



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■



Consumption in Asset Returns

Svetlana Bryzgalova
Jiantao Huang
Christian Julliard

SRC Discussion Paper No 92
November 2025

This is a revised version. The previous version was dated January 2020.



Systemic Risk Centre

Discussion Paper Series

Abstract

Using information in returns we identify the stochastic process of consumption. We find that aggregate consumption reacts over multiple quarters to innovations spanned by financial markets. This persistent component accounts for over a quarter of consumption variation. These shocks command a large and significant risk premium, driving a large share of stocks and a small yet significant fraction of bonds time series variation. Nevertheless, we find no support for stochastic volatility of consumption driving time-varying risk premia. Finally, an otherwise standard recursive utility model based on our estimated process explains equity premium and risk-free rate puzzles with low risk aversion.

JEL Classification Codes: E21, E27, G12, E43, C11.

Keywords: Consumption Dynamics, Asset Returns, Consumption-Based Asset Pricing, Term Structure.

This paper is published as part of the Systemic Risk Centre's Discussion Paper Series. The support of the Economic and Social Research Council (ESRC) in funding the SRC is gratefully acknowledged [grant number ES/R009724/1].

Svetlana Bryzgalova, London Business School and Systemic Risk Centre, London School of Economics and Political Science

Jiantao Huang, Faculty of Business and Economics, University of Hong Kong

Christian Julliard, Department of Finance, Financial Markets Group and Systemic Risk Centre, London School of Economics and Political Science, and Centre for Economic Policy Research

Published by
 Systemic Risk Centre
 The London School of Economics and Political Science
 Houghton Street
 London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© Svetlana Bryzgalova, Jiantao Huang and Christian Julliard, submitted 2025

Consumption in Asset Returns*

Svetlana Bryzgalova[†]

Jiantao Huang[‡]

Christian Julliard[§]

August 12, 2025

Abstract

Using information in returns we identify the stochastic process of consumption. We find that aggregate consumption reacts over multiple quarters to innovations spanned by financial markets. This persistent component accounts for over a quarter of consumption variation. These shocks command a large and significant risk premium, driving a large share of stocks and a small yet significant fraction of bonds time series variation. Nevertheless, we find no support for stochastic volatility of consumption driving time-varying risk premia. Finally, an otherwise standard recursive utility model based on our estimated process explains equity premium and risk-free rate puzzles with low risk aversion.

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, Jonathan Parker, 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, New York 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, and AFA 2022. We are extremely thankful for thoughtful and stimulating inputs from Leonid Kogan (the co-editor), the associated editor, and two anonymous referees. We have read The *Journal of Finance* disclosure policy and have no conflicts of interest to disclose. Any errors or omissions are the responsibility of the authors.

[†]London Business School; sbryzgalova@london.edu.

[‡]Faculty of Business and Economics, University of Hong Kong; huangjt@hku.hk.

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

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 empirical performance and policy implications of models. However, using consumption data alone, it is difficult 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 were of the type usually postulated in the literature, standard methods would fail to accurately recover the magnitude and persistence of its conditional mean and volatility – crucial ingredients for asset pricing. Instead, we develop a novel state-space approach that sheds light on the dark matter of consumption-based models. We find that consumption growth is highly predictable at business cycle frequency by shocks spanned by financial markets that account for *more than a quarter* of its variance. Furthermore, these shocks to the conditional consumption mean command a large and significant risk premium in the cross-section of asset returns. Nevertheless, we find no support for the stochastic volatilities of consumption driving time-varying risk premia. Our findings hold across a range of reasonable measures of consumption growth, choices of base assets, specifications of external predictors, assumptions about mismeasurement in consumption, and misspecification of the model.

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. This is exactly the insight that motivated the modeling of the conditional moments of consumption as a function of information in stocks and bonds, for example, in Harvey (1988) and Kandel and Stambaugh (1990). 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 innovations reflected in both consumption and financial assets. In addition, our identification strategy is confirmed

¹See, e.g., Beeler and Campbell (2012), Campbell (2017), Cochrane (2005), and Ludvigson (2013) for a review of the empirical challenges of consumption-based asset pricing.

by data on shareholders' and non-shareholders' consumption: Only the consumption of the former reacts to financial market shocks, with the aggregate data being in between the two.

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. As in Blanchard and Quah (1989), we leverage the fact that any covariance stationary process for consumption, including but not limited to those used in habits and long-run risk models, can be represented by a Moving Average (MA). As in the structural vector autoregression literature (e.g., Christiano, Eichenbaum, and Evans (2005) and Sims and Zha (2006)), we impose an economic restriction to identify the impulse response of consumption to shocks spanned by financial markets, and, hence, reliably estimate its conditional mean and persistence. This, in turn, allows us to correctly recover the volatility dynamics, because the latter, as we show, can be consistently estimated *only if* the conditional consumption mean and its predictability are properly captured. Furthermore, the MA modeling is also motivated by measurement problems in consumption data, such as time aggregation (Breeden, Gibbons, and Litzenberger (1989)) and benchmarking (Wilcox (1992)), that we formally account for in our analysis.

The crucial role of the conditional mean in driving consumption growth that we uncover (over 25% of total consumption variance) is more than twice what is normally assumed in leading macro-finance models. For instance, the conditional mean contribution to consumption variance is zero in Lucas (1987) that calculates the cost of business cycle fluctuation as well as in models that, for parsimony, do not encode any predictability in consumption (e.g., the habit framework of Campbell and Cochrane (1999) and the rare disasters model of Barro (2006)). It is instead about 12% in the long-run risk calibrations of Bansal and Yaron (2004) and Hansen, Heaton, Lee, and Roussanov (2007) (to the best of our knowledge, the largest used in the literature).

The impact of shocks to the consumption growth mean is well identified, sharply estimated, and economically large: A one-standard-deviation innovation implies a response of consumption growth of about 1% over the next two years. Moreover, it takes about two years for the conditional mean shocks to be fully reflected in the consumption process, producing substantial predictability, and no further reaction is apparent after this horizon.

The consumption shocks spanned by financial markets generate a clear business cycle pattern in the conditional expectation of consumption growth, have a large and significant risk premium in the cross-section of asset return, with an implied annualized Sharpe ratio of about 0.5. This in turn implies, as we show, that an otherwise standard Epstein and Zin (1989) preference setting calibrated with our estimated process can match the market equity risk premium with a low level of risk aversion: Point estimates range from 6.5 to 27, with values as low as 10 always within the 90% confidence intervals. Furthermore, assuming an elasticity of intertemporal substitution above one, the model can also match the level and volatility of the risk-free rate. Nevertheless, we uncover a challenge for macro-finance modeling. Often, representative agent models obtain equilibrium time-varying risk premia by postulating stochastic volatilities (SVs) in consumption, which leads to a dependence of excess returns on the conditional volatilities. We find no evidence supporting this mechanism, even after controlling for time-averaging bias in the consumption process. However, we detect a significantly positive comovement between the squared VXO index and the stochastic variance of the consumption mean shock, with a posterior mean correlation coefficient of about 45%.²

Crucially, all of our findings are obtained without restricting either the nature or number of factors driving asset returns, allowing for general and flexible SV processes in both consumption and returns (with short- and long-run SVs, as in Schorfheide, Song, and Yaron (2018), as well as the leverage effect), and hold using quarterly, monthly, and mixed-frequency estimations.

