

Consumption*

Svetlana Bryzgalova[†]

Christian Julliard[‡]

May 18, 2022

Abstract

Using information in returns we identify the stochastic process of consumption – the crucial ingredient of most macro-finance models. We find that aggregate consumption reacts over multiple quarters to innovations spanned by financial markets, and this persistent component accounts for 26% of the consumption variation. These innovations drive most of the time series variation of equity returns and are priced in the cross-sections of both bonds and stocks. The data rejects the hypothesis that the stochastic volatility of consumption is proportional to market volatility, and that either of them is priced, posing a novel challenge for consumption-based asset pricing models.

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, Dana Kiku, Martin Lettau, Lars Lochstoer, Dong Lou, Nour Meddahi, João Mergulhão, Tarun Ramadorai, Nick Roussanov, Ken Singleton, Andrea Tamoni, Annette Vissing-Jorgensen, Jessica Wachter, Amir Yaron, Irina Zviadadze, and seminar and conference participants at Berkeley Haas, Boston Fed, Federal Reserve Board of Governors, Geneva Finance Research Institute, Goethe Institute, HSE ICEF, LBS, LSE, Lund University, Northwestern Kellogg, NY Fed, Pennsylvania State University, Stanford GSB, Toulouse School of Economics, UC-San Diego, University of Milan, USC Marshall, UIUC, Wharton, WU Vienna University of Economics and Business, Duke/UNC 2016 Asset Pricing conference, EFA 2016, LEAP conference, MFA 2016 conference, SoFiE 2016 conference, 2017 FIRN Asset Pricing workshop, ESSFM Gerzensee 2019, NBER SI 2019, FARFE 2019, SFS Cavalcade 2019, SED 2021, EEA 2021. We are also thankful to Jiantao Huang for excellent research assistance. Any errors or omissions are the responsibility of the authors.

[†]Department of Finance, London Business School, and CEPR, sbryzgalova@london.edu

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

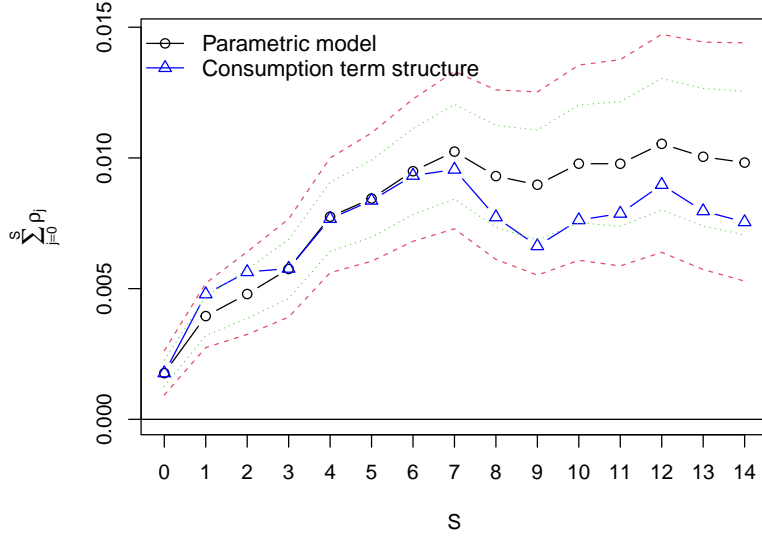
I Introduction

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 the frameworks under consideration: Chen, Dou, and Kogan (2021) refer to this source of model fragility as “dark matter.” As we show, if the stochastic process of consumption was of the type commonly postulated in the literature, standard selection procedures would fail to recover accurately *both* its degree of persistence and its volatility dynamics – the crucial ingredients for asset pricing. This paper attempts to fill the gap and shed light on the dark matter of consumption-based models.

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 in 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, we 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 establishes a new set of facts regarding the dynamics of stocks, bonds, and consumption growth. Conceptually, the method is simple: we estimate robustly the impulse response of consumption to shocks, hence we accurately identify its conditional mean and persistence. This in turn allows us to recover correctly the volatility dynamics, since the latter, as we show, can be consistently estimated *only* if the conditional consumption mean

¹See, 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.

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



The graph presents posterior means of the cumulative response function of consumption growth (black line with circles), along with the centered posterior 90% (dashed red lines) and 68% (dotted green lines) coverage regions. The blue line with triangles denotes the first principal component of $cov(r_{i,t}^{ex}, \Delta c_{t,t+1+s})$. These shocks account for 26% of consumption growth time series variation (quarterly data, 1961:Q3-2017:Q2).

and its predictability are properly captured.

Our first result is summarized in Figure 1, which 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 to three years to be fully reflected in consumption, thus producing substantial (but not excessive) predictability. The economic magnitude of this effect is large: One standard deviation shock implies a cumulated response of consumption of about 1% over the next 2 to 4 years. Most importantly, these shocks 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 the best case scenario, in leading macro-finance models.² Furthermore, these very same shocks play a

²For 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).

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.

All of our findings allow for very general stochastic volatility processes in both consumption and returns. Compared to the previous literature, we rely on general and flexible specifications, richer asset return data, and much less restrictive priors. As a result, we uncover a new fundamental challenge for theoretical models that obtain equilibrium time-varying risk premia by assuming a common stochastic volatility in consumption and returns: The data not only strongly suggests that these volatilities are distinct stochastic processes but also that they do not seem to drive time variation in excess returns.

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 consumption 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

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

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 very high-order 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. To draw the analogy with the Wold's time series decomposition, we estimate the moving average representation of consumption (or equivalently, its impulse-response) and identify its innovations with the help of asset returns. That is, instead of *ex ante* postulating, say, an AR(1) persistent component in consumption, we accurately recover its conditional mean without relying on fragile (and, as we show, untestable within popular macro-finance models) assumptions.

Importantly, as we show, unlike traditional information criteria for specification selection, our approach captures correctly the conditional mean and persistence of consumption even in the type of models that have been postulated in the literature.

The key identifying restriction we rely on is that innovations in consumption spanned by financial markets should be reflected in returns *at the same time* as they occur. This assumption is not only theoretically sound, since equilibrium prices are jump variables, but is also supported by a rich set of reduced form empirical evidence in the previous literature (that we further extended in the Appendix).

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 in its residuals. We further investigate the role of stochastic volatility by estimating a flexible state-space model that allows for (potentially) distinct and priced volatility shocks in both consumption and returns. We find very little evidence that the volatility of consumption mean shocks is the same as the volatility of financial markets and that either of these volatilities affects the time series of excess returns. Our finding is driven by the robust

approach that we use in modeling the conditional expectation of consumption growth, which uses not only information in past consumption but also asset returns data. We show that the common practice of disregarding the role of returns in driving the conditional mean of consumption generates spurious evidence of volatility clustering, as well as a spurious link between financial returns and consumption volatility.

Our second empirical 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 the term structure 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 our cross-sectional asset pricing findings by also estimating a very broad class of consumption-based pricing kernels, via Empirical Likelihood (see, e.g., Owen (2001)), *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 a 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 and show that misspecification of the consumption mean process leads to spurious evidence of volatility clustering; *c)* stock returns significantly load on these common shocks (that capture 36%–95% of their time series variation); *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)* there is very little evidence that the volatility of consumption shocks is proportional to the volatility of financial markets; *f)*

none of these volatilities drives excess returns; *g*) most of the time series variation in bonds is captured by a single factor, independent from both consumption and stock returns, which does not command a risk premium (i.e., it is unpriced); *h*) 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 average returns⁴; *i*) low frequency consumption explains up to 92% of the cross-sectional variation of bond excess returns.

The remainder of the paper is organized as follows. The next section reviews the most closely related literature. Section III shows that the type of consumption process that are often used in the macro-finance literature are unlikely to be correctly identified using canonical model specification selection. Our state-space formulation is introduced in Section IV where we also show that it can accurately recover the consumption dynamics of popular models. 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 Closely 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 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, that is, when consumption (current and expected) is low. Therefore, by its very nature, our paper is closely linked to consumption-based empirical asset pricing.⁵ Within this literature, the

⁴In our baseline specification we consider a cross-section of 46 assets given by 12 industry portfolios, 25 size and book-to-market portfolios, and nine bond portfolios, but the results appear robust to alternative specifications.

⁵E.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), and Bansal, Kiku, and Yaron (2012). See also Ludvigson (2012) for an excellent review.

underlying stochastic process of consumption 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. Different from these, we use a much broader information set, and a more flexible parametrization, for identifying shocks to the mean before assessing whether volatility is priced, and explicitly estimate the volatility processes of both consumption and asset returns (instead of proxying the former with the latter). 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. Tédongap (2014) 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. We show that a consumption growth mean misspecification leads to spurious GARCH dynamics and creates an artificial link between consumption volatility and returns.

Bandi and Tamoni (2015) and Boons and Tamoni (2015) decompose the process for consumption growth into different, frequency specific, components. They find that only the shock with a half-life of up to four years plays a significant role in explaining the cross-section of returns, thus supporting the importance of business cycle fluctuations in determining the risk premium. In contrast, Zviadadze (2021) develops a methodology for testing structural asset pricing models using the term structure of equity risk, and finds that, among the models considered in the paper, only the formulation of Drechsler and Yaron (2011), that has AR(1) consumption mean dynamics and multiple volatility shocks, generates a realistic

term structure of risk. 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 the extremely low frequency ones, and find no significant evidence of time-varying consumption volatility. They attribute the latter to a lack of testing power within their VAR-based framework. Instead, *i*) we uncover the fundamental role of persistent shocks to the consumption conditional mean, *ii*) document that asset returns and consumption have very distinct volatility processes, and *iii*) present sharp empirical evidence against consumption volatility driving excess returns.

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 parameters. The latter proposes a Bayesian strategy for identifying the deep parameters of the model using a mixed frequency setting. Instead, we do not take an *ex ante* stand on the preferences nor on the speed or the pattern (e.g., an AR(1)) of the consumption dynamic: We work with the MA representation of the consumption process and relax both functional and prior restrictions. Furthermore, we explicitly allow past returns to carry information about the consumption process. This identifies directly the underlying shocks, and the consumption response to them, through the joint behavior 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, which 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 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).⁶ Their paper, as ours, 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 and, instead, using a broad cross-section of asset returns could be much more informative. 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.⁷ These findings lead us to include, among other assets, size and value sorted portfolios of stocks and a cross-section of bond returns in our empirical analysis.

III The Challenge of Consumption Persistency

We start by showing that the type of consumption processes assumed in most macro-finance model is hard to detect, and hence test, in samples of the same size as the historical ones.

Let's consider as a working example the so-called long run risk process of Bansal and Yaron (2004). In this formulation log-consumption growth, $\Delta c_{t,t+1}$, contains a persistent AR(1) component, x_t , which is crucial to rationalise unconditional risk premia and other moments of the historical data. That is:

$$\Delta c_{t,t+1} = \mu + x_t + \sigma_t \eta_{t+1}, \quad (1)$$

$$x_{t+1} = \rho x_t + \phi_e \sigma_t e_{t+1}, \quad (2)$$

where $\eta_t, e_t, \sim iid \mathcal{N}(0, 1)$ and σ_t , depending on the calibration, is either a constant or a

⁶This approach is similar to the one in Constantinides and Ghosh (2011), which shows that, in a linearized long-run risk framework with AR dynamics, the consumption mean and volatility processes can be inverted from the market price-dividend ratio and the risk-free rate.

⁷Also, 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, but forecasts future expansion, i.e., it seems to incorporate good news about future consumption.

Stochastic Volatility (SV) process. How much of the time series variation of consumption is driven by the predictable component x_t varies in the literature. For instance, the conditional mean generates 4.5% of the consumption variance in Bansal and Yaron (2004) and, in what is to the best of our knowledge the most extreme example, 12% in the calibration of Hansen, Heaton, Lee, and Roussanov (2007).

But would a researcher be likely to detect this persistency, and identify its functional form, in samples of the same size as the historical one? We address this question formally by using the Hansen, Heaton, Lee, and Roussanov (2007) calibration to generate monthly series of consumption of the same length as the postwar sample we use in our empirical analysis (214 quarters), and formally perform ARIMA model selection for the consumption process. We consider consumption aggregated to the quarterly frequency – as in real world data – as well as the monthly observations that are normally not available for such long samples. Details of the simulation design are reported in Appendix A.1.

Table 1 reports the frequency of specification selected according to the Bayesian and Akaike Information Criteria (BIA and AIC). Strikingly, with both quarterly and monthly sampling the most often selected specification implies no predictability at all (columns (A) and (B)). And note that this result arises using the calibration of the long run risk process that has the *largest* predictability of consumption. Furthermore, the table shows that even if a researcher were to observe the conditional mean of consumption growth directly (columns (C) and (D)), canonical specification selection would fail to identify the true mean process with more than 92% probability.

Note that misidentifying the conditional mean process has also important consequences for the recovery of the volatility process. For instance, suppose that the process in equations (1)-(2) is a constant volatility one, that is, $\sigma_t = \sigma \ \forall t$. In this case, if a researcher were to conclude that there is no predictability in consumption growth (as Table 1 suggests as the most likely outcome), she would then find evidence of time-varying volatility in consumption, since its squared forecast errors would be positively autocorrelated (with the j -th autocorrelation

Table 1: ARIMA model selection of Long Run Risk consumption process

(A) BIC: $\Delta c_{t,t+1}$				(B) AIC: $\Delta c_{t,t+1}$				(C) BIC: x_{t+1}				(D) AIC: x_{t+1}			
Panel A: Quarterly frequency (214 observations)															
p	d	q	freq	p	d	q	freq	p	d	q	freq	p	d	q	freq
0	0	0	53.7%	0	0	0	25.1%	0	1	0	50.0%	0	1	1	21.5%
0	1	1	24.8%	0	1	1	14.6%	0	1	1	15.6%	0	1	0	13.7%
1	0	1	5.9%	1	0	1	13.8%	1	0	1	12.2%	1	0	1	11.6%
1	1	1	5.1%	0	0	1	6.3%	2	0	0	10.1%	2	0	0	7.6%
1	0	0	3.9 %	1	1	1	5.3%	1	1	0	5.1%	1	1	0	7.5%
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
Panel B: Monthly frequency (642 observations)															
0	0	0	55.1%	0	0	0	27.8%	0	1	0	92.3%	0	1	0	70.0%
0	1	1	31.3%	0	1	1	24.0%	1	0	0	7.2%	2	1	2	4.9%
1	0	1	4.7%	1	0	1	13.4%	1	1	0	0.3%	1	0	0	4.7%
5	1	0	1.4%	0	0	1	6.5%	1	1	1	0.1%	1	1	0	3.8%
0	0	1	0.9%	0	1	2	3.2%	2	0	0	0.1%	1	1	1	3.6%
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

Empirical frequencies of ARIMA(p,d,q) models selected by Bayesian information criterion (BIC) and Akaike information criterion (AIC) in 1,000 simulations of the Hansen, Heaton, Lee, and Roussanov (2007) long run risk specification for the consumption growth process. We only list the top five most frequent models.

proportional to ρ^{2j}). More generally, missing the true degree of predictability in the conditional mean process mechanically delivers spurious (if the true process is homoskedastic) or biased (if the true process has time varying volatility) evidence of volatility clustering.

If the true process for consumption is similar to what we rely on in macro-finance models, the above considerations imply three crucial requirements for reliable inference on the data generating process of consumption.

First, we need an estimator of the conditional mean of consumption that can capture the true degree of predictability *without* relying on fragile specification selection. This is crucial in particular if one wants to make statements about stochastic volatility in consumption (its existence, magnitude, and properties).

Second, we should not achieve identification via arbitrary (and often non-testable within a model) parametric restrictions for the consumption process, e.g., an arbitrary AR(1) process for its persistent component, or the often employed proportionality restriction between consumption and return stochastic volatilities. Given their crucial role for model predictions,

these restrictions should be tested whenever possible.

Third, ideally we would like a method that allows to learn about the consumption process by leveraging information in other variables that should be adapted to the same type of shocks that drive consumption (e.g., wealth shocks).

The above three properties are exactly what the empirical formulation that we present in the next section delivers.

IV A Model of Consumption and Returns Dynamics

In macro finance models the stochastic discount factor is typically a function of consumption growth and potentially additional variable (such as, e.g., habits, returns or wealth, leisure, leverage ratios, and aggregation weights in heterogenous agent models). Furthermore, consumption and returns both contributes to the intertemporal budget constraint. This implies that, with exclusion of knife-edge examples, there is a set of shocks with respect to which both returns and consumption growth are adapted processes.

The reason for this general feature of equilibrium models is that households react to shocks (e.g., wealth, income, or beliefs) by adjusting *both* consumption and investment decisions. Hence in principle one could leverage the information in equilibrium asset returns to learn about the shocks that drive consumption and the form of its stochastic process. This simple insight lies at the core of our empirical strategy.

In particular, 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 serially uncorrelated disturbance, w_c with variance σ_c , that is independent from financial market shocks, and a potentially autocorrelated process – a persistent component – that (potentially) depends on the current and past stocks to asset returns. For expositional simplicity, we start by focusing on a setting with constant volatilities, and generalise our framework to incorporate stochastic volatilities in all the shocks.

To avoid taking an ex ante stand on the particular time series structure of the persistent

component (or its absence), we work with its (potentially infinite) moving average representation. Obviously, by virtue of the Wold’s representation theorem, an $MA(\infty)$ modelling for the persistent component would capture the true data generating process and avoid the fallacy of model selection outlined in the previous section. Since the MA coefficients in the Wold representation are square summable, any finite order covariance stationary ARIMA can be approximated with a high order MA process, with the accuracy increasing with the MA order. Therefore, we model the (log) consumption growth process as

$$\Delta c_{t-1,t} = \mu_c + \underbrace{\sum_{j=0}^{\bar{S}} \rho_j f_{t-j}}_{MA(\bar{S})} + w_t^c, \quad (3)$$

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 normalized to have unit variance) is the fundamental innovation upon which all asset returns load contemporaneously, that is,

$$\mathbf{r}_t^e = \underbrace{\boldsymbol{\mu}_r}_{N \times 1} + \underbrace{\boldsymbol{\rho}^r}_{N \times 1} f_t + \underbrace{\mathbf{w}_t^r}_{N \times 1}, \quad (4)$$

where \mathbf{r}^e denotes a vector of log excess returns, $\boldsymbol{\mu}_r$ is a vector of expected values, $\boldsymbol{\rho}^r$ contains the asset specific loadings on the common risk factor, and \mathbf{w}_t^r is a vector of white noise shocks with diagonal⁸ covariance matrix Σ_r , which are meant to capture asset specific idiosyncratic shocks. In the above returns are modelled as reacting contemporaneously and fully to the f shocks since in equilibrium models prices are “jump” variables.

Note that the joint dynamics of consumption and returns postulated in equations (3)–(8) is consistent with the extensive preliminary evidence that we report in Appendix A.3. Therein we show, based on predictive regressions and Structural-VAR estimation, that: *i*) the consumption growth process shows significant predictability over multiple years; *ii*) this

⁸The diagonality assumption can be relaxed, as explained in Appendix A.4, and later we will allow this covariance to embed both common and idiosyncratic stochastic volatilities.

predictability is better captured by lagged returns than lagged consumption; *iii*) this predictability is generated by a shock common to returns and consumption to which the former reacts fully contemporaneously, while the latter react little contemporaneously, but substantially in the following quarters.

