Consumption in Asset Returns*

Svetlana Bryzgalova† Christian Julliard‡

March 8, 2019

Abstract

Consumption dynamics are hard to measure accurately in the data, yet they are the crucial ingredient of macro-finance asset pricing models. The central insight of these models is that, in equilibrium, both consumption and returns are largely driven by the same fundamental shocks. Therefore, we use the information in returns to identify the underlying process of consumption. We find that aggregate consumption growth reacts over multiple quarters to the innovations spanned by bond and stock returns. This persistent component: (a) is economically large i.e. it accounts for about 26% of the total variation in consumption; (b) drives most of the time series variation of stocks and a significant (yet small) fraction of bond returns; (c) is reflected in the term structure of interest rates; and (d) is priced jointly in the cross-sections of bond and stock returns. These results, stable across estimation techniques and robustness checks, pose a novel challenge for asset pricing theory.

Keywords: Consumption Dynamics, Asset Returns, Consumption-Based Asset Pricing, Term Structure.
JEL Classification Codes: E21, E27, G12, E43, C11.

---

*We benefited from helpful comments from Ravi Bansal, Douglas Breeden, Mike Chernov, Anna Cieslak, John Cochrane, Ian Dew-Bekker, Ravi Jagannathan, Martin Lettau, Lars Lochstoer, Dong Lou, Nour Meddahi, João Mergulhão, Tarun Ramadorai, Nick Roussanov, Ken Singleton, Andrea Tamoni, Annette Vissing-Jorgensen, Amir Yaron, and seminar and conference participants at Berkeley Haas, Kellogg, LSE, Stanford, LEAP conference, LBS, USC Marshall, UC San Diego, Federal Reserve Board of Governors in Washington DC, Federal Reserve Bank of Boston, Geneva Finance Research Institute, UIUC, Pennsylvania State University, WU Vienna University of Economics and Business, Duke/UNC 2016 Asset Pricing conference, EFA 2016 annual meeting, HSE ICEF, LSE, MFA 2016 conference, SoFiE 2016 conference, 2017 FIRN Asset Pricing workshop (Melbourne), Toulouse School of Economics, University of Milan. Any errors or omissions are the responsibility of the authors. Christian Julliard thanks the Economic and Social Research Council (UK) [grant number: ES/K002309/1] for financial support.

†London Business School, sbryzgalova@london.edu

‡Department of Finance, FMG, and SRC, London School of Economics, and CEPR; c.julliard@lse.ac.uk.

I Introduction

Consumption-based models have contributed to our understanding of financial markets, business cycle dynamics, and household decision making. In these settings, assumptions about consumption persistence and volatility have a profound impact on both the empirical performance and policy implications of models. However, using consumption data alone, it is hard to identify the underlying stochastic process.[1] As a result, macro-finance researchers tend to rely on assumptions that are difficult, if not impossible, to validate outside of the frameworks under consideration: Chen, Dou, and Kogan (2015) refer to this source of model fragility as 'dark matter.'

This paper tries to fill the gap in the understanding and modeling of the consumption process. Our identification strategy is rooted in the central insight of the intertemporal Euler equation of models that have consumption as one of the state variables entering the utility function: most shocks affecting the household force it to adjust *both* investment and consumption plans. In fact, asset prices that reveal information about the state variables of the economy are a feature of almost any consumption-based macro-finance model. Therefore, one can use the cross-section of returns to extract the innovations that are reflected in both consumption and financial assets. Our approach allows the *joint* consumption and return data to 'speak for itself,' and helps us establish a new set of facts regarding the dynamics of stocks, bonds, and consumption growth.

Our first result is summarized in Figure 1 that highlights the slow, economically and statistically significant, empirical cumulative response of consumption growth to the common innovations spanned by stock and bond returns. We find that these shocks take about two-three years to be fully reflected in consumption, hence producing substantial (but not excessive) predictability. The economic magnitude of this channel is large: a one standard deviation shock implies a cumulated response of consumption of about 1% over the next 2-4 years. Most importantly, they generate a clear business cycle pattern in the conditional expectation of consumption growth that accounts for more than a quarter of total consumption variance – more than twice what is normally assumed in leading macro-finance models.[2] Furthermore, these very same shocks play a fundamental role in explaining the time series of financial assets: they account for most of the variance of equities, and drive a significant yet small share of bond returns fluctuations.

[1]See e.g. Beeler and Campbell (2012), Campbell (2017), Cochrane (2007), and also Ludvigson (2012) for a review of the empirical challenges of consumption-based asset pricing.

[2]For instance, the contribution of the conditional mean to the variance of the consumption growth process is: zero in the habit framework of Campbell and Cochrane (1999), in Lucas (1987) calculation of the cost of business cycle fluctuation, and in the rare disasters model of Barro (2006), 4.5% in the long-run risk calibration of Bansal and Yaron (2004) and 12% in that of Hansen, Heaton, Lee, and Roussanov (2007).

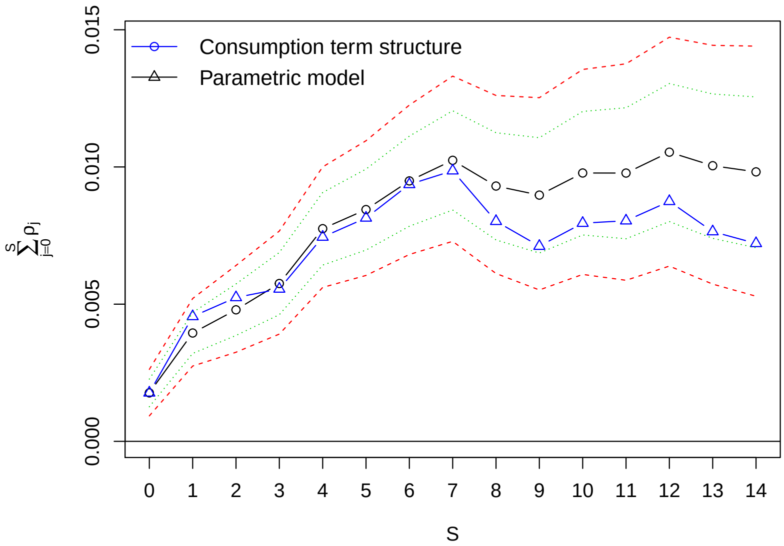

Figure 1: Cumulative response function of consumption growth to one standard deviation shock spanned by asset returns.

Posterior means (black line with circles) and centered posterior 90% (dashed red lines) and 68% (dotted green lines) coverage regions. The blue line with triangles denote the first principal component of $cov(r_{i,t}^{ex}, \Delta c_{t,t+1+S})$. Quarterly data, 1961:Q3-2017:Q2.

Although our identification approach relies only on the joint time series dynamics of consumption and asset returns, it also has powerful implications for cross-sectional asset pricing. In particular, the loadings on the common shocks provide an immediate measure of the consumption risk in stock and bond returns. We find that this measure can *jointly* explain the average term structure of the interest rates and a broad cross-section of equity risk premia (including industry portfolios). In other words, our identification strategy uncovers the fundamental link between the consumption process and asset returns, allowing us not only to accurately capture the former, but also to precisely measure its comovement with the latter. In doing so, our paper also rationalizes the empirical success of cross-sectional asset pricing models for stock returns that rely on consumption risk measured at low frequencies,[3] and establishes a complementary new finding for bonds.

To uncover, identify, and test the stochastic process of consumption growth, we employ three different empirical strategies.

First, we set up a flexible (state-space) model that extracts the common shocks to con-

[3]See, e.g. Brainard, Nelson, and Shapiro (1991), Parker and Julliard (2005), Jagannathan and Wang (2007), Malloy, Moskowitz, and Vissing-Jorgensen (2009).

sumption and asset returns and estimates their propagation pattern within the time series of consumption growth. The timing is crucial here. There is rich empirical evidence suggesting that consumption could be slow to adjust to changing economic conditions, and multiple theoretical frameworks consistent with it. Therefore, we allow consumption to react (potentially) with lags to these common shocks. In particular, we model consumption growth as the sum of two independent processes: a (potentially, since parameters are estimated) long memory moving average that (potentially) co-moves with asset returns, and a transitory component orthogonal to financial assets. Innovations to asset returns are in turn modeled as depending (potentially) on the shocks to the persistent component of consumption plus an orthogonal source. The joint system allows us to test and identify the presence of a slow moving component in consumption, and verify whether the latter is spanned by the shocks affecting financial markets. To draw the analogy with the Wold's time series decomposition, instead of ex ante postulating, say, an AR(1) process for consumption, we estimate the moving average representation and identify its innovations with the help of asset returns. The key identifying restriction is that the innovations should be reflected by financial markets *at the same time* as they occur. This assumption is not only theoretically sound, since equilibrium prices are forward-looking, but is also supported by a rich set of reduced form empirical evidence that we provide.

Our findings have important implications not only for the conditional mean of consumption growth, but for its volatility as well. We show that once the slow-moving component of consumption is properly accounted for, there is no significant evidence of volatility clustering. This implies that either aggregate consumption volatility shocks are only weakly supported by the data,or consumption-based asset pricing models need to incorporate a strong *contemporaneous* leverage effect in consumption growth, i.e. a contemporaneous correlation between mean and volatility shocks – a formulation almost never used in the literature.[4] Furthermore, we show that misspecification of the mean process makes it likely to detect time-varying volatility even in its absence.

Our second approach relies on directly assessing the main empirical prediction of our framework: a particular *term structure* of the covariances between asset returns and multi-period consumption growth. We therefore measure it directly in the data and test whether it implies a similar consumption persistence as in the state-space setting. Figure 1 confirms that the results of the two approaches are almost identical.

Finally, we confirm the asset pricing implications of our model through a cross-sectional

[4] In a typical model with stochastic volatility in consumption, the shocks to the conditional mean are not correlated with those of volatility, and even a standard lead-lag leverage effect would not be consistent with our findings. To the best of our knowledge, the only papers that consider contemporaneous leverage effects are the theoretical contribution of Tinang (2018) and the latent factor VAR model of Zviadadze (2018).

Empirical Likelihood–based estimation of a very broad class of consumption-based pricing kernels *without* imposing restrictions on the time series properties of the data. Consistent with our state-space framework, the future response of consumption growth to current asset returns accurately captures the level and the spread of the consumption risk for both stocks and bonds. Compared to standard approaches, this leads not only to sharper inference, but also remarkably better cross-sectional fit.

To summarize, we find that: *a*) consumption reacts very slowly (i.e. over a period of two to four years), but significantly, to the shocks spanned by asset returns, and this slow-moving component accounts for about 26% of its time series variation; *b*) accounting for this slow-moving dynamics, we find no significant evidence of volatility clustering in consumption, implying that models with stochastic volatility in consumption should include a contemporaneous leverage effect; *c*) stock returns significantly load on these common shocks (that capture 36%–95% of their time series variation), and display a loadings pattern that closely mimics the value and size pricing anomalies; *d*) US Treasury bond returns load significantly on the same innovations, with loadings increasing with the time to maturity, but these shocks drive no more than 3% of their time series variation; *e*) most of the time series variation in bonds is captured by a single factor, independent from both consumption and stock returns, that does not command a risk premium (i.e. it is unspanned); *f*) the total exposure to consumption risk, captured by the latent factor model and its estimated loadings, accounts for 59%-92% of the joint cross-section of stocks and bond returns;[5] *g*) low frequency consumption explains up to 92% of the cross-sectional variation in bond excess returns. Furthermore, all of the above results are confirmed by the extensive preliminary evidence that we provide.

The remainder of the paper is organized as follows. The next section reviews the related literature, while Section III provides empirical support for the state-space formulation introduced in Section IV. Our main empirical findings are presented in Section V, and Section VI concludes. Data description as well as additional methodological details, robustness checks and supplementary empirical evidence, are reported in the Appendix.

II Related literature

Our paper is related to several large strands of literature. At the core of our identification strategy lies the notion that equilibrium prices of financial assets should be determined by

[5] In our baseline specification we consider a cross-section of 46 asset given by 12 industry portfolios, 25 size and book-to-market portfolios, and 9 bond portfolio, but the results appear robust to alternative specifications.

their risk to households' marginal utilities and, in particular, current and future marginal utilities of consumption: agents are expected to demand a premium for holding assets that are more likely to yield low returns when the marginal utility of consumption is high, i.e. when consumption (current and expected) is low. Therefore, by its very nature, our paper is closely linked to consumption-based empirical asset pricing.[6] Within this literature, the underlying stochastic process of consumption growth has been the subject of a long-standing debate.

While there is ample empirical evidence suggesting that either shocks to the consumption mean or its volatility are priced in the cross-section of asset returns, there is much less agreement on what are the relative contributions of these components, the frequency of the shocks, and even the sign of their price of risk. Lettau and Ludvigson (2001a), Bansal, Dittmar, and Lundblad (2005), Malloy, Moskowitz, and Vissing-Jorgensen (2009), Savov (2011), and Kroencke (2017) focus on shocks to the first moment of consumption growth. Bansal, Kiku, Shaliastovich, and Yaron (2014) and Campbell, Giglio, Polk, and Turley (2018) argue that volatility shocks are priced, even conditional on those to the mean. Jacobs and Wang (2004), and Balduzzi and Yao (2007) use survey data to estimate the variability of idiosyncratic consumption risk across households, and find that it is priced in the cross-section of portfolios sorted by size and value. Tedongap (2007) estimates the conditional volatility of consumption through a GARCH model and finds that the value stocks are more exposed to its innovations, leading to a corresponding risk premium. Bandi and Tamoni (2015) and Boons and Tamoni (2015) go a step further and decompose the process for consumption growth into different, frequency specific, components. They find that only the shock with a half-life of about 4 years, i.e. the business cycle frequency, plays a significant role in explaining the cross-section of returns, thus supporting the importance of business cycle fluctuations in determining the risk premium. Dew-Becker and Giglio (2016) instead argue that the shocks that are most important for explaining the joint dynamics of macroeconomic fundamentals and asset returns are those of an extremely low frequency, and find no evidence of time-varying consumption volatility. In contrast, Zviadadze (2018) argues that shocks to the variance of consumption growth, and to the long-run mean of volatility, account for a large part of the time series variation of the market discount rate.

Naturally, our paper is also related to the literature on excess smoothing in consumption (e.g. Deaton (1987), Campbell and Deaton (1989)). In particular, it rationalizes the findings of Parker and Julliard (2005) that consumption risk, measured by the covariance of asset returns and cumulated future consumption growth, can explain a large fraction of the size and

[6] E.g. Breeden, Gibbons, and Litzenberger (1989), Lettau and Ludvigson (2001b), Jagannathan and Wang (2007), Piazzesi, Schneider, and Tuzel (2007), Hansen, Heaton, Lee, and Roussanov (2007), Bansal, Kiku, and Yaron (2012), and see Ludvigson (2012) for an excellent review.

value anomalies. Furthermore, our findings are consistent with the empirical evidence linking slow movements in consumption to asset returns[7] and the different economic mechanisms that can explain it. For instance, models with: constraints on information flow and rational inattention (e.g. Flavin (1993), Caballero (1995), Lynch (1996), Abel, Eberly, and Panageas (2013)), adjustment costs and consumption commitments (e.g. Marshall and Parekh (2002), Postlewaite, Samuelson, and Silverman (2008), Shore and Sinai (2010), Chetty and Szeidl (2016)), complementarities in consumption/marginal utilities (e.g. Eichenbaum, Hansen, and Singleton (1988), Startz (1989), Abel (1990), Piazzesi, Schneider, and Tuzel (2007), Yogo (2006)).

Chen and Ludvigson (2009), in the consumption habit setting, and Schorfheide, Song, and Yaron (2018), within the long-run risks framework, share our perspective of studying the process for consumption through the lens of asset returns. The former treats the functional form of the habit as unknown, and estimates it with the rest of the model's parameters. The latter proposes a bayesian strategy for identifying the deep parameters of the model using a mixed frequency setting. Contrary to most empirical work, we do not take an ex ante stand on the preferences, nor on the the speed or the pattern (e.g. an AR(1)) of the consumption dynamic. Instead, we work with the MA representation of the consumption process. This identifies directly the underlying shocks, and the consumption response to them, through the joint behaviour of consumption and asset returns. Hence, we consistently recover the horizon of the shocks spanned by asset returns and consumption growth, as well as the speed and patterns of their propagation, for a broader class of possible stochastic processes.

Our approach, that leverages the information contained in a rich cross-section of financial assets to provide insights about the underlying state variable, is also similar in spirit to a recent work by Jagannathan and Marakani (2016). They show that the price-dividend ratios of a cross-section of asset returns can be used to estimate the long-run risk process of Bansal and Yaron (2004). Their paper, as our, acknowledges that since both consumption and the real risk-free rate are measured with considerable error, it is hard to rely on market-wide indicators to infer the degree of predictability in the data. Instead, relying on a broad cross-section of asset returns provides a natural remedy to this problem. Indeed, Liew and Vassalou (2000) show that the cross-sectional value and size factors are leading indicators of future economic growth, and that the information content of these portfolios is largely independent from that of the aggregate stock market. Similarly, Ang, Piazzesi, and Min (2006) confirm that the yield curve is informative about future GDP growth.[8] These findings lead us to

[7] E.g., Daniel and Marshall (1997), Bansal, Dittmar, and Lundblad (2005), Jagannathan and Wang (2007), Hansen, Heaton, and Li (2008), Malloy, Moskowitz, and Vissing-Jorgensen (2009).

[8] Also Cochrane and Piazzesi (2005) find that a single factor (a tent-shaped linear combination of forward rates), predicts excess returns on one- to five-year maturity bonds. This factor tends to be high in recessions,

include, among other assets, size and value sorted portfolios of stocks and a cross section of bond returns in our empirical analysis.

Although the main focus of our paper is on identifying the underlying process of consumption growth, our paper is also naturally connected to the large literature on co-pricing of stocks and bonds. Our findings on the relevance of macroeconomic risk are related to a series of works that combine the affine asset pricing framework with a parsimonious mix of macro variables and bond factors for the joint pricing of bonds and stocks. In particular: Bekaert and Grenadier (1999) and Bekaert, Engstrom, and Grenadier (2010) present a linear model for the simultaneous pricing of stock and bond returns; Lettau and Wachter (2011) focus on matching an upward sloping bond yield term structure and a downward sloping equity yield curve via an affine model that incorporates persistent shocks to the aggregate dividend, inflation, risk-free rate, and price of risk processes; Koijen, Lustig, and Nieuwerburgh (2017) develop an affine model in which three factors –the level of interest rates, the Cochrane and Piazzesi (2005) factor, and the dividend-price ratio– have explanatory power for the cross-section of bonds and stock returns. Finally, our identification of an unspanned latent factor in bond returns is consistent with a similar finding in Chernov and Mueller (2012).

III Preliminary empirical evidence

To motivate the structure of the state-space model for consumption and asset returns, we first establish a set of empirical facts via model-free reduced-form approaches. We document that *a*) consumption growth is autocorrelated, *b*) not only asset returns predict future levels of consumption growth, but do it better than the past values of consumption itself, and *c*) consumption is slow to react to the innovations that it jointly spans with asset returns, while the latter react to them immediately. A detailed description of the data used throughout the paper is reported in Appendix A.1.

First, Figure 2 plots the autocorrelation function (left panel), and the p-values (right panel) of the Ljung and Box (1978) and Box and Pierce (1970) tests of joint significance of the autocorrellations, of the one quarter log consumption growth ($\Delta c_{t,t+1}$). The figure clearly shows that the autocorrelations are individually statistically significant up to the one year horizon (left panel), and jointly statistically significant (right panel) at the 1% level even after about 14 quarters (and significant at lower confidence levels at even longer horizons). That is, there is substantial persistence in the time series of consumption growth.[9]

but forecasts future expansion, i.e. it seems to incorporate good news about future consumption.