Estimation consistency of a volatility process crucially depends on the correct specification of the conditional mean, especially when the latter is persistent and difficult to detect, as postulated in leading models. For this reason, the key driver of our findings is the correct identification of the conditional mean process of consumption. Following the insight of Jagannathan and Marakani (2016), we show that using consumption data alone is not enough to reliably identify the consumption process, and we leverage instead the information in a rich cross-section of asset returns. Similar to Schorfheide, Song, and Yaron (2018),

²The VXO is an implied volatility index for the S&P 100 single-name stocks produced by the Chicago Board Options Exchange. We employ VXO, rather than the more canonical VIX, due to its longer data series and very high correlation with the latter.

we use a state-space approach to filter the latent shocks to the consumption process. But unlike them, we do not constrain the consumption mean process to a stringent parametric functional form.

Our econometric framework is flexible and comprehensive. We model consumption growth as the sum of two independent processes: a high-order MA that (potentially) co-moves with returns and a transitory component orthogonal to financial assets (and external predictors). Innovations to returns are, in turn, modeled as depending (potentially) on the shocks to the persistent component of consumption, potentially external predictors, and a set of orthogonal latent factors that drive the residual comovement among assets. As in Wold's decomposition, we estimate the MA representation of consumption (equivalently, its impulse response). Nevertheless, unlike in the existing literature, we identify the MA innovations by using a large cross-section of asset returns. These economic innovations are allowed to follow distinct SV processes, which can further act as predictors for asset returns, as in many equilibrium macro-finance models.³

With extensive simulations, we show that our method, even in samples as small as the historical sample, robustly recovers the consumption dynamics. This is a difficult task, and indeed we find that popular model selection procedures (e.g., AIC and BIC criteria) fail to identify the type of consumption processes postulated in the literature. Rather, our approach correctly recovers the conditional mean of consumption growth and, therefore, its volatility process. As we show, missed and/or misspecified predictability in the conditional mean process generates spurious volatility dynamics similar to what is documented in the existing literature.

³Our general MA formulation for the consumption process allows shocks to be potentially driven by a jump process. However, in the empirical estimation we model all innovations via continuous distributions. That is, we do not explicitly model the large jumps often used to calibrate rare disaster models. Nevertheless, note that these typical calibrations (e.g., Barro (2005), Wachter (2013)) are based on the assumption that the cumulative disasters, recorded in the data over multiple years, take place during a single year (for a detailed discussion of this counterfactual assumption, see Blanchard and Constantinides (2008) and Julliard and Ghosh (2012)). At the quarterly and monthly frequencies we focus on, properly taking into account the time series properties of disasters would imply modelling quite small jumps. For instance, the largest recorded contraction in quarterly nondurable consumption per capita during the Great Recession, -1.7% in 2008:Q4 (included in our sample), is less than 2.4 standard deviations of consumption growth. Similarly, even the cumulative contraction in U.S. consumption during the Great Depression (17%, that is, about half of the cumulative GDP drop of 31% often used in calibrations) implies average quarterly consumption contractions of about 1% per period. Given the frequency of our data, these small jumps can be well approximated by, and are hard to distinguish from, a continuous process (see, e.g., Ait-Sahalia (2004)).

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. Instead, consumption could adjust slowly to these shocks. This assumption is not only theoretically sound (see, e.g., Caballero and Simsek (2023)), 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 elaborate on in the Internet Appendix D). In addition, our formulation implies a particular term structure for the covariances between asset returns and multi-period consumption growth. We measure this term structure directly in the data and find that it yields the same size and persistence of the conditional mean process as in our state-space setting. Moreover, our estimated conditional consumption mean is a predictor of future consumption growth superior to other commonly used proxies (e.g., lagged consumption growth, its AR or VAR forecasts, lagged GDP growth, the Survey of Professional Forecaster, and the Liu and Matthies (2022) news index).

The central feature of our framework is that we do not impose stringent parametric restrictions on the functional form of the conditional consumption mean process. Our approach is by construction robust to the nature and number of factors driving asset returns, the priced or unpriced nature of the shocks, and the role of external predictors. It allows us to compare and test not only specific structural models but also entire classes of economic mechanisms. To the best of our knowledge, this yields the most comprehensive analysis of the stochastic process for consumption to date and should guide both the design and calibration of macro-finance models.

Closely Related Literature

By its very nature, our paper is closely linked to consumption-based asset pricing.⁴ Given both the saliency of consumption dynamics for equilibrium macro-finance models and the ambiguous nature of the true data-generating process, akin to the dark matter in finance (Campbell (2017) and Chen, Dou, and Kogan (2021)), much effort has been put into deter-

⁴For example, Breeden, Gibbons, and Litzenberger (1989), Lettau and Ludvigson (2001, 2001b), Jacobs and Wang (2004), Hansen, Heaton, Lee, and Roussanov (2007), Jagannathan and Wang (2007), Piazzesi, Schneider, and Tuzel (2007), Constantinides and Ghosh (2017), and Bansal, Kiku, and Yaron (2012). See also Ludvigson (2013) for an excellent review.

mining its key features. Hence, our paper is naturally related to several large strands of the literature.

Proper identification and recovery of the consumption process, as well as its link to asset returns, is known to be an elusive problem. Consumption-based asset pricing models are notoriously difficult to estimate due to the weak identification of the underlying parameters. This problem is particularly pronounced in the context of linear factor models (Burnside (2011, 2015), Kan and Zhang (1999), and Kleibergen and Zhan (2020)) that consider consumption growth as a source of priced risk in the cross-section of asset returns. The challenge of identifying the stochastic process of consumption growth has also been highlighted in the structural models of consumption-based asset pricing. In particular, Johannes, Lochstoer, and Mou (2016) and Collin-Dufresne, Johannes, and Lochstoer (2016) illustrate the problem from the perspective of a Bayesian learner and show that it has a profound impact on equilibrium consumption and return dynamics.

Our solution to the identification problem of consumption dynamics relies on insights from two strands of literature in empirical macroeconomics. First, following Blanchard and Quah (1989), we rely on the MA representation of the stochastic process. Second, as in the structural VAR literature (e.g., Christiano, Eichenbaum, and Evans (2005) and Sims and Zha (2006)), we use the intertemporal Euler equation to formulate a constraint on the joint dynamics of asset returns and consumption and identify their common shocks and propagation mechanisms. This allows us to trace out the response of consumption to shocks spanned by financial market as in Parker (2001), but without assuming that return innovations are fully captured by the market index and that consumption reacts only with a lag to financial shocks.

Using data from the Panel Study of Income Dynamics (PSID), we show that the dynamics uncovered in aggregate data also hold for shareholders but – importantly – not for non-shareholders, further supporting our identification strategy and rationalizing the findings of Malloy, Moskowitz, and Vissing-Jorgensen (2009). Finally, our approach is also supported by the evidence in Ang, Piazzesi, and Min (2006), who show that financial markets are informative about macro dynamics, and Liew and Vassalou (2000), who find that the equity value and size factors are leading indicators of economic growth.