Obviously, we cannot feasibly use an infinite number of lags in the MA component of consumption growth. At the same time, since the MA representation of the persistent component is meant to only approximate the true latent dynamic, model selection would not be appropriate and, as outlined in Section III, possibly unreliable. Note also that employing an excessively high number of lags in the MA does *not* affect the consistency of the estimation, but only its efficiency. Consequently in our empirical analysis we rely on a conservative approach and: *i*) in the baseline estimation we use up to 3.5 years of lagged quarterly shocks in the conditional mean of consumption since this is almost twice as much the degree of persistency that we find in the preliminary evidence in Appendix A.3 and it covers the span of predictability uncovered in the previous literature (see, e.g., Liew and Vassalou (2000), Parker and Julliard (2005), Bandi and Tamoni (2015)); *ii*) we show in the simulation in Section IV.2 that the approach is very robust to the precise number of included MA lags; *iii*) we confirm that all of our empirical results are virtually identical even when including up to 12.5 years of past quarterly shocks (i.e., 50 lags).⁹

The dynamic system in equations (3)–(4) can be reformulated as a state-space model, and Bayesian posterior inference can be conducted to estimate both the unknown parameters $(\mu_c, \boldsymbol{\mu}_r, \{\rho_j\}_{j=0}^{\bar{S}}, \boldsymbol{\rho}^r, \sigma_c^2, \Sigma_r)$ 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.4.

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 (3), the shocks would be underidentified, making it difficult to give any particular rotation

⁹Additional results for alternative lag lengths are available upon request.

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 the Uhlig (2005) identification in Structural-VARs. In particular, our approach is akin to constructing the Generalised Impulse Response Function of consumption and financial markets building upon the insights of Koop, Pesaran, and Potter (1996) and Pesaran and Shin (1998).

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 as¹⁰

$$\frac{\partial \mathbb{E} [\Delta c_{t-1,t+S}]}{\partial f_t} = \sum_{j=0}^S \rho_j, \quad (5)$$

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

Note that equations (3)–(4) 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 characterized by the loadings of the dynamic system on the factor f_t , as follows:

$$Cov(\Delta c_{t-1,t+S}; \mathbf{r}_t^e) \equiv \sum_{j=0}^S \rho_j \boldsymbol{\rho}^r. \quad (6)$$

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

¹⁰This immediately follows from the observation that, since $\Delta c_{t-1,t+S} \equiv \sum_{j=0}^S \Delta c_{t-1+j,t+j} \equiv \ln(C_{t+S}/C_{t-1})$, we have $[\Delta c_{t-1,t}, \Delta c_{t-1,t+1}, \dots, \Delta c_{t-1,t+S}]' \equiv \Gamma [\Delta c_{t-1,t}, \Delta c_{t,t+1}, \dots, \Delta c_{t-1+S,t+S}]'$, where Γ is a lower triangular square matrix of ones (of dimension S).

The particular one-factor *term structure* in the exposures of asset returns to consumption growth in equation (6) also provides an alternative method, which does not rely on the parametric likelihood and directly stems from the empirical covariances, for recovering moving average component in consumption. We discuss this further in Section V.2.

Finally, note that the formulation in equations (3)–(4) 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

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + \sum_{j=0}^{\bar{S}} \theta_j g_{t-j} + w_t^c; \quad \text{and} \quad (7)$$

$$\mathbf{r}_t^e = \underbrace{\boldsymbol{\mu}_r}_{N \times 1} + \underbrace{\boldsymbol{\rho}^r}_{N \times 1} f_t + \left[\underbrace{\boldsymbol{\theta}^b}_{N_b \times 1}, \underbrace{\mathbf{0}'_{N-N_b}}_{N \times 1} \right]' g_t + \underbrace{\mathbf{w}_t^r}_{N \times 1}, \quad (8)$$

where N_b is the number of bonds and they are ordered first in the vector \mathbf{r}_t^e , and $\boldsymbol{\theta}^b \in \mathbb{R}^{N_b}$ 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

$$Cov(\Delta c_{t-1,t+S}; \mathbf{r}_t^e) \equiv \sum_{j=0}^S \rho_j \boldsymbol{\rho}^r + \left[\boldsymbol{\theta}^b, \mathbf{0}'_{N-N_b} \right]' \sum_{j=0}^S \theta_j. \quad (9)$$

Two observations regarding our parametric framework deserve mentioning. First, both the one-factor (equations (3)–(4)) and two-factor (equations (7)–(8)) models are strongly overidentified. Second, estimation of the model assuming constant volatility is generally consistent even in the presence of time-varying volatility in the true processes; hence, our formulation is robust along this dimension. We address this issue formally in Section V.1.2, where we show that misspecification of the consumption mean process leads to spurious evidence of consumption volatility clustering, and in Section V.3, where we generalize our state-space model to allow for stochastic volatilities affecting all the shocks in the system.

IV.1 Estimation

We can rewrite the dynamic model in equations (3)–(4) in state-space form, assuming Gaussian innovations, as

$$\mathbf{z}_t = \mathbf{F}\mathbf{z}_{t-1} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}_{\bar{S}+1}; \Psi), \quad (10)$$

$$\mathbf{y}_t = \boldsymbol{\mu} + \mathbf{H}\mathbf{z}_t + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}_{N+1}, \Sigma). \quad (11)$$

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

$$\Psi := \underbrace{\begin{bmatrix} 1 & \mathbf{0}_{\bar{S}}' \\ \mathbf{0}_{\bar{S}} & \mathbf{0}_{\bar{S} \times \bar{S}} \end{bmatrix}}_{(\bar{S}+1) \times (\bar{S}+1)}, \quad \mathbf{F} := \underbrace{\begin{bmatrix} \mathbf{0}_{\bar{S}}' & 0 \\ I_{\bar{S}} & \mathbf{0}_{\bar{S}} \end{bmatrix}}_{(\bar{S}+1) \times (\bar{S}+1)}, \quad \Sigma := \underbrace{\begin{bmatrix} \sigma_c^2 & \mathbf{0}_N' \\ \mathbf{0}_N & \Sigma_r \end{bmatrix}}_{(N+1) \times (N+1)}, \quad \mathbf{H} := \underbrace{\begin{bmatrix} \rho_0 & \rho_1 & \dots & \rho_{\bar{S}} \\ \boldsymbol{\rho}^r & \mathbf{0}_N & \dots & \mathbf{0}_N \end{bmatrix}}_{(N+1) \times (\bar{S}+1)}, \quad (12)$$

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

Similarly, the dynamic system in equations (7)–(8) can be represented in the state-space form (10)–(11) with $\mathbf{z}_t := [f_t, \dots, f_{t-\bar{S}}, g_t, \dots, g_{t-\bar{S}}]'$; $\mathbf{v}_t := [f_t, \mathbf{0}_{\bar{S}}', g_t, \mathbf{0}_{\bar{S}}']' \sim \mathcal{N}(\mathbf{0}_{\bar{S}+1}; \Psi)$; Ψ and \mathbf{F} being block diagonal with blocks repeated twice and given, respectively, by the first two matrices in equation (12); and with space equation matrix of coefficients given by

$$\mathbf{H} := \underbrace{\begin{bmatrix} \rho_0 & \dots & \dots & \rho_{\bar{S}} & \theta_0 & \dots & \dots & \theta_{\bar{S}} \\ \rho_1^r & 0 & \dots & 0 & \theta_1^b & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \rho_{N_b}^r & 0 & \dots & 0 & \theta_{N_b}^b & 0 & \dots & 0 \\ \dots & 0 & \dots & 0 & 0 & 0 & \dots & 0 \\ \rho_N^r & 0 & \dots & 0 & \dots & 0 & \dots & 0 \end{bmatrix}}_{(N+1) \times 2(\bar{S}+1)}. \quad (13)$$

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

$$\mathbf{y}_t | \mathcal{I}_{t-1}, \boldsymbol{\mu}, \mathbf{H}, \Psi, \Sigma, \mathbf{z}_t \sim \mathcal{N}(\boldsymbol{\mu} + \mathbf{H}\mathbf{z}_t; \Sigma), \quad (14)$$

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,¹¹ the specification above for the dynamics of consumption and asset returns implies, in the presence of only one latent factor (f_t) common to both assets and consumption,

$$\mathbb{E}[\mathbf{R}_t^e] = \alpha + \left(\sum_{j=0}^S \rho_j \boldsymbol{\rho}^r \right) \lambda_f, \quad (15)$$

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 expression above 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 (7)–(8), the implied cross-sectional model of returns is

$$\mathbb{E}[\mathbf{R}_t^e] = \alpha + \left(\sum_{j=0}^S \rho_j \boldsymbol{\rho}^r \right) \lambda_f + \left[\boldsymbol{\theta}^b, \mathbf{0}_{N-N_b}' \right]' \sum_{j=0}^S \theta_j \lambda_g, \quad (16)$$

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

¹¹See equation (28) and discussion therein.

Equation (15) (and similarly equation (16)), 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.4.

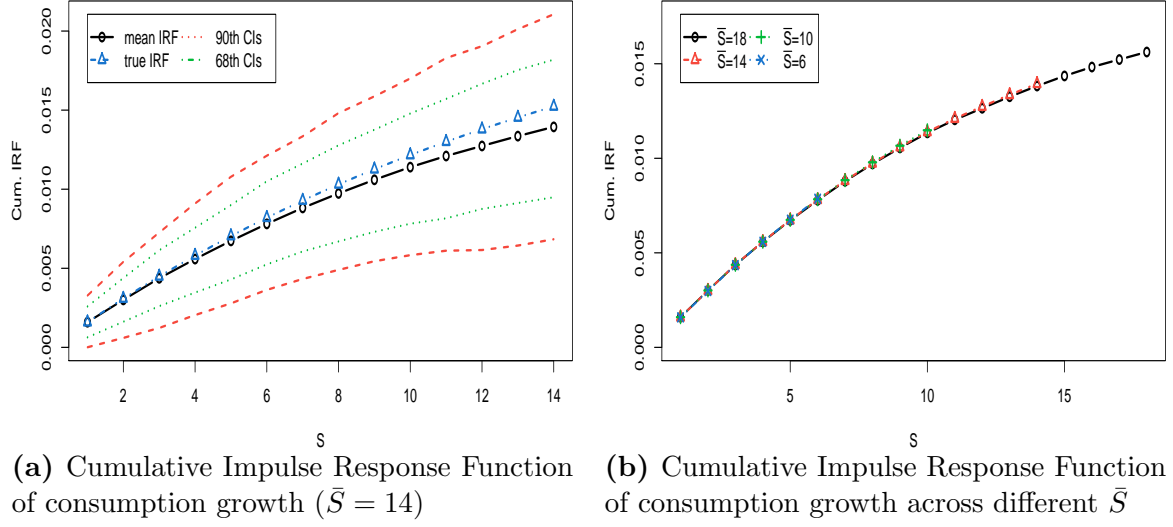
Furthermore, we confirm the main empirical results obtained with the above state-space framework, for both the impulse response function of consumption and the cross-sectional pricing patterns, using a principal component based approach and an Empirical Likelihood-type estimation. The details of these nonparametric procedures are discussed in Sections V.2 and A.5, respectively.

IV.2 The challenge of consumption persistency redux

A natural question is whether our state-space representation of consumption and returns in equation (3)–(4) is able to recover the consumption process when, as shown in Section III, standard methods fail. For doing so, we use once more the simulated Hansen, Heaton, Lee, and Roussanov (2007) long run risk consumption process of Section III, calibrate the asset loadings on f to the values observed in the historical data, and apply our state-space estimation to it. The results are summarized in Figure 2.

As shown in panel (a) of Figure 2, our state-space model with a long consumption MA does an outstanding job in capturing the effect of time t shocks on subsequent consumption growth: the difference between the mean estimate (across simulated time series) and the true is extremely small, and the variability across simulated samples is also quite small. If anything, we observe a small attenuation bias in the long run, implying that our approach is

Figure 2: Cumulative response function of consumption growth to one standard deviation shock to the conditional mean of consumption growth in 1,000 simulations



Panel (a) plots the mean, 5th, 14th, 84th, 95th percentiles of cumulative IRF in 1,000 simulations. The model is estimated under the assumption that $\bar{S} = 14$. Panel (b) plots the average cumulative impulse response function (IRF) of consumption growth to one standard deviation shock to the conditional mean of consumption growth, in 1,000 simulations, using MA representation with different maximum number of lags: $\bar{S} = 6, 10, 14, 18$.

in fact conservative in estimating the true extent of consumption predictability. That is, the conditional mean of consumption is accurately captured by the state-space representation method. In addition, Figure A5 in Appendix A.7 shows that our estimation recovers very precisely the loading of asset returns on the shocks to the conditional mean of consumption.

Furthermore, as shown in panel (b) of Figure 2, the IRF estimates are almost identical using different orders for the MA component. The only difference is that with longer MA we can trace the effect further in the future. Hence, a finite order MA gives a conservative measure of the long run effect of time t shocks.

Recall also that there is a one-to-one mapping between IRFs (or equivalently, MA representation) and variance decomposition. Hence, our accurate IRF estimates imply that we can perform, as we do in later sections, accurate variance decomposition for both consumption and returns.

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 (3)–(4) (as well as equations (7)–(8)) 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 parametrizes 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 A.5 we rely on a standard nonparametric technique, Empirical Likelihood,¹² 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 (3)–(4) 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 $\bar{S} = 14$ quarters for a number of reasons.

First, equation (3) implies a certain autocorrelation structure of the nondurable con-

¹²For an excellent review and comparison with other moment-based estimation approaches, see Kitamura (2006).

consumption growth through the combination of the common factor lags and loadings. Hence, the value of \bar{S} should be high enough to capture most of the time series autocorrelation in the consumption growth. Figures A1–A3 of Appendix A.3 shows that most of the dependence is conservatively captured by the first 14 lags; hence, this value is a natural choice for the lag truncation. Second, an extensive literature (see, e.g., Liew and Vassalou (2000), Parker and Julliard (2005), Bandi and Tamoni (2015)) shows that such lag length more than covers the predictability in consumption. Third, our results remain robust to including many more lags. For robustness, we have experimented with up to $\bar{S} = 50$, and the results (available upon request) remain virtually identical. And finally, note that using more lags than potentially needed does not affect the consistency of our estimates but only their efficiency (this is confirmed by the fact that using many more lags our point estimates are identical, yet with larger confidence bands). Hence, instead of trying to select a lower dimensional MA specification, we take a robust stance and use a long lag window.¹³ This in turn allows us, in Section V.3, to compare classes of macro-finance models rather than specific standalone specifications.

V.1.1 The consumption mean process

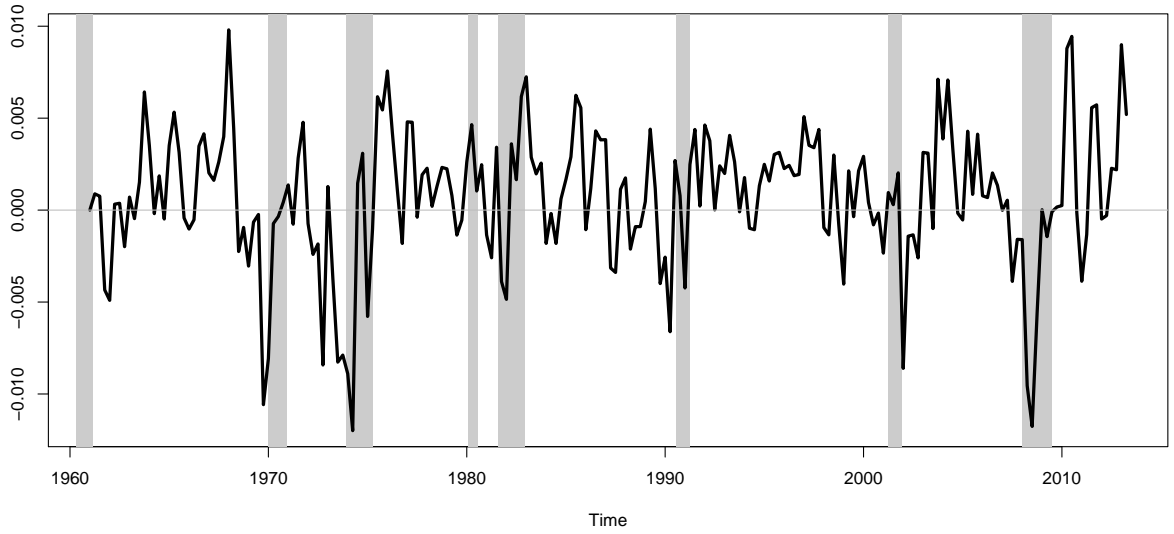
The consumption growth representation in equations (3) and (7) is similar to the moving average decomposition and allows us to infer the dynamics of multi-period consumption growth ($\Delta c_{t,t+1+S}$) in response to a common and/or a bond-specific shock. Figure 1 in the Introduction depicts the (cumulated) 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

¹³An alternative robust approach would be to compute Bayesian posterior probabilities for a wide range of lag lengths and then perform a Bayesian Model Averaging estimate (see, e.g., Raftery, Madigan, and Hoeting (1997), and Hoeting, Madigan, Raftery, and Volinsky (1999)) of the conditional mean of consumption, therefore obtaining estimates that are statistically optimal along several dimensions (see the discussion in Bryzgalova, Huang, and Julliard (2021)).

pronounced, leveling off at around $S = 10$. At this horizon, the effect is economically very large: The cumulative response of consumption growth to a one standard deviation shock is about 1%.

Note that allowing for a bond-specific latent factor (equations (7)–(8)) leaves the consumption loadings on f shocks virtually unchanged, and consumption does not load significantly on the bond-specific factor g (see, respectively, Figures A6 and A7 in Appendix A.7).

Figure 3: Moving average component of consumption growth



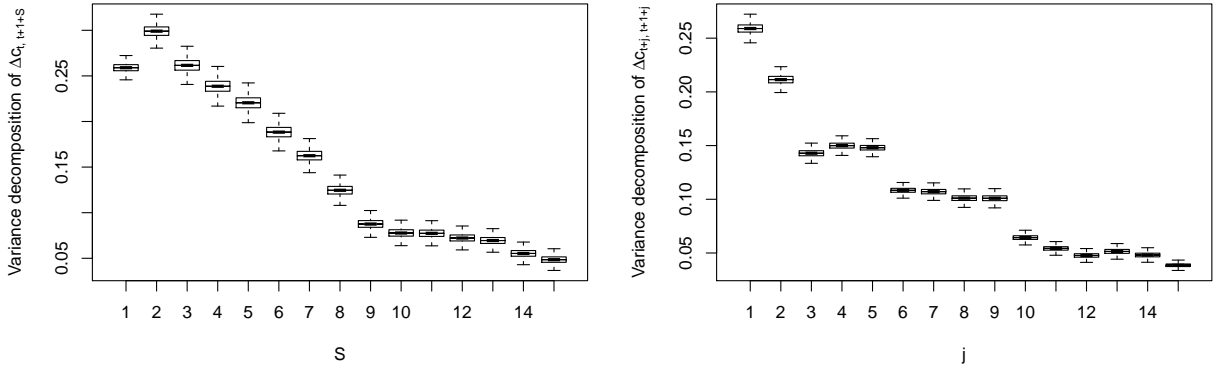
Posterior mean of the moving average f_t component of consumption growth. Grey areas denote NBER recessions.

Figure 3 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: For example, this is five to six 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 4 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 toward 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 three to four years from now. As shown in Figure A7 of the Appendix, adding a bond-specific factor has a minimum impact on the explanatory power of the model for future consumption growth.

Figure 4: Share of consumption growth variance driven by its moving average component.



