74
Value	-0.46	0.37	0.22	0.66	0.42
Growth	-0.39	-0.90	0.14	0.14	0.07
Variance explained	0.83	0.09	0.06	0.02	0.00
Cumulative variance	0.83	0.92	0.98	1.00	1.00





Figure A.1: Market Price-Consumption Cycle: Full Sample vs Real-Time.

This figure shows the market price-consumption cycle estimated in-sample over the full 1947:Q1–2024:Q4 period together with its real-time counterpart. Quarterly observations.

A.2 Characteristics-Sorted Portfolios



(a) Size

(b) Value

Figure A.2: Price Deviations: Size and Value

This figure shows the real-time price-consumption cycles for different portfolios. The sample period starts in 1947:Q1 and we use the first 30 years as a burn-in sample. Estimates span the sample 1977:Q1 to 2024:Q4. Quarterly observations.

A.3 Real-Time Consumption




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

This figure replicates Figure 5 using *vintage* real PCE, lagged by one quarter, to compute price deviations.




Figure A.4: IS vs OOS R^2 : Vintage Consumption

This figure replicates Figure 6 using *vintage* real PCE, lagged by one quarter, to compute price deviations.

A.4 Alternative Consumption Measures




Figure A.5: IS vs OOS R^2 : NIPA ND&S

This figure replicates Figure 6 using real NIPA non-durable and services consumption data to compute price deviations.

A.5 Alternative Price Measures




Figure A.6: IS vs OOS R^2 : Excess Total Return Price

This figure replicates Figure 6 using cumulative excess returns (including dividends) to compute price-consumption cycles.






Figure A.7: IS vs OOS R^2 : Real Total Return Price

This figure replicates Figure 6 using cumulative real returns (including dividends) to compute price-consumption cycles.


A.6 Sub-Samples




(a) From 1990



(b) From 1995



(c) From 2000



(d) From 2005

Figure A.8: IS vs OOS R^2 : Different Samples

This figure replicates Figure 6 using different sample periods.

A.7 Further Control Variables

Table A.2 shows that the predictive content of the price-consumption cycle is distinct from standard interest-rate state variables. The PCC coefficient remains remarkably stable, between -0.099 and -0.105 , and statistically significant after controlling separately for the Treasury bill rate, term spread, default return spread, and default yield spread. By contrast, none of these yield-based predictors significantly forecasts equity excess returns, either alone or alongside the PCC. Thus, the PCC does not simply repackage variation in interest rates or the yield curve; rather, it captures variation in expected equity excess returns beyond that contained in standard yield-based predictors.

Table A.2: Multivariate Predictive Regressions: PCC and Yield-Based Predictors

This table presents results on aggregate market return predictability. We regress the one-quarter-ahead log market excess return on several aggregate predictors. *tbl* is the three-month Treasury bill rate; *tms* is the term spread, measured as the difference between the long-term government bond yield and the Treasury bill rate; *dfr* is the default return spread, measured as the difference between long-term corporate and government bond returns; and *dfy* is the default yield spread, measured as the difference between BAA- and AAA-rated corporate bond yields. Values in parentheses are Newey–West standard errors. The sample period is 1977:Q1–2024:Q4. Quarterly observations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PCC	−0.098*** (0.030)					−0.099*** (0.034)	−0.099*** (0.030)	−0.100*** (0.027)	−0.105*** (0.031)
<i>tbl</i>		−0.155 (0.176)				0.020 (0.160)			
<i>tms</i>			0.085 (0.328)				−0.135 (0.313)		
<i>dfr</i>				0.056 (0.383)				−0.083 (0.358)	
<i>dfy</i>					0.379 (1.640)				−1.080 (1.568)
Observations	191	191	191	191	191	191	191	191	191
Adjusted R ²	0.052	−0.001	−0.005	−0.005	−0.005	0.047	0.047	0.048	0.050

A.8 Economic Restrictions Implied by the Campbell–Shiller Identity

The log-linearized CS identity implies the following expression:

$$d_t - p_t = r_{t+1} - \Delta d_{t+1} - c + \rho (d_{t+1} - p_{t+1}), \quad (\text{A.1})$$

where d_t is log dividend, p_t is log price, $c = \log(1 + \exp\{\mathbb{E}[p - d]\}) - \rho \mathbb{E}[p - d]$ and $\rho = \frac{\exp\{\mathbb{E}[p - d]\}}{1 + \exp\{\mathbb{E}[p - d]\}}$.