We follow the logic of Jagannathan and Marakani (2016) and leverage a large cross-section of equity and bond returns to amplify the signal common to consumption and asset returns. The high-dimensional nature of the problem, therefore, naturally calls for the use of Bayesian filtering techniques. Similar to Schorfheide, Song, and Yaron (2018) and Zviadadze (2021), we leverage state-space modeling to infer the latent dynamics of consumption and returns. Unlike them, however, we do not rely on the restrictive assumption of a low-order autoregressive structure (e.g., AR(1)), because this would constrain the *entire* long-run dynamics of the conditional mean of consumption (Beeler and Campbell (2012)). This element is key, because misspecifying the conditional mean gives rise, as we show, to spurious volatility dynamics in consumption.

Our study is also related to a recent strand of papers that employ Bayesian tools to evaluate the assumptions, long considered to be salient, yet largely untestable by the frequentist literature. Giannone, Lenza, and Primiceri (2021) rely on the spike-and-slab priors to test the cornerstone assumption of sparsity in a variety of applications. Bryzgalova, Huang, and Julliard (2023) leverage the hierarchical structure implied by the no-arbitrage constraint in cross-sectional asset pricing to parse the factor zoo and test the underlying assumptions of uniqueness and factor sparsity of linear models. Our flexible modeling allows us to compare entire classes of data-generating processes for consumption and returns, postulated (and not) in the prior literature. As a result, we are able to test the key assumptions – usually taken for granted – commonly used to generate time-varying risk premia in macro-finance models.

The low power of frequentist approaches in consumption-based models often forces researchers to rely on proxies of consumption volatility. Bansal, Kiku, Shaliastovich, and Yaron (2014) and Campbell, Giglio, Polk, and Turley (2018) use realized volatility of returns as a proxy for that of consumption growth, motivated by the market clearing condition of a closed Lucas tree economy with no labor or public spending. Dew-Becker and Giglio (2016) further acknowledge the lack of sufficient power in testing whether stochastic volatility drives time-varying risk premia. Our framework, instead, allows us to estimate the volatilities of consumption and assets and test their persistence, equality, and other features. We find that the data reject the notion of equality between consumption and return volatilities and indicate that neither of them drives time variation in conditional risk premia.

Our work is also connected to a large body of literature on robust inference in potentially misspecified models (Hansen and Sargent (2001)). Similar to Giglio and Xiu (2021), we allow for multiple latent and observable factors to drive a cross-section of asset returns. Our specification for the consumption dynamics is also agnostic about the exact functional form consistent with the underlying MA decomposition, treating the consumption data-generating process as unknown. In the consumption habit setting, a similar approach has been undertaken by Chen and Ludvigson (2009), who, rather than postulating a particular functional form of habit, provide its general nonparametric estimate.

Finally, our paper also contributes to the literature that studies the link between low-frequency consumption movements and financial returns (Lettau and Ludvigson (2001, 2001b), Parker and Julliard (2005), Jagannathan and Wang (2007), Ortú, Tamoni, and Tebaldi (2013), Dew-Becker and Giglio (2016), and Bandi and Tamoni (2023)). We identify the stochastic process of consumption by (primarily) tracing out the term structure of its response to shocks spanned by financial markets. In doing so, we document that it is the same consumption shocks, weakly identified in the short run (as in Kleibergen and Zhan (2020)), but economically and statistically salient at business cycle frequency. Furthermore, we show that this mechanism is unlikely to be driven by measurement error, time-averaging and benchmarking of consumption data (Hansen and Singleton (1983), Breeden, Gibbons, and Litzenberger (1989), and Wilcox (1992)), as our results also hold with mixed-frequency estimation, monthly data, formally accounting for observational error, and in a broad spectrum of robustness exercises.

Although our method cannot distinguish the economic mechanisms driving the persistency of the conditional consumption mean (e.g., adjustment costs in consumption and complementary factors in the marginal utility, or constraint on information processing), we do pin down the process for consumption (and asset returns) that a valid macro-finance model should either deliver in equilibrium or that a researcher should at least use for a realistic calibration of the data-generating process.

The remainder of the paper is organized as follows. Section I demonstrates that the types 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 II, where we also show that it can accurately recover the consumption dynamics of popular models, as well as more general formulations. Our main empirical findings are presented in Section III, Section IV outlines their main implications for equilibrium asset pricing models, and Section V concludes.

I The Challenge of Consumption Persistency

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

Consider, for 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 in rationalizing unconditional risk premia and other moments of the historical data. That is,

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

where $\eta_t, e_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and σ_t , depending on the calibration, is either a constant or a SV process. The amount of time-series variation of consumption that is driven by the predictable component x_t varies in the literature. For instance, the conditional mean generates about 12% of the quarterly consumption variance in Bansal and Yaron (2004) and 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 sample? We address this question formally by using the Hansen, Heaton, Lee, and Roussanov (2007) calibration to generate a 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 generally much noisier in reality. We report the details of the simulation design in Internet Appendix A.

Table 1 presents the frequency of specification selected according to the Bayesian and Akaike Information Criteria (BIC and AIC). Strikingly, with both quarterly and monthly sampling, the most often selected specification implies no predictability at all (Columns

Table 1: ARIMA model selection of the 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)											
<i>p</i>	<i>d</i>	<i>q</i>	freq	<i>p</i>	<i>d</i>	<i>q</i>	freq	<i>p</i>	<i>d</i>	<i>q</i>	freq
0	0	0	53.7%	0	0	0	25.1%	0	1	0	50.0%
0	1	1	24.8%	0	1	1	14.6%	0	1	1	15.6%
1	0	1	5.9%	1	0	1	13.8%	1	0	1	12.2%
1	1	1	5.1%	0	0	1	6.3%	2	0	0	10.1%
1	0	0	3.9 %	1	1	1	5.3%	1	1	0	5.1%
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Panel B: Monthly frequency (642 observations)											
0	0	0	55.1%	0	0	0	27.8%	0	1	0	92.3%
0	1	1	31.3%	0	1	1	24.0%	1	0	0	7.2%
1	0	1	4.7%	1	0	1	13.4%	1	1	0	0.3%
5	1	0	1.4%	0	0	1	6.5%	1	1	1	0.1%
0	0	1	0.9%	0	1	2	3.2%	2	0	0	0.1%
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

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 list only the top five most frequent models.

(A) and (B)). Moreover, 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.⁵

Note that misidentifying the conditional mean process also has important consequences for recovering the volatility process. For instance, suppose that the process in equation (1) has a constant volatility; 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, because its squared forecast errors would be positively autocorrelated (with the j -th autocorrelation 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.

⁵Similarly, Colacito and Croce (2011) show, with a Monte Carlo simulation exercise, that recovering the consumption dynamic of equation (1) using a cross-section of countries is a daunting task when using only consumption data, even in a long time series sample.

If the true process for consumption is similar to what we rely on in macro-finance models, the considerations above 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, for example, 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 in model predictions, these restrictions should be tested whenever possible. Third, ideally, we would use a method that allows us 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). These three properties above are exactly what our empirical formulation delivers. We present it in the next section.