[9]Note that, even in the seminal examination of the random walk hypothesis of Hall (1978), the presence

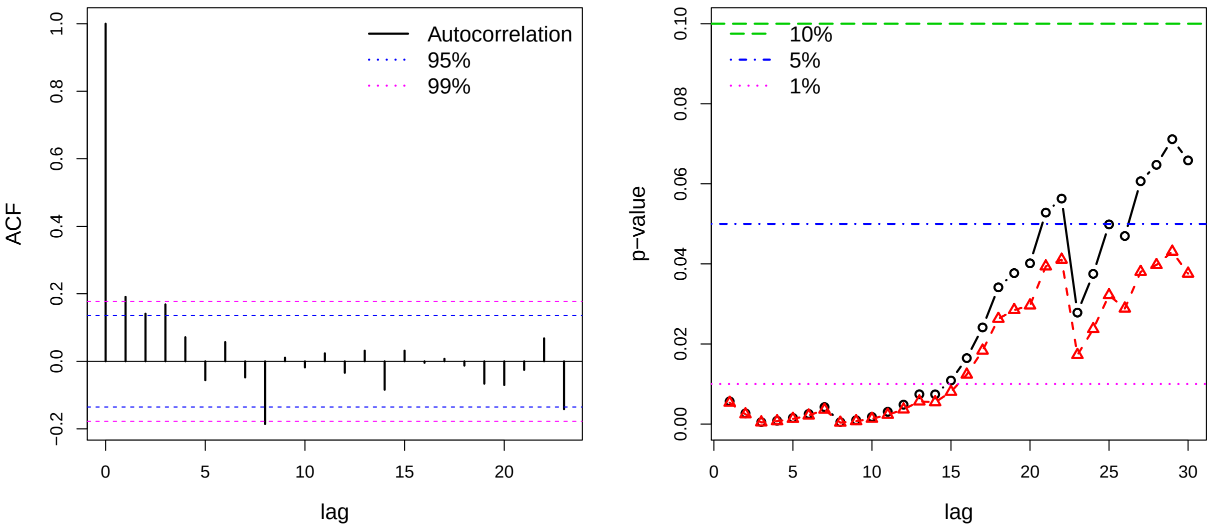

Figure 2: Autocorrelation structure of consumption growth.

Left panel: autocorrelation function of consumption growth ($\Delta c_{t,t+1}$) with 95% and 99% confidence bands. Right panel: $p-$values of Ljung and Box (1978) (triangles) and Box and Pierce (1970) (circles) tests.

Second, we run multivariate linear predictive regressions of cumulated log consumption growth $\Delta c_{t,t+1+S}$, for several values of S, on the first eight principal components of time t asset returns.[10] Figure 3 depicts summary statistics for these predictive regressions at different horizons (S). In particular, the left panel plots the time series adjusted R^2 of these regressions, and the right panel the p-value of the F-test of joint significance of the regressors for this and some additional specifications.

Several observations are in order. At $S = 0$ the time series adjusted R^2 is quite large being about 6.3%. Moreover, the regressors are jointly statistically significant (the p-value of the F-test is less than 1%). For $S > 0$, since $\Delta c_{t,t+1+S} \equiv \Delta c_{t,t+1} + \Delta c_{t+1,t+1+S}$, if asset returns did not predict the autocorrelated component of future consumption growth, the adjusted R^2 should actually decrease monotonically in S, as depicted by the red dashed lines with triangles in the left panel of Figure 3. Instead, for $S > 0$, the figure shows no such decrease in the data (black line with circles) – in fact, predictability increases at intermediate horizons. Moreover, the regressors are jointly statistically significant for any horizon up to 12 quarters following the returns.[11]

of predictability in consumption growth could not be rejected.

[10]We use the first eight principal components of the 25 size and book-to-market Fama-French portfolios, 12 industry portfolios, and 9 bond portfolios, because they explain about 95% of the asset returns variance. Using fewer, or more, principal components, or even directly the asset returns series, we have obtained very similar results to the ones reported.

[11]Liu and Matthies (2018) also document the existence of long-run predictability in consumption using a

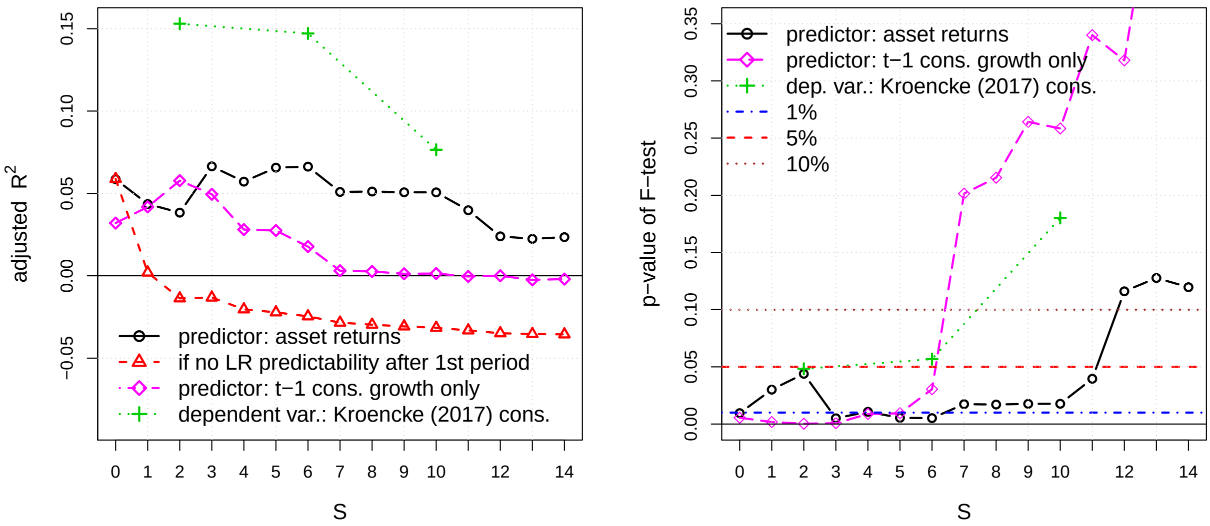

Figure 3: Predictive regressions of $\Delta c_{t,t+1+S}$ on time t asset returns and consumption.

Predictive regressions of $\Delta c_{t,t+1+S}$ on the first eight principal components of asset returns between time $t-1$ and t for different values of S. Left panel: adjusted R^2 (black solid line with circles) and theoretical adjusted R^2 (red dashed line with triangles) if all the predictability was driven by the first period, and purple dashed line with rhombi stands for the adjusted R^2 when only $t-1$ consumption growth is used as a predictor. Dotted line with pluses corresponds to using asset returns to predict the unfiltered consumption growth of Kroencke (2017). Right panel: p-value of the F-test of joint significance of the covariates, as well as the 10% (dotted line), 5% (dashed line), and 1% (dot-dashed line) significance thresholds.

Could one achieve the same level of predictability by using just consumption data, either due to a persistent component (independent of returns) propagating through the actual consumption growth (as e.g. an AR(1)), or through accumulated non-classical measurement errors that display a certain degree of persistence?[12] This is unlikely. The purple lines in Figure 3 depict the degree of predictability obtained using just lagged consumption growth, $\Delta c_{t-1,t}$. While highly significant at the horizon for up to six quarters, using lagged consumption as a predictor is inferior to extracting information from asset returns: not only this variable fails to capture the long range of true predictability, but even at the short horizon it is almost always underperforming stock and bond returns.

Measurement errors in consumption are unlikely to yield such a persistent level of predictability either. While non-classical errors could possibly contribute to a wide range of statistical artifacts, most of their impact should either disappear within a horizon of about one year (should it be related to seasonal smoothing), or be much smaller in magnitude. In order to test this conjecture, we repeat the same predictive exercise with the unfiltered

news-based measure of economic growth prospects.

[12] Seasonal smoothing of consumption levels often leads to countercyclical measurement errors in the growth rates.

consumption data of Kroencke (2017)[13] (green dotted lines in Figure 3). Should the predictability result be an accidental by-product of a countercyclical measurement error due to smoothing, it must go away when using the unfiltered data. If anything, the figure shows, the power of asset returns to forecast consumption becomes even more apparent. Unfortunately, since only yearly data is available for unfiltered consumption, the sample is naturally shorter, which increases standard errors and leads to the feasible use of only three predictive horizons within our time window. However, even taking these limitations into account, asset returns still remain significant predictors of future consumption.

The above results highlight that not only there is substantial predictability in consumption growth, but also it is best captured by asset returns.

Third, the state-space representation of the slow consumption adjustment process presented in the next section postulates the presence of (potentially) long-run shocks in the consumption growth process to which asset returns react only contemporaneously. To verify this conjecture, we recover the long-run impact of common innovations to financial market returns and nondurable consumption using a simple bivariate structural vector autoregression (S-VAR) for a broad market excess return index and consumption growth.[14] We achieve identification via long-run restrictions á la Blanchard and Quah (1989). That is, we distinguish a fundamental long-run shock, that can have a long-run impact on both market return and consumption, and a transitory shock that is restricted not to have a long-run impact on the former.[15].

Figure 4 displays the cumulated impulse response functions to a one standard deviation long-run S-VAR shock, and highlights fundamentally different responses of asset returns and consumption growth to this shock. Asset returns (left panel) fully adjust to the shock contemporaneously, with no statistically significant additional response in the periods following the initial shock. Instead, consumption growth (right panel) reacts gradually to the shock over several quarters, with the cumulated effect reaching its peak only 5-7 quarters after the initial shock. These patterns are consistent with the moving average process we postulate in the next section. Moreover, we later show that the above reaction of consumption to the long-run shocks is extremely similar to the one obtained in our state-space model.

[13] We are grateful to Tim Kroenke for making the data available on his website.

[14] We construct the excess return index as the first principal component of a cross-section that includes excess returns (with respect to the three-month Treasury bill) of the 25 Fama-French size and book-to-market portfolio, 12 industry portfolio, and Treasury securities with maturities of 6 months, 1, 2, 3, 4, 5, 6, 7, and 10 years.

[15] The details of the estimation procedure are presented in Appendix A.2

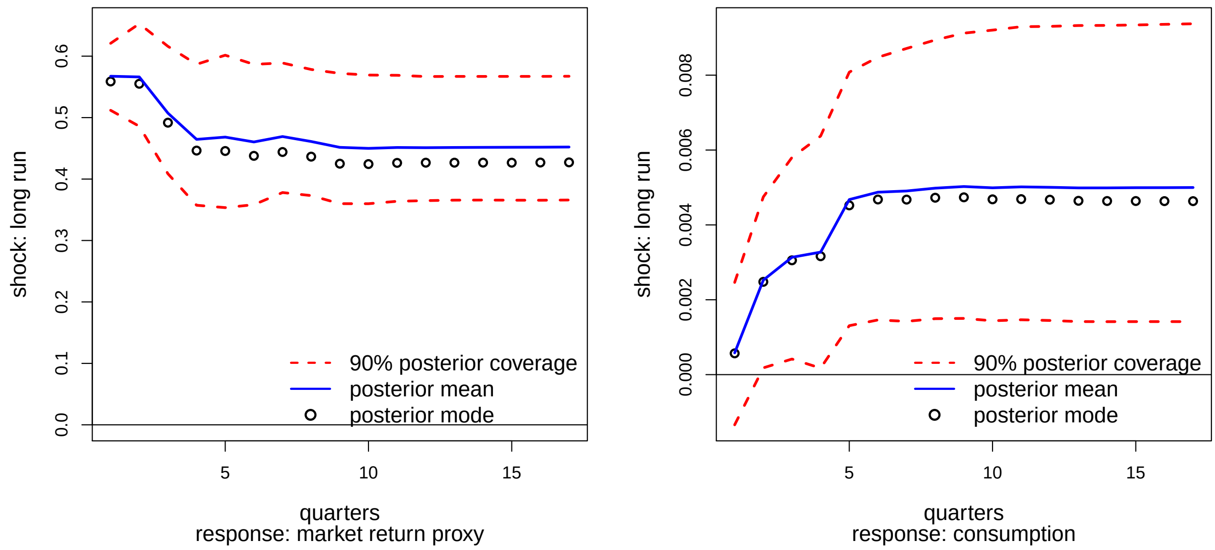

Figure 4: Cumulated response functions to a long-run shock.

Left panel depicts the cumulated response function for the excess return index from the composite cross-section of stock and bond returns, while the right panel plots the cumulated impulse response function of consumption growth. The graphs include posterior mean (continuous blue line), mode (black circles), and centered 90% coverage region (red dashed lines).

IV A Model of Consumption and Returns Dynamics

To model parametrically the reaction of consumption to the same shocks that are spanned by asset returns, we postulate that the consumption growth process can be decomposed in two terms: a white noise disturbance, w_c with variance σ_c^2, that is independent from financial market shocks, plus an autocorrelated process – a slow adjusting component – that (potentially) depends on the current and past stocks to asset returns. To avoid taking an ex ante stand on the particular time series structure of the slow adjustment component (or its absence), we work with its (potentially infinite) moving average representation. That is we model the (log) consumption growth process as:

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + w_t^c, \tag{1}$$

where $\bar{S}$ is a large positive integer (potentially equal to $+\infty$), μ_c is the unconditional mean, the ρ_j coefficients are square summable, and most importantly f_t (a white noise process normalised to have unit variance) is the fundamental innovation upon which all asset returns load contemporaneously, i.e.

$$\underset{N\times 1}{\mathbf{r}_t^e} = \underset{N\times 1}{\boldsymbol{\mu}_r} + \underset{N\times 1}{\boldsymbol{\rho}^r} f_t + \underset{N\times 1}{\mathbf{w}_t^r}, \tag{2}$$

where $\mathbf{r}^e$ denotes a vector of log excess returns, $\boldsymbol{\mu}_\tau$ is a vector of expected values, $\boldsymbol{\rho}^r$ contains the asset specific loadings on the common risk factor, $\mathbf{w}_t^r$ is a vector of white noise shocks with diagonal[16] covariance matrix Σ_r, that are meant to capture asset specific idiosyncratic shocks.

The dynamic system in equations (1)-(2) can be reformulated as a state-space model, and Bayesian posterior inference can be conducted to estimate both the unknown parameters (μ_c, $\boldsymbol{\mu}_r$, $\{\rho_j\}_{j=0}^{\bar{S}}$, $\boldsymbol{\rho}^r$, σ_c^2, Σ_r) and the time series of the unobservable common factor of consumption and asset returns ($\{f_t\}_{t=1}^T$). This estimation procedure is described below and additional details are presented in Appendix A.3.

A crucial point that allows us to achieve identification of the shocks, is the lead-lag structure of the consumption process and its (potential) link to asset returns. Without equation (1), the shocks would be underidentified, making it difficult to give any particular rotation a structural interpretation. Another interpretation of this estimation approach is that of uncovering the shocks that drive financial returns through the impulse response function on consumption, in the spirit of Uhlig (2005) identification in Structural-VARs. In fact, it is easy to see that the ρ_j coefficients identify the impulse response function of multiperiod consumption growth to the shock f_t as[17]

$$\frac{\partial \mathbb{E}\left[\Delta c_{t-1,t+S}\right]}{\partial f_t} = \sum_{j=0}^{S} \rho_j \tag{3}$$

where $\rho_{j>\bar{S}} := 0$.

Note that equations (1)-(2) rely *only* on the time series properties of asset returns and consumption. Therefore, nothing in the formulation of the joint system requires the shocks to be priced in the cross-section of returns, or equivalently expected asset returns to align with their exposure to consumption (albeit this is what we would expect in a consumption-based asset pricing model). However, this becomes a testable implication, since the covariance between asset returns and consumption growth over one or several periods is fully characterised by the loadings of the dynamic system on the factor f_t:

$$Cov\left(\Delta c_{t-1,t+S}; \mathbf{r}_t^e\right) \equiv \sum_{j=0}^{S} \rho_j \boldsymbol{\rho}^r. \tag{4}$$

This implies that the time series estimates of the latent factor loadings ($\hat{\rho}_j$ and $\hat{\boldsymbol{\rho}}^r$) can be

[16] The diagonality assumption can be relaxed as explained in Appendix A.3

[17] This immediately follows from the observation that, since $\Delta c_{t-1,t+S} \equiv \sum_{j=0}^{S} \Delta c_{t-1+j,t+j} \equiv \ln\left(C_{t+S}/C_{t-1}\right)$, we have $\left[\Delta c_{t-1,t}, \Delta c_{t-1,t+1}, ..., \Delta c_{t-1,t+S}\right]' \equiv \Gamma\left[\Delta c_{t-1,t}, \Delta c_{t,t+1}, ..., \Delta c_{t-1+S,t+S}\right]'$ where Γ is a lower triangular square matrix of ones (of dimension S).

used to assess whether the slow consumption adjustment component has explanatory power for the cross-section of risk premia (via, for instance, simple cross-sectional regressions of returns on these estimated covariances).

The particular one factor *term structure* in the exposures of asset returns to consumption growth in equation (4) also provides an alternative method, that does not rely on the parametric likelihood and directly stems from the empirically measured covariances, for recovering moving average component in consumption. We discuss this further in Section V.2.

Finally, note that the formulation in equations (1)-(2) can be generalized to allow for a bond-specific latent factor (g_t), to which consumption potentially reacts slowly over time. The dynamic system in this case becomes:

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + \sum_{j=0}^{\bar{S}} \theta_j g_{t-j} + w_t^c; \tag{5}$$

$$\underset{N\times 1}{\mathbf{r}_t^e} = \underset{N\times 1}{\boldsymbol{\mu}_r} + \underset{N\times 1}{\boldsymbol{\rho}^r} f_t + \left[\begin{array}{cc} \underset{N_b\times 1}{\boldsymbol{\theta}'^b}, & \mathbf{0}'_{N-N_b} \end{array} \right]' g_t + \underset{N\times 1}{\mathbf{w}_t^r}; \tag{6}$$

where N_b is the number of bonds and they are ordered first in the vector $\mathbf{r}_t^e$, $\boldsymbol{\theta}^b \in \mathbb{R}^{N_b}$ contains the bond loadings on the factor g_t – a white noise process with variance normalized to one. In this case the implied covariance of consumption and returns becomes:

$$Cov\left(\Delta c_{t-1,t+S}; \mathbf{r}_t^e\right) \equiv \sum_{j=0}^{S} \rho_j \boldsymbol{\rho}^r + \left[\begin{array}{cc} \boldsymbol{\theta}'^b, & \mathbf{0}'_{N-N_b} \end{array} \right]' \sum_{j=0}^{S} \theta_j. \tag{7}$$

Two observations regarding our parametric framework deserve mentioning. First, both the one-factor (equations (1)-(2)) and two-factor (equations (5)-(6)) models are strongly overidentified. Second, estimation of the model would generally remain consistent even in the presence of time varying volatility in the true processes (with the only exception of a *contemporaneous* 'leverage effect'), hence our formulation is robust along this dimension. We address this issue formally in Section V.1.2.