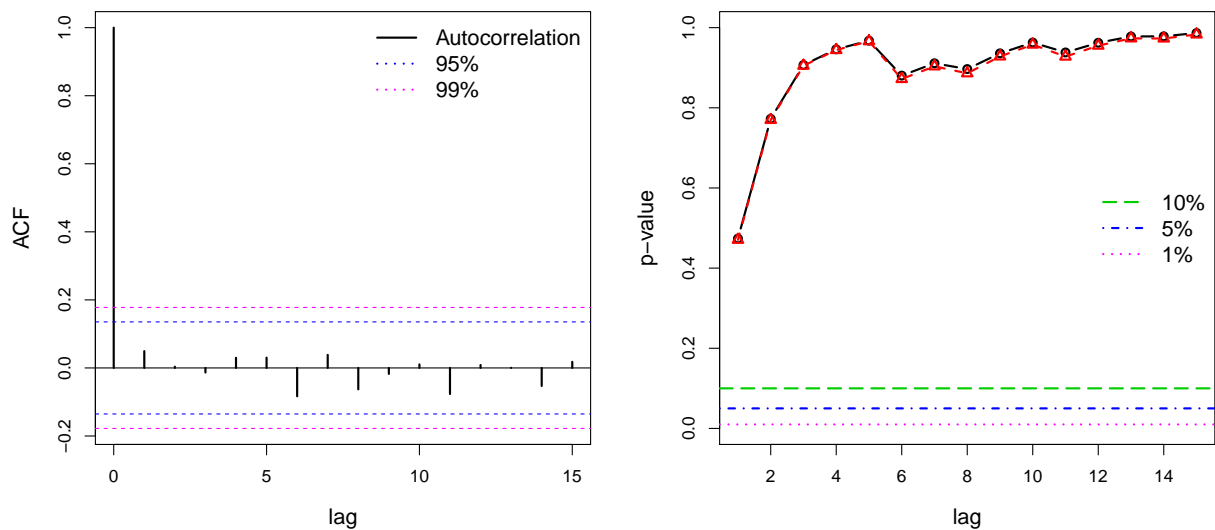
Box-plots (posterior 95% coverage area) of the percentage of time series variances of consumption growth explained by the MA component. Left panel: Cumulated consumption growth $\Delta c_{t,t+s}$. Right panel: One period consumption growth $\Delta c_{t-1+j,t+j}$.

V.1.2 Clustering and predictability of consumption volatility

Note that the estimation of the conditional mean of consumption growth in equations (3) and (7) 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 $f_{\tau \leq t}$ and $g_{\tau \leq t}$.¹⁴

¹⁴That 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 5: Autocorrelation structure of consumption growth squared forecast errors.

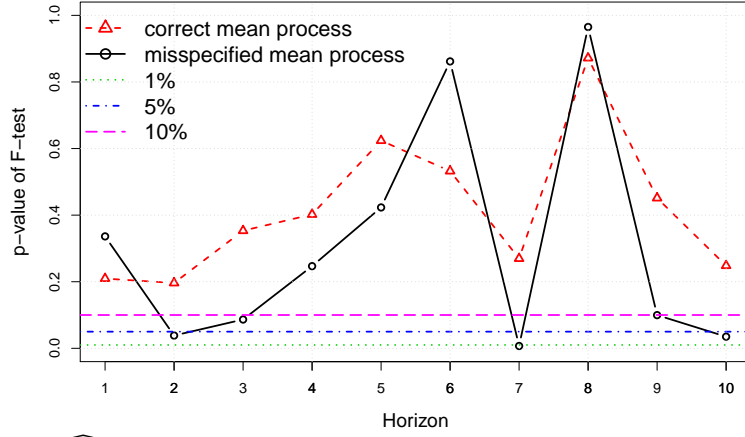


Left panel: Autocorrelation function of $\widehat{Var}_t(\Delta c_{t,t+1})$ with 95% and 99% confidence bands. Right panel: p -values of Ljung and Box (1978) (red triangles) and Box and Pierce (1970) (black circles) tests.

Figure 5 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, in line with the previous literature (see, e.g., Liew and Vassalou (2000)), predictability of the first conditional moment of consumption growth in Appendix A.3. The p -values of the F -tests for these predictability regressions are depicted by the dashed red line with triangles in Figure 6. The p -values (which 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 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

Figure 6: Predictability of consumption squared forecast errors.



Predictive regressions of $\widehat{Var}_{t+h}(\Delta c_{t+h,t+h+1})$ 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.

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 $\widehat{Var}_t(\Delta c_{t,t+1})$ would be mechanically different from zero. For instance, the k -th autocorrelation of $(\Delta c_{t,t+1} - \mu_c)^2$, for $k \leq \bar{S}$, is proportional to

$$Cov \left(\left(\sum_{j=k}^{\bar{S}} \rho_j f_{t-j} + \theta_j g_{t-j} \right)^2, \left(\sum_{j=k}^{\bar{S}} \rho_{j-k} f_{t-j} + \theta_{j-k} g_{t-j} \right)^2 \right) \neq 0. \quad (17)$$

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

First, we run again predictability regressions of $\widehat{Var}_{t+h}(\Delta c_{t+h,t+h+1})$ on the first eight principal components of asset returns (the same predictive variables as in Section A.3) but, as a misspecification benchmark, we construct this measures assuming a constant conditional mean for consumption growth. Summary statistics for these regressions are depicted by the black continuous line with circles in Figure 6. 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. Instead, with the robust MA specification of the mean process, there is no evidence of predictability in the volatility proxy of consumption growth (dashed red line with triangles of Figure 6).¹⁵

Table 2: Estimates of GARCH (1,1) for the innovations of different models

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

	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: $\mu_t = \mu_0 + \sum_{i=1}^S \rho_i f_{t+1-i}$			
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 the conditional mean. Models are estimated using QMLE, and robust t-statistics are constructed using Newey and West (1987) standard errors.

Second, in Table 2 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. This formulation (AR(1)-GARCH(1,1)), is exactly the one that has been often used in the literature (see, e.g., Bansal, Khatchatrian, and Yaron (2005) and Tédongap (2014)) to provide evidence of time varying volatility in consumption.¹⁶ In this case, there is statistically significant evidence of volatility clustering (with a half-life of about 5–6 quarters). However, as shown in Appendix A.3, lagged consumption alone does not capture the full extent of consumption predictability. Therefore, in Panel B we add

¹⁵This is in line with the evidence of Dew-Becker and Giglio (2016, Appendix D) who, proxying consumption volatility with the realized vol of the S&P500 index (as often assumed in the prior literature), find no predictability.

¹⁶We provide an extensive analysis of Stochastic Volatility processes in consumption and asset returns, of the type also commonly used in the literature, in Section V.3.

to the AR(1), as drivers of the conditional mean, the same principal components of asset returns that predict consumption (see Appendix A.3). The resulting change is striking: Once we better control for consumption predictability, the evidence in 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 (as shown in Section IV.2), and lead to very similar implications for volatility – something clearly contradicted by the table.

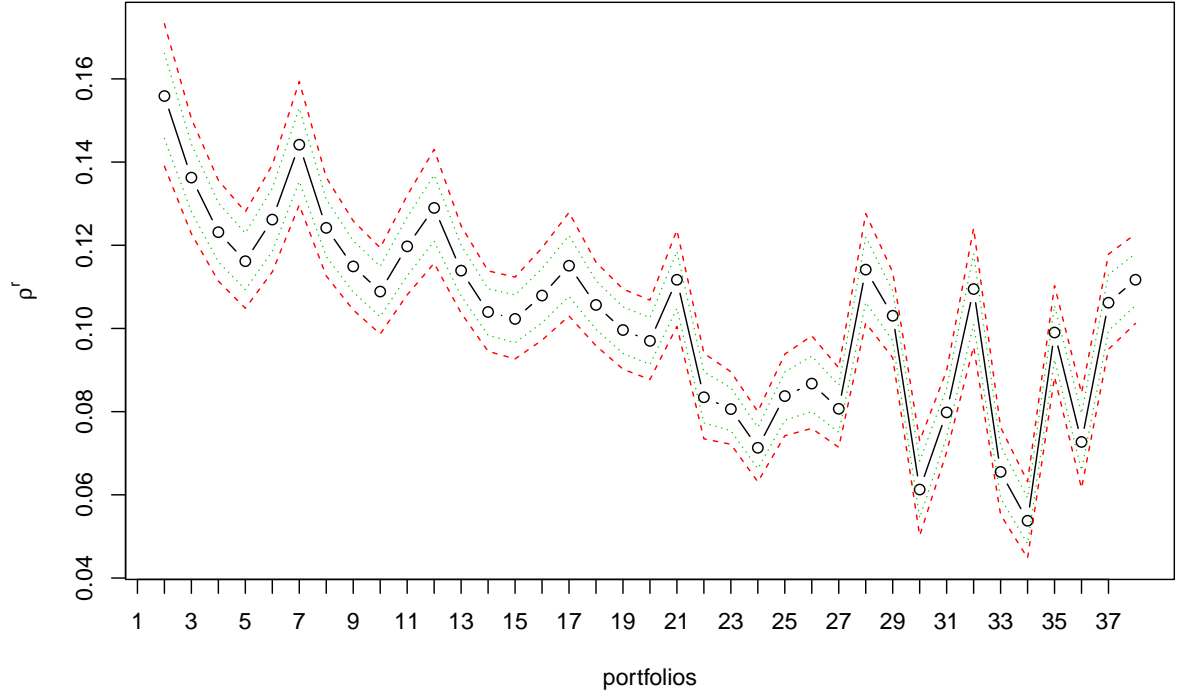
V.1.3 Time series properties of stocks and bonds

We now turn to the time series properties of stocks and bonds implied by our model in equations (3)–(4) (and (7)–(8)). The loadings of equity portfolios on the latent factor f_t are depicted in Figure 7.

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 (7)–(8) (see Figure A8 in the Appendix).

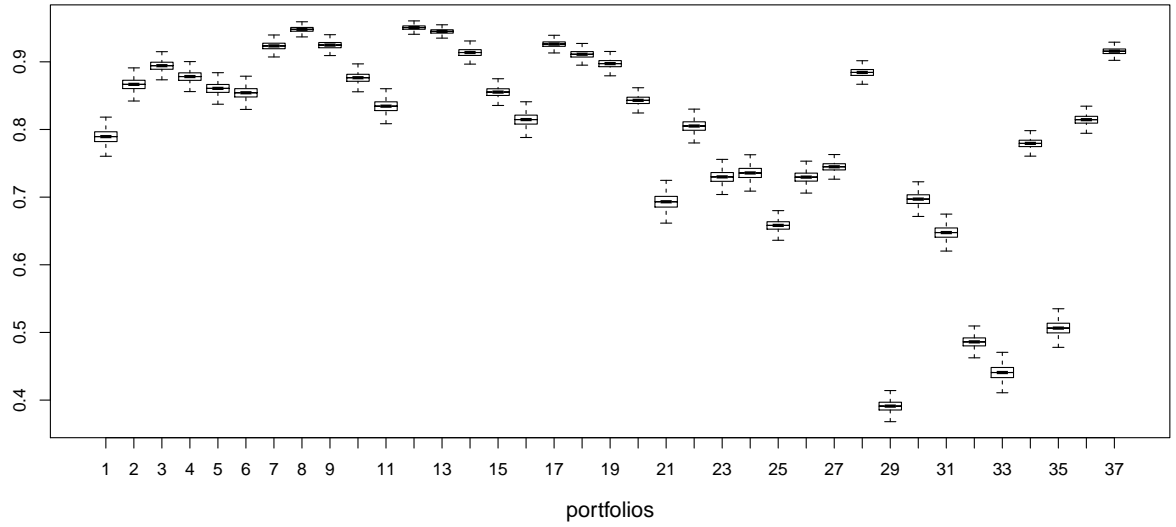
These loadings are not only statistically, but also economically significant as shown in Figure 8: The common factor f explains on average 79% of the time series variation of stock

Figure 7: Common factor loadings (ρ^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 centered 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 8: 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.

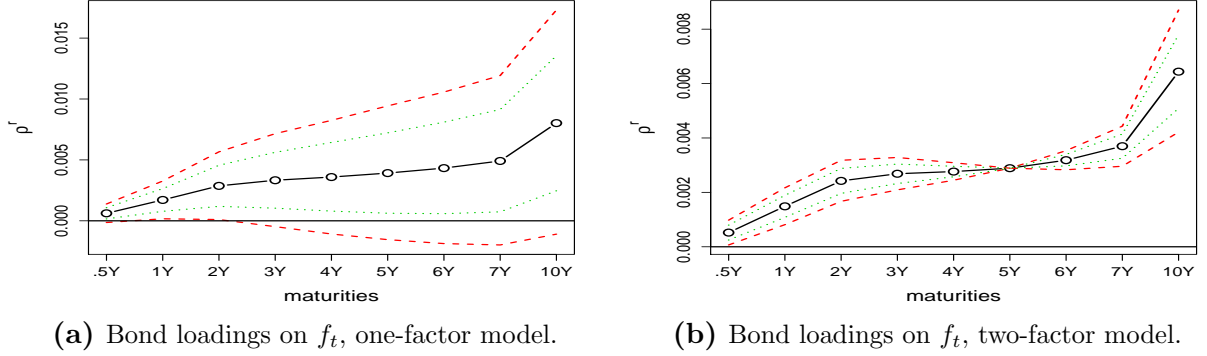
returns, ranging from 36% to nearly 95% for individual portfolios. Moreover, this explanatory power in our model is produced by a single consumption-based factor, as opposed to some of the alternative successful specifications that typically rely on three or more explanatory variables. As shown in Figure A9 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 Figures 9a and 9b 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 10. These loadings are highly statistically significant and increase steeply and monotonically with maturity, revealing a very pronounced term structure pattern.

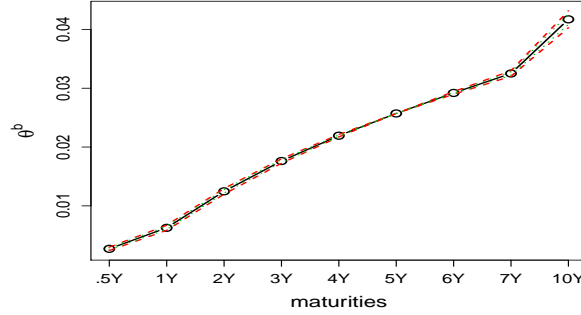
Finally, Figure 11 reports the share of time series variation of bond returns explained by the f_t shocks (left panel), and the f_t and g_t shocks (right panel), and highlights the importance of allowing for a bond-specific factor to characterize the time series of bond returns. The common factor f_t 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 six-month bond returns, its performance rapidly improves with maturity and results in a nearly perfect fit for the time horizon of about five years.

Figure 9: Bond loadings on common factor f_t .



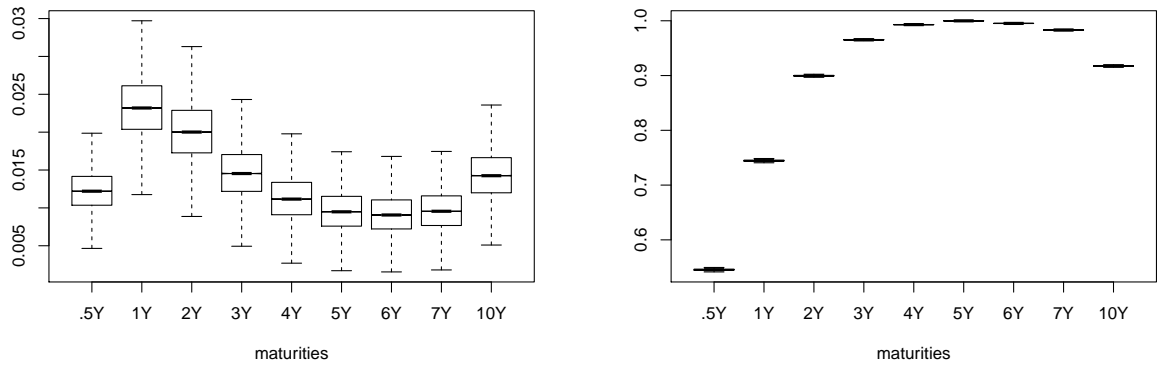
Posterior means of the bonds factor loadings on f_t (black continuous line with circles) and centered posterior 90% (red dashed line) and 68% (green dotted line) coverage regions in the one-latent-factor model.

Figure 10: Bond loadings (θ^b) on the bond-specific factor (g_t).



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

Figure 11: Variance decomposition box-plots of bond returns.



(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.

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.

V.1.4 What drives the consumption shocks spanned by financial markets?

Figure 3 revealed a clear business cycle fluctuation within the moving average component of consumption growth, but what exactly is being captured by those shocks in financial markets? Figure 12 looks at the (posterior) correlation of the consumption mean process shocks, 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.

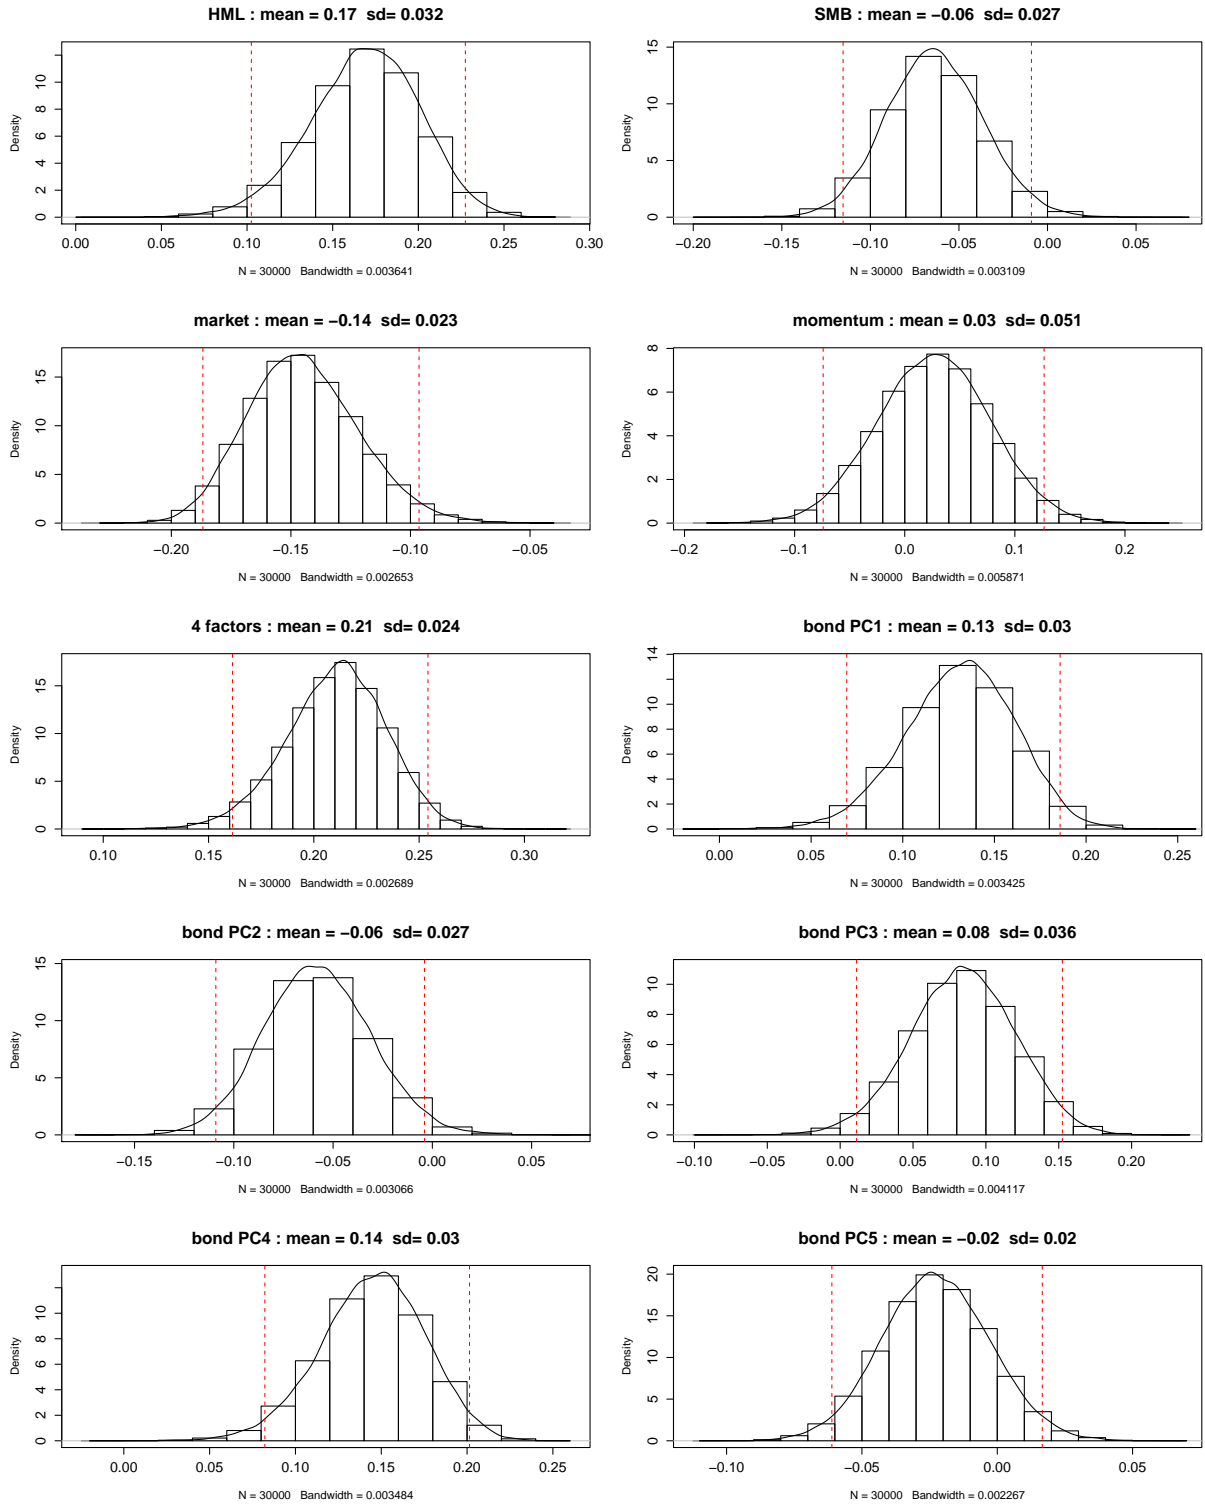
Focusing on equity risk factors, the conditional mean of consumption shows a strong correlation with both the HML factor (a proxy for the value premium) and the overall market excess return, while the correlation with SMB (a proxy for the size premium), is small (albeit significant), while there is virtually no co-movement with the momentum factor. The negative co-movement 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 in line with the finding of Liew and Vassalou (2000) which shows that HML forecasts good economic states, and with equilibrium models with capital irreversibility (as, in e.g. Seru, Papanikolaou, Kogan, and Stoffman (2017)).