II 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 variables (e.g., habits, returns on wealth, leisure, leverage ratios, and aggregation weights in heterogeneous agent models). Furthermore, consumption and returns both contribute to the intertemporal budget constraint. This implies that, with the exclusion of exceptional 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, and 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 , which is independent of financial market shocks, and a possibly autocorrelated process – a persistent component – that possibly depends on the current and past shocks to asset returns. For expositional simplicity, we start by focusing on a setting with constant volatilities and generalize 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 (possibly infinite) MA representation. Obviously, by virtue of Wold's representation theorem, an $MA(\infty)$ model for the persistent component would capture the true data-generating process and avoid the fallacy of model selection outlined in the previous section. Because the MA coefficients in the Wold representation are square summable, any finite order covariance stationary ARIMA can be approximated using a high-order MA process, with the accuracy increasing with the MA order. Therefore, we model the (log) consumption growth process as

$$\Delta c_{t-1,t} = \mu_c + \underbrace{\sum_{j=0}^{\bar{S}} \rho_j f_{t-j}}_{MA(\bar{S})} + w_t^c, \quad (2)$$

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.

This representation encompasses most leading consumption-based asset pricing models. In particular, setting $\rho_j = w_t^c = 0 \forall j > 0$ and letting f_t to be an i.i.d. Gaussian process yields the formulation of Campbell and Cochrane (1999). The two-state Markov chain process of Mehra and Prescott (1985) is equivalent to an AR(1) process, and, hence, $MA(\bar{S} = \infty)$ and $w_t^c = 0$. In the rare disaster formulation of Barro (2005), $\rho_j = 0 \forall j > 0$, f_t is an i.i.d. jump variable, and w_t^c are i.i.d. shocks. Wachter (2013) shares the same formulation but with varying jump intensity. The AR(1) process in Vasicek (1977) maps into $MA(\bar{S} = \infty)$ and $w_t^c = 0$. Finally, adding an uncorrelated w_t^c shock to the Vasicek (1977) formulation results in the consumption process of Campbell (1999) and Bansal and Yaron (2004). Nevertheless, in

our empirical estimation, we restrict all innovations to follow Gaussian processes. Hence, we do not formally model large and infrequent jumps as considered in the rare disaster literature (e.g., Wachter (2013) and Julliard and Ghosh (2012)).

Consistent with these models, we assume that returns load on f_t contemporaneously:

$$\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}_f^r f_t + \boldsymbol{\rho}_g^r \mathbf{g}_t + \mathbf{w}_t^r, \quad f_t \perp \mathbf{g}_t, \quad \mathbf{g}_t \sim \mathcal{N}(\mathbf{0}_{K-1}, \mathbf{I}_{K-1}), \quad (3)$$

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, \mathbf{g}_t is the additional $K - 1$ latent factors driving only asset returns but not appearing in the consumption growth equation, and \mathbf{w}_t^r is uncorrelated with all other shocks and has a diagonal covariance matrix $\boldsymbol{\Sigma}_r$.⁶ The key identification assumption in equation (3) is that f_t is the only common component between consumption growth and asset returns and is orthogonal to the remaining factors \mathbf{g}_t . In the formulation above, returns are modeled as reacting contemporaneously and fully to the f shocks since prices are “jump” variables in equilibrium models. Note that, as shown below, equations (2)–(3) can also accommodate external predictor variables and, as we demonstrate later on, a direct dependency on the stochastic volatility processes of *all* shocks.

The joint dynamics of consumption and returns postulated in equations (2)–(3) are consistent with the extensive preliminary evidence that we report in Internet Appendix D. Therein, based on predictive regressions, we show that: *i*) the consumption growth process shows significant predictability over multiple years and *ii*) this predictability is better captured by lagged returns than by lagged consumption.

Obviously, we cannot feasibly use an infinite number of lags in the MA component of consumption growth. At the same time, because the MA representation of the persistent component is meant to approximate only the true latent dynamic, model selection would not be appropriate and, as outlined in Section I, would possibly be 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

⁶The diagonality and time-invariance assumption can be relaxed, as explained in Internet Appendix B.2, and we do so in empirical applications. We will later allow this covariance to embed both common and idiosyncratic stochastic volatilities. Note also that the violation of the diagonality restriction affects only efficiency, not consistency, of the estimation of the conditional mean processes for consumption and returns.

quarterly shocks in the conditional mean of consumption because this is almost twice as much the degree of persistency that we find in the preliminary evidence in Internet Appendix D, and it covers the span of predictability uncovered in the previous literature (see, e.g., Liew and Vassalou (2000), Parker and Julliard (2005), and Bandi and Tamoni (2023)); *ii*) we show in the simulation in Section II.1 that the approach is robust to the precise number of included MA lags; and *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).

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

A crucial point that allows us to identify the shocks is the lead-lag structure of the consumption process and its possible link to asset returns. Without equation (2), the shocks would be under-identified, making it difficult to give any particular rotation a structural interpretation. Another interpretation of this estimation approach is that of uncovering the shocks that drive financial returns through the impulse response function on consumption, in the spirit of the Uhlig (2005) identification in structural-VARs. In particular, our approach is akin to constructing the Generalized 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 (IRF) of multiperiod consumption growth to the shock f_t as⁷

$$\frac{\partial \mathbb{E} [\Delta c_{t-1,t+S}]}{\partial f_t} = \sum_{j=0}^S \rho_j, \quad \rho_{j>\bar{S}} := 0. \quad (4)$$

The framework described in equations (2)–(3) is closely related to Giglio and Xiu (2021), who propose a three-pass estimator of factors' risk premia. Specifically, their paper projects the nontradable factor onto the space of asset returns' latent factors. Their building block is that the projection of any variable onto the asset space is rotation invariant; hence, they can

⁷This 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).

recover latent factors of asset returns via Principal Component Analysis (PCA) and regress the nontradable factor on the first several principal components (PCs). Similar to their paper, we extract the common component in nontradable factors and asset returns relying on the rotation invariance of latent factors. We show in Section II.2 that if the true f_t is a linear combination of large principal components of \mathbf{r}_t^e , our framework can identify it.

Furthermore, our paper nests in and improves the framework in Giglio and Xiu (2021). We allow the lead-lag relationship between nontradable factors and asset returns, whereas Giglio and Xiu (2021) study only the contemporaneous covariance structure between these variables. Because most nontradable factors, especially macro variables, are persistent to a certain extent, the flexibility of our approach can improve the identification of the spanned component between nontradable economic fundamentals and asset returns.

Finally, note that the formulation in equations (2)–(3) can be generalized to *i*) allow for a bond-specific latent factor (b_t) to which consumption could react slowly over time and *ii*) external predictor variables, \mathbf{x} , predetermined at time t , as follows:

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + \sum_{j=0}^{\bar{S}} \theta_j b_{t-j} + \boldsymbol{\gamma}_x^c \mathbf{x}_{t-1} + w_t^c \quad \text{and} \quad (5)$$

$$\mathbf{r}_t^e = \underbrace{\boldsymbol{\mu}_r}_{N \times 1} + \underbrace{\boldsymbol{\rho}^r}_{N \times 1} f_t + \left[\underbrace{(\boldsymbol{\theta}^b)'}_{N_b \times 1}, \underbrace{\mathbf{0}'_{N-N_b}}_{N_b \times 1} \right]' b_t + \underbrace{\boldsymbol{\rho}_g^r \mathbf{g}_t}_{N \times 1} + \boldsymbol{\gamma}_x^r \mathbf{x}_{t-1} + \underbrace{\mathbf{w}_t^r}_{N \times 1}, \quad (6)$$

where N_b is the number of bonds that 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 b_t – a white noise process with variance normalized to one.