IV.1 Estimation

We can rewrite the dynamic model in equations (1)-(2) in state-space form, assuming Gaussian innovations, as

$$\mathbf{z}_t = \mathbf{F}\mathbf{z}_{t-1} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}\left(\mathbf{0}_{\bar{S}+1}; \Psi\right); \tag{8}$$

$$\mathbf{y}_t = \boldsymbol{\mu} + \mathbf{H}\mathbf{z}_t + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}\left(\mathbf{0}_{N+1}; \Sigma\right). \tag{9}$$

where $\mathbf{y}_t := [\Delta c_t, \mathbf{r}_t^{e\prime}]$, $\mathbf{z}_t := [f_t, ..., f_{t-\bar{S}}]'$, $\boldsymbol{\mu} := [\mu_c, \boldsymbol{\mu}_r']'$, $\mathbf{v}_t := \left[f_t, \mathbf{0}_{\bar{S}}'\right]'$, $\mathbf{w}_t := [w_t^c, \mathbf{w}_t'^r]'$,

$$\Psi := \underbrace{\begin{bmatrix} 1 & \mathbf{0}_{\bar{S}}' \\ \mathbf{0}_{\bar{S}} & \mathbf{0}_{\bar{S}\times\bar{S}} \end{bmatrix}}_{(\bar{S}+1)\times(\bar{S}+1)}, \quad \mathbf{F} := \underbrace{\begin{bmatrix} \mathbf{0}_{\bar{S}}' & 0 \\ I_{\bar{S}} & \mathbf{0}_{\bar{S}} \end{bmatrix}}_{(\bar{S}+1)\times(\bar{S}+1)}, \quad \Sigma := \underbrace{\begin{bmatrix} \sigma_c^2 & \mathbf{0}_N' \\ \mathbf{0}_N & \Sigma_r \end{bmatrix}}_{(N+1)\times(N+1)}, \quad \mathbf{H} := \underbrace{\begin{bmatrix} \rho_0 & \rho_1 & \cdots & \rho_{\bar{S}} \\ \boldsymbol{\rho}^r & \mathbf{0}_N & \ldots & \mathbf{0}_N \end{bmatrix}}_{(N+1)\times(\bar{S}+1)}, \tag{10}$$

and $I_{\bar{S}}$ and $\mathbf{0}_{\bar{S}\times\bar{S}}$ denote, respectively, an identity matrix and a matrix of zeros of dimension $\bar{S}$.

Similarly, the dynamic system in equations (5)-(6) can be represented in the state-space form (8)-(9) with: $\mathbf{z}_t := [f_t, ..., f_{t-\bar{S}}, g_t, ..., g_{t-\bar{S}}]'$; $\mathbf{v}_t := \left[f_t, \mathbf{0}_{\bar{S}}', g_t, \mathbf{0}_{\bar{S}}',\right]' \sim \mathcal{N}\left(\mathbf{0}_{\bar{S}+1}; \Psi\right)$; Ψ and $\mathbf{F}$ being block diagonal with blocks repeated twice and given, respectively, by the first two matrices in equation (10); and with space equation matrix of coefficients given by

$$\mathbf{H} := \underbrace{\begin{bmatrix} \rho_0 & \ldots & \ldots & \rho_{\bar{S}} & \theta_0 & \ldots & \ldots & \theta_{\bar{S}} \\ \rho_1^r & 0 & \ldots & 0 & \theta_1^b & 0 & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\ \rho_{N_b}^r & 0 & \ldots & 0 & \theta_{N_b}^b & 0 & \ldots & 0 \\ \ldots & 0 & \ldots & 0 & 0 & 0 & \ldots & 0 \\ \rho_N^r & 0 & \ldots & 0 & \ldots & 0 & \ldots & 0 \end{bmatrix}}_{(N+1)\times 2(\bar{S}+1)}. \tag{11}$$

The above state-space system implies the following conditional likelihood for the data:

$$\mathbf{y}_t | \mathcal{I}_{t-1}, \boldsymbol{\mu}, \mathbf{H}, \Psi, \Sigma, \mathbf{z}_t \sim \mathcal{N}\left(\boldsymbol{\mu} + \mathbf{H}\mathbf{z}_t; \Sigma\right) \tag{12}$$

where $\mathcal{I}_{t-1}$ denotes the history of the state and space variables until time $t-1$. Hence, under a diffuse (Jeffreys') prior and conditional on the history of $\mathbf{z}_t$ and $\mathbf{y}_t$, and given the diagonal structure of Σ, we have the standard Normal-inverse-Gamma posterior distribution for the parameters of the model (see e.g. Bauwens, Lubrano, and Richard (1999)). Moreover, the posterior distribution of the unobservable factors $\mathbf{z}_t$ conditional on the data and the parameters, can be constructed using a standard Kalman filter and smoother approach (see, e.g., Primiceri (2005)).

When combined with a log linearized consumption Euler equation for a very broad class of asset pricing models,[18] the above specification for the dynamics of consumption and asset returns implies, in the presence of only one latent factor (f_t) common to both assets and

[18] See equation (20) and discussion therein.

consumption

$$\mathbb{E}\left[\mathbf{R}_t^e\right] = \alpha + \left(\sum_{j=0}^{S} \rho_j \boldsymbol{\rho}^r\right) \lambda_f \tag{13}$$

where $\mathbf{R}_t^e \in \mathbb{R}^N$ denotes the vector of returns in excess of the risk-free rate, λ_f is a positive scalar variable that captures the price of risk associated with the exposure to the slow consumption adjustment risk, and $\alpha \in \mathbb{R}^N$ is the vector of average mispricings. If consumption fully captures the risk of asset returns, the above expression should hold with $\alpha = \mathbf{0}_N$, otherwise α should capture the covariance between the omitted risk factors and asset returns.

Similarly, if we allow for a bond specific latent factor (g_t), as in equations (5)-(6), the implied cross-sectional model of returns is

$$\mathbb{E}\left[\mathbf{R}_t^e\right] = \alpha + \left(\sum_{j=0}^{S} \rho_j \boldsymbol{\rho}^r\right) \lambda_f + \left[\begin{array}{cc} \boldsymbol{\theta}'^b, & \mathbf{0}'_{N-N_b} \end{array}\right]' \sum_{j=0}^{S} \theta_j \lambda_g \tag{14}$$

with the additional testable restriction $\lambda_f = \lambda_g$.

Equation (13) (and similarly equation (14)), conditional on the data and the parameters of the state-space model, defines a conventional cross-sectional regression, hence the parameters α, λ_f and λ_g can be estimated via the standard Fama and MacBeth (1973) procedure. Therefore, not only can we compute posterior means and confidence bands for both the coefficients of the state space model and for the unobservable factor's time series, but we can also compute means and confidence bands for the Fama and MacBeth (1973) estimates. That is, we can jointly assess the ability of the slow consumption adjustment risk of explaining both the time series and the cross-section of asset returns with a simple Gibbs sampling approach described in detail in Appendix A.3.

Furthermore, we confirm the main empirical results obtained with the above state-space framework, for both the impulse response function of consumption and the cross-sectional pricing patterns, using a principal component based approach and an Empirical Likelihood type estimation. The details of these nonparametric procedures are discussed in Sections V.2 and V.3 respectively.

V Empirical Evidence

Our empirical analysis is based on both parametric and nonparametric inference, and both Bayesian and frequentist inference, therefore ensuring that the results presented are robust to the methodology employed. The main approach, in Section V.1, combines the state-space system implied by equations (1)-(2) (as well as equations (5)-(6)) with standard Bayesian filtering techniques to recover the unobservable latent consumption factor (f_t) and the other model parameters. We analyze loading patterns, variance decompositions, and impulse response functions. Furthermore, since our framework parametrises the time series dynamics without imposing cross-sectional asset pricing restrictions *a priori*, we separately test the latter using Fama and MacBeth (1973) cross-sectional regressions.

In Section V.2 we examine the term structure implications of asset exposure to consumption growth at different horizons to recover the moving average parameters of the consumption process. Finally, in Section V.3 we rely on a standard nonparametric technique, Empirical Likelihood,[19] to document that a) the slow-moving component of consumption growth is priced in the cross-section of bond returns, b) this slow-moving component provides substantially better identification of the assets' exposure to consumption growth, revealing precise estimates of both the level and the spread of the co-movement, and c) these patterns are uniform across different cross-sections of the assets.

V.1 Evidence from the state-space model

While our model in equations (1)-(2) allows for a potentially infinite number of lags for the consumption process, in order to proceed with the actual estimation one has to choose a particular value of $\bar{S}$. For the rest of the section we use $\bar{S} = 14$ for a number of reasons.

First, equation (1) implies a certain autocorrelation structure of the nondurable consumption growth through the combination of the common factor lags and its loadings. Hence, the value of $\bar{S}$ should be high enough to capture most of the time series autocorrelation in the consumption growth. Figures 2-4 of the previous sections suggest that most of the dependence is captured by the first 14 lags, hence this value is a natural choice for the lag truncation. Second, intuitively most of the pricing impact from the consumption adjustment is probably taking place within the business cycle frequency, consistent with a number of recent empirical studies (e.g. Bandi and Tamoni (2015)). Therefore, $\bar{S} = 14$ quarters is a rather conservative choice, since it provides a 3.5 year window to capture most of the interaction between the consumption and returns. Finally, our results remain robust to including additional lags.[20]

V.1.1 The consumption mean process

The consumption growth representation in equations (1) and (5) is similar to the moving average decomposition, and allows us to model the dynamics of multi-period consumption

[19] For an excellent review and comparison with other moment-based estimation approaches, see Kitamura (2006).

[20] For robustness, we have experimented with up to $\bar{S} = 50$, and the results, available upon request, remained virtually identical.

growth ($\Delta c_{t,t+1+S}$) in response to a common and/or a bond-specific shock. Figure 5 depicts the (cumulated) loadings of consumption on the latent factor f as a function of the horizon S. At $S = 0$, the case of a standard consumption-based asset pricing model, the moving average component of consumption virtually does not load on the common factor. Instead, as S increases, the impact of the common factor becomes more and more pronounced, levelling off at around $S = 11$. At this horizon, the effect is economically very large: the cumulative response of consumption growth to a one standard deviation shock is about 1%.

Interestingly, the pattern of estimated loadings is similar (at least at the horizons we consider) to the one implied by the moving average representation of the consumption process in Bansal and Yaron (2004),[21] and reveals a comparable degree of persistency at the 3-4 years frequency. Note that allowing for a bond specific latent factor (as in equations (5)-(6)) leaves the consumption loadings on the f_t shocks virtually unchanged and consumption does not load significantly on the bond specific factor g_t (see, respectively, figures A1 and A2 in Appendix A.6).

Figure 6 shows the (posterior mean of the) MA component of consumption, based on our filtered f_t innovations. The slow moving component within consumption aptly captures the business cycle fluctuations and has a pronounced exposure to recession risk. Furthermore, this component generates economically large swings in *quarterly* consumption growth, with contractions and expansions of about 1% being not uncommon.

But how much of the total consumption volatility can this slow moving component explain? More than a quarter – which is very large (and sharply estimated) compared to the leading asset pricing frameworks: e.g. this is 5-6 times larger than in long-run risk framework of Bansal and Yaron (2004), while this quantity should actually be equal to zero in the the rare disasters models (e.g. Barro (2006)) and in the habit setting of Campbell and Cochrane (1999). Figure 7 demonstrates that the common factor is responsible for roughly 26% of the variation in the one-period nondurable consumption growth, 33% of the two-period consumption growth, and so on, followed by a slow decline towards just above 5% for the 15-period growth. Interestingly, the model retains significant predictive power (albeit, much lower) even for the one-period consumption that will occur 3-4 year from now. As shown in Figure A2 of the Appendix, adding a bond specific factor has a minimum impact on the explanatory power of the model for future consumption growth.

[21] For more details on the construction of the MA representation for the Bansal and Yaron (2004) framework see Appendix A.4.

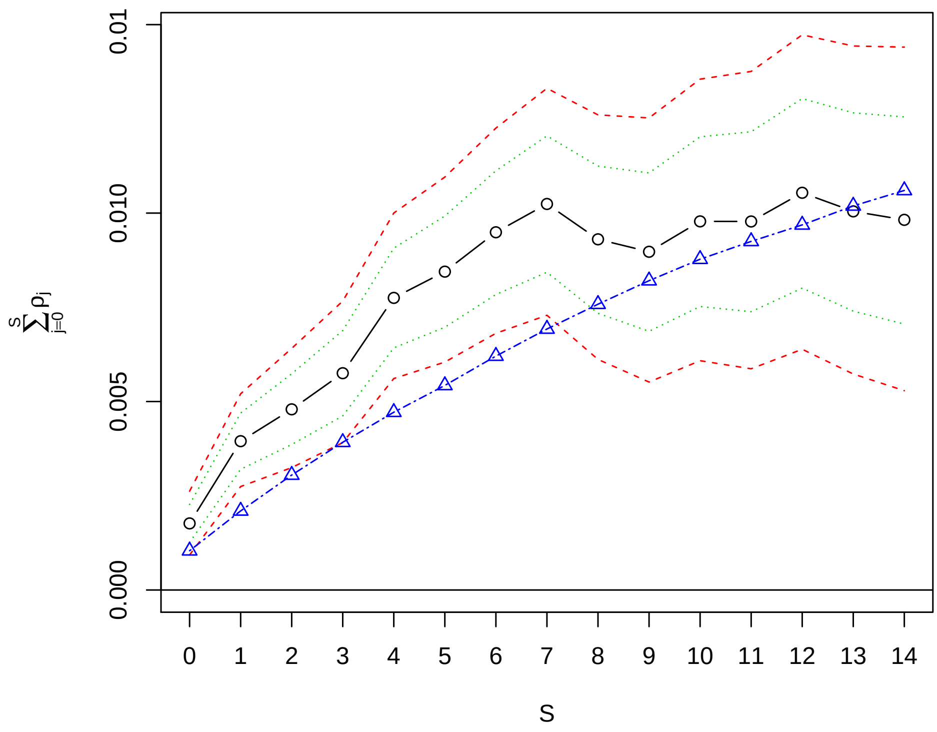

Figure 5: Cumulated consumption adjustment response to the common factor (f_t) shock.

Posterior means (black continuous line with circles) and centred posterior 90% (red dashed line) and 68% (green dotted line) coverage regions. Estimation based on the one-factor model in equations (1) and (2). Blue triangles denote Bansal and Yaron (2004) implied values.

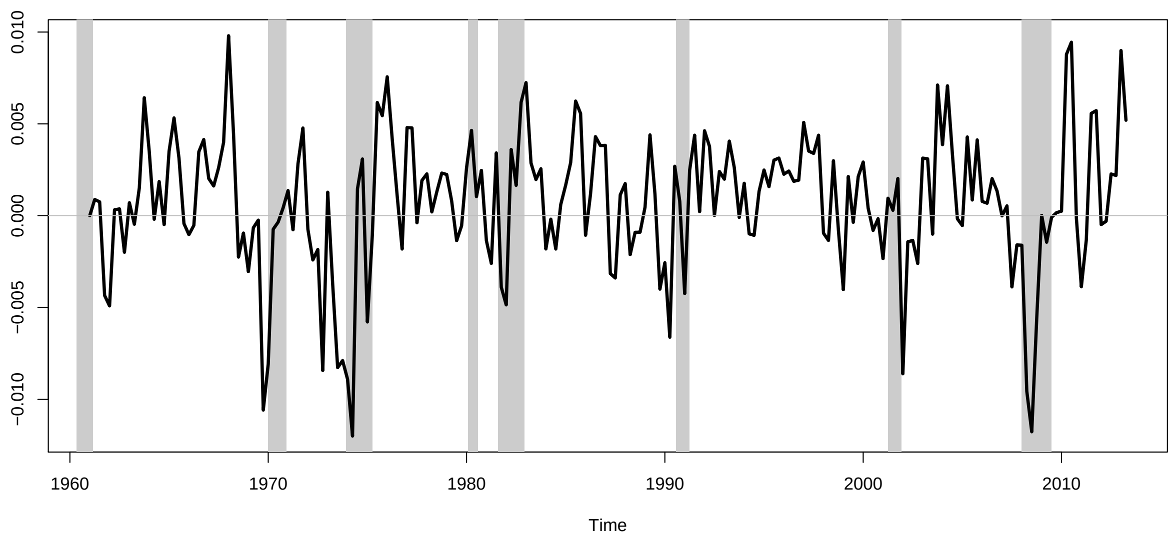

Figure 6: Moving average of f_t component of consumption growth

Posterior mean of the moving average f_t component of consumption growth. Grey areas denote NBER recessions.

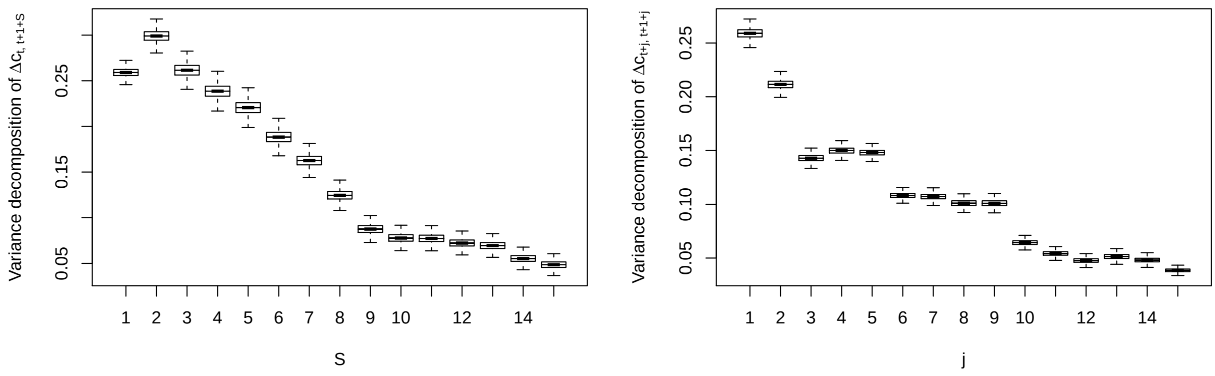

Figure 7: Share of consumption growth variance driven by its conditional mean (the MA component $\sum_{j=0}^{\bar{S}} \rho_j f_{t-j}$).

Box-plots (posterior 95% coverage area) of the percentage of time series variances of consumption growth explained by the MA component. Left panel: cumulated consumption growth $\Delta c_{t,t+S}$. Right panel: one period consumption growth $\Delta c_{t-1+j,t+j}$.

V.1.2 Clustering and predictability of consumption volatility

Note that the estimation of the conditional mean of consumption growth in equations (1) and (5) is generally consistent even in the presence of time varying consumption volatility. Therefore, the presence of volatility clustering can be assessed by analyzing the serial correlation of the squared one-step ahead forecast errors (see, e.g., Engle (1982)) of the consumption growth process. That is, by examining the autocorrelation and predictability of $\widehat{Var}_t(\Delta c_{t,t+1}) := \left(\Delta c_{t,t+1} - \widehat{\mathbb{E}}_t[\Delta c_{t,t+1}]\right)^2$, where the conditional mean is computed at each t using the estimated ρ_j and θ_j coefficients and latent state variables $f_{\tau \le t}$ and $g_{\tau \le t}$.[22]

Figure 8 reports the autocorrelation function (left panel), as well as the p-values of the Ljung and Box (1978) and Box and Pierce (1970) tests (right panel) of joint significance of the autocorrelations of $\widehat{Var}_t(\Delta c_{t,t+1})$ and shows no evidence of volatility clustering in the consumption growth process. Nevertheless, conditional consumption volatility might still, in principle, be correlated with financial asset returns. We test this hypothesis by running linear predictive regressions of $\widehat{Var}_{t+h}(\Delta c_{t+h,t+h+1})$, at several horizons h, on the time t first eight principal components of stock and bond returns. Note that this is the same test used to establish predictability of the first conditional moment of consumption growth in Section III. The p-values of the F-tests for these predictability regressions are depicted by the dashed red line with triangles in Figure 9. The p-values (that range from 0.2826 to 0.922) show that asset returns are not significant predictors of future consumption volatility. For this feature of the data to be revealed, it is key to properly account for the conditional mean of the

[22] That is, $\widehat{\mathbb{E}}_t[\Delta c_{t,t+1}]$ is the posterior mean of $\mu_c + \sum_{j=1}^{\bar{S}} \rho_j f_{t+1-j} + \sum_{j=1}^{\bar{S}} \theta_j g_{t+1-j}$ at each t.