Focusing on bonds, 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 – that is, the so called slope 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), and Hamilton and Kim (2002)).¹⁷

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

¹⁷Also, 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.

Figure 12: Common innovations, popular risk factors, 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.

V.1.5 The price of consumption risk

The latent factor model in equations (7)–(8) 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 strong implied factor structure in the returns, since that could lead to spuriously high significance levels and quality of fit, significantly complicating the overall model assessment. However, as Figure 12 indicates, the f shocks of consumption do not heavily load on any of the main principal components of returns. Furthermore, our state-space framework 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)), which could also lead to the unjustifiably high significance levels.

Table 3 summarizes the cross-sectional pricing performance of our parametric model of consumption on a mixed cross-section of nine 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 specification, 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.

Table 3: 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 [0.0053, .0063]	16.00 [9.54, 28.78]			.59 [.56, .62]
(2)		22.00 [13.12, 39.67]			.91 [.90, .92]
Two latent factor specification					
(3)	.0059 [.0051, .0065]	15.79 [9.10, 29.71]			.59 [.56, .62]
(4)		21.75 [12.52, 40.89]			0.91 [0.90, 0.92]
(5)	.0065 [−243.4, 235.1]	15.11 [8.64, 28.46]	37.89 [−243.3, 235.1]		.59 [0.56, 0.61]
(6)		21.73 [12.51, 40.85]	−37.01 [−1152, 1198]		.92 [.91, .92]
(7)	.0057 [.0040, .0071]			15.95 [9.12, 30.57]	.59 [.56, .62]
(8)				21.69 [12.52, 40.79]	.91 [.89, .92]

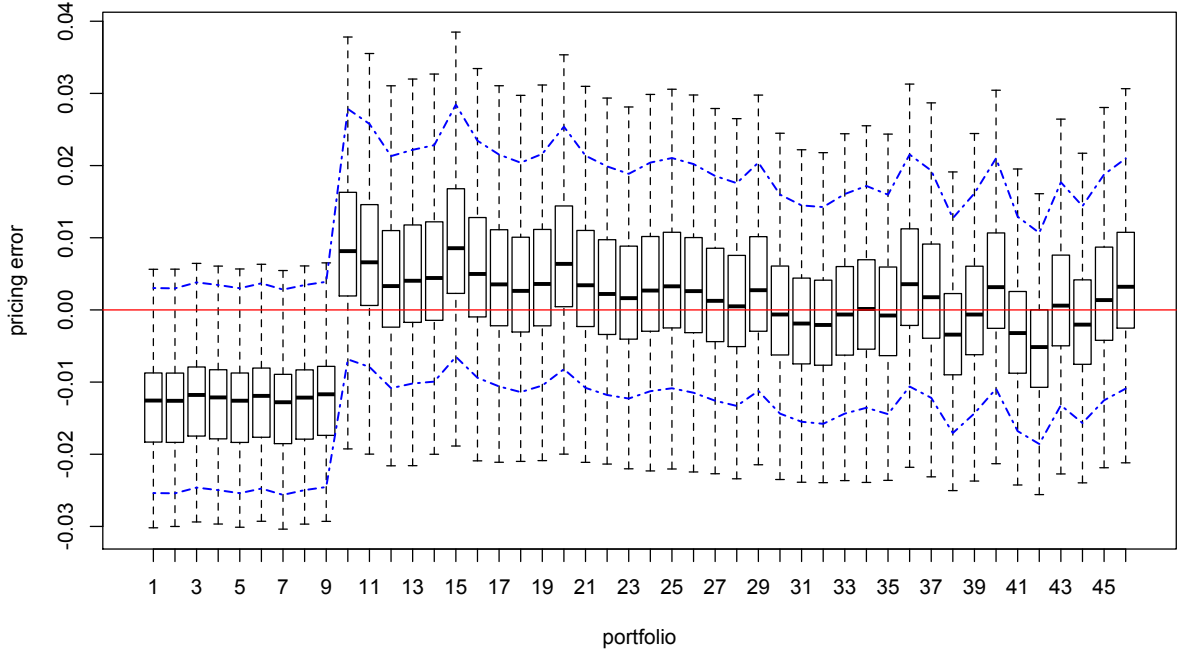
The table presents posterior means and centered 95% posterior coverage (in square brackets) of the Fama and MacBeth (1973) cross-sectional regression of excess returns on $\sum_{j=0}^S \rho_j \boldsymbol{\rho}^r$ (with associated coefficient λ_f) and $[\boldsymbol{\theta}^b, \mathbf{0}'_{N-N_b}]' \sum_{j=0}^S \boldsymbol{\theta}_j$ (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, and bond portfolios.

The average pricing error is small, about 60 b.p. per quarter, but statistically significant in most specifications. Figure 13 summarizes the posterior for the individual pricing errors for the specification in row (1) of Table 3: They are all individually statistically insignificant even at the 10% level. The figure further suggests that the cross-sectional pricing could benefit from a bond-specific intercept – that is, a level factor for bonds.

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

In addition, we confirm these cross-sectional findings using the a semiparametric estimation via Empirical Likelihood, as reported in Appendix A.5–A.5.1.

Figure 13: 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 3. 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 (3)–(4) allow us to pin down the parameters of the joint consumption-returns process with a high degree of precision. However, this comes at the 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 (3)–(4) imply a very particular pattern in the covariances of asset returns with

multi-period consumption growth, that is, for any asset i ,

$$\begin{aligned}
cov(r_{i,t+1}^{ex}, \Delta c_{t,t+1}) &= \rho_i^r \rho_0, \\
cov(r_{i,t+1}^{ex}, \Delta c_{t,t+2}) &= \rho_i^r (\rho_0 + \rho_1), \\
&\dots \\
cov(r_{i,t+1}^{ex}, \Delta c_{t,t+k}) &= \rho_i^r \left(\sum_{j=1}^k \rho_{k-j} \right).
\end{aligned} \tag{18}$$

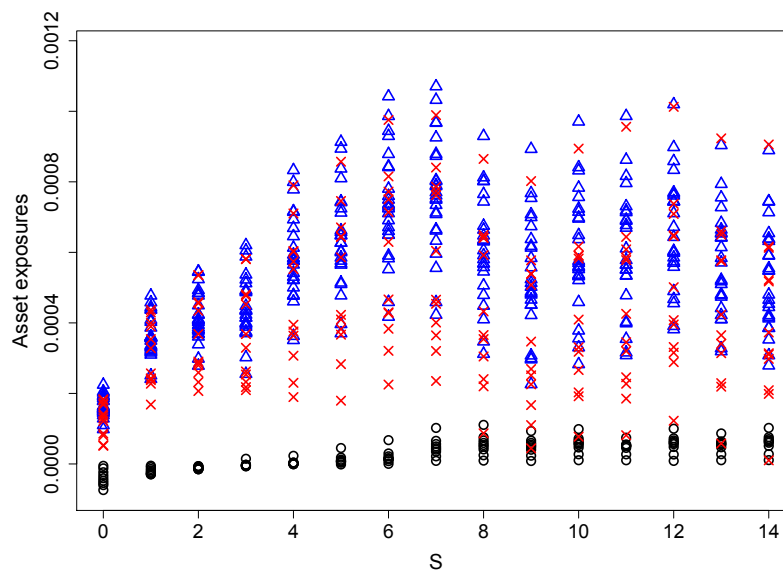
Therefore, the term structure of asset exposures to consumption risk is driven by a single common component: $(\rho_0, \rho_0 + \rho_1, \dots, \sum_{j=1}^k \rho_{k-j})'$ – that is, the cumulative response function of consumption to an f_t shock. This property is not affected by the potential presence 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 (3) 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 $cov(r_{i,t+1}^{ex}, \Delta c_{t,t+k})$ 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.¹⁸

A second testable implication of equation (18) 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 14, which depicts $Cov(\Delta c_{t,t+1+S}, r_{j,t+1}^e)$ 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, 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),

¹⁸In 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.

Figure 14: Cross-sectional spread of exposure to slow consumption adjustment risk



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

which uses $\Delta c_{t,t+1+S}$ to price time t asset returns.

V.3 The role of stochastic volatility

Albeit our state-space model for consumption and returns in equations (3)–(4) delivers consistent estimates of the conditional mean parameters even in the presence of time-varying volatility for the error terms, the literature has often focused on models with a joint stochastic volatility process for consumption and asset return process. This approach has been common since it provides time-varying risk premia for equilibrium asset returns in representative agent models.

As shown in Section V.1.2, ad hoc specifications for the mean process (e.g., constant or AR(1)), that miss the true degree of predictability in the consumption process, mechanically invalidate the identification of the consumption vol process and its predictability. This is precisely why we rely on a long MA representation, that can correctly captures the conditional mean dynamics irrespectively of its true functional form, and therefore provides a

robust way to assess the nature and properties of the consumption volatility without taking a stand on the data generating process of the conditional mean.

That is, our state-space representation provides a reliable framework to tackle inference on a class of structural models, rather than a particular parametrization.

We generalize our state-space model in equations (3)–(4) to allow for stochastic volatility in *all* the error terms. Furthermore, we allow the volatilities to be priced, that is, the process for log excess returns is now

$$\underset{N \times 1}{\mathbf{r}_t^e} = \underset{N \times 1}{\boldsymbol{\mu}_r} + \underset{N \times 1}{\boldsymbol{\rho}^r} f_t + \underset{N \times 1}{\boldsymbol{\beta}_f} \sigma_{f,t-1}^2 + \underset{N \times 1}{\boldsymbol{\beta}_r} \sigma_{r,t-1}^2 + \underset{N \times 1}{\mathbf{w}_t^r}, \quad (19)$$

where $\sigma_{r,t-1}^2$ is a common market volatility process that affects the volatility of all excess return shocks (\mathbf{w}^r), and $\sigma_{f,t-1}^2$ is the volatility of the shocks (f) to the conditional mean of consumption growth (in equation (3)). The above formulation is very general and encompasses the standard models from the literature as particular cases. For instance, setting $\sigma_{f,t} \equiv \sigma_{r,t}$, one obtains the canonical “long-run risk” formulation of Bansal and Yaron (2004).

We follow the past literature (e.g., Hull and White (1987) and Chesney and Scott (1989)) by assuming the following distributions for w_t^c , f_t and \mathbf{w}_t^r :

$$w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \exp(h_{ct})), \quad (20)$$

$$f_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \exp(h_{ft})), \quad (21)$$

$$\mathbf{w}_t^r \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{rt}), \quad \boldsymbol{\Sigma}_{rt} = \text{diag}\{\sigma_{r,1t}^2, \dots, \sigma_{r,Nt}^2\}, \quad (22)$$

where $\sigma_{r,it}^2$ is the stochastic volatility of the i -th asset that is driven by both idiosyncratic volatility shocks ($e_{i,rt}$) and the common volatility process ($\sigma_{r,t}^2 \equiv \exp(h_{rt})$), that is,

$$\sigma_{r,it}^2 = \exp\{\kappa_0^{(i)} + \kappa_1^{(i)} h_{rt} + e_{i,rt}\}. \quad (23)$$

In what follows we refer for simplicity to $\sigma_{r,t}^2 \equiv \exp(h_{rt})$ as the “market” variance, albeit

in our formulation of asset return dynamics the total volatility of market returns is a linear combination of $\sigma_{r,t}$ and $\sigma_{f,t}$.

In order to identify the model, without loss of generality we normalize h_{ft} and the common market volatility component h_{rt} to have zero mean and unit variance. The stochastic volatility processes h_{ct} , h_{ft} , and h_{rt} are all modeled as independent AR(1) processes in the unrestricted formulation (see Appendix A.6 for further details including the MCMC algorithm needed to evaluate this model). Furthermore, we also consider several restricted versions of the above specification that correspond to the processes employed in the previous literature.

Using this general formulation, the estimated parameters for the consumption ($\{\rho_j\}_{j=0}^S$) and asset returns ($\boldsymbol{\rho}^r$) loadings on the common shock f stay virtually unchanged (see figure A10 in the Appendix). Similarly, the variance decompositions for consumption and returns as well as the cross-sectional price of risk are almost identical to the ones presented in the previous sections. While this is not surprising (since the formulation (3)–(4) is consistent for mean equation parameters independently from the volatility process), this finding is quite reassuring. In other words, it is clear that the returns significantly load on the shocks to the consumption mean that drive more than a quarter of the consumption growth variance. That is, all our results seem robust to the modeling choice of the volatility processes.

What do we learn about the volatility itself for both consumption and returns? Figure 15 reports the estimated volatility processes (posterior median and 95% credible intervals) under a diffuse prior for the autoregressive coefficients of the processes (see Appendix A.6 for details and very similar estimation results obtained under alternative priors). Several observations are in order. First, there is clear evidence of time-varying volatility in stock returns (Panel C), while there is a whole range of constant volatility levels for the shocks to consumption (shaded areas in Panels A and B) that are within the posterior confidence bands. Note also that all our volatilities are quite sharply estimated, even in the case of consumption shocks. Therefore, the evidence in favor of stochastic volatility in consumption growth is weak at best. Second, all but two asset volatilities have significant loadings (κ_1^i) on the common

financial market volatility process. Third, the common financial market volatility clearly increases during recessions and market crashes, while for the other two volatility processes there is no such a clear pattern. Forth, there seems to be little co-movement between market and consumption shocks volatilities.

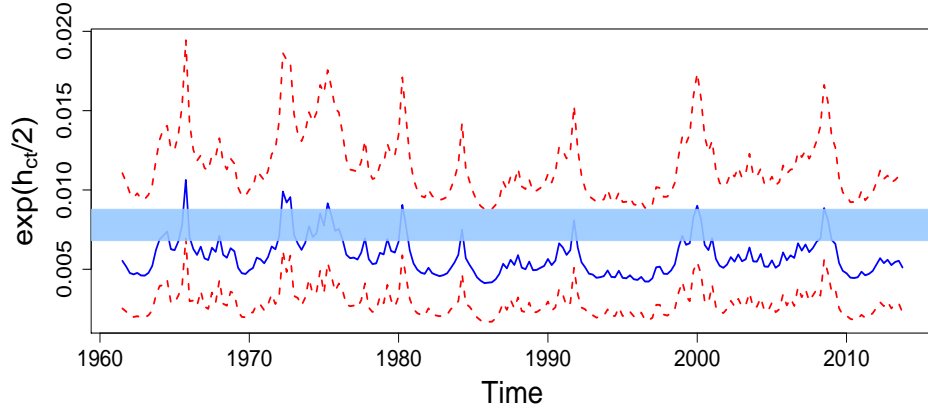
We further explore the correlation between the estimated volatilities and (the square of) the VXO volatility index (a proxy for the underlying volatility process commonly used in the literature¹⁹) in Table 4. Both financial market vol and the vol of the consumption conditional mean shocks have non-trivial correlations with the volatility index (.40 for the former and .48 for the latter, at the median, in Panel A), while being very weakly correlated with each other (the median correlation is only .15, Panel B).

Under the null of the stochastic process for asset returns that we postulate (equations (19), (20)-(23) and normalizations therein) these correlations give an estimate of the variance decomposition for the S&P100 (the index underling the VXO). Hence, Panel A of Table 4 implies that about 40% of the (option-implied) variance of the S&P100 index is generated by f_t – the shocks to the conditional mean of consumption. This is a strong external verification of our (very general) modeling choice for the joint dynamics of consumption and returns. Note that in the canonical long run risk framework (Bansal and Yaron (2004)), the pairwise correlations in Panel B (i.e., between the stochastic volatilities of consumption shocks and market returns) should all be equal to one, and these would be expected to have close to unit correlation with the VXO^2 .

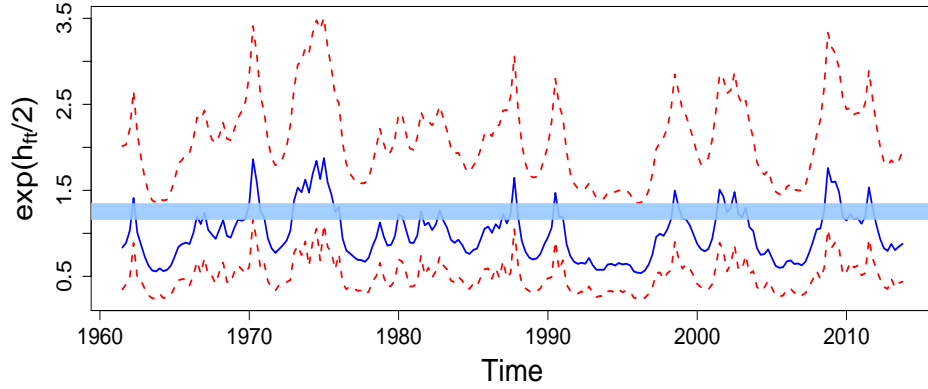
Finally, we check whether the volatility process of asset returns and consumption mean shocks are significant drivers of excess returns dynamics. Figure 16 reports the posterior distributions of the loadings in equation (19). Strikingly, none of the excess return series loads significantly on the volatility of consumption growth shocks, and only a few have marginally significant loadings on the common financial market vol process. This finding poses a challenge to the literature that has modeled time-varying risk premia by assuming

¹⁹We use the VXO index, instead of VIX, due to the longer time series available for the former.