Two observations regarding our parametric framework warrant mentioning. First, both the one-factor (equations (2)–(3)) and two-factor (equations (5)–(6)) models are overidentified. Second, the 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 III.2, where we show that misspecification of the consumption mean process leads to spurious evidence of consumption volatility clustering, and in Section III.6, where we generalize our state-space model to allow for stochastic volatilities affecting all the shocks in the system.

II.1 The challenge of consumption persistency redux


A natural question is whether our state-space representation of consumption and returns in equations (2)–(3) is able to recover the consumption process when standard methods fail (as shown in Section I). To do so, we again use the simulated Hansen, Heaton, Lee, and Roussanov (2007) long-run risk consumption process described in Section I, calibrate the asset loadings on f to the values observed in the historical data, and apply our state-space estimation to it. Note that this simulation is particularly challenging for two reasons. First, the simulated time-series sample is small, with only 200 quarterly observations. Second, consumption and return data are generated at a monthly frequency and then aggregated to the quarterly frequency. Hence, the simulated data are affected by time-averaging, making the true conditional mean much harder to recover, as illustrated in Schorfheide, Song, and Yaron (2018).

As shown in Panel (a) of Figure 1, our state-space model with a long consumption MA accurately captures the effect of time- t shocks on subsequent consumption growth: The difference between the mean estimate (across simulated time-series) and the true value is extremely small, as is the variability across simulated samples. If anything, we observe a small attenuation bias in the long run, implying that our approach is, 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 IA.19 in Internet Appendix P demonstrates that our estimation precisely recovers the loading of asset returns on the shocks to the conditional mean of consumption.

Furthermore, as shown in Panel (b) of Figure 1, the IRF estimates are almost identical when using different orders for the MA component. The only difference is that with a longer MA, we can trace the effect further in the future. Hence, a finite order MA yields 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 accurate variance decomposition for both consumption and returns, as we do in later sections.

Figure 1: Cumulative impulse 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, 16th, 84th, and 95th percentiles of cumulative IRF in 1,000 simulations. The model is estimated under the assumption that $\bar{S} = 14$. Red dots denote pseudo true values, blue solid line denotes the mean estimates across simulations, while shaded areas denote 68% and 90% confidence intervals. The simulated time-series sample size is 200 (quarters). Panel (b) plots the average cumulative 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: $\bar{S} = 6, 10, 14$, and 18.

II.2 Identifying the conditional dynamic of $\Delta c_{t-1,t}$

Section II.1 confirms that our framework can accurately recover the conditional dynamics of the consumption process in a single-factor model, where the shock to the consumption conditional mean is the only common driver in consumption and asset returns. However, some may be concerned that our approach mechanically identifies f_t as the largest principal component of returns. To address this concern, we consider a more general simulation setup in which asset returns are driven by their PCs and the conditional mean of consumption growth is determined by large or small PCs or even their linear combinations.

Specifically, we assume the following multifactor model for asset returns \mathbf{r}_t^e :

$$\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r \mathbf{u}_t + \mathbf{w}_t^r, \quad \mathbf{u}_t \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_K, \mathbf{I}_K), \quad (\boldsymbol{\rho}^r)' \boldsymbol{\rho}^r = \mathbf{I}_K, \quad \mathbf{w}_t^r \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_N, \boldsymbol{\Sigma}_r), \quad (7)$$

where \mathbf{u}_t denotes pseudo-true latent factors, $\boldsymbol{\rho}^r$ are their factor loadings, $\boldsymbol{\mu}_r$ is the vector of unconditional mean returns, and \mathbf{w}_t^r are idiosyncratic shocks with a diagonal covariance matrix. $\Delta c_{t-1,t}$ is simulated using equation (2).


We consider four scenarios: (I) $\boldsymbol{\rho}^r$ contains the eigenvectors of the first and second PCs of

returns ($K = 2$), and $f_t = u_{1t}$, (II) ρ^r contains the eigenvectors of the first and second PCs of returns ($K = 2$), and $f_t = u_{2t}$, (III) ρ^r contains the eigenvectors of the first and second PCs of returns ($K = 2$), and $f_t = (u_{1t} + u_{2t})/\sqrt{2}$, and (IV) ρ^r contains the eigenvectors of the first five PCs ($K = 5$), and $f_t = (u_{1t} + u_{4t})/\sqrt{2}$. In addition, we estimate the eigenstructure of returns using observed return data of 37 stock portfolios and calibrate them such that $(\rho_1^r)' \rho_1^r = 16(\rho_2^r)' \rho_2^r = 25(\rho_3^r)' \rho_3^r = 36(\rho_4^r)' \rho_4^r = 144(\rho_5^r)' \rho_5^r$. Hence, the first latent factor is a dominant one. In all simulations, we calibrate ρ^c using the estimates obtained in Figure 3, and the variance of w_{ct} such that the time-series R^2 explained by the MA component of f_t is 23–24% in the consumption equation, consistent with our empirical evidence presented later.

Figure 2 presents the distribution of correlations between the estimated MA components and their pseudo-true values, that is, $\text{Corr}(\sum_{j=0}^{\bar{S}} \rho_j^c f_{t-j}, \sum_{j=0}^{\bar{S}} \hat{\rho}_j^c \hat{f}_{t-j})$. If we perfectly recover the MA component, the correlation coefficients should equal 1. It is worth noting that the identification of $\sum_{j=0}^{\bar{S}} \rho_j^c f_{t-j}$ depends on whether (1) we recover f_t and (2) we precisely estimate ρ_j^c . In experiment I, f_t is the strong factor (the first PC of asset returns), so identifying f_t is relatively simple. The box plot confirms this: Our estimates identify virtually the entire MA component in consumption growth, with a mean correlation coefficient of 0.89 (and a 90% confidence interval of [0.76, 0.93]), while the correlation between the true and filtered f shocks is 98%.

Moreover, our identification strategy does not mechanically recover the first principal component of \mathbf{r}_t^e . Experiment II assumes that f_t is the second PC of asset returns, whose variance is $\frac{1}{16}$ of the first PC, meaning that it drives on average only 4.4% of the time series variation of returns. However, our method can still recover the MA component in this challenging case. In particular, the second column of Figure 2 shows that the mean correlation is about 0.85, while the mean correlation between the true and filtered shocks is 93%, quite close to the value of experiment I. This implies that in this experiment, in a sample of the same size as the historical one, the shocks recovered by our method are not contaminated by the first principal component in a meaningful way. That is, our method does not deliver filtered consumption shocks that mechanically explain a large part of the time series variation of asset returns.

Figure 2: Correlation between estimated and pseudo-true MA components.



The graph presents the distribution of correlations between the estimated MA components and their pseudo-true values, that is, $\text{Corr}(\sum_{j=0}^{\bar{S}} \rho_j^c f_{t-j}, \sum_{j=0}^{\bar{S}} \hat{\rho}_j^c \hat{f}_{t-j})$. The model is estimated under the assumption that $\bar{S} = 14$. The simulated time-series sample size is 200 (quarters), under the following four simulation setups:

Experiment I: ρ^r contains the eigenvectors of the first two PCs of asset returns ($K = 2$), and $f_t = u_{1t}$.