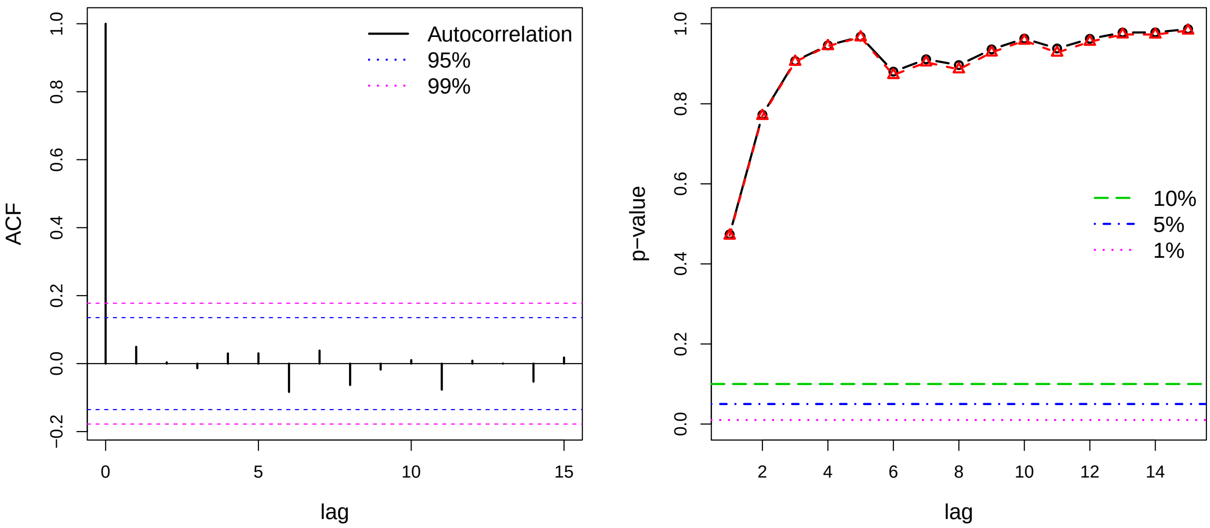

Figure 8: Autocorrelation structure of consumption growth squared forecast errors.

Left panel: autocorrelation function of $\widehat{Var}_t\left(\Delta c_{t,t+1}\right)$ with 95% and 99% confidence bands. Right panel: p-values of Ljung and Box (1978) (red triangles) and Box and Pierce (1970) (black circles) tests.

consumption process.

Indeed, given our finding of a common latent factor driving both asset returns and consumption, if one were to erroneously model the conditional mean of consumption growth, one would be likely to find spurious evidence of volatility clustering. For instance, erroneously modeling the conditional mean of consumption as being constant, the autocorrelation of $\widehat{Var}_t\left(\Delta c_{t,t+1}\right)$ would be mechanically different from zero. For instance, the k-th autocorrelation of $\left(\Delta c_{t,t+1} - \mu_c\right)^2$, for $k \leq \bar{S}$, is proportional to

$$Cov\left(\left(\sum_{j=k}^{\bar{S}} \rho_j f_{t-j} + \theta_j g_{t-j}\right)^2, \left(\sum_{j=k}^{\bar{S}} \rho_{j-k} f_{t-j} + \theta_{j-k} g_{t-j}\right)^2\right) \neq 0. \tag{15}$$

To verify that a misspecification of the consumption mean process leads to spurious evidence of time varying volatility we perform two exercises.

First, in Table 1 we estimate a GARCH(1,1) for consumption volatility under different assumptions about the mean process. Panel A considers the common AR(1) specification for consumption growth. In this this case there is statistically significant evidence of volatility clustering (with a half-life of about 5-6 quarters). However, from Section III we know that lagged consumption alone does not capture the full extent of consumption predictability. Therefore, in Panel B we add to the AR(1), as drivers of the conditional mean, the same principal components of asset returns that predict consumption in Section III. The resulting change is striking: once we better control for consumption predictability, the evidence in

Table 1: Estimates of GARCH (1,1) for the innovations of different models

$$\Delta c_{t+1} = \mu_t + \epsilon_{t+1}$$
$$\sigma_{t+1}^2 = \omega + \alpha \epsilon_t^2 + \beta \sigma_t^2$$

	ω	α	β
Panel A: $\mu_t = \mu_0 + \mu_1^c \Delta c_t$			
Estimate	7.279×10^{-6}	0.141	0.719
t-stat	[1.912]	[1.664]	[8.928]
Panel B: $\mu_t = \mu_0 + \mu_1^c \Delta c_t + \sum_{i=1}^{8} \mu_i^r r_{PC_i,t}^{ex}$			
Estimate	4.019×10^{-5}	0.136	6.845×10^{-3}
t-stat	[2.424]	[1.297]	[0.022]
Panel C: $\mu_t = \mu_0 + \sum_{i=1}^{\bar{S}} \rho_i f_{t+1-i}$			
Estimate	3.502×10^{-5}	7.678×10^{-2}	3.421×10^{-2}
t-stat	[1.363]	[0.957]	[0.053]

The table presents GARCH estimates for consumption growth volatility with different models for of the conditional mean. The model is estimated using QMLE and t-statistics are constructed using Newey and West (1987) standard errors.

favor of time varying volatility vanishes. Finally, in Panel C we use the conditional mean of our moving average specification (without the contemporaneous shock) evaluated at its posterior mean. Two observations are in order. First, as in Panel B, there is neither statistical nor economic evidence of volatility clustering (the implied half-life of the volatility shocks is about one quarter i.e. the same as for an *i.i.d.* process). Second, the sharply different results in Panels A and C suggest that the AR(1) approximation of the conditional mean is not innocuous for the identification of the volatility process. Moreover, if the AR(1) were the true process, our MA($\bar{S}$) specification should closely approximate it, and lead to very similar implications for volatility – something clearly contradicted by the table.

Second, we run predictability regressions of $\widehat{Var}_t(\Delta c_{t,t+1})$ on the first eight principal components of asset returns (the same predictive variables as in Section III). We construct the consumption volatility proxy both under the null of our state-space framework and, as a misspecification benchmark, under constant conditional mean for consumption growth. Summary statistics for these regressions are depicted by the black continuous line with circles in Figure 9. The figure shows that the misspecification of the mean process generates spurious predictability of consumption volatility in 50% of the horizons considered.

That is, modeling the mean of consumption growth without exploiting the information in asset returns and the flexibility of the MA representation, leads to spurious evidence of time-varying volatility of consumption growth.

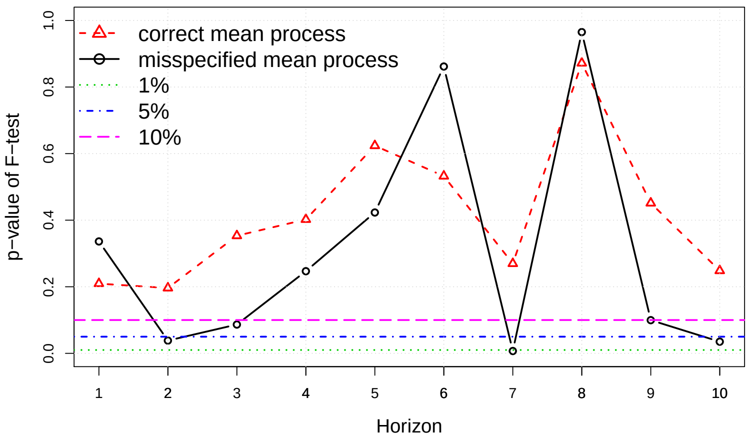

Figure 9: Predictability of consumption squared forecast errors.

Predictive regressions of $\widehat{Var}_{t+h}\left(\Delta c_{t+h,t+h+1}\right)$ on the time t first eight principal components of asset returns at several horizons h. p-value of the F-test of joint significance of the covariates as well as the 10%, 5%, and 1% significance thresholds (respectively, horizontal dotted, dashed and dot-dashed lines). The dashed red line with triangles denotes statistics for the correctly specified conditional mean for the consumption growth process, while the black continuous line with circles corresponds to the assumption of a constant conditional mean.

V.1.3 Time series properties of stocks and bonds

We now turn to the time series properties of stocks and bonds implied by our model in equations (1)–(2) (and (5)–(6)). The loadings of equity portfolios on the latent factor f_t are depicted in Figure 10.

The size and book-to-market sorted portfolios are ordered first (e.g. portfolio 2 is the smallest decile of size and the second smallest decile of book-to-market ratio), followed by the 12 industry portfolios (portfolios 26-39 in the graph). All the portfolios have significant and positive exposures to the common factor. Furthermore, there is an easily recognizable pattern in the factor loadings, in line with the size and value anomalies. This provides some preliminary evidence that the f_t shocks play an important role in explaining the cross-sectional dispersion of stock returns. The findings remain unchanged when a bond-specific factor is added to the model as in equations (5)–(6) (see Figure A3 in the Appendix).

These loadings are not only statistically, but also economically significant as shown in Figure 11: the common factor f explains on average 79% of the time series variation of stock returns, ranging from 36% to nearly 95% for individual portfolios. Moreover, this explanatory power in our model is produced by a single consumption-based factor, as opposed to some of the alternative successful specifications that typically rely on 3 or more explanatory variables. As shown in Figure A4 in the Appendix, adding a bond specific factor leaves the variance decomposition of stock returns virtually unaffected.

The loadings of the bond portfolios on the common consumption factor f_t are reported in

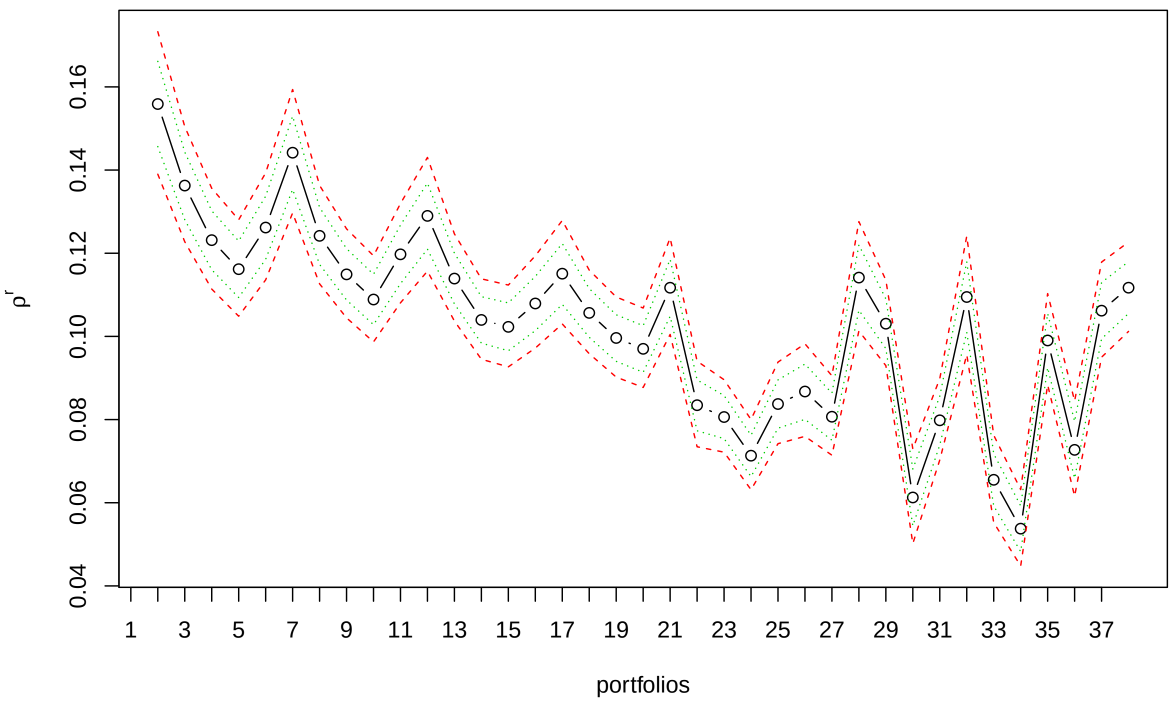

Figure 10: Common factor loadings ($\boldsymbol{\rho}^r$) of the stock portfolios in the one-factor model.

Posterior means of the stocks factor loadings on f_t (black continuous line with circles) and centred posterior 90% (red dashed line) and 68% (green dotted line) coverage regions in the one latent factor model. Ordering of portfolios: 25 Fama and French (1992) size and book-to-market sorted portfolios and 12 industry portfolios.

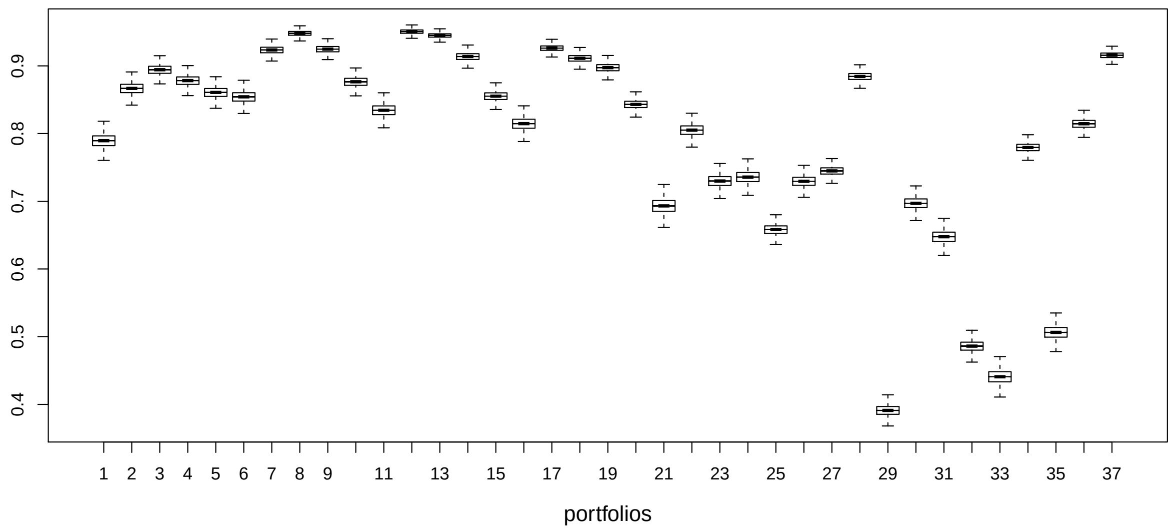

Figure 11: Share of stock portfolios' return variance explained by the f component.

Box-plots (posterior 95% coverage area) of the percentage of time series variances of individual stock portfolio returns explained by the f component in the one factor model. Ordering of portfolios: 25 Fama and French (1992) size and book-to-market sorted portfolios and 12 industry portfolios.

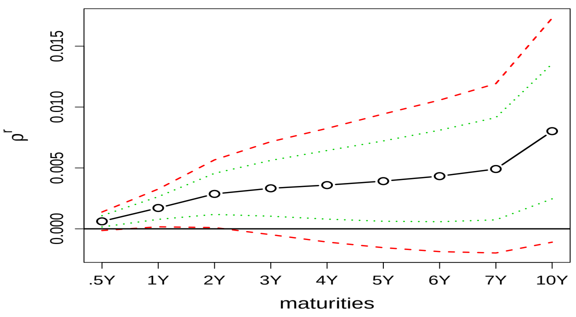

(a) Bond loadings on f_t, one-factor model.

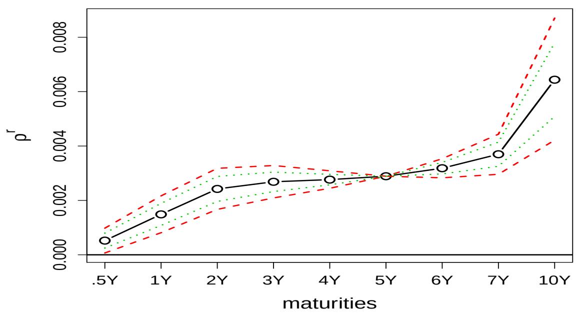

(b) Bond loadings on f_t, two-factor model.

Figure 12: Bond loadings on common factor f_t.

Posterior means of the bonds factor loadings on f_t (black continuous line with circles) and centred posterior 90% (red dashed line) and 68% (green dotted line) coverage regions in the one latent factor model.

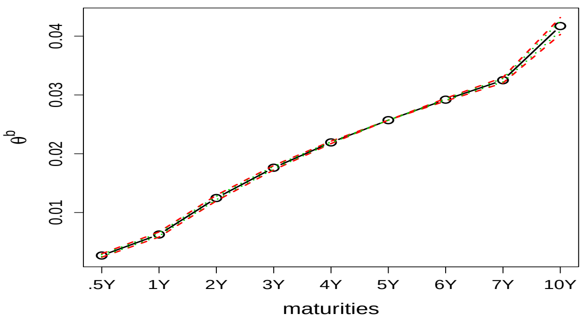

Figure 13: Bond loadings ($\boldsymbol{\theta}^b$) on the bond-specific factor (g_t).

Posterior means (black continuous line with circles) and centred posterior 90% (red dashed line) and 68% (green dotted line) coverage regions, of bond loadings on the bond specific factor g_t.

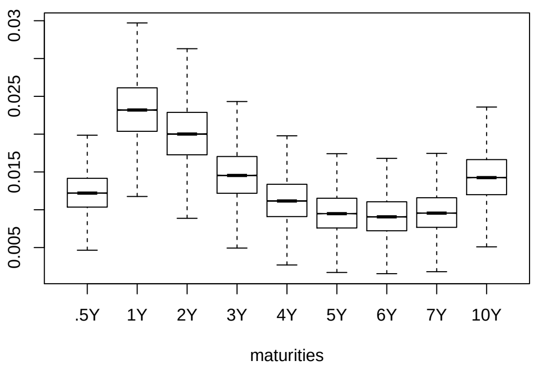

(a) Percentage of time series variances of bond returns explained by f_t, one-factor specification.

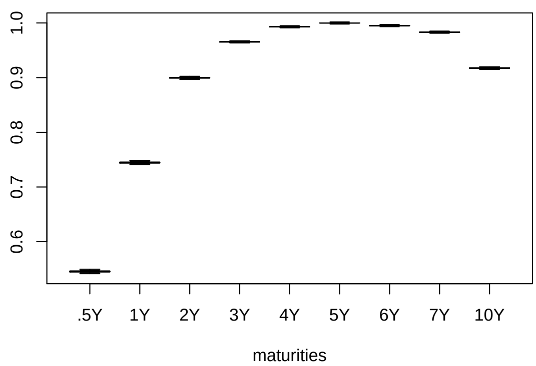

(b) Percentage of time series variances of bond returns jointly explained by f_t and g_t components.

Figure 14: Variance decomposition box-plots of bond returns.

Box-plots (posterior 95% coverage area) of the percentage of time series variances of bond returns explained by the f_t (left panel), and f_t and g_t (right panel) shocks.