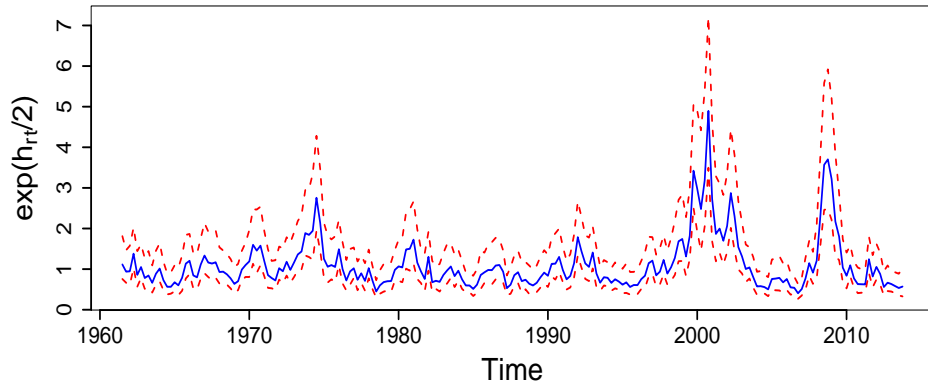
Figure 15: Filtered stochastic volatilities of consumption and returns



Panel A: Log volatility of the idiosyncratic shock to consumption (w_t^c) of equation (3)



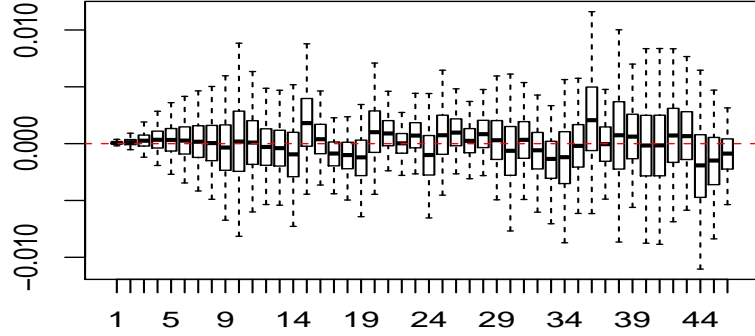
Panel B: Log volatility of the shock to the conditional mean of consumption growth (f_t).



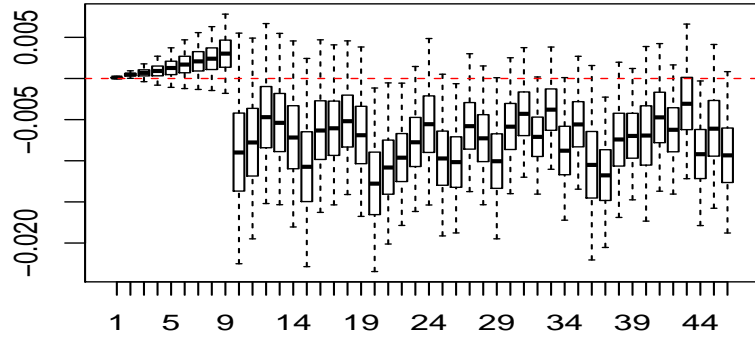
Panel C: Common log volatility of asset return (h_{rt} of equation (23)).

Estimated stochastic volatilities of the model in equations (3), (19), (20)–(23), and Section A.6 under a diffuse prior for the autoregressive volatility coefficients. Solid blue lines depict the posterior median of the log volatility, while dotted red lines denote 2.5% and 97.5% credible intervals. Shaded areas reflect constant volatility levels that would not be rejected given the credible intervals.

Figure 16: Loadings of excess returns on consumption and returns volatilities



Panel A: posterior distribution of excess return loadings (β_f) on the variance of shocks to the conditional consumption growth mean ($\sigma_{f,t-1}^2$) in equation (19).



Panel B: posterior distributions of excess return loadings (β_r) on the common financial return variance ($\sigma_{r,t-1}^2$) in equation (19).

Box-plots of the posterior distributions of the loadings of portfolio excess returns on the variance of shocks to the conditional consumption growth and the common financial returns variance. 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.

Table 4: Correlations among stochastic volatility processes and VXO index.

	Mean	2.5%	5%	50%	95%	97.5%
Panel A: correlations of vol processes with VXO ² index						
$cor(\sigma_{ct}^2, VXO_t^2)$	0.15	-0.07	-0.04	0.15	0.36	0.40
$cor(\sigma_{ft}^2, VXO_t^2)$	0.40	0.16	0.20	0.40	0.58	0.61
$cor(\sigma_{rt}^2, VXO_t^2)$	0.48	0.32	0.35	0.48	0.59	0.61
Panel B: Pairwise correlations of vol processes						
$cor(\sigma_{ct}^2, \sigma_{ft}^2)$	0.06	-0.11	-0.08	0.05	0.24	0.29
$cor(\sigma_{ct}^2, \sigma_{rt}^2)$	0.15	-0.02	-0.00	0.13	0.34	0.38
$cor(\sigma_{ft}^2, \sigma_{rt}^2)$	0.16	0.02	0.03	0.15	0.31	0.35

The table summarises posterior mean, 2.5%, 5%, 50%, 95% and 97.5% quantiles of correlation among σ_{ct}^2 , σ_{ft}^2 , σ_{rt}^2 and the VXO index under a diffuse prior for the autoregressive coefficients of the vol processes.

Table 5: Model Comparison Using Log Bayes Factors

	Models								
	I	II	III	IV	V	VI	VII	VIII	IX
Log of BF:	0	-24.72	-33.03	-44.4	-75.1	-90.72	-144.11	-160.58	-203.26
Post. Prob.:	100%	0%	0%	0%	0%	0%	0%	0%	0%

Model I: w_t^c , f_t and w_t^r follow SV processes, and $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \mathbf{w}_t^r$

Model II: w_t^c , f_t and w_t^r follow SV processes, $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \mathbf{w}_t^r$, and $h_{ft} = h_{rt}$

Model III: w_t^c and w_t^r follow SV processes, $f_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \mathbf{w}_t^r$

Model IV: f_t and w_t^r follow SV processes, $w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_c^2)$, and $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \mathbf{w}_t^r$

Model V: w_t^r follows SV process, $w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_c^2)$, $f_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \mathbf{w}_t^r$

Model VI: w_t^c , f_t and w_t^r follow SV processes, $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \boldsymbol{\beta}_f \sigma_{f,t-1}^2 + \mathbf{w}_t^r$, and $h_{ft} = h_{rt}$

Model VII: w_t^r follows SV process, $w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_c^2)$, $f_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \boldsymbol{\beta}_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r$

Model VIII: w_t^c , f_t and w_t^r follow SV processes, and $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \boldsymbol{\beta}_f \sigma_{f,t-1}^2 + \boldsymbol{\beta}_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r$

Model IX: f_t and w_t^r follow SV processes, $w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_c^2)$, and $\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \boldsymbol{\beta}_f \sigma_{f,t-1}^2 + \boldsymbol{\beta}_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r$

The table summarizes the model comparison for restricted and unrestricted versions of the specification in equations (3), (19), (20)–(23), and Section A.6. We approximate the Bayes Factor using the Schwartz criterion. We use model I as benchmark and calculate the (log) odds of each model compared to model I. A negative number implies that the chosen model is less likely than model I conditional on the observed data. The model posterior probabilities are computed under the prior of the specifications being all equally likely.

time variation in the common volatility process of consumption and asset returns. We formally test this mechanism below.

Table 5 reports Bayes factors and posterior probabilities for several restricted and unrestricted versions of the specification in equations (3), (19), (20)–(23), and Section A.6.

In particular, we test both commonality in the various volatility processes (for returns and consumption mean shocks), as well as their impact on excess returns. Note that the posterior model probabilities (and Bayes factors) are particularly appropriate for this type of test because they yield valid model selection even over the space of misspecified models. That is, they select the model that has the highest probability of being the true data generating process (not just the model with the highest likelihood, see, e.g., Schervish (1995)).

The data strongly favors a specification in which *i*) volatilities do *not* affect excess returns (Models I to V) and *ii*) returns and consumption have *distinct* volatility processes – the posterior probability of such a formulation (Model I) is (almost) 100%. This constitutes an important challenge for a large class of asset pricing models that obtain time-varying risk premia via stochastic volatility.

The above result is salient because assuming common volatilities (up to a factor of proportionality) in consumption and returns has been customary in the literature. For this reason, empirical analyses typically proxy this common volatility by using the market realized vol (see, e.g., Bansal, Kiku, Shaliastovich, and Yaron (2014), and Campbell, Giglio, Polk, and Turley (2018)) as a priced state variable. Instead, similar to Dew-Becker and Giglio (2016), we do not find evidence that volatility is priced but, different from them, our result does not seem to be driven by a potential lack of power since we are directly comparing parametric alternatives and, most importantly, the data overwhelmingly prefers the Model I specification – a specification with distinct, and unpriced, vol processes for consumption and returns. That is, our analysis does shed light on the dark matter of consumption based asset pricing models.

To summarize, even allowing for flexible models of time-varying volatility, we confirm all the results obtained in the previous sections: The conditional mean of the consumption process reacts over multiple quarters to shocks spanned by financial returns, and this persistent component is economically very large and a significant driver of both the time series of consumption and returns and the cross-section of risk premia.

Additionally, we find very little evidence that the volatility of consumption mean shocks is the same as the volatility of financial markets, and that these volatilities are reflected in the time series of excess returns. This finding poses a fundamental challenge for theoretical models that assume such commonality and dynamics in order to generate time-varying risk premia. Note that our findings have been obtained using a significantly more general modeling of both mean and volatility processes, and much less stringent prior assumptions, than in the previous empirical literature.

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 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.

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 stock returns, as well as the positive (unconditional) slope of the term structure of bond excess returns. Our model explains 36%–95% of the time series variation in stock returns, and 57%–90% of the joint cross-sectional variation in stocks and bonds.

Furthermore, we uncover a fundamental challenge for theoretical models that obtain equilibrium time-varying risk premia by assuming a common stochastic volatility in consumption

and returns. The data strongly suggest that these are not only distinct stochastic processes, but also that they do not seem to drive time variation in excess returns.

Compared to the previous literature, our findings are obtained using not only more general and flexible specifications, but also less restrictive priors and much richer asset return data. We also provide a large set of complementary evidence that does not rely on our parametric formulation, yet strongly supports all the main time series and cross-sectional findings.

Our findings have first order implications not only for macro-finance models but also, and arguably more importantly, for the assessment of the costs of business cycle fluctuations and optimal fiscal and monetary policies. We defer the study of these effects to future work.

References

- ABEL, A. B. (1990): “Asset Prices Under Habit Formation and Catching Up with the Joneses,” *American Economic Review*, 80, 38–42.
- 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., V. KHATCHATRIAN, AND A. YARON (2005): “Interpretable Asset Markets?,” *European Economic Review*, 49(3), 531–560.
- 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*. Oxford 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.
- BLANCHARD, O. J., AND D. QUAH (1989): “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *The American Economic Review*, 79, 655–673.
- 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.
- BRYZGALOVA, S., J. HUANG, AND C. JULLIARD (2021): “Bayesian Solutions for the Factor Zoo: We Just Run Two Quadrillion Models,” mimeo.
- 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), 205–251.
- 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 (2021): “Measuring the ‘Dark Matter’ in Asset Pricing Models,” *Journal of Finance*, forthcoming.
- 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.
- CHESNEY, M., AND L. SCOTT (1989): “Pricing European Currency Options: A Comparison of the Modified Black-Scholes Model and a Random Variance Model,” *Journal of Financial and Quantitative Analysis*, 24(3), 267–284.
- 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), 519–543.
- CONSTANTINIDES, G. M., AND A. GHOSH (2011): “Asset Pricing Tests with Long-run Risks in Consumption Growth,” *The Review of Asset Pricing Studies*, 1(1), 96–136.
- CSISZAR, I. (1975): “I-Divergence Geometry of Probability Distributions and Minimization Problems,” *Annals of Probability*, 3, 146–158.
- DEW-BECKER, I., AND S. GIGLIO (2016): “Asset Pricing in the Frequency Domain: Theory and Empirics,” *Review of Financial Studies*, 29, 2029–2068.
- DRECHSLER, I., AND A. YARON (2011): “What’s Vol Got to Do with It,” *Review of Financial Studies*, 24(1), 1–45.

- 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.
- 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, R., B. SACK, 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, 340–360.
- 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., 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.
- HOETING, J. A., D. MADIGAN, A. E. RAFTERY, AND C. T. VOLINSKY (1999): “Bayesian Model Averaging: A Tutorial,” *Statistical Science*, 14(4), 382–401.
- HULL, J. C., AND A. WHITE (1987): “The Pricing of Options on Assets with Stochastic Volatilities,” *Journal of Finance*, 42(2), 281–300.
- JACOBS, K., AND K. K. WANG (2004): “Idiosyncratic Consumption Risk and the Cross-Section of Asset Returns,” *Journal of Finance*, 59, 2211–2252.
- 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), 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.
- KIM, S., N. SHEPHARD, AND S. CHIB (1998): “Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models,” *The Review of Economic Studies*, 65(3), 361–393.
- KITAMURA, Y. (2006): “Empirical Likelihood Methods in Econometrics: Theory and Practice,” Cowles Foundation Discussion Papers 1569, Cowles Foundation, Yale University.
- KOOP, G., M. PESARAN, AND S. M. POTTER (1996): “Impulse Response Analysis in Nonlinear Multivariate Models,” *Journal of Econometrics*, 74(1), 119–147.
- 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), 815–849.

- (2001b): “Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying,” *Journal of Political Economy*, 109, 1238–1286.
- 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.
- 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.
- MENZLY, L., T. SANTOS, AND P. VERONESI (2004): “Understanding Predictability,” *Journal of Political Economy*, 112(1), 1–47.
- NEWHEY, W., AND R. SMITH (2004): “Higher-Order Properties of GMM and Generalized Empirical Likelihood Estimators,” *Econometrica*, 72, 219–255.
- NEWHEY, W. K., AND K. D. WEST (1987): “A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- OWEN, A. B. (2001): *Empirical Likelihood*. Chapman and Hall.
- PARKER, J. A., AND C. JULLIARD (2005): “Consumption Risk and the Cross-Section of Expected Returns,” *Journal of Political Economy*, 113(1).
- PESARAN, H., AND Y. SHIN (1998): “Generalized Impulse Response Analysis in Linear Multivariate Models,” *Economics Letters*, 58(1), 17–29.
- PIAZZESI, M., M. SCHNEIDER, AND S. TUZEL (2007): “Housing, Consumption and Asset Pricing,” *Journal of Financial Economics*, 83(3), 531–569.
- PRIMICERI, G. E. (2005): “Time Varying Structural Vector Autoregressions and Monetary Policy,” *Review of Economic Studies*, 72(3), 821–852.
- RAFTERY, A. E., D. MADIGAN, AND J. A. HOETING (1997): “Bayesian Model Averaging for Linear Regression Models,” *Journal of the American Statistical Association*, 92(437), 179–191.
- SAVOV, A. (2011): “Asset Pricing with Garbage,” *Journal of Finance*, 66(1), 177–201.
- SCHERVISH, M. J. (1995): *Theory of statistics*, Springer series in statistics. Springer-Verlag.
- SCHORFHEIDE, F., D. SONG, AND A. YARON (2018): “Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach,” *Econometrica*, 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.
- SIMS, C. A., AND T. ZHA (1999): “Error Bands for Impulse Responses,” *Econometrica*, 67(5), 1113–1155.
- 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.

- TÉDONGAP, R. (2014): “Consumption Volatility and the Cross-Section of Stock Returns,” *Review of Finance*, 19(1), 367–405.
- 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. (2021): “Term Structure of Risk in Expected Returns,” *The Review of Financial Studies*, forthcoming.

A Appendix

A.1 Simulation design

We assume the same data generating process of the log consumption growth 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:

$$\begin{aligned}\Delta c_{t,t+1} &= \mu + x_t + \sigma_t \eta_{t+1}, \\ x_{t+1} &= \rho x_t + \phi_e \sigma_t e_{t+1}, \\ \sigma_{t+1}^2 &= \sigma^2(1 - v_1) + v_1 \sigma_t^2 + \sigma_w \sigma_t \omega_{t+1},\end{aligned}$$

where $\eta_{t+1}, e_t, \omega_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. The calibrated monthly parameter values are: $\mu = 0.0015, \rho = 0.979, \phi_e = 0.044, \sigma = 0.0078, v_1 = 0.987, \sigma_w = 0.00029487$.

We use this particular LRR calibration because it ensures the non-negativity of the volatility process and, most importantly, gives the best chance of detecting the predictability of consumption using the simple ARIMA selection methods commonly used in empirical work: the contribution of the conditional mean to the variance of the consumption growth in this calibration is about 12% – the largest share among leading LRR calibrations.

When also simulating return data, we further assume that log excess returns (\mathbf{r}_{t+1}^e) follow

$$\underset{N \times 1}{\mathbf{r}_{t+1}^e} = \underset{N \times 1}{\boldsymbol{\mu}_r} + \underset{N \times 1}{\boldsymbol{\rho}^r} e_{t+1} + \underset{N \times 1}{\mathbf{w}_{t+1}^r},$$

where $\boldsymbol{\mu}_r$ is a vector of average monthly log excess returns, $\boldsymbol{\rho}^r$ is a vector of returns' loadings on the contemporaneous shock, e_{t+1} , which drives the conditional mean of $\Delta c_{t,t+1}$, and $\mathbf{w}_{t+1}^r \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \boldsymbol{\Sigma}_r)$. For simplicity, we assume that $\boldsymbol{\Sigma}_r$ is diagonal and constant over time in the simulations. Values of $\boldsymbol{\mu}_r$, $\boldsymbol{\rho}^r$ and $\boldsymbol{\Sigma}_r$ are chosen to be their sample estimates in the main text.²⁰ After simulating monthly sequences of $(\Delta c_{t,t+1}, \mathbf{r}_{t+1}^e)$, we aggregate them

²⁰More precisely, we obtain the estimates of $\boldsymbol{\mu}_r$, $\boldsymbol{\rho}^r$ and $\boldsymbol{\Sigma}_r$ using quarterly data. In the simulation of

into quarterly sequences by summing all three monthly observations within a quarter. We simulate 1,000 sample paths with 214 quarterly observations. Since we approximate the log consumption growth by an MA(14) process, there are 200 effective quarterly observations.

A.2 Data description

Bond holding returns are calculated on a quarterly basis using the zero coupon yield data constructed by Gurkaynak, Sack, and Wright (2007)²¹ 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: six months, 1, 2, 3, 4, 5, 6, 7, and 10 years, which gives us a set of nine 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 the 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) expenditure on nondurable goods per capita available from the National Income and Product Accounts (NIPA). As in e.g. Parker and Julliard (2005), we do not include services in our baseline definition of consumption, since these are likely to mechanically bias both the persistence of the consumption proxy (due to, e.g., utilities and health care) and the co-movement with market returns (due to, e.g., financial services and insurance). Our results are robust to the usage of alternative measures and refinements of the consumption proxy (as, e.g., the exclusion of shoes and clothing, as in Lettau and Ludvigson (2001a, 2001b), due to their semi-durable nature). We use the end-of-period timing convention and assume that all of the expenditure occurs at the

monthly returns, we divide all above three parameters by 3.

²¹The data is regularly updated and available at <http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html>

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 realized. 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.

The VXO volatility index is publicly available since 1986:Q1.