Experiment II: ρ^r contains the eigenvectors of the first two PCs of asset returns ($K = 2$), and $f_t = u_{2t}$.

Experiment III: ρ^r contains the eigenvectors of the first two PCs ($K = 2$), and $f_t = (u_{1t} + u_{2t})/\sqrt{2}$.

Experiment IV: ρ^r contains the eigenvectors of the first five PCs ($K = 5$), and $f_t = (u_{1t} + u_{4t})/\sqrt{2}$.

Another concern of our approach is the single-factor structure in the consumption growth equation; that is, we allow for only one f_t . To address this critique, experiment III assumes that the true f_t is a linear combination of the top two PCs of \mathbf{r}_t^e . The simulation results are assuring: The mean correlation between the estimated and true MA components is 0.88, and the 90% confidence interval is similar to that of experiment I (while the mean correlation between true and filtered MA shocks is 96%).

Finally, the fourth experiment simulates returns using a five-factor model, and f_t consists of the first and fourth PCs. Despite the challenge that this scenario presents, we are still able to identify the MA component with a similar correlation coefficient as in the other three experiments (and a mean correlation of filtered MA shocks with the true ones of about 92%). However, the confidence interval is wider than that of experiment III, in which we consider a two-factor model. As shown in later sections, adding more latent factors in returns only increases estimation noise, thus leading to wider confidence intervals.

Overall, there are three main takeaways from the simulation study. First, our method does not mechanically identify the first principal component of asset returns as the driver

of the conditional mean shocks. Second, we can identify f_t even when it is a small principal component or a linear combination of several latent factors of asset returns. That is, the filtered MA innovations delivered by our method do not mechanically explain a large share of the time series variation of asset returns. Finally, it is essential to consider both single-factor and multifactor models. If we do not identify additional valuable information in consumption growth using multifactor models, the simple single-factor model is preferred because we can achieve sharper estimates and statistical inference.


III Empirical Evidence

In this section, we bring the state-space formulation in equations (2)–(3), as well as its generalization in equations (5)–(6), to the data to recover the conditional consumption mean process (subsections III.1, III.3–III.5). Furthermore, our analysis identifies the forms of volatility processes that are consistent with the data, identifies the degree of commonality in the consumption and returns stochastic volatility processes, and determines whether they drive time-varying risk premia (subsections III.2 and III.6).

We first focus on the conditional consumption growth mean dynamics in subsections III.1, its implications for consumption volatility dynamics (subsection III.2), and the loadings of asset returns on the shocks spanned by consumption (subsection III.3). In these subsections, we focus on the most robust specification in that we consider only one latent factor (f), do not include external predictors, leave the covariance structure of returns unrestricted, and do not impose a particular dynamic on the volatility processes. We begin with this streamlined formulation because it is consistent for the estimation of the conditional mean process of consumption and its forecast errors under general conditions and, hence, it allows us to analyze the properties of consumption volatility without taking a stand on its true process (see, e.g., Engle (1982)).

We then confirm that the results are stable if we consider external predictors for consumption and asset returns (subsection III.4) and multiple latent factors driving asset returns (subsection III.5), if we explicitly model stochastic volatility of all shocks (subsection III.6), and if we consider different MA lengths and cross-sections of asset returns (subsection III.7).

Figure 3: Cumulative impulse response function of consumption growth to a one-standard-deviation shock spanned by asset returns.



Posterior means of the cumulative response function of consumption growth (solid line with circles), with the centered posterior 90% (dotted lines) and 68% (dashed lines) coverage regions. Red line with triangles denotes the first principal component of $\text{cov}(r_{i,t}^{ex}, \Delta c_{t,t+1+S})$. Quarterly data, 1963:Q3-2019:Q4. Green line with triangles is the simulated cumulative impulse responses assuming that (1) monthly consumption growth is independent over time and has a contemporaneous correlation of 0.20 with the monthly f_t shock, and (2) monthly consumption data are benchmarked to the annual data (see Internet Appendix H for details).

We also show that the key stylized facts we uncover in the data are confirmed nonparametrically without the additional assumptions needed for the state-space model estimation.

III.1 The consumption mean process

The consumption growth representation in equations (2) and (5) is similar to the MA 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 bond-specific shock. Figure 3 depicts the (cumulated) loadings of consumption on the latent factor f as a function of the horizon S . At $S = 0$, as in the case of a standard consumption-based asset pricing model, the MA component of consumption virtually does not load on the common factor. Instead, as S increases, the impact of the common factor becomes increasingly pronounced, leveling off at approximately $S = 10$. At this horizon, the effect is economically substantial: The cumulative response of consumption growth to a one-standard-deviation shock is about 1%. Nevertheless, this large effect is *not* excessive: As shown in Internet Appendix E, it does not violate the bound on the long-run standard deviation of consumption growth obtained in Dew-Becker (2017).

Importantly, the finding in Figure 3 reconciles seemingly discordant empirical claims in the literature. For instance, Kleibergen and Zhan (2020) argue that the correlation between stock returns and consumption growth is small and not significantly different from zero. In contrast, Parker and Julliard (2005) find a large and significant long-run response of consumption to asset return shocks. Our MA representation sheds light on this apparent contradiction by pinning down the slow propagation mechanism of shocks spanned by financial markets into consumption: It is the same consumption shock that is both weakly identified in the short horizon yet strong and evident at business cycle frequency. Furthermore, the tight confidence intervals in the figure emphasize that our parametric state-space approach has a higher power of detecting the link between financial markets and consumption compared to the nonparametric approaches in the previous literature that yield high sampling uncertainty.

Note that allowing for a bond-specific latent factor (equations (5)–(6)) leaves the consumption loadings on f shocks virtually unchanged, and consumption does not significantly load on the bond-specific factor b_t . See, respectively, Figure IA.21 in Internet Appendix P.

The relatively tight restrictions on the parametric model in equations (2)–(3) 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 (weak) constraints on the data-generating process, some of which may, in principle, not hold in the data. However, the strongest prediction of our parametric setting – the term structure of asset exposure to consumption risk – can be tested without all the ancillary assumptions of the state-space model.

To see this, note that equations (2)–(3) imply a particular pattern in the covariances of asset returns with multi-period consumption growth; that is, for any asset i ,

$$\begin{aligned}
\text{cov}(r_{i,t+1}^{ex}, \Delta c_{t,t+1}) &= \rho_i^r \rho_0, \\
\text{cov}(r_{i,t+1}^{ex}, \Delta c_{t,t+2}) &= \rho_i^r (\rho_0 + \rho_1), \quad \text{and} \\
&\dots \\
\text{cov}(r_{i,t+1}^{ex}, \Delta c_{t,t+k}) &= \rho_i^r \left(\sum_{j=1}^k \rho_{k-j} \right).
\end{aligned} \tag{8}$$

Therefore, the term structure of asset exposures to consumption risk is driven by a single

common component: $(\rho_0, \rho_0 + \rho_1, \dots, \sum_{j=1}^k \rho_{k-1})'$, 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 (2) 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 $\text{cov}(r_{i,t+1}^{ex}, \Delta c_{t,t+k})$ at different horizons k .