Iterating Equation (A.1) forward, write:

$$d_t - p_t = \sum_{j=1}^k \rho^{j-1} (r_{t+j} - \Delta d_{t+j}) + \rho^k (d_{t+k} - p_{t+k}), \quad (\text{A.2})$$

$$= r_t^k - \Delta d_t^k + (d/p_t)^k. \quad (\text{A.3})$$

This expression readily implies the following relationship of covariances

$$\mathbb{C}(d_t - p_t, d_t - p_t) = \mathbb{C}(d_t - p_t, r_t^k) - \mathbb{C}(d_t - p_t, \Delta d_t^k) + \mathbb{C}(d_t - p_t, d/p_t^k).$$

In terms of (univariate regression) β s, the restriction on the covariances becomes

$$1 = \beta_{r,dp}^k - \beta_{\Delta d,dp}^k + \beta_{d/p,dp}^k. \quad (\text{A.4})$$

Panel (a) of Figure A.9 confirms that most variation in the price-dividend ratio reflects discount-rate variation, with $\beta_{r,dp}^k$ accounting for the dominant share.

An immediate implication of the same logic applies: if a variable x_t , *orthogonal* with respect to the dividend-to-price ratio, were to forecast long-run excess returns r_t^∞ , the same variable would also have to forecast long-run *excess* dividend growth Δd_t^∞ . These forecasts would offset each other so that, given $d_t - p_t$, the forecast of the entire right hand side of the present value identity is not altered (c.f., Equation (A.3) with $k \rightarrow \infty$). This is easy to see. Write

$$\mathbb{C}(x_t, d_t - p_t) = \mathbb{C}(x_t, r_t^k) - \mathbb{C}(x_t, \Delta d_t^k) + \mathbb{C}\left(x_t, (d/p_t)^k\right),$$

where $\mathbb{C}(x_t, d_t - p_t) = 0$, due to the *orthogonality* between x_t and $d_t - p_t$. In terms of (univariate regression) β s, the previous restriction on the covariances now becomes:

$$0 = \beta_{r,x}^k - \beta_{\Delta d,x}^k + \beta_{d/p,x}^k. \quad (\text{A.5})$$

This restriction is solely a by-product of (1) the CS identity and (2) the orthogonality of the predictor.

Panel (b) shows that variation in the price-consumption cycle primarily reflects shifts




Figure A.9: Campbell–Shiller Decomposition

This figure plots the term structure of direct multi-horizon predictive coefficients for the variance decomposition associated with the dividend-to-price ratio. The predictive slopes are obtained from weighted long-horizon regressions. The coefficients are associated with the log return (r), log dividends growth (Δd), and log dividend-to-price ratio (dp). The forecasting variable is dp in Panel (a) and the price-consumption cycle orthogonalized with respect to dp in Panel (b). “sum” denotes the value of the variance decomposition. H represents the number of years ahead. The sample period is 1977:Q1–2024:Q4. Quarterly observations.

in near-term expected returns, $\beta_{r,x}^k < 0$, accompanied by a smaller movement in expected dividend growth, $\beta_{\Delta d,x}^k < 0$. The remaining variation is accounted for by future returns and dividend growth through $\beta_{d/p,x}^k$.