Figures 12a and 12b for, respectively, the one and two latent factor specifications. Both sets of estimates show an upward sloping term structure of the loadings, and the point estimates are very similar in the two specifications, with the main difference being that, allowing for a bond specific factor (g_t) delivers much sharper estimates of the loadings on the common factor f_t. The magnitude of these loadings is considerably smaller than that of stocks – a feature that, as shown later, will allow us to price jointly the cross-sections of stocks and bonds. While these numbers may not seem as impressive as those for the cross-section of stocks, the pattern is highly persistent and significant, confirming a common factor structure between nondurable consumption growth and asset returns.

The loadings on the bond specific factor g_t are reported in Figure 13. These loadings are highly statistically significant, and increase steeply and monotonically with maturity, revealing a very pronounced term structure pattern.

Finally, Figure 14 reports the share of time series variation of bond returns explained by the f_t shocks (left panel), and the f_t and g_t shocks (right panel), and highlights the importance of allowing for a bond specific factor to characterize the time series of bond returns. The common factor f_t accounts for a small (about 1.5%), but statistically significant, proportion of the time series variation in bond returns. The bond-specific factor, in turn, manages to capture most of the residual time series variation in returns. While the model captures just about 55% of the variation in the 6-month bond returns, its performance rapidly improves with maturity and results in a nearly perfect fit for the time horizon of about five years.

V.1.4 What drives the consumption mean process?

Figure 6 revealed a clear business cycle fluctuation within the moving average component of consumption growth, but what exactly is being captured by those shocks? Figure 15 looks at the (posterior) correlation of the consumption mean process (i.e. the MA of f_t) with traditional asset pricing factors and principal components of bond returns. Interestingly, consumption growth is characterized by a complicated mix of exposures to several popular proxies for risk.

Focusing on equity risk factors, the conditional mean of consumption shows a strong correlation with both the HML factor (a proxy for the value premium) and the overall market excess return, while the the correlation with SMB (a proxy for the size premium), is small (albeit significant), while there is virtually no comovement with the momentum factor. The negative comovement with the market is consistent with a reduction of the overall risk premium in good consumption states (as e.g. in habit, and many other, macro-finance models). The positive correlation with the value factor is both in line with the finding of

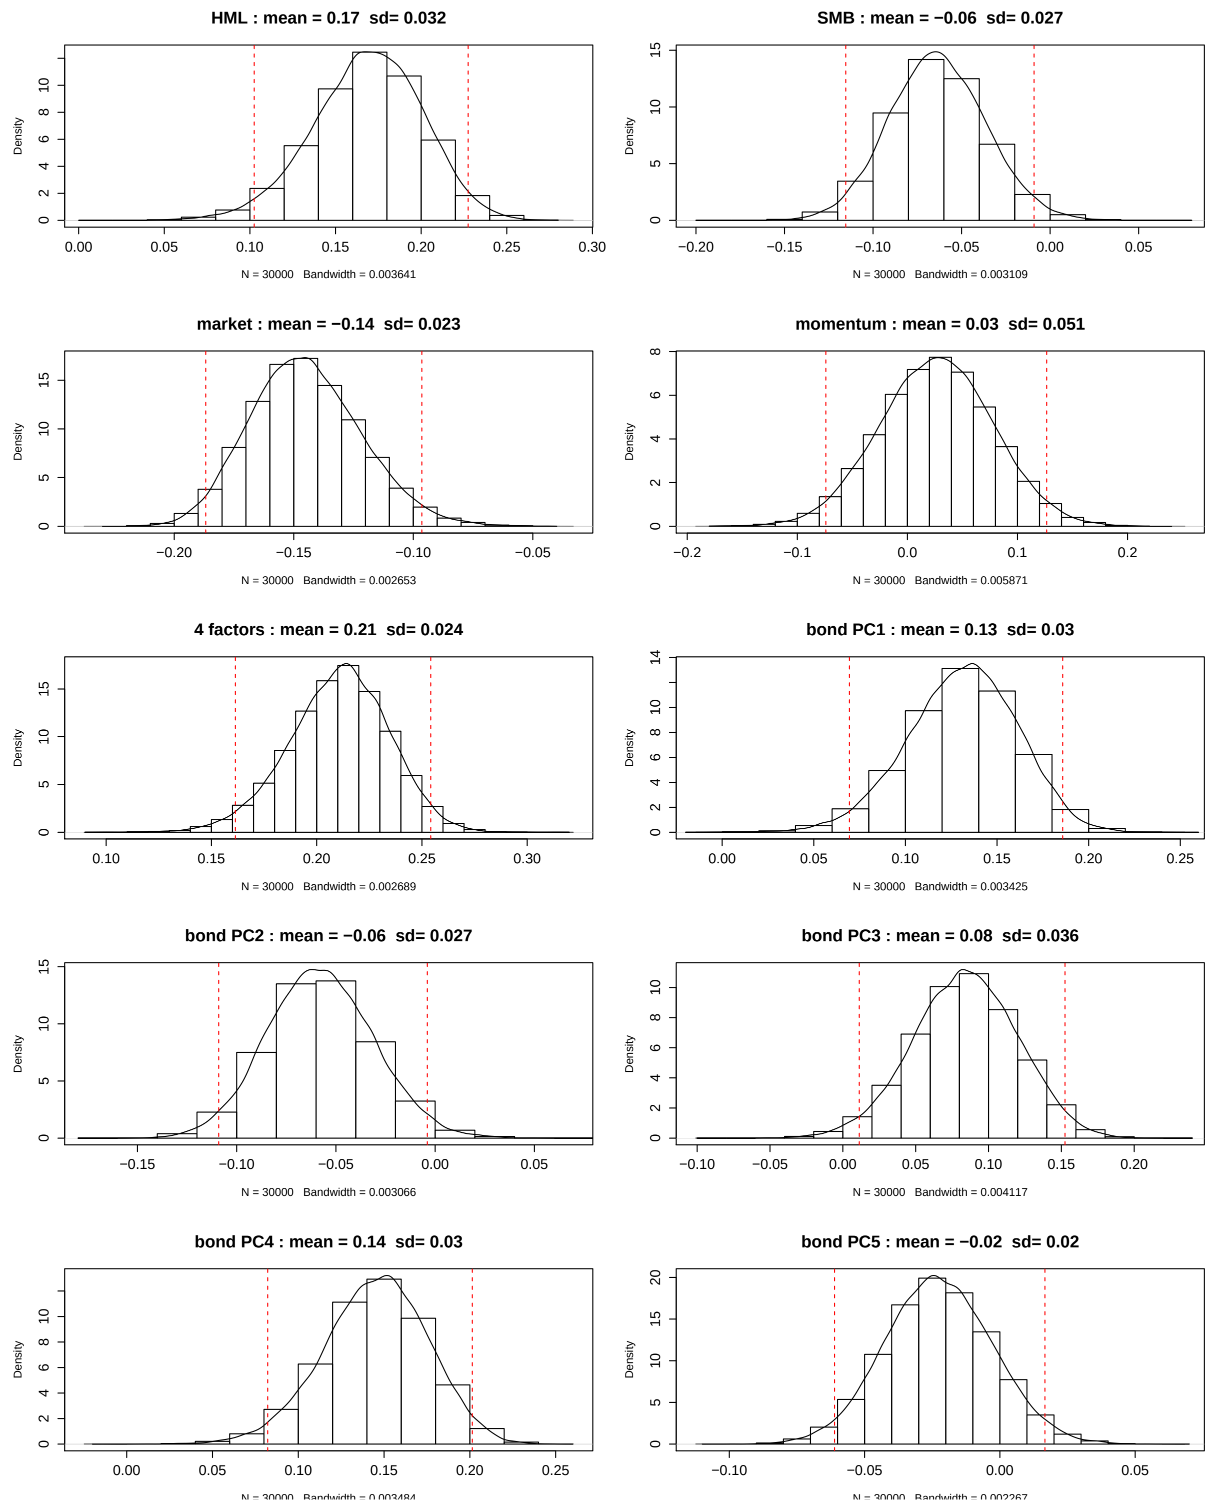

Figure 15: Common innovations, popular risk factor, and principal components of bond returns.

Posterior distributions of the correlation coefficients between the MA innovations f_t and: the HML and SMB factors (first row); market excess returns and momentum factor (second row); a linear combination of HML, SMB, market, and momentum factors (left panel on third row); the first five principal components (PC1 to PC5) of bond excess returns (remaining panels). Dashed red lines indicate centered 90% posterior coverage areas. Posterior means and standard deviations of the correlation are reported at top of each sub-graph.

Liew and Vassalou (2000) that HML forecasts good economic states and with equilibrium models with capital irreversibility (as e.g. Seru, Papanikolaou, Kogan, and Stoffman (2017)).

Focusing on the bond risk factors, we find that the conditional consumption mean is positively, and strongly, correlated with the first principal component (PC) of bond excess returns, and weakly correlated the the second and third PCs. Since the first PC of bond excess returns corresponds to the second PC of bond returns – i.e. the so called factor – our finding supports the large empirical evidence connecting the slope of the term structure of interest rates to both the consumption process (e.g. Harvey (1988)) and future economic activity (e.g. Stock and Watson (1989), Hamilton and Kim (2002)).[23]

Next, we investigate whether this risk is actually *priced* in the cross-section of assets.

V.1.5 The price of consumption risk

The latent factor model in equations (5)–(6) naturally measures the covariance between asset returns and consumption growth at different horizons. We use this implication to test whether consumption risk, as captured by our framework, is priced in the cross-section of asset returns.

Following the critique of Lewellen, Nagel, and Shanken (2010), we use a mixed cross-section of assets to ensure that there is no dominating implied factor structure in the returns. Indeed, if that was the case, it could lead to spuriously high significance levels, quality of fit, and significantly complicate the overall model assessment. However, as Figure 15 indicated, the MA(f_t) component in consumption growth does not heavily load on any of the main principal components of returns. Furthermore, our state-space frameworks allows us to create simultaneous posterior confidence bands for any cross-sectional asset pricing statistics, including the factor risk premia, pricing errors, and measures of fit. Note also that since both stocks and bonds have significant loadings on the common factor (and in the case of bonds, also on the bond-specific one), we do not face the problem of *irrelevant*, or *spurious* factors (Kan and Zhang (1999)), that could also lead to the unjustifiably high significance levels.

Table 2 summarizes the cross-sectional pricing performance of our parametric model of consumption on a mixed cross-section of 9 bond portfolios, 25 Fama-French portfolios sorted by size and book-to-market, and 12 industry portfolios. For each of the specifications, we recover the full posterior distribution of the factor loadings, and estimate the associated risk premia using Fama-MacBeth (1973) cross-sectional regressions. Regardless of the specifica-

[23] Also, the significant loading on the third PC of excess returns supports the finding of Adrian, Crump, and Moench (2013) that the fourth PC of yields is one of the significant determinant for the cross-section of bond returns.

Table 2: Cross-Sectional Regressions with State-Space Loadings

Row:	α	λ_f	λ_g	$\lambda_f = \lambda_g$	$\bar{R}^2$
One latent factor specification					
(1)	.0058	16.00			.59
	[0.0053, .0063]	[9.54, 28.78]			[.56, .62]
(2)		22.00			.91
		[13.12, 39.67]			[.90, .92]
Two latent factor specification					
(3)	.0059	15.79			.59
	[.0051, .0065]	[9.10, 29.71]			[.56, .62]
(4)		21.75			0.91
		[12.52, 40.89]			[0.90, 0.92]
(5)	.0065	15.11	37.89		.59
	[−243.4, 235.1]	[8.64, 28.46]	[−243.3, 235.1]		[0.56, 0.61]
(6)		21.73	−37.01		.92
		[12.51, 40.85]	[−1152, 1198]		[.91, .92]
(7)	.0057			15.95	.59
	[.0040, .0071]			[9.12, 30.57]	[.56, .62]
(8)				21.69	.91
				[12.52, 40.79]	[.89, .92]

The table presents posterior means and centred 95% posterior coverage (in square brackets) of the Fama and MacBeth (1973) cross-sectional regression of excess returns on $\sum_{j=0}^{S} \rho_j \boldsymbol{\rho}^r$ (with associated coefficient λ_f) and $\left[\ \boldsymbol{\theta}'^b, \ \ \mathbf{0}'_{N-N_b} \ \right]' \sum_{j=0}^{S} \theta_j$ (with associated coefficient λ_g). The column labeled $\lambda_f = \lambda_g$ reports restricted estimates. Cross-section of assets: 25 Fama and French (1992) size and book-to-market portfolio; 12 industry sorted portfolios; 9 bond portfolios.

tion, there is strong support in favor of the persistent shocks to consumption being a priced risk in the cross-section of stocks and bonds: the associated cross-sectional slope is always positive and highly statistically significant, and the $\bar{R}^2$ varies from 59% to 92%, depending on whether the intercept is included in the model. While allowing for a common intercept in the estimation substantially lowers cross-sectional fit, 95% posterior coverage remains very tight, providing a reliable indication about the model performance.

The average pricing error is small, about 60 b.p. per quarter, but statistically significant in most specifications. Figure 16 summarises the posterior for the individual pricing errors for the specification in row (1) of Table 2: they are all individually statistically insignificant even at 10% level. The figure further suggests that the cross-sectional pricing could benefit from a bond-specific intercept.

While the common factor f plays an important role in explaining the cross-section of both stock and bond returns, the bond factor g is not priced. The bond-specific factor is *unspanned*, because consumption does not significantly load on it. This factor is nevertheless essential for explaining most of the time series variation in bond returns.

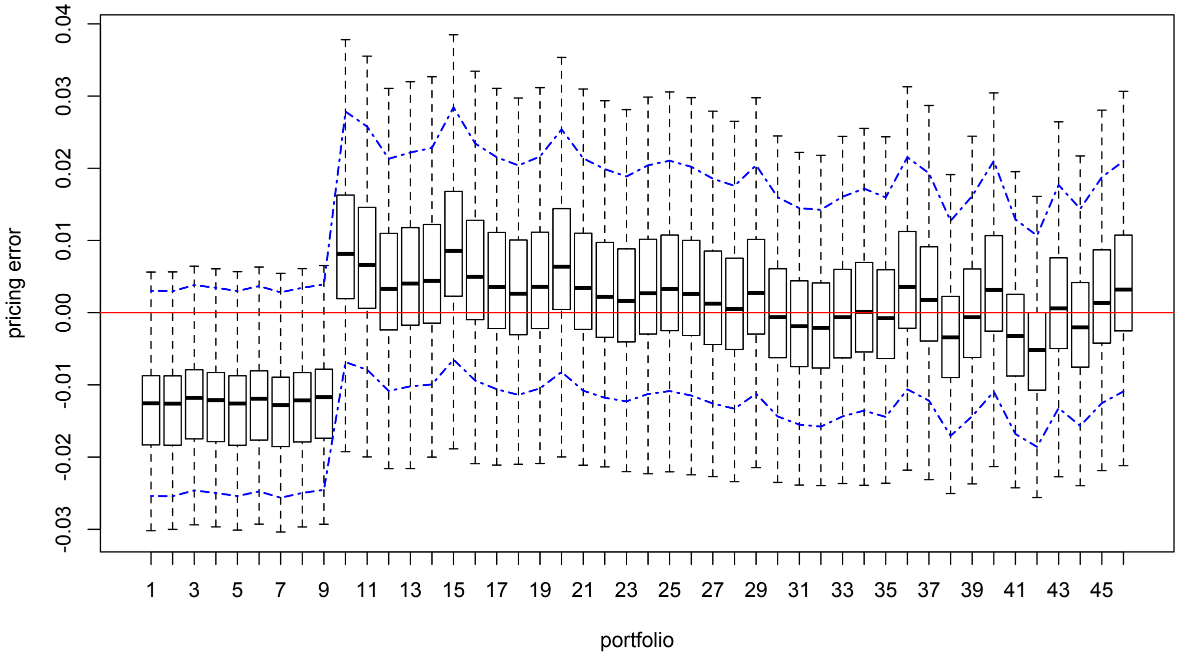

Figure 16: Posterior distributions of cross-sectional pricing errors.

Box-plots of the posterior distributions of pricing errors for the cross-sectional specification in row (1) of Table 2. Portfolios are ordered with bonds first (1 to 9), Fama-French 25 size and book-to-market second (10 to 34), and industry portfolios last. Blue dash-dot lines denote centered 90% posterior coverage areas.

V.2 The term structure of consumption exposure

The relatively tight restrictions on the parametric model in equations (1)-(2) allow us to pin down the parameters of the joint consumption-returns process with a high degree of precision. However, this comes at a price of imposing constraints on the data generating process, some of which may in principle not hold in the data. In this section we aim to test the strongest prediction of our parametric setting – the term structure of asset exposure to consumption risk – while relaxing the ancillary assumptions needed to estimate our state-space model.

Equations (1)–(2) imply a very particular pattern in the covariances of asset returns with multiperiod consumption growth, i.e. for any asset i

$$\begin{aligned}
cov(r^{ex}_{i,t+1}, \Delta c_{t,t+1}) &= \rho^r_i \rho_0 \\
cov(r^{ex}_{i,t+1}, \Delta c_{t,t+2}) &= \rho^r_i (\rho_0 + \rho_1) \\
&\dots \\
cov(r^{ex}_{i,t+1}, \Delta c_{t,t+k}) &= \rho^r_i \left(\sum_{j=1}^{k} \rho_{k-1} \right)
\end{aligned} \tag{16}$$

Therefore, the term structure of asset exposures to consumption risk is driven by a single common component: $(\rho_0, \rho_0 + \rho_1, \ldots, \sum_{j=1}^{k} \rho_{k-1})'$ – i.e. the cumulative response function of consumption to an f_t shock. This property is not affected by the potential presence of cross-sectional correlations between stocks and bonds, or additional factors driving stocks and bonds that are orthogonal to consumption. Therefore, if the time-varying dynamics of consumption growth in equation (1) describes well the data generating process, we should be able to recover the same pattern of loadings by simply extracting the first uncentered principal component of $cov(r^{ex}_{i,t+1}, \Delta c_{t,t+k})$ at different horizons k.

Figure 1 in the Introduction illustrates our findings. Remarkably, the loadings on the first PC almost perfectly match the cumulated response function from the state-space model, therefore identifying the same persistent time-varying mean for consumption growth.[24]

Figure 17: Cross-sectional spread of exposure to slow consumption adjustment risk

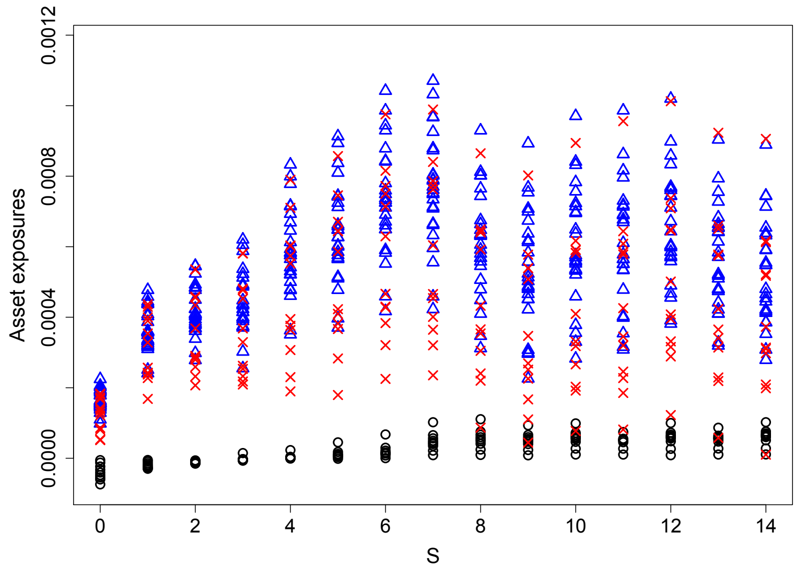

Panels present the spread of consumption betas measured as $Cov(\Delta c_{t,t+1+S}, r^e_{j,t+1})$ for different horizon S (0-15) and asset j: 9 bonds (circle), 6 Fama-French size and book-to-market (triangle), and 12 Industry (cross) portfolios.