A.3 Preliminary empirical evidence

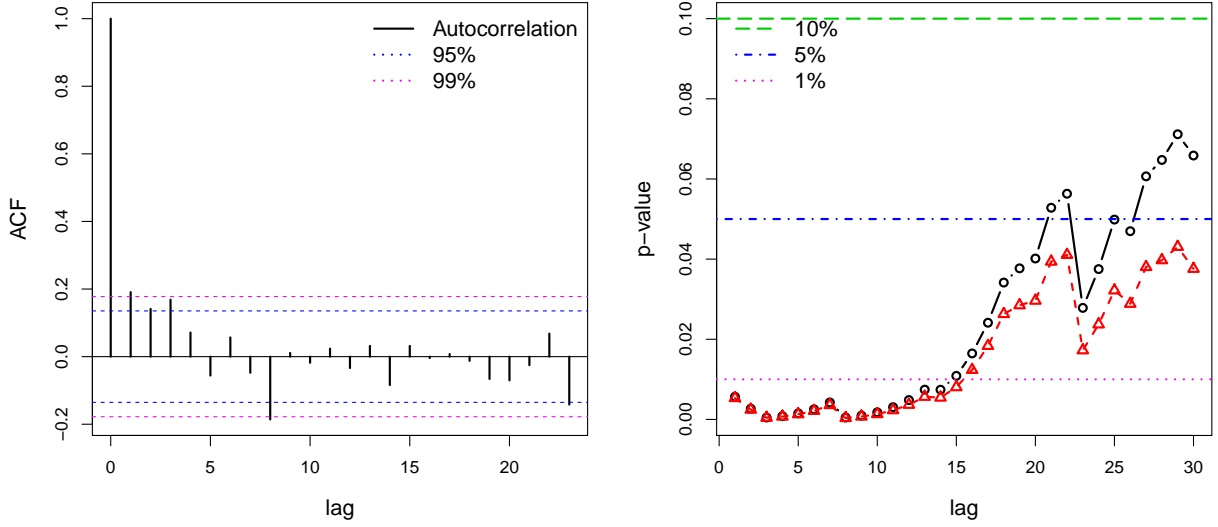
To motivate the structure of the state-space model for consumption and asset returns, we first establish a set of empirical facts via model-free reduced-form approaches. We document that *a*) consumption growth is autocorrelated, *b*) not only asset returns predict future levels of consumption growth 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 is reported in Appendix A.2.

First, Figure A1 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 autocorrelations, 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.²²

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

²²Note that, even in the seminal examination of the random walk hypothesis of Hall (1978), the presence of predictability in consumption growth could not be rejected.

Figure A1: 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.

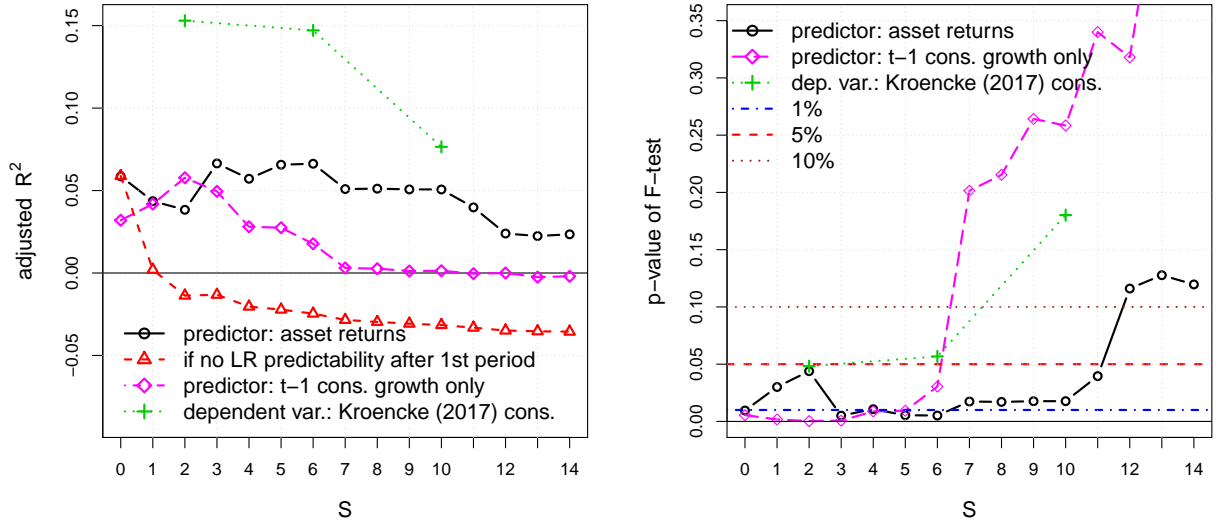
t asset returns.²³ Figure A2 depicts summary statistics for these predictive regressions at different horizons (S). In particular, the left panel plots the time series adjusted R^2 of these regressions, and the right panel the p -value of the F -test of joint significance of the regressors for this and 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 A2. 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.²⁴

²³We use the first eight principal components of the 25 size and book-to-market Fama-French portfolios, 12 industry portfolios, and nine 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.

²⁴Liu and Matthies (2018) also document the existence of long-run predictability in consumption using a

Figure A2: 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?²⁵ This is unlikely. The purple lines in Figure A2 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 does this variable fail 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 news-based measure of economic growth prospects.

²⁵Seasonal smoothing of consumption levels often leads to countercyclical measurement errors in the growth rates.

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 consumption data of Kroenke (2017)²⁶ (green dotted lines in Figure A2). 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, as 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 results above highlight that not only there is substantial predictability in consumption growth, but also it is best captured by asset returns.

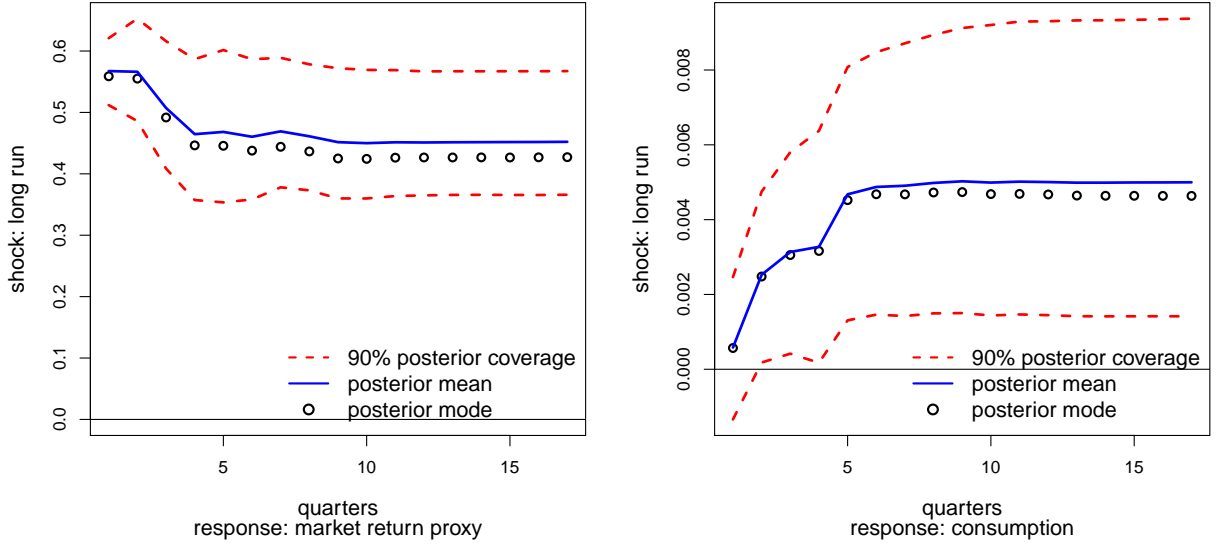
Third, the state-space representation of the slow consumption adjustment process that we use postulates the presence of (potentially) long-run shocks in the consumption growth process to which asset returns react only contemporaneously. To verify this conjecture, we recover the long-run impact of common innovations to financial market returns and non-durable consumption using a simple bivariate structural vector autoregression (S-VAR) for a broad market excess return index and consumption growth.²⁷ We achieve identification via long-run restrictions á la Blanchard and Quah (1989). That is, we distinguish a fundamental long-run shock, which can have a long-run impact on both asset prices and consumption *levels*, and a transitory shock that is restricted not to have a long-run impact on the latter. The details of the estimation procedure are presented in Appendix A.3.1.

Figure A3 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

²⁶We are grateful to Tim Kroenke for making the data available on his website.

²⁷We 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 portfolios, 12 industry portfolios, and Treasury securities with maturities of six months, 1, 2, 3, 4, 5, 6, 7, and 10 years.

Figure A3: Cumulated response functions to a long-run shock.



The 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).

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 five to seven quarters after the initial shock. These patterns are consistent with the moving average process we postulate in Section IV. Moreover, as shown by the estimates of the state-space model in the main text, the above reaction of consumption to the long-run shocks is extremely similar to the one obtained with our state-space model.

A.3.1 S-VAR Identification via Long-Run Restrictions

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

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

where c is a vector of constants, L denotes the lag operator, $\Gamma(L) \equiv \Gamma_1 + \Gamma_2 L + \dots + \Gamma_p L^{p-1}$, and each Γ_j is a two-by-two matrix. In order to identify the S-VAR using long-run restrictions, we follow Blanchard and Quah (1989) and work with the moving average representation

$$X_t = \kappa + A(L) \begin{bmatrix} \varepsilon_t^{sr} \\ \varepsilon_t^{lr} \end{bmatrix} \quad (24)$$

where κ is a vector of constants, $A(L) \equiv A_0 + A_1 L + A_2 L^2 + \dots A_\infty L^\infty \equiv [\Gamma_0 + L\Gamma(L)]^{-1}$, 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 $\{\cdot\}_{1,2}$ 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 variables in levels (i.e., consumption level or asset prices). This restriction also implies that $A(1) \equiv \sum_{j=0}^{\infty} A_j$ should be a lower triangular matrix.

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

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

where $B(L) = B_0 + B_1 L + \dots + B_p L^p$ and $\Omega = \Gamma_0^{-1} (\Gamma_0^{-1})'$ 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 $DD' = [I - B(1)]^{-1} \Omega [I - B(1)]^{-1'}$. 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 $\hat{\Gamma}_0^{-1} = [I - \hat{B}(1)] \hat{D}$, $\hat{\Gamma}(L) = -\hat{\Gamma}_0 \hat{B}(L)$ 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 five 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.4 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 (14) 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 by²⁸

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

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 (12). Therefore, the missing part necessary for taking draws via MCMC using a Gibbs sampler is the conditional distributions of \mathbf{z}_t . Since

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

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

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

²⁸Relaxing the diagonality assumption the posterior distribution of Σ^{-1} is a Wishart centered at the OLS estimates.

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

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

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:

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

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), 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)

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

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 (15) and (16)) 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.5 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,²⁹ all imply a consumption Euler equation of the form of

$$C_t^{-\phi} = \mathbb{E}_t \left[C_{t+1}^{-\phi} \tilde{\psi}_{t+1} R_{j,t+1} \right] \quad (25)$$

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.³⁰

Note that equation (25) 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 multi-period forward-looking SDF.

Hence, the Euler equation can be equivalently rewritten as

$$\mathbf{0}_N = \mathbb{E} \left[\left(\frac{C_{t+1+S}}{C_t} \right)^{-\phi} \psi_{t+1+S} \mathbf{R}_{t+1}^e \right] = \mathbb{E} \left[M_{t+1}^S \mathbf{R}_{t+1}^e \right], \quad (26)$$

²⁹See, 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), and Lustig and Nieuwerburgh (2005).

³⁰E.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.

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

$$\mathbb{E} [\mathbf{R}_{t+1}^e] = -\frac{\text{Cov} (M_{t+1}^S; \mathbf{R}_{t+1}^e)}{\mathbb{E} [M_{t+1}^S]}, \quad (27)$$

where $M_{t+1}^S := (C_{t+1+S}/C_t)^{-\phi} \psi_{t+1+S}$. 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: M_{t+1}^S for every $S \geq 0$. Log-linearizing the above expression, we have the linear factor model

$$\mathbb{E} [\mathbf{r}_{t+1}^e] = [\phi \text{Cov} (\Delta c_{t,t+1+S}; \mathbf{r}_{t+1}^e) - \text{Cov} (\log \psi_{t+1+S}; \mathbf{r}_{t+1}^e)] \lambda_S, \quad (28)$$

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 (26).

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)).³¹

Since the relevance of the multi-period consumption for the cross-section of stocks has

³¹Consider the following transformation of the Euler equation:

$$\mathbf{0} = \mathbb{E} [M_t^S \mathbf{R}_t^e] \equiv \int \left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\phi} \psi_{t+S} \mathbf{R}_t^e dP = \int \left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\phi} \mathbf{R}_t^e d\Psi = \mathbb{E}^\Psi \left[\left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\phi} \mathbf{R}_t^e \right], \quad (29)$$

where P is the unconditional physical probability measure and Ψ is an (absolutely continuous) probability measure such that $\frac{d\Psi}{dP} = \frac{\psi_{t+S}}{\psi}$, where $\bar{\psi} = \mathbb{E} [\psi_{t+S}]$. Both ϕ and Ψ can then be recovered via the Empirical Likelihood approach. Section A.5.1 describes the estimation procedure in detail.

Table A1: Cross-section of bond returns and consumption risk

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

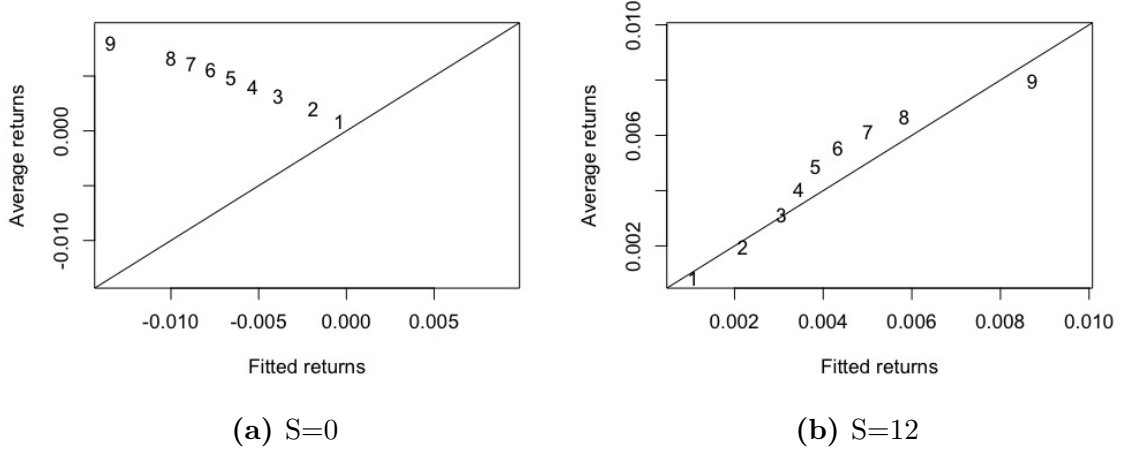
The table reports the pricing of nine excess bond holding returns for various values of the horizon S , allowing for an intercept. Standard errors are reported in parentheses and p-values in brackets. Estimation is performed using Empirical Likelihood and allowing for a common pricing error, α , such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

already been highlighted by Parker and Julliard (2005), we first focus on the cross-section of bonds only. Table A1 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, 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 $S = 14$, and 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).³²

Most importantly, for $S \gg 0$, 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 14). 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 14), resulting in lower standard errors and a better quality of fit. In fact, for large S , as shown in Figure A4, 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 it actually reverses the order of the portfolios: That is, 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

³²Note that the LR-test indicates for all S , as one would expect, that there is a statistically significant component ψ_t in the pricing kernel.

Figure A4: Slow consumption adjustment factor and the cross-section of bond returns



The figures show average and fitted returns on the cross-section of nine 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 their correlation with the consumption growth over 13 quarters, starting from the contemporaneous one.

risk is increased, the quality of fit substantially improves, leading to an R-squared of 90% 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 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 A2 presents cross-sectional estimates for various portfolios of stocks and bonds as S varies. Three patterns, common to also Table A1, are evident. First, cross-sectional fit is generally higher for $S \gg 0$. Second, the point estimates of ϕ are much smaller as S increases. 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 14: as S increases, both the level and the spread of asset loadings on consumption growth become

Table A2: Expected excess returns and consumption risk

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

larger and better identified. This, in turn, delivers a better fit, a lower curvature for the consumption growth, and a tighter identification of the latter. For robustness, Appendix A.8 provides similar empirical evidence for the alternative model specifications that include asset class-specific intercepts.

A.5.1 Empirical Likelihood estimation

Empirical Likelihood provides a natural framework for recovering parameter estimates and probability measure Ψ defined by equation (29), by minimizing Kullback-Leibler Information

Criterion (KLIC), as follows

$$(\hat{\Psi}, \hat{\phi}) = \arg \min_{\Psi, \phi} D(P||\Psi) \equiv \arg \min_{\Psi} \int \ln \frac{dP}{d\Psi} dP \quad \text{s.t.} \quad \mathbf{0} = \mathbb{E}^{\Psi} \left[\left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\phi} \mathbf{R}_t^e \right], \quad (30)$$

From Csiszar (1975) duality result we have

$$\hat{\psi}_t = \frac{1}{T \left(1 + \hat{\lambda}(\theta)' \left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\hat{\phi}} \mathbf{R}_t^e \right)} \quad \forall t = 1..T, \quad (31)$$

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

$$\hat{\phi} = \arg \min_{\phi \in \mathbb{R}} - \sum_{t=1}^T \ln \left(1 + \hat{\lambda}(\phi)' \left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\phi} \mathbf{R}_t^e \right) \quad (32)$$

$$\hat{\lambda}(\phi) = \arg \min_{\lambda \in \mathbb{R}^n} - \sum_{t=1}^T \ln \left(1 + \lambda(\phi)' \left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\phi} \mathbf{R}_t^e \right) \quad (33)$$

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 (33); then minimize the value of the dual objective function w.r.t. $\phi(\hat{\lambda})$, following equation (32).

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 A generalized model with stochastic volatilities

To complete the specification in equations (3), (19), and (20)–(23) we need to formalize the autoregressive volatility processes and the prior formulations. We do so in what follows, and we also provide the sampling algorithm for the generalized state-space model for consumption, returns, and their volatilities.

A.6.1 Stochastic Volatility of w_t^c

Let

$$y_{ct}^* = \log((w_t^c)^2) = h_{ct} + \log(\epsilon_{ct}^2),$$

$$h_{ct} = \psi_c + \delta_c(h_{c,t-1} - \psi_c) + \sigma_{\eta c}\eta_{ct},$$

where $\epsilon_{ct} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $\eta_{ct} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and they are independent. Let $e_{ct} = \log(\epsilon_{ct}^2) \sim \log(\chi^2(1))$. We estimate the model following Kim, Shephard, and Chib (1998). More specifically, we approximate e_{ct} using a mixture of Gaussian distributions,

$$e_{ct} = \sum_{j=1}^7 q_j \mathcal{N}(m_j - 1.2704, v_j^2), \quad (34)$$

or, equivalently,

$$e_{ct} \mid S_{ct} = j \sim \mathcal{N}(m_j - 1.2704, v_j^2), \quad (35)$$

where $\sum_{j=1}^7 q_j = 1$ and values of $\{m_j\}_{j=1}^7$ and $\{v_j\}_{j=1}^7$ could be found in Table 4 in Kim, Shephard, and Chib (1998). Below is the algorithm of estimating the unknown parameters $\{h_{ct}, S_{ct}, \psi_c, \delta_c, \sigma_{\eta c}\}$:

Step 1: initialize $\{S_{ct}\}_{t=1}^T$, ψ_c , δ_c , $\sigma_{\eta c}$.

Step 2: sample $\{h_{ct}\}_{t=1}^T$ from $h_c \mid y_c^*, S_c, \psi_c, \delta_c, \sigma_{\eta c}$ (Kalman Smoother step).