B Theoretical Model: Additional Implications

B.1 Time-Varying Consumption β and Cointegration

The estimated time-varying beta estimated in equation (32) is a valid *conditional short-run projection coefficient*. However, the estimated time-varying beta is not the long-run coefficient relating the levels of asset prices and consumption. To see this, note first that

$$p_t - \phi x_t = \theta + \delta_t .$$

Since x_t is $I(1)$ and δ_t is $I(0)$, p_t and x_t are cointegrated with cointegrating coefficient ϕ . Moreover, because

$$\begin{aligned} c_t &= x_t + s_t \\ s_t &= \sigma_{s,t-1} \eta_t \end{aligned}$$

with s_t stationary, we also have

$$\begin{aligned} p_t - \phi c_t &= \theta + \phi x_t + \delta_t - \phi(x_t + s_t) \\ &= \theta + \delta_t - \phi s_t \end{aligned}$$

Therefore p_t and c_t are themselves cointegrated with the *constant* long-run coefficient ϕ .

This creates an important distinction:

$$\boxed{\text{Long-run coefficient} = \phi} .$$

whereas

$$\boxed{\text{Short-run coefficient} = \beta_t = \phi \frac{\sigma_{\ell,t}^2}{\sigma_{\ell,t}^2 + \sigma_{s,t}^2}} .$$

Consequently, equation (32) should not in general be interpreted as a time-varying cointegrating coefficient. A sufficiently long regression in levels may be dominated by the common stochastic trend x_t and therefore identify ϕ , even when the short-run conditional projection coefficient β_t moves substantially over time.

A related issue arises from the intercept in equation (32),

$$\alpha_{0,t-1} = \theta + (\phi - \beta_{t-1}) x_{t-1} .$$

Since x_{t-1} is nonstationary, the model does not naturally generate a simple deterministic intercept or linear time trend when $\beta_t \neq \phi$. The time-varying intercept is precisely what allows the short-run conditional slope to differ from the constant long-run coefficient.

C Alternative model based on mispricing

Section 6 provides a simple model of time-varying risk that reconciles our empirical findings. To demonstrate that time-varying risk is not the only possible explanation, we also provide an alternative model based on biased expectations for the long-term component of consumption. We intentionally keep nearly all of the equations unchanged from Section 6 to highlight how biased expectations can play the same role as time-varying risk. Specifically, the only changes are that we remove the time-variation in the volatility of long-term shocks to consumption, $\sigma_{\ell,t}$, and instead introduce a sentiment variable S_t which impacts the agent's expectations.

There is a macroeconomic fundamental x_t and an observable sentiment variable S_t that objectively follow

$$x_t = x_{t-1} + \sigma_\ell \varepsilon_t \quad (\text{C.1})$$

$$S_t = \rho S_{t-1} + \sigma_S \varepsilon_{S,t} . \quad (\text{C.2})$$

The agent believes that x_t follows

$$x_t = x_{t-1} + S_{t-1} + \sigma_\ell \varepsilon_t . \quad (\text{C.3})$$

The only difference from equations (18)-(19) is that the agent believes x_t is affected by an AR(1) sentiment S_{t-1} rather than AR(1) volatility $\sigma_{\ell,t-1}^2$. While this is not needed for our results, we allow for non-zero correlation in the shocks, $\sigma_{x,S} \equiv \text{Cov}(\sigma_\ell \varepsilon_t, \sigma_S \varepsilon_{S,t})$. A natural case would be $\sigma_{x,S} > 0$, which means the agent is overly optimistic after periods of objectively good shocks to the macroeconomic fundamental.

The remainder of the model is identical to Section 6. Consumption depends on the macroeconomic fundamental plus a short-lived shock,

$$c_t = x_t + \sigma_{s,t-1} \eta_t \quad (\text{C.4})$$

$$\sigma_{s,t}^2 = \bar{\sigma}_s^2 + \rho_s (\sigma_{s,t-1}^2 - \bar{\sigma}_s^2) + \sigma_{w,s} w_{s,t} . \quad (\text{C.5})$$

and the risky asset pays dividends which are a levered version of the macroeconomic fundamental plus noise,

$$d_t = \phi x_t + \sigma_d \eta_{d,t} .$$

We assume all shocks are independent Gaussian white noise, with the aforementioned exception of allowing correlation between ε_t and $\varepsilon_{S,t}$. The representative agent has a log SDF