A second testable implication of equation (16) is that, given our state space results, the covariance between asset returns and multi-period consumption growth should display an increase in both its level and cross-sectional dispersion. This conjecture is supported by figure 17 that depicts $Cov(\Delta c_{t,t+1+S}, r^e_{j,t+1})$ for various assets j and horizons S. As we move away from the standard case of $S = 0$, two observations immediately arise. First, there is a substantial increase in the average exposure of asset returns to consumption growth. Second,

[24]In the figure the level of the first PC is normalized to have the same origin as the ρ_0 estimated from the state-space formulation.

there is a strong 'fanning out' effect, observed for the higher values of the consumption horizon S. This spread in covariances rationalizes the finding of Parker and Julliard (2005) that uses $\Delta c_{t,t+1+S}$ to price time t asset returns.

V.3 Semi-parametric inference

Representative agent based consumption asset pricing models with either CRRA, Epstein and Zin (1989), or habit based preferences, as well as several models of complementarity in the utility function, consumption commitment, and models with either departures from rational expectations, or robust control, or ambiguity aversion, and even some models with solvency constraints,[25] all imply a consumption Euler equation of the form

$$C_t^{-\phi} = \mathbb{E}_t \left[C_{t+1}^{-\phi} \tilde{\psi}_{t+1} R_{j,t+1} \right] \tag{17}$$

for any gross asset return j including the risk-free rate R_{t+1}^f, and where $\mathbb{E}_t$ is the rational expectation operator conditional on information up to time t, C_t denotes flow consumption, $\tilde{\psi}_{t+1}$ depends on the particular form of preferences (and expectation formation mechanism) considered, and the ϕ parameter is a function of the underlying preference parameters.[26]

Note that equation (17) above also implies that $C_t^{-\phi} = \mathbb{E}_t \left[C_{t+1+S}^{-\phi} \psi_{t+1+S} \right]$ where $\psi_{t+1+S} := R_{t+1,t+1+S}^f \prod_{j=0}^{S} \tilde{\psi}_{t+1+j}$ is a multiplicative component of a multiperiod forward-looking SDF. Hence, the Euler equation can be equivalently rewritten as

$$\mathbf{0}_N = \mathbb{E} \left[\left(\frac{C_{t+1+S}}{C_t} \right)^{-\phi} \psi_{t+1+S} \mathbf{R}_{t+1}^e \right] = \mathbb{E} \left[M_{t+1}^S \mathbf{R}_{t+1}^e \right] \tag{18}$$

where $\mathbf{R}^e \in \mathbb{R}^N$ is a vector of excess returns, and $M_{t+1}^S := \left(C_{t+1+S}/C_t \right)^{-\phi} \psi_{t+1+S}$. Using the definition of covariance, we can rewrite the above equation as a model of expected returns

$$\mathbb{E} \left[\mathbf{R}_{t+1}^e \right] = -\frac{Cov \left(M_{t+1}^S; \mathbf{R}_{t+1}^e \right)}{\mathbb{E} \left[M_{t+1}^S \right]}. \tag{19}$$

where $M_{t+1}^S := \left(C_{t+1+S}/C_t \right)^{-\phi} \psi_{t+1+S}$. That is, under the null of the model being correctly specified, there is an entire family of SDFs that can be equivalently used for asset pricing:

[25]See, e.g.: Bansal and Yaron (2004); Abel (1990), Campbell and Cochrane (1999), Constantinides (1990), Menzly, Santos, and Veronesi (2004); Piazzesi, Schneider, and Tuzel (2007), Yogo (2006); Basak and Yan (2010), Hansen and Sargent (2010); Chetty and Szeidl (2016); Ulrich (2010); Lustig and Nieuwerburgh (2005).

[26]E.g., ϕ would denote relative risk aversion in the CRRA framework, while it would be a function of both risk aversion and elasticity of intertemporal substitution with Epstein and Zin (1989) recursive utility.

M_{t+1}^S for every $S \geq 0$. Log-linearizing the above expression, we have the linear factor model

$$\mathbb{E}\left[\mathbf{r}_{t+1}^e\right] = \left[\phi Cov\left(\Delta c_{t,t+1+S}; \mathbf{r}_{t+1}^e\right) - Cov\left(\log \psi_{t+1+S}; \mathbf{r}_{t+1}^e\right)\right] \lambda_S \tag{20}$$

where λ_S is a positive scalar. The above pricing restriction, ignoring the ψ factor, is *exactly* the one that we have tested in section V.1.5 under the null of our state-space model. We now instead want to avoid 1) using linearized relationships, 2) imposing our parametric model for consumption and asset returns and 3) ignoring the ψ component of the SDF. That is, we want to tackle directly the pricing restriction in Euler equation (18).

Since the stochastic discount factor M_{t+1}^S can be decomposed into the product of consumption growth over several periods (C_{t+1+S}/C_t) and an additional, potentially unobservable, component, we can use an Empirical Likelihood–based approach to assess the ability of low frequency consumption to price returns without taking a stand on the actual model (i.e. on the ψ_t component). The EL-based inference allows to separate the unobservable part of the SDF from that related to consumption growth, treating ψ_t like a nuisance parameter that is concentrated out (see Ghosh, Julliard, and Taylor (2016)).[27]

Since the relevance of the multiple period consumption for the cross-section of stocks has already been highlighted by Parker and Julliard (2005), we first focus on the cross-section of bonds only. Table 3 summarizes the ability of the consumption component of M_{t+1}^S to explain the cross-section of bond returns for various values of S. When $S = 0$ we have the standard CRRA consumption-CAPM, where the expected returns are driven only by their contemporaneous correlation with consumption growth. The output reflects the well-known failure of the classical model to capture the cross-section of bond returns: according to the LR-test, the model is rejected in the data, and the cross-sectional adjusted R-squared is negative and large, and the curvature parameter is very large and imprecisely estimated. Increasing the span of consumption growth, S, to 2 or more quarters drastically changes the picture: the level of cross-sectional fit increases dramatically, up to 92% for $S = 14$, the point estimates of ϕ, which in the case of additively separable CRRA utility corresponds to the Arrow-Pratt relative risk-aversion coefficient) are drastically reduced (hence more in line with economic theory).[28]

[27]Consider the following transformation of the Euler equation:

$$\mathbf{0} = \mathbb{E}\left[M_t^S \mathbf{R}_t^e\right] \equiv \int \left(\frac{C_{t+S}}{C_{t-1}}\right)^{-\phi} \psi_{t+S} \mathbf{R}_t^e dP = \int \left(\frac{C_{t+S}}{C_{t-1}}\right)^{-\phi} \mathbf{R}_t^e d\Psi = \mathbb{E}^{\Psi}\left[\left(\frac{C_{t+S}}{C_{t-1}}\right)^{-\phi} \mathbf{R}_t^e\right] \tag{21}$$

where P is the unconditional physical probability measure, Ψ is an (absolutely continuous) probability measure such that $\frac{d\Psi}{dP} = \frac{\psi_{t+S}}{\bar{\psi}}$ where $\bar{\psi} = \mathbb{E}\left[\psi_{t+S}\right]$. Both ϕ and Ψ can then be recovered via Empirical Likelihood approach. Section A.5 describes the estimation procedure in detail.

[28]Note that the LR-test indicates for all S, as one would expect, that there is a statistically significant

Table 3: Cross-section of bond returns and consumption risk

Horizon S (Quarters)	$R^2_{adj}(\%)$ (1)	α (2)	ϕ (3)	LR-test (4)
0	-668	0.0007	105	12.4191
		(0.0003)	(27.8)	[0.0004]
1	-138	0.0010	88	33.3935
		(0.0004).	(24.5)	[0.0000]
2	62	0.0024	98	101.9077
		(0.0006)	(24.7)	[0.0000]
3	81	0.0013	61	47.5635
		(0.0004)	(17.2)	[0.0000]
4	51	0.0007	50	29.2312
		(0.0003)	(13.9)	[0.0000]
5	4	0.0009	46	32.8188
		(0.0003)	(10.6)	[0.0000]
6	10	0.0009	45	29.8647
		(0.0003)	(10.0)	[0.0000]
7	58	0.0006	38	21.6852
		(0.0003)	(9.7)	[0.0000]
8	65	0.0008	40	26.2551
		(0.0002)	(9.9)	[0.0000]
9	62	0.0008	54	25.6001
		(0.0002)	(10.4)	[0.0000]
10	74	0.0008	40	20.9373
		(0.0002)	(12.1)	[0.1911]
11	74	0.0009	45	25.6532
		(0.0002)	(14.1)	[0.0000]
12	90	0.0008	77	27.8346
		(0.0002)	(16.3)	[0.0000]
13	82	0.0007	85	26.9717
		(0.0002)	(17.4)	[0.0000]
14	92	0.0007	74	25.8651
		(0.0002)	(19.7)	[0.0000]

The table reports the pricing of 9 excess bond holding returns for various values of the horizon S, and allowing for an intercept. Standard errors are reported in parentheses and p-values in brackets. Estimation is performed using Empirical Likelihood and allowing for a common pricing error, α, such that $\mathbf{0} = \mathbb{E}\left[M^S_{t+1}\left(\mathbf{R}^e_{t+1} - \alpha\right)\right]$. The fit measured as: $R^2_{adj} = 1 - \frac{n-2}{n-1}\widehat{Var}_c\left(\frac{1}{T}R_{i,t+1} - \hat{\alpha} - \frac{\widehat{Cov}\left((C_{t+1+S}/C_t)^{-\hat{\phi}}, \mathbf{R}^e_{t+1}\right)}{\widehat{\mathbb{E}}\left[(C_{t+1+S}/C_t)^{-\hat{\phi}}\right]}\right) \Big/ \widehat{Var}_c\left(\frac{1}{T}R_{i,t+1}\right)$ where $\widehat{Var}_c(\cdot)$ denotes the sample cross-sectional variance.

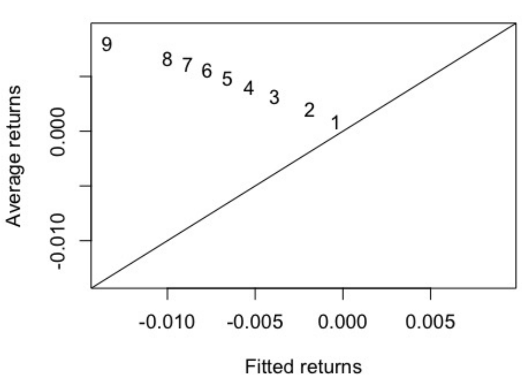

(a) S=0

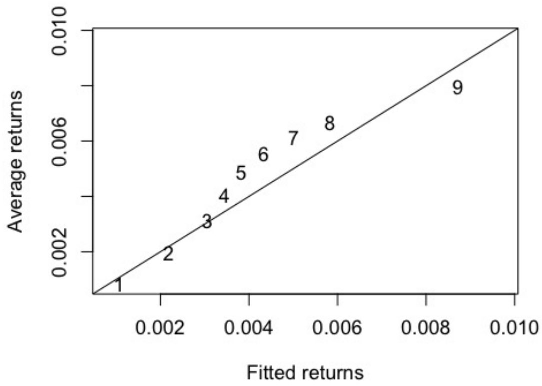

(b) S=12

Figure 18: Slow consumption adjustment factor and the cross-section of bond returns

The figures show average and fitted returns on the cross-section of 9 bond portfolios, sorted by maturity. The model is estimated by Empirical Likelihood for various values of consumption horizon S. $S = 0$ corresponds to the standard consumption-based asset pricing model; $S = 12$ corresponds to the use of ultimate consumption risk, where the cross-section of returns is driven by the their correlation with the consumption growth over 13 quarters, starting from the contemporaneous one.

Most importantly, fo $S >> 0$, ϕ is much more precisely estimated. The large standard error associated with this parameter for the standard consumption-based model ($S = 0$) is due to the fact that the level and spread of the contemporaneous correlation between asset returns and consumption growth is rather low (see Figure 17). This in turn leads to substantial uncertainty in parameter estimation. As the number of quarters used to measure consumption risk increases, the link between returns and the slow moving component of the consumption becomes more pronounced (as in Figure 17), resulting in lower standard errors and a better quality of fit. In fact for large S, as shown in Figure 18, the model-implied average excess returns are very close to the actual ones, in drastic contrast to the $S = 0$ case. The contemporaneous correlation between bond returns and consumption growth (Panel (a), $S = 0$) is so low that not only it delivers a very poor fit, but actually reverses the order of the portfolios: i.e. the fitted average return from holding long-term bonds is smaller than that of the short term ones. Instead, when the horizon used to measure consumption risk is increased, the quality of fit substantially improves, leading to an R-squared of 85% for $S = 12$ in Panel (b). Note that, in light of our parametric evidence, this improvement in fit is driven by the fact that the cumulated response of consumption over many quarters following the returns captures the slow consumption response to time t wealth shocks (to which asset returns react immediately).

The ability of slow consumption adjustment risk to capture a large proportion of the

component ψ_t in the pricing kernel.

Table 4: Expected excess returns and consumption risk

Horizon S (Quarters)	$R^2_{adj}(\%)$ (1)	α (2)	ϕ (3)	LR-test (4)	$R^2_{adj}(\%)$ (5)	α (6)	ϕ (7)	LR-test (8)
	Cross-sectional estimates							
	Panel A: 9 Bonds and Fama-French 6 portfolios				*Panel B: 9 Bonds, Fama-French 6, and Industry 12 portfolios*			
0	33	-0.0002	113	6.7888	38	-0.0002	91	9.0198
		(0.0003)	(27.3)	[0.0092]		(0.0002)	(22.8)	[0.0027]
10	76	0.0000	38	28.8862	55	0.0006	16	23.6425
		(0.0004)	(8.4)	[0.0000]		(0.0002)	(4.1)	[0.0000]
11	75	0.0001	53	32.1007	48	0.0005	12	18.9782
		(0.0002)	(10.8)	[0.0000]		(0.0002)	(3.3)	[0.0000]
12	91	0.0005	45	29.4202	50	0.0005	12	17.7822
		(0.0003)	(13.4)	[0.0000]		(0.0002)	(3.2)	[0.0000]
	Panel C: 9 Bonds and Fama-French 25 portfolios				*Panel D: 9 Bonds, Fama-French 25, and Industry 12 portfolios*			
0	66	-0.0002	88	6.3405	39	0.0000	32	4.5356
		(0.0002)	(19.7)	[0.0118]		(0.0002)	(16.6)	[0.0332]
10	79	0.0002	23	11.0360	45	0.0002	10	11.3431
		(0.0002)	(3.9)	[0.0009]		(0.0002)	(2.6)	[0.0007]
11	80	0.0002	22	9.8914	42	0.0002	9	9.5157
		(0.0002)	(3.7)	[0.0017]		(0.0002)	(2.3)	[0.0020]
12	69	0.0002	17	6.7009	38	0.0002	8	8.9695
		(0.0002)	(3.1)	[0.0096]		(0.0002)	(2.1)	[0.0027]

The table reports the pricing of excess returns of stocks and bonds, allowing for no intercept. Standard errors are reported in parentheses and p-values in brackets. Estimation is done using empirical likelihood.

cross-sectional variation in returns is not a feature of the bond market alone: it works equally well on the joint cross-section of stocks and bonds, providing a simple and parsimonious one factor model for co-pricing securities in both asset classes. Table 4 presents cross-sectional estimates for various portfolios of stocks and bonds as S varies. Three patterns, common to also table 3, are evident. First, cross-sectional fit is generally higher for $S >> 0$. Second, the point estimates of ϕ are much smaller as S increase. Finally, cumulated consumption growth delivers much sharper estimates of the curvature parameter (with standard errors often an order of magnitude smaller). Again, this is in line with the evidence in Figure 17: as S increases, both the level and the spread of asset loadings on consumption growth become larger and better identified. This, in turn, delivers better fit, a lower curvature for the consumption growth, and a tighter identification of the latter. For robustness, Appendix A.7 provides similar empirical evidence for the alternative model specifications that include asset class-specific intercepts.

VI Conclusions

We identify the stochastic process of consumption growth using the information contained in financial returns. Our strategy relies on the central insight of the intertemporal Euler equation of models that have consumption as one of the state variables entering the utility function: most shocks affecting the household force it to adjust both investment and consumption plans.

We show that a flexible parametric model with common factors driving asset dynamics and consumption identifies the slow varying conditional mean of the consumption growth. This component is persistent at the business cycle frequency, and is economically large, capturing more than a quarter of the time series variation of consumption growth. This indicates that the shocks spanned by financial markets are first order drivers of consumption risk. Controlling for the time variation in the conditional mean of consumption growth, we do not find evidence of volatility clustering in consumption. One implication of this is that models with stochastic volatility in consumption should probably include a strong contemporaneous leverage effect.

Turning to the asset pricing implications, we find that both stocks and bonds load significantly on the innovations spanned by consumption growth. This generates sizeable risk premia and dispersion in returns, consistent with the size and value anomalies, as well as with the positive slope of the term structure of bond excess returns. As a result, our model explains between 36% and 95% of the time series variation in returns and between 57% and 90% of the joint cross-sectional variation in stocks and bonds.

We also provide a large set of complementary evidence that, despite not relying on our parametric formulation, strongly supports all the main time series and cross-sectional findings.

References

Abel, A. B. (1990): "Asset Prices Under Habit Formation and Catching Up with the Joneses," *American Economic Review*, 80, 38–42.

Abel, A. B., J. C. Eberly, and S. Panageas (2013): "Optimal Inattention to the Stock Market With Information Costs and Transactions Costs," *Econometrica*, 81(4), 1455–1481.

Adrian, T., R. K. Crump, and E. Moench (2013): "Pricing the Term Structure with Linear Regressions," *Journal of Financial Economics*, 110(1), 110–138.

Ang, A., M. Piazzesi, and W. Min (2006): "What does the yield curve tell us about GDP growth?," *Journal of Econometrics*, 131, 359–403.

Balduzzi, P., and T. Yao (2007): "Testing Heterogeneous Agent Models: an Alternative Aggregation Approach," *Journal of Monetary Economics*, 54, 369–412.

Bandi, F., and A. Tamoni (2015): "Business-Cycle Consumption Risk and Asset Prices," available at SSRN: http://ssrn.com/abstract=2337973.

Bansal, R., R. F. Dittmar, and C. T. Lundblad (2005): "Consumption, Dividends, and the Cross Section of Equity Returns," *The Journal of Finance*, 60(4), pp. 1639–1672.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2014): "Volatility, the Macroeconomy and Asset Prices," *Journal of Finance*, (69), 2471–2511.

Bansal, R., D. Kiku, and A. Yaron (2012): "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 1, 183–221.

Bansal, R., and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59(4), 1481–1509.

Barro, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121(3), 823–866.

Basak, S., and H. Yan (2010): "Equilibrium Asset Prices and Investor Behaviour in the Presence of Money Illusion," *Review of Economic Studies*, 77(3), 914–936.

Bauwens, L., M. Lubrano, and J.-F. Richard (1999): *Bayesian Inference in Dynamic Econometric Models*. Oxofrd University Press, Oxford.

Beeler, J., and J. Y. Campbell (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1, 141–182.

Bekaert, G., E. Engstrom, and S. R. Grenadier (2010): "Stock and Bond Returns with Moody Investors," *Journal of Empirical Finance*, 17(5), 867–894.