The red line with triangles in Figure 3 illustrates our findings. Remarkably, the loadings on the first PC of consumption term structure exposure almost exactly match the cumulated response function from the state-space model, therefore identifying the same persistent time-varying mean for consumption growth.⁸

Recall that the identifying restriction for the finding above is that, for a representative agent, the intertemporal Euler equation determining equilibrium prices implies that both returns and consumption react to a common set of innovations (e.g., wealth shocks and changes in expectations). But is this really what we uncover? For an agent who does not participate in the equity market, such an Euler equation argument does not hold. Hence, we can verify this mechanism by using micro-data on the consumption of stockholders and non-stockholders.

Figure 4 uses the consumption series constructed in Malloy, Moskowitz, and Vissing-Jorgensen (2009) over the sample of 1982:Q1 to 2004:Q3 and reestimates our state-space model separately for these two types of households. Although the confidence bands are wider than in our baseline estimation due to the short sample size, the message is clear: Consumption growth of households that participate in the financial market (left panel) reacts significantly, and gradually over multiple quarters, to shocks spanned by financial assets, whereas the consumption of other households (right panel) does not. This further confirms the power and the economic mechanism underlying our identification strategy: Some of the shocks affecting the households that participate in financial markets are reflected in *both* equilibrium asset prices and consumption. Hence, we use this insight to uncover their joint dynamics and compare them with those postulated in popular macro-finance models. Note also that the estimated effect of a one-standard-deviation shock for shareholders' consumption is actually larger, by a factor of 2–3, than the effect uncovered in aggregate data –

⁸In 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 4: Cumulative impulse response functions of shareholders’ and non-shareholders’ consumption growth to a one-standard-deviation shock spanned by asset returns.



Posterior means of the cumulative response functions of consumption growth, centered posterior 90% (dotted lines) and 68% (dashed lines) coverage regions. Panels (a) and (b) use shareholders and non-shareholders’ consumption growth, respectively, following Malloy, Moskowitz, and Vissing-Jorgensen (2009). Data are downloaded from Tobias Moskowitz’s website. We use the “Dn1” and “Ds1” variables as proxies for non-shareholders’ and shareholders’ quarterly consumption growth. Data sample: 1982:Q1 to 2004:Q3.

exactly as one would expect, given that the aggregate series also contains the consumption of non-shareholders.


Figure 5 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.

III.1.1 Measurement issues of consumption

There are two possible measurement issues in quarterly data that might bias our estimates in Figure 3: time averaging and benchmarking of the quarterly consumption series. We now turn to evaluate their potential impact.

The evidence of persistency of the conditional consumption mean process that we document in Figure 3 might be contaminated by the well-known time-aggregation “bias” (see, e.g., Breeden, Gibbons, and Litzenberger (1989)). That is, consumption growth measured from quarter t to $t + 1$ is in fact a moving average of true consumption growth from $t - 1$ to $t + 1$ (as shown in Equation (IA.11) of the Internet Appendix). We formally address this

Figure 5: Moving average component of consumption growth.



Posterior mean of the moving average f_t component of consumption growth. Grey areas denote NBER recessions. Estimate based on the single-factor model in equations (2)–(3), with $\bar{S} = 14$. The cross-section of base assets includes 25 size- and value-sorted portfolios, 12 industry portfolios, and nine bond portfolios.


issue in two ways.

First, in Section G of the Internet Appendix, we introduce a mixed-frequency state-space formulation that can be overlayed on our estimation framework to directly model the time aggregation of quarterly consumption. Therein, we also show that, under the null of serially uncorrelated monthly consumption growth, we should not observe significant impulse responses of consumption after two quarters, even in the presence of time-aggregation bias. Furthermore, we estimate a latent factor model for monthly excess returns and allow quarterly consumption growth (observed only at the end of each quarter) to react slowly to the monthly asset return shocks as in Equation (2).

Second, as an additional way of addressing the time-aggregation bias, in Section F of the Internet Appendix, we reestimate our dynamic model for consumption and asset returns using monthly consumption and return data. Note that monthly consumption requires considering the well-known issue of measurement error in these data, which is directly modeled in our extended state-space formulation (see equation (IA.8) of the Internet Appendix).

Figure 6 presents the cumulative responses of quarterly consumption growth to one-standard-deviation monthly asset return shocks. We report the mixed-frequency state-space estimates in Panel (a), where the impulse response of consumption continues to increase until about 24 months, with a cumulative effect of roughly 0.015. Panel (b), instead, depicts the cumulative impulse function obtained with the monthly consumption data and mapped

Figure 6: Cumulative impulse response function of quarterly consumption growth to a one-standard-deviation shock spanned by monthly asset returns.



Posterior means of the cumulative impulse response function of quarterly consumption growth (solid lines with circles) to monthly asset return shocks, with centered posterior 90% (dotted lines) and 68% (dashed lines) coverage regions. The x-axis denotes the number of *monthly* lags. Panel A uses the mixed-frequency state-space model in Section G of the Internet Appendix. Panel B is based on the impulse responses of monthly consumption growth to monthly asset return shocks, converted into the CIRFs of quarterly consumption growth to monthly asset return shocks (see equation (IA.15) of the Internet Appendix). Green lines display the simulated cumulative impulse responses, assuming that (1) monthly consumption growth is independent over time and has a contemporaneous correlation of 0.20 with the monthly f_t shock, and (2) monthly consumption data are benchmarked to the annual data (see Internet Appendix H for details).

into quarterly consumption responses. The estimates in Panels (a) and (b) are strikingly similar, with wider confidence intervals in the case of monthly consumption data (due to the substantial measurement error). Furthermore, as Figure IA.3 of the Internet Appendix shows, monthly consumption growth still slowly adjusts to asset return shocks: The f_t shocks explain more than 11% of the time-series variation of *monthly* consumption growth.

But how would the time-aggregated quarterly consumption growth react to monthly asset return shocks if there were no predictability in monthly consumption? We answer this question in Figure IA.11 of the Internet Appendix by conducting a counterfactual exercise. Specifically, we impose that monthly consumption growth correlates only with contemporaneous asset return shocks. This implies that the time-aggregated quarterly consumption growth comoves with only contemporaneous and lagged asset returns up to four months.⁹ In this setting, we find a much smaller cumulative impulse response of quarterly consumption growth that peaks at less than 0.004 (one quarter of what we show in Figure 6). Hence,

⁹That is, we impose $\bar{S} = 0$ in equation (IA.14) of the Internet Appendix.

the time aggregation of consumption is unlikely to fully explain the predictability that we observe in quarterly consumption growth.

In addition to time aggregation, the serial dependence of quarterly consumption growth could be biased by benchmarking (see Wilcox (1992) and Triplett (1997)). Specifically, monthly and quarterly estimates of consumption are based on the Monthly Retail Trade Survey, which is of lower quality than the annual survey. Because the four quarters of consumption never precisely match the annual measure, the monthly and quarterly estimates are ex-post revised to benchmark to the annual estimates.