Step 3: sample S_{ct} from $p(S_{ct} \mid y_{ct}^*, h_{ct})$, where

$$\begin{aligned} p(S_{ct} = j \mid y_{ct}^*, h_{ct}) &\propto q_j \times p(y_{ct}^* \mid h_{ct}, S_{ct} = j) \\ &\propto q_j \times f_N(y_{ct}^* \mid h_{ct} + m_j - 1.2704, v_j^2). \end{aligned}$$

Step 4: update ψ_c , δ_c , $\sigma_{\eta c}$. We assign an inverse-gamma prior for $\sigma_{\eta c}^2$, i.e. $\sigma_{\eta c}^2 \sim \Gamma^{-1}(\frac{\sigma_0}{2}, \frac{s_\sigma}{2})$, hence:

$$\begin{aligned}
p(\sigma_{\eta c}^2 \mid y_c^*, S_c, \psi_c, \delta_c, h_c) &\propto p(h_c \mid \psi_c, \delta_c, \sigma_{\eta c}^2) \pi(\sigma_{\eta c}^2) \propto p(h_{c,1} \mid \psi_c, \delta_c, \sigma_{\eta c}^2) \prod_{t=0}^{T-1} p(h_{c,t+1} \mid \psi_c, \delta_c, \sigma_{\eta c}^2) \pi(\sigma_{\eta c}^2) \\
&\propto (\sigma_{\eta c}^2)^{-\frac{T}{2}} \exp \left\{ -\frac{\sum_{t=1}^{T-1} [h_{c,t+1} - \psi_c - \delta_c(h_{c,t} - \psi_c)]^2 + (h_{c,1} - \psi_c)^2 (1 - \delta_c^2)}{2\sigma_{\eta c}^2} \right\} (\sigma_{\eta c}^2)^{-\frac{\sigma_0}{2}-1} \exp \left\{ -\frac{s_\sigma}{2\sigma_{\eta c}^2} \right\}.
\end{aligned}$$

Therefore the conditional posterior is

$$\sigma_{\eta c}^2 \mid y_c^*, S_c, \psi_c, \delta_c, h_c \sim \Gamma^{-1} \left(\frac{T + \sigma_0}{2}, \frac{s_\sigma + \sum_{t=1}^{T-1} [h_{c,t+1} - \psi_c - \delta_c(h_{c,t} - \psi_c)]^2 + (h_{c,1} - \psi_c)^2 (1 - \delta_c^2)}{2} \right). \quad (36)$$

Let $\delta_c = 2\phi - 1$, and $\phi \sim \text{Beta}(\phi^{(1)}, \phi^{(2)})$. Therefore, the prior distribution of δ_c is $\pi(\delta_c) \propto (1 + \delta_c)^{\phi^{(1)}-1} (1 - \delta_c)^{\phi^{(2)}-1}$. The posterior distribution of δ_c is

$$\begin{aligned}
p(\delta_c \mid \psi_c, h_c, \sigma_{\eta c}^2) &\propto \pi(\delta_c) p(h_c \mid \psi_c, \delta_c, \sigma_{\eta c}^2) \\
p(h_c \mid \psi_c, \delta_c, \sigma_{\eta c}^2) &\propto (1 - \delta_c^2)^{\frac{1}{2}} \exp \left\{ -\frac{\sum_{t=1}^{T-1} [h_{c,t+1} - \psi_c - \delta_c(h_{c,t} - \psi_c)]^2 + (h_{c,1} - \psi_c)^2 (1 - \delta_c^2)}{2\sigma_{\eta c}^2} \right\},
\end{aligned}$$

$$\log p(h_c \mid \psi_c, \delta_c, \sigma_{\eta c}^2) \propto -\frac{\sum_{t=1}^{T-1} [h_{c,t+1} - \psi_c - \delta_c(h_{c,t} - \psi_c)]^2 + (h_{c,1} - \psi_c)^2 (1 - \delta_c^2)}{2\sigma_{\eta c}^2} + \frac{1}{2} \log(1 - \delta_c^2).$$

The above function is concave in δ_c for all values of $\phi^{(1)}$ and $\phi^{(2)}$. Hence, δ_c can be sampled using a reject-accept algorithm. Let

$$\hat{\delta}_c = \frac{\sum_{t=1}^{T-1} [h_{c,t+1} - \psi_c][h_{c,t} - \psi_c]}{\sum_{t=1}^{T-1} [h_{c,t} - \psi_c]^2}, \quad V_{\delta_c} = \frac{\sigma_{\eta c}^2}{\sum_{t=1}^{T-1} [h_{c,t} - \psi_c]^2}.$$

We first sample a proposal δ_c^* from a normal distribution $\mathcal{N}(\hat{\delta}_c, V_{\delta_c})$ and accept the new value δ_c^* with probability $\min\{1, \exp(g(\delta_c^*) - g(\delta_c^{i-1}))\}$, where

$$g(\delta_c) = \log(\pi(\delta_c)) - \frac{(h_{c,1} - \psi_c)^2 (1 - \delta_c^2)}{2\sigma_{\eta c}^2} + \frac{1}{2} \log(1 - \delta_c^2).$$

Finally, we assign a diffuse prior for ψ_c .

$$p(\psi_c \mid h_c, \delta_c, \sigma_{\eta_c}^2) \propto p(h_c \mid \delta_c, \psi_c, \sigma_{\eta_c}^2) \\ \propto \exp \left\{ -\frac{\sum_{t=1}^{T-1} [h_{c,t+1} - \psi_c - \delta_c(h_{ct} - \psi_c)]^2 + (h_{c,1} - \psi_c)^2(1 - \delta_c^2)}{2\sigma_{\eta_c}^2} \right\}.$$

Therefore, we obtain the posterior distribution of ψ_c as follows

$$\psi_c \mid h_c, \delta_c, \sigma_{\eta_c}^2 \sim \mathcal{N}(\hat{\psi}_c, \sigma_{\psi}^2),$$

where $\sigma_{\psi}^2 = \frac{\sigma_{\eta_c}^2}{[(T-1)(1-\delta_c)^2 + (1-\delta_c^2)]}$ and $\hat{\psi}_c = \sigma_{\psi}^2 \left\{ \frac{(1-\delta_c^2)h_{c,1}}{\sigma_{\eta_c}^2} + \frac{(1-\delta_c)\sum_{t=1}^{T-1} [h_{c,t+1} - \delta_c h_{ct}]}{\sigma_{\eta_c}^2} \right\}$.

Step 5: go back to step 2 until convergence.

A.6.2 Stochastic Volatility of f_t

The consumption mean shock f_t follows a normal distribution with the stochastic volatility process given by

$$y_{ft}^* = \log(f_t^2) = h_{ft} + \log(\epsilon_{ft}^2),$$

$$h_{ft} = \delta_f h_{f,t-1} + \sqrt{1 - \delta_f^2} \eta_{ft},$$

where $\epsilon_{ft} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $\eta_{ft} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and they are independent. Let $e_{ft} = \log(\epsilon_{ft}^2) \sim \log(\chi^2(1))$. Note that the above process for h_{ft} is just an AR(1) normalized to have unconditional zero mean and unit variance. As before, we approximate e_{ft} using a mixture of Gaussian distributions, that is, $e_{ft} = \sum_{j=1}^7 q_j \mathcal{N}(m_j - 1.2704, v_j^2)$. The sampling algorithm is given below.

Step 1: initialize $\{S_{ft}\}_{t=1}^T, \delta_f$.

Step 2: sample $\{h_{ft}\}_{t=1}^T$ from $h_f \mid y_f^*, S_f, \delta_f$ (Kalman smoother step).

Step 3: sample S_{ft} from $p(S_{ft} \mid y_{ft}^*, h_{ft})$.

Step 4: update δ_f using the Metropolis algorithm. Similar to the prior distribution of δ_c , we assume $\delta_f = 2\phi - 1$, and $\phi \sim \text{Beta}(\phi^{(1)}, \phi^{(2)})$. The posterior distribution of δ_f is

$$p(\delta_f \mid h_f, \psi_f, \gamma_f = 0) \propto \pi(\delta_f) p(h_f \mid \psi_f, \delta_f, \gamma_f = 0) \\ \propto (1 + \delta_f)^{\phi^{(1)}-1} (1 - \delta_f)^{\phi^{(2)}-1} (1 - \delta_f^2)^{-\frac{T-1}{2}} \exp \left\{ -\frac{\sum_{t=2}^T (h_{ft} - \delta_f h_{f,t-1})^2}{2(1 - \delta_f^2)} - \frac{h_{f1}^2}{2} \right\}.$$

We draw δ_f using a Metropolis algorithm:

- Step M1. initialize δ_f^0 .
- Step M2. draw δ_f^* from a normal distribution $\mathcal{N}(\delta_f^{(i-1)}, c_{mh}^2)^{33}$.
- Step M3. calculate $\rho(\delta_f^*, \delta_f^{(i-1)}) = \min\{1, \frac{p(\delta_f^* | h_f)}{p(\delta_f^{(i-1)} | h_f)}\}$;
- Step M4. set $\delta_f^{(i)} = \delta_f^{(i-1)}$ with probability $1 - \rho(\delta_f^*, \delta_f^{(i-1)})$ and $\delta_f^{(i)} = \delta_f^*$ with probability $\rho(\delta_f^*, \delta_f^{(i-1)})$.

Step 5: go back to Step 2 until convergence.

A.6.3 Stochastic Volatility of asset returns

Let

$$\mathbf{y}_{rt}^* = \begin{matrix} \boldsymbol{\kappa}_0 \\ N \times 1 \end{matrix} + \begin{matrix} \boldsymbol{\kappa}_1 \\ N \times 1 \end{matrix} h_{rt} + \mathbf{e}_{rt},$$

$$h_{rt} = \delta_r h_{r,t-1} + \sqrt{1 - \delta_r^2} \eta_{rt},$$

where

$$\mathbf{y}_{rt}^* = \begin{pmatrix} \log(w_{1t}^r)^2 \\ \log(w_{2t}^r)^2 \\ \vdots \\ \log(w_{Nt}^r)^2 \end{pmatrix}, \mathbf{e}_{rt} = \begin{pmatrix} e_{1,rt} \\ e_{2,rt} \\ \vdots \\ e_{N,rt} \end{pmatrix} = \begin{pmatrix} \log(\epsilon_{1t}^2) \\ \log(\epsilon_{2t}^2) \\ \vdots \\ \log(\epsilon_{Nt}^2) \end{pmatrix},$$

³³ c_{mh} determines the step size in the Metropolis algorithm. We aim at choosing c_{mh} such as the frequency of accepting a new δ_f is around 50%.

and where the ϵ_{it} shocks are independent across different assets. In the above model, we assume that one hidden state, h_{rt} , drives the common component of asset-specific stochastic volatilities. In order to identify the model, we normalize this component h_{rt} to have zero mean and unit variance. In order to simplify the estimation, we further exclude κ_0 by demeaning \mathbf{y}_{rt}^* to have the same sample mean as \mathbf{e}_{rt} . Therefore, the model is simplified as

$$\kappa_0 = \mathbf{y}_{rt}^* - \bar{\mathbf{y}}_{rt}^*,$$

$$\bar{\mathbf{y}}_{rt}^* = \kappa_1 h_{rt} + \mathbf{e}_{rt},$$

$$h_{rt} = \delta_r h_{r,t-1} + \sqrt{1 - \delta_r^2} \eta_{rt},$$

for $i = 1, 2, \dots, N$, $e_{i,rt} = \sum_{j=1}^7 q_j \mathcal{N}(m_j - 1.2704, v_j^2)$, or equivalently, $e_{i,rt} \mid S_{i,rt} = j \sim \mathcal{N}(m_j - 1.2704, v_j^2)$. Therefore, $\bar{y}_{i,rt}^* \mid \kappa_1^{(i)}, h_{rt}, S_{i,rt} = j \stackrel{\text{iid}}{\sim} \mathcal{N}(\kappa_1^{(i)} h_{rt} + m_j - 1.2704, v_j^2)$.

Let

$$\bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)} = \begin{pmatrix} \bar{y}_{r1}^{*,(i)} - m_{i1} + 1.2704 \\ \vdots \\ \bar{y}_{rT}^{*,(i)} - m_{iT} + 1.2704 \end{pmatrix}, \mathbf{V}_{\mathbf{r}}^{(i)} = \begin{pmatrix} v_{i1}^2 & & \\ & \ddots & \\ & & v_{iT}^2 \end{pmatrix}, \mathbf{H}_{\mathbf{r}} = \begin{pmatrix} h_{r1} \\ \vdots \\ h_{rT} \end{pmatrix},$$

where $m_{it} = m_j$ and $v_{it}^2 = v_j^2$ if $S_{it} = j$.

Assuming a Jeffreys (diffuse) prior for $\kappa_1^{(i)}$, we have

$$\begin{aligned} p(\kappa_1^{(i)} \mid \bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)}, \mathbf{H}_{\mathbf{r}}, \{S_{i,rt}\}_{t=1}^T) &\propto p(\bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)} \mid \kappa_1^{(i)}, \mathbf{H}_{\mathbf{r}}, \{S_{i,rt}\}_{t=1}^T) \\ &\propto \exp\left\{-\frac{1}{2}(\bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)} - \mathbf{H}_{\mathbf{r}}\kappa_1^{(i)})^\top (\mathbf{V}_{\mathbf{r}}^{(i)})^{-1} (\bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)} - \mathbf{H}_{\mathbf{r}}\kappa_1^{(i)})\right\}, \end{aligned}$$

$$\kappa_1^{(i)} \mid \bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)}, \mathbf{H}_{\mathbf{r}}, \{S_{i,rt}\}_{t=1}^T \sim \mathcal{N}(\hat{\kappa}_1^{(i)}, [\mathbf{H}_{\mathbf{r}}^\top (\mathbf{V}_{\mathbf{r}}^{(i)})^{-1} \mathbf{H}_{\mathbf{r}}]^{-1}),$$

where $\hat{\kappa}_1^{(i)} = [\mathbf{H}_{\mathbf{r}}^\top (\mathbf{V}_{\mathbf{r}}^{(i)})^{-1} \mathbf{H}_{\mathbf{r}}]^{-1} \mathbf{H}_{\mathbf{r}}^\top (\mathbf{V}_{\mathbf{r}}^{(i)})^{-1} \bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)}$.

Finally, we also need to update $\{S_{i,rt}\}_{t=1}^T$ as follows

$$\begin{aligned} p(S_{i,rt} = j \mid \bar{y}_{rt}^{*,(i)}, h_{rt}, \kappa_1^{(i)}) &\propto q_j \times p(\bar{y}_{rt}^{*,(i)} \mid h_{rt}, S_{i,rt} = j, \kappa_1^{(i)}) \\ &\propto q_j \times f_N(\bar{y}_{rt}^{*,(i)} \mid \kappa_1^{(i)} h_{rt} + m_j - 1.2704, v_j^2) \end{aligned}$$

The posterior draws of h_{rt} and δ_r are then obtained in a similar way as before.

A.6.4 The consumption growth equation coefficients

Since shocks in equation (1) and (2) are uncorrelated, we can sample the unknown parameters equation by equation. Let's introduce some notations.

$$\boldsymbol{\rho}^c = \begin{pmatrix} \mu_c \\ \rho_0 \\ \vdots \\ \rho_{\bar{s}} \end{pmatrix}, \boldsymbol{\Delta C} = \begin{pmatrix} \Delta C_{0,1} \\ \Delta C_{1,2} \\ \vdots \\ \Delta C_{T-1,T} \end{pmatrix}, \boldsymbol{\Sigma}_w^c = \begin{pmatrix} \sigma_{c1}^2 & 0 & \dots & 0 \\ 0 & \sigma_{c2}^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{cT}^2 \end{pmatrix}, \mathbf{X}^c = \begin{pmatrix} 1 & f_1 & \dots & f_{1-\bar{s}} \\ 1 & f_2 & \dots & f_{2-\bar{s}} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & f_T & \dots & f_{T-\bar{s}} \end{pmatrix}.$$

Then the posterior distribution of $\boldsymbol{\rho}^c$ under a flat prior is ³⁴

$$\boldsymbol{\rho}^c \mid \boldsymbol{\Delta C}, \mathbf{X}^c, \boldsymbol{\Sigma}_w^c \sim \mathcal{N}(\hat{\boldsymbol{\rho}}_{gls}, [(\mathbf{X}^c)^\top (\boldsymbol{\Sigma}_w^c)^{-1} \mathbf{X}^c]^{-1}),$$

$$\hat{\boldsymbol{\rho}}_{gls} = [(\mathbf{X}^c)^\top (\boldsymbol{\Sigma}_w^c)^{-1} \mathbf{X}^c]^{-1} (\mathbf{X}^c)^\top (\boldsymbol{\Sigma}_w^c)^{-1} \boldsymbol{\Delta C}.$$

³⁴**Proof.** The likelihood function of the data is

$$\boldsymbol{\Delta C} \mid \boldsymbol{\rho}^c, \mathbf{X}^c, \boldsymbol{\Sigma}_w^c \sim \mathcal{N}(\mathbf{X}^c \boldsymbol{\rho}^c, \boldsymbol{\Sigma}_w^c),$$

Applying the diffuse prior for $\boldsymbol{\rho}^c$, the posterior distribution of $\boldsymbol{\rho}^c$ is

$$\begin{aligned} p(\boldsymbol{\rho}^c \mid \boldsymbol{\Delta C}, \mathbf{X}^c, \boldsymbol{\Sigma}_w^c) &\propto p(\boldsymbol{\Delta C} \mid \boldsymbol{\rho}^c, \mathbf{X}^c, \boldsymbol{\Sigma}_w^c) \\ &\propto \exp\left\{-\frac{1}{2}(\boldsymbol{\Delta C} - \mathbf{X}^c \boldsymbol{\rho}^c)^\top (\boldsymbol{\Sigma}_w^c)^{-1} (\boldsymbol{\Delta C} - \mathbf{X}^c \boldsymbol{\rho}^c)\right\} \\ &\propto \exp\left\{-\frac{1}{2}(\boldsymbol{\rho}^c - \hat{\boldsymbol{\rho}}_{gls})^\top (\mathbf{X}^c)^\top (\boldsymbol{\Sigma}_w^c)^{-1} \mathbf{X}^c (\boldsymbol{\rho}^c - \hat{\boldsymbol{\rho}}_{gls})\right\}. \end{aligned}$$

■

A.6.5 The excess returns equation coefficients

Let

$$\boldsymbol{\rho}_i^r = \begin{pmatrix} \mu_{ri} \\ \rho_i^r \\ \beta_{fi} \\ \beta_{ri} \end{pmatrix}, \quad \mathbf{X}_i^r = \begin{pmatrix} 1 & f_1 & \sigma_{f,0}^2 & \sigma_{r,0}^2 \\ 1 & f_2 & \sigma_{f,1}^2 & \sigma_{r,1}^2 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & f_T & \sigma_{f,T-1}^2 & \sigma_{r,T-1}^2 \end{pmatrix}, \quad \text{and} \quad \boldsymbol{\Sigma}_{wr,i} = \begin{pmatrix} \sigma_{r,i1}^2 & 0 & \dots & 0 \\ 0 & \sigma_{r,i2}^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{r,iT}^2 \end{pmatrix}.$$

The conditional posterior under a flat prior is then

$$\boldsymbol{\rho}_i^r \mid \mathbf{R}_i, \mathbf{X}_i^r, \{\sigma_{ft}^2\}_{t=0}^{T-1}, \{\sigma_{r,it}^2\}_{t=0}^T \sim \mathcal{N}(\hat{\boldsymbol{\rho}}_{i,gl}^r, [(\mathbf{X}_i^r)^\top \boldsymbol{\Sigma}_{wr,i}^{-1} \mathbf{X}_i^r]^{-1}),$$

$$\hat{\boldsymbol{\rho}}_{i,gl}^r = [(\mathbf{X}_i^r)^\top \boldsymbol{\Sigma}_{wr,i}^{-1} \mathbf{X}_i^r]^{-1} (\mathbf{X}_i^r)^\top \boldsymbol{\Sigma}_{wr,i}^{-1} \mathbf{R}_i.$$