Bekaert, G., and S. R. Grenadier (1999): "Stock and Bond Pricing in an Affine Economy," NBER Working Papers 7346, National Bureau of Economic Research, Inc.

Blanchard, O. J., and D. Quah (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *The American Economic Review*, 79, 655–73.

Boons, M., and A. Tamoni (2015): "Horizon-Specific Macroeconomic Risks," available at SSRN: http://ssrn.com/abstract=2516251.

Box, G. E. P., and D. A. Pierce (1970): "Distribution of Residual Autocorrelations in Autoregressive-Integrated Moving Average Time Series Models," *Journal of the American Statistical Association*, 65(332), pp. 1509–1526.

Brainard, W. C., W. R. Nelson, and M. D. Shapiro (1991): "The Consumption Beta Explains Expected Returns at Long Horizons," Yale University Manuscript.

Breeden, D. T., M. R. Gibbons, and R. H. Litzenberger (1989): "Empirical Test of the Consumption-Oriented CAPM," *The Journal of Finance*, 44(2), 231–262.

Caballero, R. J. (1995): "Near Rationality, Heterogeneity, and Aggregate Consumption," *Journal of Money, Credit, and Banking*, 27(1), 29–48.

Campbell, J. Y. (2017): *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton University Press, NJ.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–51.

Campbell, J. Y., and A. Deaton (1989): "Why is Consumption so Smooth?," *Review of Economic Studies*, 56(3), 357–73.

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley (2018): "An Intertemporal CAPM with Stochastic Volatility," *Journal of Financial Economics*, 128, 207–233.

Chen, H., W. Dou, and L. Kogan (2015): "Measuring the 'Dark Matter' in Asset Pricing Models," Available at SSRN: https://ssrn.com/abstract=2326753.

Chen, X., and S. C. Ludvigson (2009): "Land of Addicts? An Empirical Investigation of Habit-based Asset Pricing Models," *Journal of Applied Econometrics*, 24(7), 1057–1093.

Chernov, M., and P. Mueller (2012): "The Term Structure of Inflation Expectations," *Journal of Financial Economics*, 106(2), 367 – 394.

Chetty, R., and A. Szeidl (2016): "Consumption Commitments and Habit Formation," *Econometrica*, 84, 855–890.

Cochrane, J. H. (2007): "Financial Markets and the Real Economy," in *Handbook of the Equity Premium*, ed. by R. Mehra, chap. 7, pp. 237–325. Elsevier Press.

Cochrane, J. H., and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95(1), 138–160.

Constantinides, G. M. (1990): "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98(2), 519–43.

Csiszar, I. (1975): "I-Divergence Geometry of Probability Distributions and Minimization Problems," *Annals of Probability*, 3, 146–158.

Daniel, K. D., and D. Marshall (1997): "The Equity Premium Puzzle and the Risk-Free Rate Puzzle at Long Horizons," *Macroeconomic Dynamics*, 1(2), 452–84.

Deaton, A. (1987): "Life-Cycle Models of Consumption: Is the Evidence Consistent With the Theory," in *Advances in Econometrics, Fifth World Congress*, ed. by T. F. Bewley, vol. 2. Cambridge University Press.

Dew-Becker, I., and S. Giglio (2016): "Asset Pricing in the Frequency Domain: Theory and Empirics," *Review of Financial Studies*, 29, 2029–68.

Eichenbaum, M., L. P. Hansen, and K. Singleton (1988): "A Times Series Analysis of Representative Agent Models of Consumption and Leisure," *Quarterly Journal of Economics*, 103, 51–78.

Engle, R. F. (1982): "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987–1007.

Epstein, L. G., and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–968.

Fama, E. F., and K. R. French (1992): "The Cross-Section of Expected Stock Returns," *The Journal of Finance*, 47, 427–465.

Fama, E. F., and J. MacBeth (1973): "Risk, Return and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81, 607–636.

Flavin, M. (1993): "The Excess Smoothness of Consumption: Identification and Interpretation," *The Review of Economic Studies*, 60(3), 651–666.

Ghosh, A., C. Julliard, and A. Taylor (2016): "What is the Consumption-CAPM missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models," *Review of Financial Studies*, 30, 442–504.

Gurkaynak, Refet, B. S., and J. H. Wright (2007): "The US Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, 54(8), 2291–2304.

Hall, R. E. (1978): "The Stochastic Implications of the Life Cycle Permanent Income Hypothesis," *Journal of Political Economy*, 86(6), 971–87.

Hamilton, J. D., and D. H. Kim (2002): "A Re-examination of the Predictability of the Yield Spread for Real Economic Activity," *Journal of Money, Credit, and Banking*, 34, 340–60.

Hansen, L. P., J. Heaton, J. Lee, and N. Roussanov (2007): "Intertemporal Substitution and Risk Aversion," in *Handbook of Econometrics*, ed. by J. Heckman, and E. Leamer, vol. 6 of *Handbook of Econometrics*, chap. 61. Elsevier.

Hansen, L. P., J. C. Heaton, and N. Li (2008): "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy*, 116(2), 260–302.

Hansen, L. P., and T. J. Sargent (2010): "Fragile Beliefs and the Price of Uncertainty," *Quantitative Economics*, 1(1), 129–162.

Harvey, C. R. (1988): "The Real Term Structure and Consumption Growth," *Journal of Financial Economics*, 22(2), 305 – 333.

Jacobs, K., and K. K. Wang (2004): "Idiosyncratic Consumption Risk and the Cross-Section of Asset Returns," *Journal of Finance*, 59, 2211–52.

Jagannathan, R., and S. Marakani (2016): "Price-Dividend Ratio Factor Proxies for Long-Run Risks," *The Review of Asset Pricing Studies*, 5, 1–47.

Jagannathan, R., and Y. Wang (2007): "Lazy Investors, Discretionary Consumption, and the Cross-Section of Stock Returns," *The Journal of Finance*, 62(4), pp. 1623–1661.

Kan, R. M., and C. Zhang (1999): "Two-pass Tests of Asset Pricing Models with Useless Factors," *Journal of Finance*, 54, 204–35.

Kitamura, Y. (2006): "Empirical Likelihood Methods in Econometrics: Theory and Practice," Cowles Foundation Discussion Papers 1569, Cowles Foundation, Yale University.

Koijen, R. S., H. Lustig, and S. V. Nieuwerburgh (2017): "The Cross-Section and Time-Series of Stock and Bond Returns," *Journal of Monetary Economics*, 88, 50–69.

Kroencke, T. A. (2017): "Asset Pricing Without Garbage," *Journal of Finance*, 72(2), 47–98.

Lettau, M., and S. Ludvigson (2001a): "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, 56(3), 815–49.

——— (2001b): "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying," *Journal of Political Economy*, 109, 1238–1286.

Lettau, M., and J. A. Wachter (2011): "The Term Structures of Equity and Interest Rates," *Journal of Financial Economics*, 101(1), 90–113.

Lewellen, J., S. Nagel, and J. Shanken (2010): "A Skeptical Appraisal of Asset Pricing Tests," *Journal of Financial Economics*, 96(2), 175–194.

Liew, J., and M. Vassalou (2000): "Can Book-to-market, Size and Momentum be Risk Factors that Predict Economic Growth?," *Journal of Financial Economics*, 57(2), 221–245.

Liu, Y., and B. Matthies (2018): "Long Run Risk: Is It There?," mimeo.

Ljung, G. M., and G. E. P. Box (1978): "On a Measure of Lack of Fit in Time Series Models," *Biometrika*, 65(2), 297–303.

Lucas, R. E. (1987): *Models of Business Cycles.* Blackwell, Oxford.

Ludvigson, S. C. (2012): "Advances in Consumption-Based Asset Pricing: Empirical Tests," in *Handbook of Economics and Finance*, ed. by G. M. Constantinides, M. Harris, and R. M. Stulz, vol. 2, chap. 12, pp. 799–906. Elsevier Press.

Lustig, H. N., and S. G. V. Nieuwerburgh (2005): "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective," *The Journal of Finance*, 60(3), pp. 1167–1219.

Lynch, A. W. (1996): "Decision Frequency and Synchronization Across Agents: Implications for Aggregate Consumption and Equity Return," *Journal of Finance*, 51(4), 1479–97.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jorgensen (2009): "Long-Run Stockholder Consumption Risk and Asset Returns," *The Journal of Finance*, 64(6), 2427–2479.

Marshall, D. A., and N. G. Parekh (2002): "Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles?," *The Journal of Finance*, 54(2), 623–654.

Menzly, L., T. Santos, and P. Veronesi (2004): "Understanding Predictability," *Journal of Political Economy*, 112(1), 1–47.

Newey, W., and R. Smith (2004): "Higher Order Properties of GMM and Generalized Empirical Likelihood Estimators," *Econometrica*, 72, 219–255.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–08.

Parker, J. A., and C. Julliard (2005): "Consumption Risk and the Cross-Section of Expected Returns," *Journal of Political Economy*, 113(1).

Piazzesi, M., M. Schneider, and S. Tuzel (2007): "Housing, Consumption and Asset Pricing," *Journal of Financial Economics*, 83(3), 531–569.

Postlewaite, A., L. Samuelson, and D. Silverman (2008): "Consumption Commitments and Employment Contracts," *The Review of Economic Studies*, 75(2), 559–578.

Primiceri, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72(3), 821–852.

Savov, A. (2011): "Asset Pricing with Garbage," *Journal of Finance*, 66(1), 177–201.

Schorfheide, F., D. Song, and A. Yaron (2018): "Identifying Long-Run Risks: A Bayesian Mixed-Frency Approach," 86, 617–654.

Seru, A., D. Papanikolaou, L. Kogan, and N. Stoffman (2017): "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, 132(2), 665–712.

Shore, S. H., and T. Sinai (2010): "Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?," *The Review of Economics and Statistics*, 92(2), 408–424.

Sims, C. A., and T. Zha (1999): "Error Bands for Impulse Responses," *Econometrica*, 67(5), 1113–1155.

Startz, R. (1989): "The Stochastic Behavior of Durable and Nondurable Consumption," *Review of Economics and Statistics*, 71, 356–63.

Stock, J. H., and M. W. Watson (1989): "New Indexes of Coincident and Leading Indicators," In: Blanchard, O., Fischer, S. (Eds.), 1989 NBER Macroeconomics Annual. MIT Press, Cambridge, MA.

Tedongap, R. (2007): "Consumption Volatility and the Cross-Section of Stock Returns," working paper.

Tinang, J. (2018): "Macro Uncertainty and the Term Structure of Risk Premium," mimeo.

Uhlig, H. (2005): "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, 52, 381–419.

Ulrich, M. (2010): "Observable Long-Run Ambiguity and Long-Run Risk," Columbia University Manuscript.

Yogo, M. (2006): "A Consumption-Based Explanation of Expected Stock Returns," *Journal of Finance*, 61(2), 539–580.

Zviadadze, I. (2018): "Term Structure of Risk in Expected Returns," mimeo.

A Appendix

A.1 Data Description

Bond holding returns are calculated on a quarterly basis using the zero coupon yield data constructed by Gurkaynak and Wright (2007)[29] from fitting the Nelson-Siegel-Svensson curves daily since June 1961, and excess returns are computed subtracting the return on a three-month Treasury bill. We consider the set of the following maturities: 6 months, 1, 2, 3, 4, 5, 6, 7, and 10 years, which gives us a set of 9 bond portfolios.

We consider several portfolios of stock returns. The baseline specification relies, in addition to the bond portfolios, on the 25 size and book-to-market Fama-French portfolios (Fama and French (1992)), and 12 industry portfolios, available from Kenneth French data library. We consider monthly returns from July, 1961 to July, 2017, and accumulate them to form quarterly returns, matching the frequency of consumption data. Excess returns are then formed by subtracting the corresponding return on the three-month Treasury bill.

Consumption flow is measured as real (chain-weighted) consumption expenditure on non-durable goods per capita available from the National Income and Product Accounts (NIPA). We use the end-of-period timing convention and assume that all of the expenditure occurs at the end of the period between t and $t+1$. We make this (common) choice because under this convention the entire period covered by time t consumption is part of the information set of the representative agent before time $t+1$ returns are realised. All the returns are made real using the corresponding consumption deflator.

Overall, this gives us consumption growth and matching real excess quarterly holding returns on 46 portfolios, from the third quarter of 1961 to the second quarter of 2017.

A.2 S-VAR Identification via Long-Run Restrictions

Consider the structural vector autoregression of order p for the vector of variables X_t (given by the quarterly consumption growth and market returns):

$$\Gamma_0 X_t + \Gamma(L) X_{t-1} = c + \varepsilon_t, \;\; \varepsilon_t \sim iid\mathcal{N}(0, I),$$

where c is a vector of constants, L denotes the lag operator, $\Gamma(L) \equiv \Gamma_1 + \Gamma_2 L + ... + \Gamma_p L^{p-1}$, and each Γ_j is a two by two matrix. In order to identify the S-VAR using long-run restrictions

[29] The data is regularly updated and available at:
http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html

we follow Blanchard and Quah (1989) and work with the moving average representation

$$X_t = \kappa + A(L) \begin{bmatrix} \varepsilon_t^{sr} \\ \varepsilon_t^{lr} \end{bmatrix} \tag{22}$$

where κ is a vector of constants, $A(L) \equiv A_0 + A_1 L + A_2 L^2 + ...A_\infty L^\infty \equiv [\Gamma_0 + L\Gamma(L)]^{-1}$, ε_t^{sr} and ε_t^{lr} denote, respectively, short- and long-run Gaussian shocks with covariance matrix normalized to be equal to the identity matrix.

The two types of shocks are identified imposing the restriction $\sum_{j=0}^{\infty} \{A_j\}_{1,2} = 0$ where $\{.\}_{1,2}$ returns the $(1,2)$ element of the matrix. That is, the short-run shock has no long-run effect on at least one of the elements of the vector Y. This restriction also implies that $A(1) \equiv \sum_{j=0}^{\infty} A_j$ should be a lower triangular matrix.

The S-VAR coefficient can be easily recovered from the reduced form VAR

$$X_t = \gamma + B(L) X_{t-1} + v_t,\ v_t \sim \mathcal{N}(0, \Omega)$$

where $B(L) = B_0 + B_1 L + ... + B_p L^p$ and $\Omega = \Gamma_0^{-1} \left(\Gamma_0^{-1}\right)'$ can be estimated via OLS.

Given the restrictions $\sum_{j=0}^{\infty} \{A_j\}_{1,2} = 0$, it follows that $D := [I - B(1)]^{-1} \Gamma_0^{-1}$ should be a lower triangular matrix. Note also that $DD' = [I - B(1)]^{-1} \Omega [I - B(1)]^{-1\prime}$. Hence, an estimate of the DD' matrix, $\widehat{DD'}$, can be constructed from the reduced form OLS estimates $\hat{B}(L)$ and $\hat{\Omega}$, and imposing the lower triangular structure on D, we can estimate $\hat{D}$ from the Choleski decomposition of $\widehat{DD'}$. Finally, we can recover the S-VAR parameters from $\hat{\Gamma}_0^{-1} = \left[I - \hat{B}(1)\right] \hat{D}$, $\hat{\Gamma}(L) = -\hat{\Gamma}_0 \hat{B}(L)$ and $\hat{c} = \hat{\Gamma}_0 \hat{\gamma}$. Impulse response functions and their confidence regions can then be constructed following Sims and Zha (1999).

The reduced form version of the VAR is estimated with 5 lags in order to allow for possible seasonality issues and a potentially rich dynamic. Nevertheless, the shape of the estimated impulse response functions remains very stable across lag length specifications.

A.3 State Space Estimation and Generalizations

Let $\Pi' := [\boldsymbol{\mu}, \mathbf{H}]$, $x_t' := [1, \mathbf{z}_t']$. Under a (diffuse) Jeffreys' prior the likelihood of the data in equation (12) implies the posterior distribution

$$\Pi' | \Sigma, \{\mathbf{z}_t\}_{t=1}^T, \{\mathbf{y}_t\}_{t=1}^T \sim \mathcal{N}\left(\hat{\Pi}_{OLS}'; \Sigma \otimes (x'x)^{-1}\right)$$

where x contains the stacked regressors, and the posterior distribution of each element on the main diagonal of Σ is given by[30]

$$\sigma_j^2 \Big| \{\mathbf{z}_t\}_{t=1}^T, \{\mathbf{y}_t\}_{t=1}^T \sim \text{Inv-}\Gamma\left(\left(T - m_j - 1\right)/2, T\hat{\sigma}_{j,OLS}^2/2\right)$$

where m_j is the number of estimated coefficients in the j-th equation. That is, the conditional posterior has a Normal-inverse-Γ structure. Moreover, $\mathbf{F}$ and Ψ have a Dirac posterior distribution at the points defined in equation (10). Therefore, the missing part necessary for taking draws via MCMC using a Gibbs sampler, is the conditional distributions of $\mathbf{z}_t$. Since

$$\begin{array}{c} \mathbf{y}_t \\ \mathbf{z}_t \end{array} \Bigg| \mathcal{I}_{t-1}, \mathbf{H}, \Psi, \Sigma \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu} \\ \mathbf{F}\mathbf{z}_{t-1} \end{bmatrix}; \begin{bmatrix} \Omega & \mathbf{H} \\ \mathbf{H}' & \Psi \end{bmatrix}\right),$$

where $\Omega := Var_{t-1}\left(\mathbf{y}_t\right) = \mathbf{H}\Psi\mathbf{H}' + \boldsymbol{\Sigma}$, this can be constructed, and values can be drawn, using a standard Kalman filter and smoother approach. Let

$$\mathbf{z}_{t|\tau} := E\left[\mathbf{z}_t | \mathbf{y}^\tau, \mathbf{H}, \Psi, \Sigma\right]; \quad \mathbf{V}_{t|\tau} := Var\left(\mathbf{z}_t | \mathbf{y}^\tau, \mathbf{H}, \Psi, \Sigma\right).$$

where $\mathbf{y}^\tau$ denotes the history of $\mathbf{y}_t$ until τ. Then, given $\mathbf{z}_{0|0}$ and $\mathbf{V}_{0|0}$, the Kalman filer delivers:

$$\begin{aligned} \mathbf{z}_{t|t-1} &= \mathbf{F} z'_{t-1|t-1}; \;\; \mathbf{V}_{t|t-1} = \mathbf{F}\mathbf{V}_{t-1|t-1}\mathbf{F}' + \Psi; \;\; \mathbf{K}_t = \mathbf{V}_{t|t-1}\mathbf{H}'\left(\mathbf{H}\mathbf{V}_{\mathbf{t|t-1}}\mathbf{H}' + \Sigma\right)^{-1} \\ \mathbf{z}_{t|t} &= \mathbf{z}_{t|t-1} + \mathbf{K}_t\left(\mathbf{y}_t - \boldsymbol{\mu} - \mathbf{H}\mathbf{z}_{t|t-1}\right); \;\; \mathbf{V}_{t|t} = \mathbf{V}_{t|t-1} - \mathbf{K}_t\mathbf{H}\mathbf{V}_{t|t-1}. \end{aligned}$$

The last elements of the recursion, $\mathbf{z}_{T|T}$ and $\mathbf{V}_{T|T}$, are the mean and variance of the normal distribution used to draw $\mathbf{z}_T$. The draw of $\mathbf{z}_T$ and the output of the filter can then be used for the first step of the backward recursion, which delivers the $\mathbf{z}_{T-1|T}$ and $\mathbf{V}_{T-1|T}$ values necessary to make a draw for $\mathbf{z}_{T-1}$ from a gaussian distribution. The backward recursion can be continued until time zero, drawing each value of $\mathbf{z}_t$ in the process, with the following updating formulae for a generic time t recursion:

$$\mathbf{z}_{t|t+1} = \mathbf{z}_{t|t} + \mathbf{V}_{t|t}\mathbf{F}'\mathbf{V}_{t+1|t}^{-1}\left(\mathbf{z}_{t+1} - \mathbf{F}\mathbf{z}_{t|t}\right); \;\; \mathbf{V}_{t|t+1} = \mathbf{V}_{t|t} - \mathbf{V}_{t|t}\mathbf{F}'\mathbf{V}_{t+1|t}^{-1}\mathbf{F}\mathbf{V}_{t|t}.$$

Hence parameters and states can be drawn via the Gibbs sampler using the following algorithm:

1. Start with a guess $\tilde{\Pi}'$ and $\tilde{\Sigma}^{-1}$ (e.g. the frequentist maximum likelihood estimates),

[30] Relaxing the diagonality assumption the posterior distribution of Σ^{-1} is a Wishart centered at the OLS estimates.

and use it to construct initial draws for $\boldsymbol{\mu}$ and $\mathbf{H}$. Using also $\mathbf{F}$ and Ψ, draw the $\mathbf{z}_t$ history using the Kalman recursion above with (Kalman step)

$$\mathbf{z}_t \sim \mathcal{N}\left(\mathbf{z}_{t|t+1}; \mathbf{V}_{t|t+1}\right).$$

2. Conditioning on $\{\mathbf{z}_t\}_{t=1}^T$ (drawn at the previous step) and $\{\mathbf{y}_t\}_{t=1}^T$ run *OLS* imposing the zero restrictions and get $\hat{\Pi}'_{OLS}$ and $\hat{\Sigma}_{OLS}$, and draw $\tilde{\Pi}'$ and $\tilde{\Sigma}^{-1}$ from the Normal-inverse-Γ (N-i-Γ step). Use these draws as the initial guess for the previous point of the algorithm, and repeat.