$$m_{t+1} = -r^f - \frac{\gamma^2}{2} \sigma_{c,t}^2 - \gamma \sigma_{c,t} \varepsilon_{c,t+1} \quad (\text{C.6})$$

where $\sigma_{c,t-1} \varepsilon_{c,t} \equiv \sigma_\ell \varepsilon_t + \sigma_{s,t-1} \eta_t$ and $\sigma_{c,t-1}^2 \equiv \sigma_\ell^2 + \sigma_{s,t-1}^2$.

We guess that the log price for the risky asset is

$$p_t = \theta + \lambda S_t + \phi x_t \quad (\text{C.7})$$

and use the Campbell–Shiller log-linearized return identity,

$$r_{t+1} = \kappa + \Delta d_{t+1} + \varphi (p_{t+1} - d_{t+1}) - p_t + d_t . \quad (\text{C.8})$$

Let $E_t^* [\cdot]$ denote the agent's subjective expectations. The values of (λ, θ) are determined by the agent's Euler equation,

$$E_t^* [\exp (m_{t+1} + r_{t+1})] = 1$$

which gives

$$\lambda = \frac{\phi}{1 - \varphi \rho} \quad (\text{C.9})$$

$$\theta = \frac{\kappa - r^f + \frac{1}{2} (1 - \varphi)^2 \sigma_d^2 + \frac{1}{2} \varphi^2 \lambda^2 \sigma_S^2 - (\gamma - \frac{1}{2} \phi) \phi \sigma_\ell^2 - (\gamma - \phi) \varphi \lambda \sigma_{x,S}}{1 - \varphi} . \quad (\text{C.10})$$

Define $\delta_t \equiv \lambda S_t$ as our AR(1) deviation. Then, we have the same solution for log prices

$$p_t = \theta + \phi x_t + \delta_t \quad (\text{C.11})$$

as Section 6. This means the intuition for our empirical results continues to hold in this behavioral version of the model. A conditional projection

$$p_t - \delta_t = \alpha_{0,t-1} + \beta_{t-1} c_t + \sigma_{u,t-1} u_t \quad (\text{C.12})$$

gives

$$\alpha_{0,t-1} = \theta + (\phi - \beta_{t-1}) x_{t-1} \quad (\text{C.13})$$

$$\beta_{t-1} = \frac{\phi}{1 + \sigma_{s,t-1}^2 / \sigma_\ell^2} \quad (\text{C.14})$$

$$\sigma_{u,t-1}^2 = \phi^2 \frac{\sigma_\ell^2 \sigma_{s,t-1}^2}{\sigma_\ell^2 + \sigma_{s,t-1}^2} . \quad (\text{C.15})$$

Thus, the value of β_{t-1} varies over time, declining during periods of volatile short-lived shocks to consumption. The objective expected price growth is

$$E_t [\Delta p_{t+1}] = - (1 - \rho) \delta_t , \quad (\text{C.16})$$

the objective expected return is

$$E_t[r_{t+1}] = \bar{r} - (1 - \varphi\rho) \delta_t , \quad (\text{C.17})$$

and the objective expected consumption growth is

$$E_t[\Delta c_{t+1}] = -\sigma_{s,t-1} \eta_t . \quad (\text{C.18})$$

Thus, the deviation δ_t predicts future price growth and future returns with geometrically decaying coefficients, and the deviation does not predict consumption growth. The price-dividend ratio is

$$p_t - d_t = \theta + \delta_t - \sigma_d \eta_{d,t} . \quad (\text{C.19})$$

Because of the $\sigma_d \eta_{d,t}$ component, the price-dividend ratio will be a worse predictor of future returns than δ_t , and δ_t will negatively predict future dividend growth after controlling for the price-dividend ratio,

$$E_t[\Delta d_{t+1}] = -\theta - \delta_t + (p_t - d_t) . \quad (\text{C.20})$$