A.6.6 Drawing the conditional consumption mean shocks

To sample $\{f_t\}_{t=1}^T$, let

$$\mathbf{y}_t = \begin{pmatrix} \Delta c_{t-1,t} \\ \mathbf{r}_t^e \end{pmatrix}, \quad \tilde{\boldsymbol{\mu}} = \begin{pmatrix} \mu_c \\ \boldsymbol{\mu}_r + \boldsymbol{\beta}_f \sigma_{f,t-1}^2 + \boldsymbol{\beta}_r \sigma_{r,t-1}^2 \end{pmatrix}, \quad \mathbf{H} = \begin{pmatrix} \rho_0 & \rho_1 & \dots & \rho_{\bar{S}} \\ \boldsymbol{\rho}^r & \mathbf{0}_N & \dots & \mathbf{0}_N \end{pmatrix}$$

$$\mathbf{w}_t = \begin{pmatrix} w_t^c \\ \mathbf{w}_t^r \end{pmatrix} \sim \mathcal{N}(\mathbf{0}_{N+1}, \boldsymbol{\Sigma}_t), \quad \mathbf{z}_t = \begin{pmatrix} f_t \\ f_{t-1} \\ \vdots \\ f_{t-\bar{S}} \end{pmatrix}, \quad \boldsymbol{\Sigma}_t = \begin{pmatrix} \sigma_{ct}^2 & 0 & \dots & 0 \\ 0 & \sigma_{r,1t}^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{r,Nt}^2 \end{pmatrix}.$$

Hence, the joint distribution of observables and f shocks is

$$\begin{pmatrix} \mathbf{y}_t \\ \mathbf{z}_t \end{pmatrix} \mid \mathcal{I}_{t-1}, \tilde{\boldsymbol{\mu}}, \mathbf{H}, \boldsymbol{\Sigma}_t, \sigma_{ft}^2 \sim \mathcal{N} \left(\begin{pmatrix} \tilde{\boldsymbol{\mu}} \\ \mathbf{F} \mathbf{z}_{t-1} \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Omega}_t & \mathbf{H} \\ \mathbf{H}^\top & \boldsymbol{\Psi}_t \end{pmatrix} \right), \quad (37)$$

$$\text{where } \boldsymbol{\Psi}_t = \begin{pmatrix} \sigma_{ft}^2 & \mathbf{0}_{\bar{S}}^\top \\ \mathbf{0}_{\bar{S}} & \mathbf{0}_{\bar{S} \times \bar{S}} \end{pmatrix}, \mathbf{F} := \underbrace{\begin{bmatrix} \mathbf{0}'_{\bar{S}} & 0 \\ I_{\bar{S}} & \mathbf{0}_{\bar{S}} \end{bmatrix}}_{(\bar{S}+1) \times (\bar{S}+1)}, \boldsymbol{\Omega}_t = \mathbf{H} \boldsymbol{\Psi}_t \mathbf{H}^\top + \boldsymbol{\Sigma}_t.$$

We then use the Kalman smoother to draw f_t , following the same procedure as in Appendix A.3.

A.6.7 Model Comparison

The model comparison for the restricted and unrestricted specification in Table 5 of Section V.3 is performed using Bayes Factors (i.e., the the marginal likelihoods of the various models) and posterior probabilities. Based on equation (37), the likelihood function of the observed data $\{\mathbf{y}_t\}_{t=1}^T$ (after integrating out f_t) is

$$\begin{aligned} p(\mathbf{Y} \mid \tilde{\boldsymbol{\mu}}, \mathbf{H}, \{\boldsymbol{\Sigma}_t\}_{t=1}^T, \{\sigma_{ft}^2\}_{t=1}^T) &= \prod_{t=1}^T f_N(\mathbf{y}_t \mid \tilde{\boldsymbol{\mu}}, \boldsymbol{\Omega}_t) \\ &= (2\pi)^{-\frac{T(N+1)}{2}} \prod_{t=1}^T |\boldsymbol{\Omega}_t|^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(\mathbf{y}_t - \tilde{\boldsymbol{\mu}})^\top \boldsymbol{\Omega}_t^{-1}(\mathbf{y}_t - \tilde{\boldsymbol{\mu}})\right\} \\ &= (2\pi)^{-\frac{T(N+1)}{2}} \prod_{t=1}^T |\mathbf{H} \boldsymbol{\Psi}_t \mathbf{H}^\top + \boldsymbol{\Sigma}_t|^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(\mathbf{y}_t - \tilde{\boldsymbol{\mu}})^\top (\mathbf{H} \boldsymbol{\Psi}_t \mathbf{H}^\top + \boldsymbol{\Sigma}_t)^{-1}(\mathbf{y}_t - \tilde{\boldsymbol{\mu}})\right\}. \end{aligned}$$

A full Bayesian analysis requires us to specify a proper prior for $\boldsymbol{\mu}$, \mathbf{H} and the parameters underlying stochastic volatility processes. The difficulty is that there is no closed-form solution for the marginal likelihood of data. Furthermore, a flat prior for $(\boldsymbol{\mu}, \mathbf{H}, \boldsymbol{\beta}_f, \boldsymbol{\beta}_r)$ is improper, hence, the marginal likelihood of the data is unnormalized and there would be an undetermined constant term in model comparison. And even if we were to assign a proper prior, the numerical integration would be very imprecise due to the high dimensionality of the parameter and hidden state spaces. Therefore, we follow the literature and approximate

the marginal likelihood of the data using the Schwartz criterion (i.e., a Laplace, or particular second order approximation of the marginal likelihood), as follows

$$\log(BF_{1,2}) \approx \log p(\mathbf{Y} \mid \hat{\theta}(\mathbb{M}_1)) - \log p(\mathbf{Y} \mid \hat{\theta}(\mathbb{M}_2)) - \frac{d_1 - d_2}{2} \log(T), \quad (38)$$

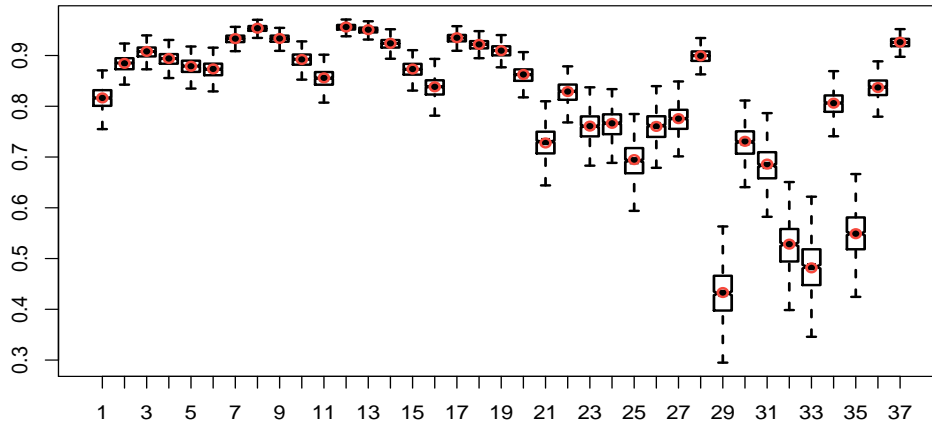
where \mathbf{Y} is the observed data, T is the sample size, \mathbb{M}_1 and \mathbb{M}_2 represent model 1 and 2, $\hat{\theta}(\mathbb{M}_1)$ and $\hat{\theta}(\mathbb{M}_2)$ are posterior mean of parameter θ under model 1 and 2, and d_1 and d_2 are model dimensions.³⁵ The Bayes Factor in equation (38) ignores the prior distribution, hence we do not need to change our current improper prior. Note that the model selection based on the above is analogous to likelihood ratio testing (the LR statistic is proportional to the first two terms in the equation (38)) or BIC based model selection.

Posterior model probabilities are then computed using the above approximation of the Bayes Factor and equal prior probability for all specifications; for example, the posterior probability of model 1 is computed as $\frac{BF_{1,i}}{\sum_j BF_{j,i}}$, where the identity of the reference model i is irrelevant.

³⁵Note that the vector of “parameters” encompasses both the “frequentist” parameters $(\mu_c, \boldsymbol{\mu}_r, \boldsymbol{\beta}_f, \boldsymbol{\beta}_r, \boldsymbol{\rho}^c, \boldsymbol{\rho}^r, \psi_c, \delta_c, \sigma_{\eta c}, \delta_f, \boldsymbol{\kappa}_0, \boldsymbol{\kappa}_0, \delta_r)$ and the latent states $(\{f_t, \sigma_{ct}^2, \sigma_{ft}^2, \sigma_{rt}^2\}_{t=1}^T)$.

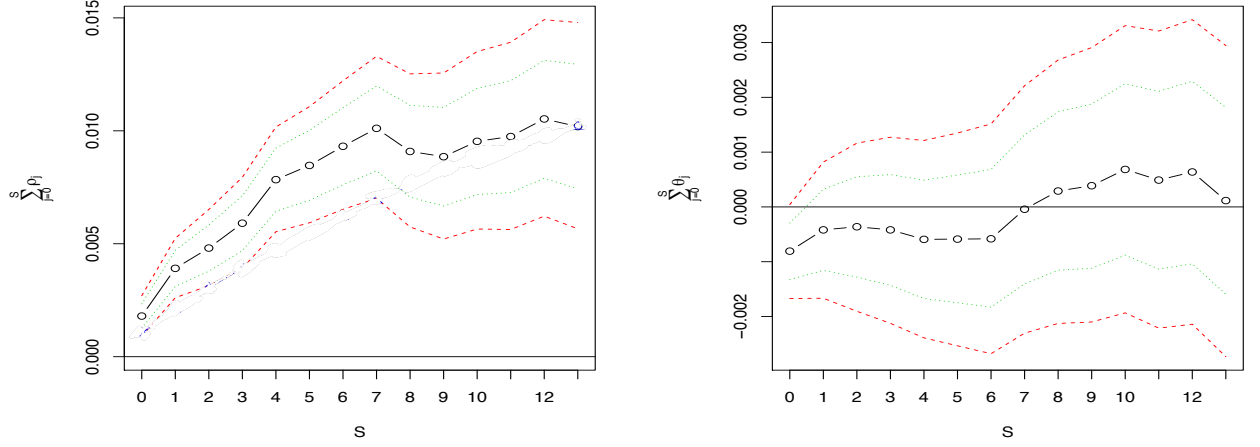
A.7 Additional figures

Figure A5: Variance decomposition of asset returns (average of 1,000 simulations)



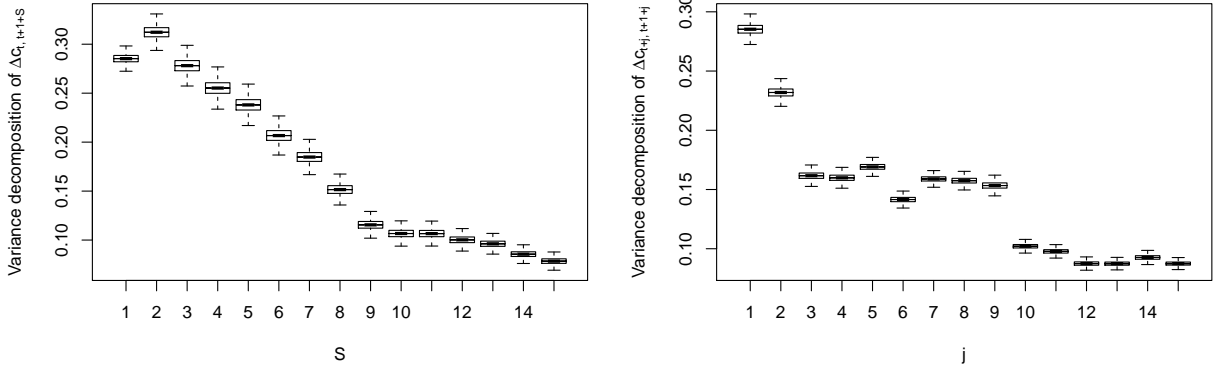
Box-plots (95% percentiles) of the percentage of time series variances of individual stock portfolio returns explained by the f component in the one factor model, as estimated by the state-space model. Red circles denote true calibrated values.

Figure A6: Consumption growth response to the latent factors f_t and g_t shocks.



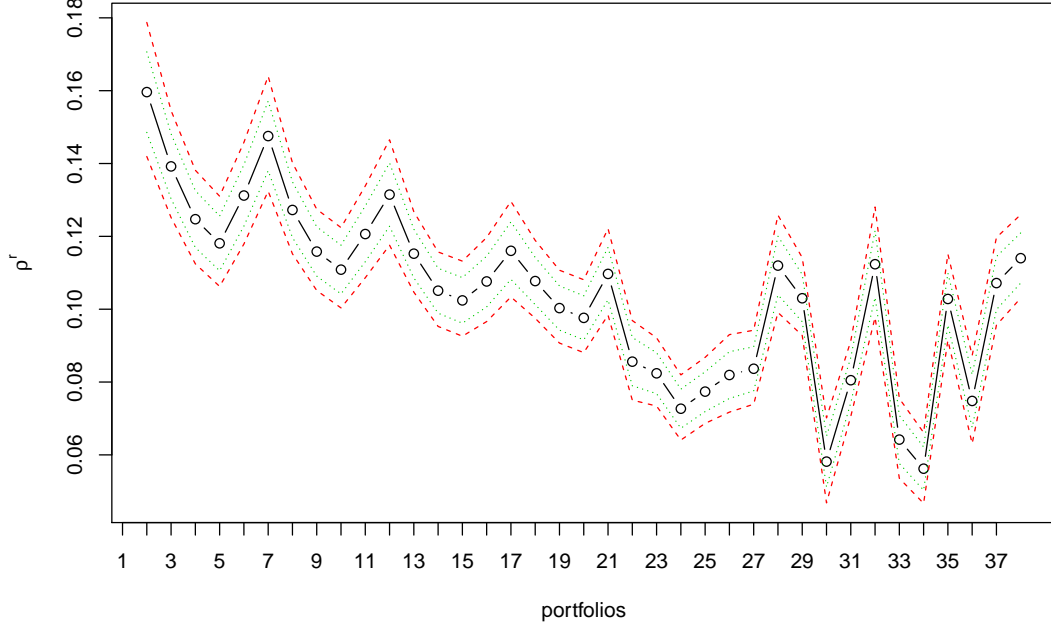
Posterior means (continuous line with circles) and centered posterior 90% (dashed line) and 68% (dotted line) coverage regions. Estimation based on the two-factor model in equations (7) and (8). Left panel: Cumulated consumption response to common factor, f_t , shocks. Right panel: Cumulated consumption response to bond factor (g_t) shocks. Triangles denote a potential AR(1), à la Bansal and Yaron (2004), calibrated to match our estimates.

Figure A7: Variance of consumption growth explained by the MA components f and g .



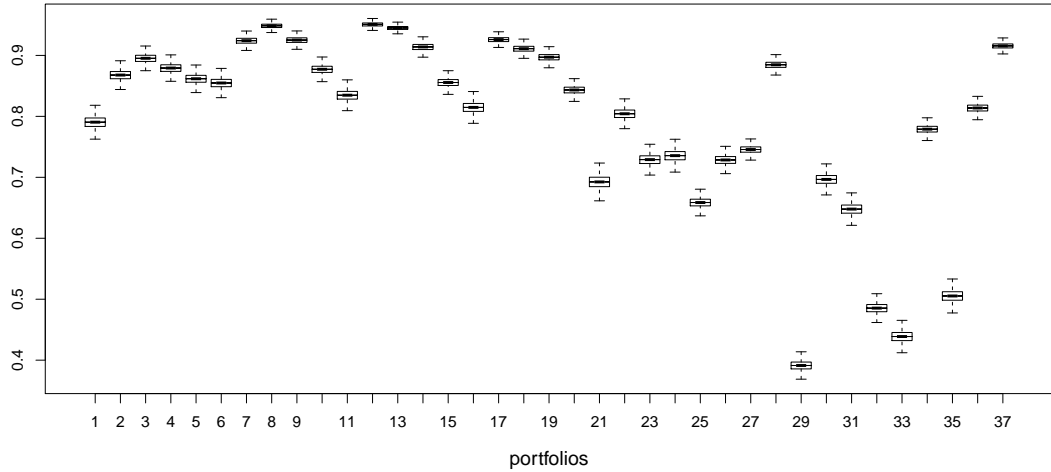
Box-plots (posterior 95% coverage area) of the percentage of time series variances of consumption growth explained by the MA components f and g . Left panel: Cumulated consumption growth $\Delta c_{t,t+1+S}$. Right panel: One period consumption growth $\Delta c_{t+j,t+1+j}$.

Figure A8: Common factor loadings (ρ^r) of the stock portfolios in the two-factor model.



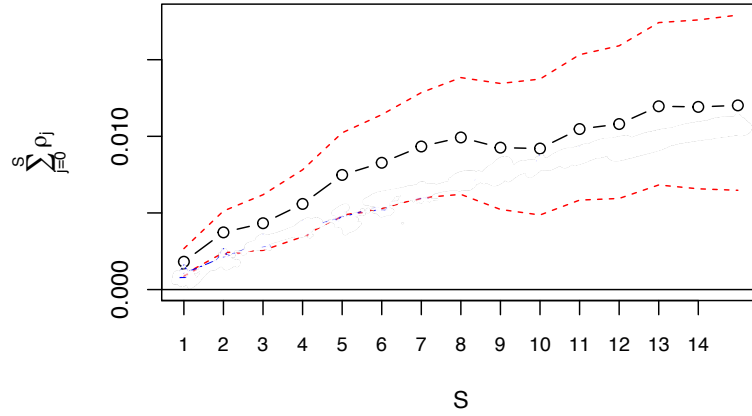
The graph presents posterior means of the stocks factor loadings on f_t (continuous line with circles) and centered posterior 90% (dashed line) and 68% (dotted line) coverage regions in the two latent factors model. Ordering of portfolios: 25 Fama and French (1992) size and book-to-market sorted portfolios (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 A9: 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.

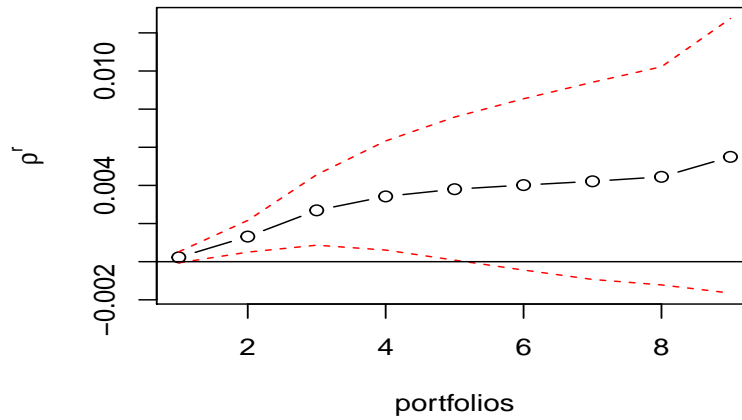
Figure A10: Cumulated consumption response to, and asset loadings on, f shocks.



Panel A: Cumulative response function of consumption growth to one standard deviation shock to f_t .



Panel B: Common factor loadings (ρ^r) of the stock portfolios on f_t .



Panel C: Bond loadings on f_t .

Posterior means (black line with circles) and centered posterior 90% (dashed red lines) coverage regions obtained using the one-factor model with unrestricted stochastic volatilities. Quarterly data, 1961:Q3-2017:Q2.

A.8 Additional tables

Table A3: Expected Excess Returns and Consumption Risk, 1967:Q3-2017:Q2

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