Computing posterior confidence intervals for the cross-sectional performance of the model, conditional on the data, is relatively simple since, conditional on a draw of the time series parameters, estimates of the risk premia (λ's in equations (13) and (14)) are just a mapping obtainable via the linear projection of average returns on the asset loadings in $\mathbf{H}$. Hence, to compute posterior confidence intervals for the cross-sectional analysis, we repeat the cross-sectional estimation for each posterior draw of the time series parameters, and report the posterior distribution of the cross-sectional statistics across these draws..

A.4 The Moving Average Representation of The Long Run Risk Process

We assume the same data generating process as in Bansal and Yaron (2004), with the only exception that we introduce a square-root process for the variance, as in Hansen, Heaton, Lee, and Roussanov (2007), that is:

$$\Delta c_{t,t+1} = \mu + x_t + \sigma_t \eta_{t+1}; \;\; x_{t+1} = \rho x_t + \phi_e \sigma_t e_{t+1}; \;\; \sigma_{t+1}^2 = \sigma^2(1-\nu_1) + \nu_1 \sigma_t^2 + \sigma_w \sigma_t w_{t+1},$$

where η_t, e_t, w_t, $\sim$ iid$\mathcal{N}(0,1)$. The calibrated *monthly* parameter values are: $\mu = 0.0015$, $\rho = 0.979$, $\phi_e = 0.044$, $\sigma = 0.0078$, $\nu_1 = 0.987$, $\sigma_w = 0.00029487$. To extract the quarterly frequency moving average representation of the process, we proceed in two steps. First, we simulate a long sample (five million observations) from the above system treating the given parameter values as the truth. Second, we aggregate the simulated data into quarterly observation and we use them to estimate, via MLE, the moving average representation of consumption growth in equation (1).

A.5 Empirical Likelihood Estimation

Empirical Likelihood provides a natural framework for recovering parameter estimates and probability measure Ψ defined by equation (21), by minimising Kullback-Leibler Information Criterion (KLIC):

$$(\hat{\Psi}, \hat{\phi}) = \underset{\Psi,\phi}{\arg\min}\ D(P||\Psi) \equiv \underset{\Psi}{\arg\min} \int \ln \frac{dP}{d\Psi} dP \quad \text{s.t.} \quad \mathbf{0} = \mathbb{E}^{\Psi}\left[\left(\frac{C_{t+S}}{C_{t-1}}\right)^{-\phi} \mathbf{R}_t^e\right] \tag{23}$$

From Csiszar (1975) duality result we have:

$$\hat{\psi}_t = \frac{1}{T\left(1 + \hat{\boldsymbol{\lambda}}(\theta)' \left(\frac{C_{t+S}}{C_{t-1}}\right)^{-\hat{\phi}} \mathbf{R}_t^e\right)} \quad \forall t = 1..T, \tag{24}$$

where $\hat{\phi}$ and $\hat{\boldsymbol{\lambda}} \equiv \hat{\boldsymbol{\lambda}}(\hat{\phi}) \in \mathbb{R}^n$ are the solution to the dual optimisation problem:

$$\hat{\phi} = \underset{\phi\in\mathbb{R}}{\arg\min}\ -\sum_{t=1}^{T} \ln\left(1 + \hat{\boldsymbol{\lambda}}(\phi)' \left(\frac{C_{t+S}}{C_{t-1}}\right)^{-\phi} \mathbf{R}_t^e\right) \tag{25}$$

$$\hat{\boldsymbol{\lambda}}(\phi) = \underset{\boldsymbol{\lambda}\in\mathbb{R}^n}{\arg\min}\ -\sum_{t=1}^{T} \ln\left(1 + \boldsymbol{\lambda}(\phi)' \left(\frac{C_{t+S}}{C_{t-1}}\right)^{-\phi} \mathbf{R}_t^e\right) \tag{26}$$

The dual problem is usually solved via the combination of internal and external loops: (Kitamura (2006)): first, for each ϕ find the optimal values of the Lagrange multipliers λ, as in equation (26); then minimize the value of the dual objective function w.r.t. $\phi(\hat{\lambda})$, following equation (25).

Empirical likelihood estimator is known not only for its nonparametric likelihood interpretation, but also for its convenient asymptotic representation and properties (Newey and Smith (2004)).

A.6 Additional Figures

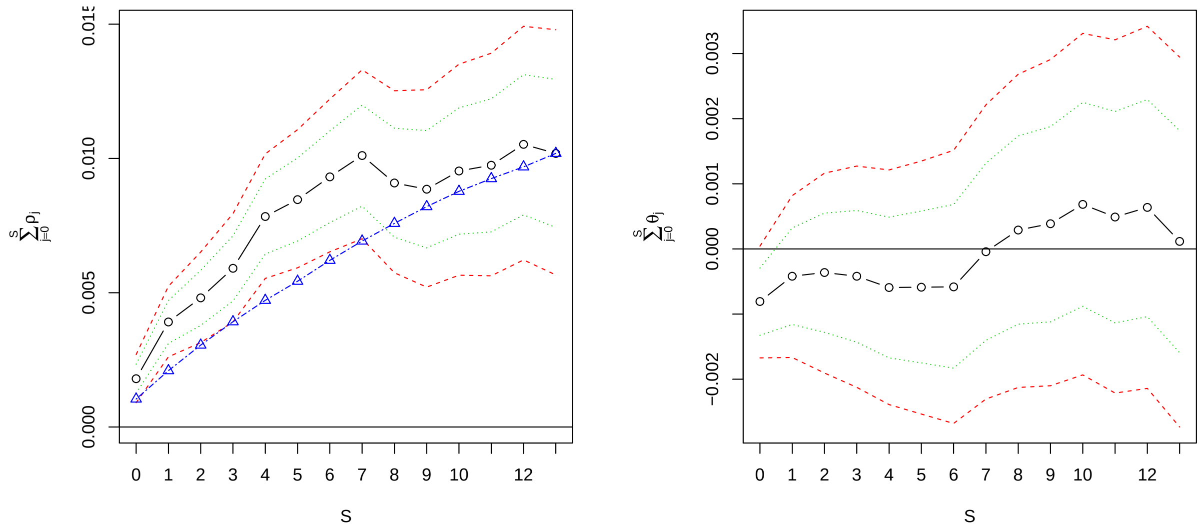

Figure A1: Slow consumption adjustment response to the latent factors f_t and g_t shocks.

Posterior means (continuous line with circles) and centred posterior 90% (dashed line) and 68% (dotted line) coverage regions. Estimation based on the two-factor model in equations (5) and (6). Left panel: cumulated consumption response to common factor, f_t, shocks. Right panel: cumulated consumption response to e to bond factor, g_t, shocks. Triangles denote Bansal and Yaron (2004) implied values.

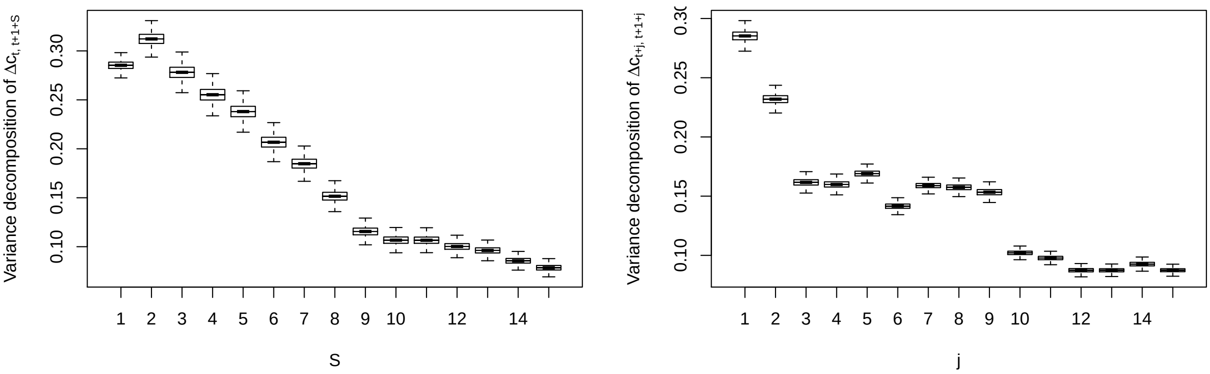

Figure A2: Variance of consumption growth explained by the MA components f and g.

Box-plots (posterior 95% coverage area) of the percentage of time series variances of consumption growth explained by the MA components f and g. Left panel: cumulated consumption growth $\Delta c_{t,t+1+S}$. Right panel: one period consumption growth $\Delta c_{t+j,t+1+j}$.

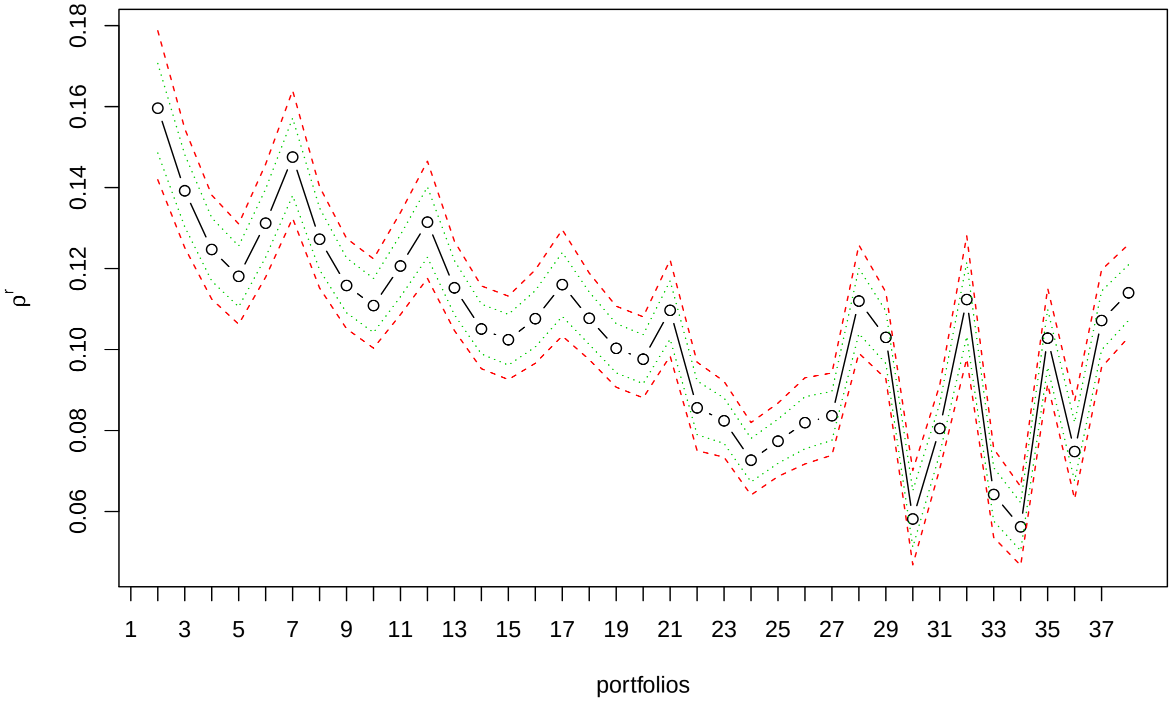

Figure A3: Common factor loadings ($\boldsymbol{\rho}^r$) of the stock portfolios in the two-factor model.

The graph presents posterior means of the stocks factor loadings on f_t (continuous line with circles) and centred posterior 90% (dashed line) and 68% (dotted line) coverage regions in the two latent factors model. Ordering of portfolios: 25 Fama and French (1992) size and book-to-market sorted portfolios (e.g. portfolio 2 is the smallest decile of size and the second smaller decile of book-to-market ratio), and 12 industry portfolios.

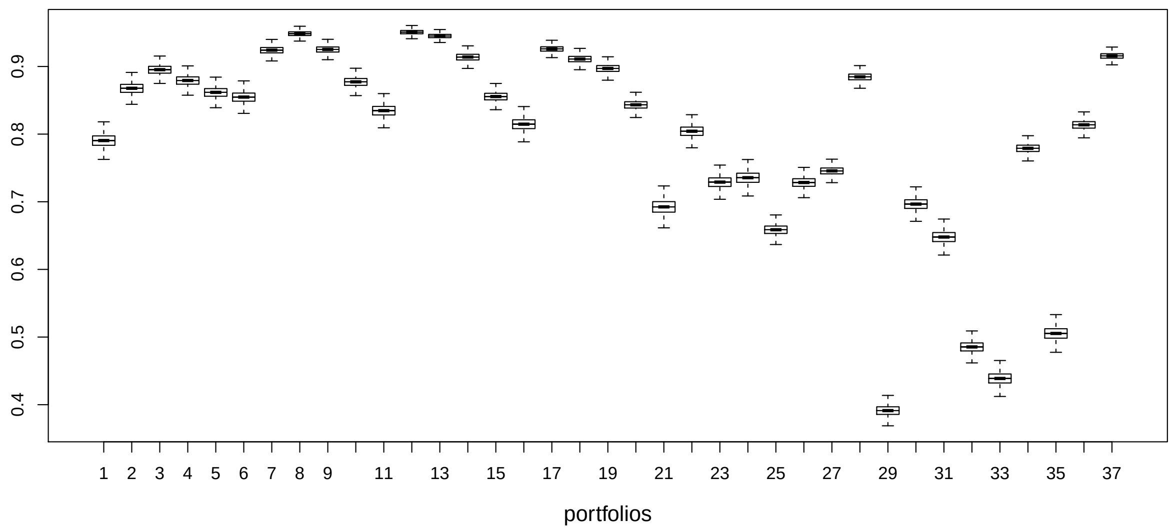

Figure A4: Share of stock portfolios' return variance explained by the f component in the two-factor model.

Box-plots (posterior 95% coverage area) of the percentage of time series variances of individual stock portfolio returns explained by the f component in the two-factor model. Ordering of portfolios: 25 Fama and French (1992) size and book-to-market sorted portfolios (e.g. portfolio 2 is the smallest decile of size and the second smaller decile of book-to-market ratio), and 12 industry portfolios.

A.7 Additional Tables

Table A1: Expected Excess Returns and Consumption Risk, 1967:Q3-2017:Q2

	No intercept			Stocks/bonds-specific intercept				
Horizon S	$R^2_{adj}(\%)$	ϕ	LR-test	$R^2_{adj}(\%)$	α_b	α_s	ϕ	LR-test
(Quarters)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: 9 Bonds and Fama-French 6 portfolios								
0	-30	114	6.3875	-2	-0.0004	0.0152	80	3.7466
		(27.2)	[0.0115]		(0.0003)	(0.0069)	(33.6)	[0.0529]
10	76	38	29.1022	65	-0.0001	0.0089	29	19.8292
		(8.3)	[0.0000]		(0.0003)	(0.0052)	(7.4)	[0.0000]
11	75	53	32.1557	64	0.0001	0.0036	51	21.6484
		(10.7)	[0.0000]		(0.0002)	(0.0047)	(11.8)	[0.0000]
12	87	55	28.7694	83	0.0005	0.0061	43	19.4720
		(12.5)	[0.0000]		(0.0003)	(0.0058)	(13.4)	[0.0000]
Panel B: 9 Bonds and Fama-French 25 portfolios								
0	5	92	5.6714	58	0.0000	0.0220	-40	1.2112
		(17.7)	[0.0172]		(0.0002)	(0.0040)	(16.0)	[0.2711]
10	78	21	10.6559	61	-0.0001	0.0146	21	6.7906
		(3.7)	[0.00112]		(0.0002)	(0.0039)	(4.1)	[0.0092]
11	79	21	9.3890	73	-0.0001	0.0146	17	4.2215
		(3.6)	[0.0022]		(0.0002)	(0.0040)	(3.6)	[0.0399]
12	66	16	6.4810	72	-0.0001	0.0163	10	3.9525
		(3.0)	[0.0109]		(0.0002)	(0.0040)	(2.9)	[0.0468]
Panel C: 9 Bonds, Fama-French 6, and Industry 12 portfolios								
0	43	88	5.1343	-4	-0.0001	0.0141	69	12.2465
		(22.6)	[0.0235]		(0.0002)	(0.0046)	(25.33)	[0.0005]
10	55	16	18.8887	33	0.0003	0.0174	16	23.1406
		(3.9)	[0.0000]		(0.0002)	(0.0039)	(4.4)	[0.0000]
11	50	13	14.5781	40	0.0004	0.0179	14	20.1932
		(3.4)	[0.0000]		(0.0002)	(0.0039)	(4.1)	[0.0000]
12	49	12	13.4594	46	0.0004	0.0174	13	18.0017
		(3.2)	[0.0000]		(0.0002)	(0.0040)	(3.9)	[0.0000]
Panel D: 9 Bonds, Fama-French 25, and Industry 12 portfolios								
0	32	19	3.9034	56	-0.0002	0.0170	20	1.0981
		(16.6)	[0.0482]		(0.0002)	(0.0036)	(18.2)	[0.2947]
10	45	8	10.9676	65	-0.0001	0.0166	7	6.4872
		(2.6)	[0.0009]		(0.0002)	(0.0033)	(2.6)	[0.0109]
11	37	8	9.0100	64	-0.0002	0.0166	6	5.2788
		(2.3)	[0.0026]		(0.0002)	(0.0033)	(2.3)	[0.0216]
12	37	8	8.3904	62	-0.0001	0.0167	5	4.3966
		(2.2)	[0.0037]		(0.0002)	(0.0033)	(2.1)	[0.0361]

The table reports the pricing of excess returns of stocks and bonds, allowing for separate asset class-specific intercepts. Standard errors are reported in parentheses and p-values in brackets. Estimation is done using Empirical Likelihood.