What would our state-space estimation method yield if consumption were IID but time-aggregated and benchmarked? We answer this question by simulation (see details in Internet Appendix H). In particular, we add to Figures 3 and 6 the simulated CIRFs (green lines), assuming that i) monthly consumption growth is independent over time and has a contemporaneous correlation of 0.20 with the monthly f_t shock¹⁰ and ii) monthly consumption data are benchmarked to the annual consumption data. Two observations are in order.


First, even if the true monthly consumption growth were serially independent, we would still detect slow responses of quarterly consumption growth to asset return shocks due to time aggregation and benchmarking. Furthermore, the asset return shock would predict quarterly consumption growth up to four quarters ahead. However, the cumulative impulse responses would be flat after the benchmarking horizon of four quarters.

Second, nevertheless, time aggregation and benchmarking of consumption data cannot fully explain the predictability of consumption growth that we estimate using our state-space formulation. In the observed data, as shown in Figures 3 and 6, the cumulative responses increase until the seventh quarter and are much larger than those predicted by time-aggregation bias and benchmarking (the green lines in the figures). Even when calibrating a much larger contemporaneous correlation between monthly consumption growth and f_t (e.g., 0.4 in Figure IA.15 of the Internet Appendix), our state-space model estimates of the cumulative responses are much more persistent and sizable than those obtained in IID time-aggregated and benchmarked consumption data.

Overall, the large degree of consumption predictability that we uncover cannot be rationalized by time aggregation and benchmarking of consumption data.

¹⁰In the real data, this contemporaneous correlation is about 0.17. Figure IA.15 of the Internet Appendix explores different correlation coefficients of between 0.10 and 0.40.

Figure 7: 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. The state-space models are estimated at both quarterly (blue bars) and mixed (red bars) frequencies. Plots report the unadjusted R-squared. Left panel: Cumulated consumption growth $\Delta c_{t-1,t+S}$. Right panel: One-period consumption growth $\Delta c_{t-1+j,t+j}$. Estimates based on single-factor model, with $\bar{S} = 14$. The cross-section of base assets includes 25 size- and value-sorted portfolios, 12 industry portfolios, and nine bond portfolios.

III.1.2 Predicting consumption

But how much of the total consumption volatility can this slow-moving component explain? The answer is more than a quarter, which is large (and sharply estimated) compared to the leading asset pricing frameworks: For example, this is more than twice as large as in the long-run risk framework of Bansal and Yaron (2004).¹¹ Figure 7 demonstrates that the common factor is responsible for roughly 23% of the variation in the one-period nondurable consumption growth, 26% of the two-period consumption growth, and so on, followed by a slow decline toward slightly more than 6% for the 12-period growth.¹² Furthermore, as shown in the figure, the predictability of consumption is even *higher* at the longer horizons when we estimate the mixed-frequency state-space model to correct the time aggregation. In this case, the MA component can explain about 12% of the quarterly consumption growth at the three-year horizon. Interestingly, the model retains significant predictive power (albeit lower), even for the one-period consumption three years ahead.

But does the conditional mean component uncovered by our method do a good job at predicting consumption? Does it survive horse races with the canonical predictors? To

¹¹This quantity is zero in frameworks that, for parsimony, do not model any predictability in consumption (e.g., Campbell and Cochrane (1999) and Barro (2006)).

¹²As shown in Figure IA.22 of the Internet Appendix, adding a bond-specific factor has a minimum impact on the explanatory power of the model for future consumption growth.

answer these questions, we explore whether the future realized consumption growth at short to medium horizons could be explained by the estimated conditional mean implied by our MA model. Specifically, given the estimated MA coefficients $\{\hat{\rho}_j\}_{j=0}^{\bar{S}}$, we can estimate the conditional mean of future one-period consumption growth as

$$\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}] = \mu_c + \sum_{s=j}^{\bar{S}} \hat{\rho}_s f_{t+j-s}, \quad (9)$$

where j ranges from 1 to 12 quarters. To avoid forward-looking bias and preserve consistency of our predictive regressions, we employ a (one-side) Kalman filter (rather than a Kalman smoother) to extract the latent state f_t . We next regress $\Delta c_{t+j-1,t+j}$ on $\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$,

$$\Delta c_{t+j-1,t+j} = \beta_{0,j} + \beta_{1,j} \widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}] + \beta_{2,j} \text{controls}_t + e_{t+j-1,t+j}^c, \quad (10)$$

and test whether the slope coefficient, $\beta_{1,j}$, is significantly different from one—the value that we should obtain in a long sample with a correctly specified model. As we do not reestimate the MA coefficients to fit future consumption growth, this exercise is an additional test of the validity of our method.

Table 2 reports the predictive regression results. In Panel A, we regress future (one-period) consumption growth (from one to 12 quarters ahead) on our estimated conditional means. The predictive R^2 s are extremely close to those obtained in Panel B of Figure 7, with the predictability after two years declining faster than shown in the figure. This small decline in predictability is not surprising, given the one-sided nature of the Kalman filter compared with the Kalman smoother in the figure. Moreover, the slope coefficients associated with the conditional mean are extremely close to, and statistically indistinguishable from, one—as should be under the null of our specification. As a benchmark, in Table IA.III of the Internet Appendix, we repeat the same predictability analysis using as a proxy for the conditional mean the VAR-based estimate of Parker (2001). At the rather short horizon (1–2 quarters), the VAR has performance similar to our state-space formulation, but at longer horizons, its performance quickly deteriorates, with R^2 s essentially equal to zero and slope coefficients much smaller than one after one year.

Panels B to D of Table 2 conduct horse races with the canonical predictors: Lagged consumption and GDP growth, as well as a Kalman filter estimate from a standard AR(1)

latent expected growth model for consumption growth.¹³ The results therein confirm that our conditional estimate captures information not fully spanned by these canonical predictors. First, the slope coefficient associated with the conditional mean is once again close to and indistinguishable from one. Second, the (unadjusted) R^2 s are only marginally higher than those in Panel A, highlighting the fact that most of the information contained in these variables is already captured by our conditional consumption mean process. Third, the added predictors never drive out the statistical significance of our conditional mean, and are themselves statistically different from zero only in rare cases (3 out of 36, i.e., at the expected Type I error frequency) and only at the short horizon (one or two quarters).

As additional horse races, in Tables IA.VII and IA.VIII of the Internet Appendix, we consider the mean consumption forecasts from the Survey of Professional Forecasters (SPF) as well as the news-based index of Liu and Matthies (2022).¹⁴

The SPF forecasts are available for only one- to four-quarter horizons and over a much shorter sample than our baseline (starting from 1981:Q3). Panel A of Table IA.VII reveals that the SPF predictive ability is much lower (about half in most cases) than the one of our conditional mean estimates in Panel A of Table 2. Furthermore, when used in conjunction with our conditional mean (Panel B of Table IA.VII), the SPF forecasts do not drive out the statistical significance of the former (albeit at the one-quarter horizon, but not others, the coefficient is statistically smaller than one). Moreover, SPF forecasts only marginally improve the predictive power achievable with our conditional mean model. Interestingly, while the predictive ability of SPF forecasts three and four quarters ahead is small (see Panel A of Table IA.VII), the one of our conditional mean model is substantial, with predictive R^2 s of about 10% (see Panel A of Table 2). The robustness of our conditional model to the inclusion of SPF forecasts is remarkable since professional forecasters are likely to use the information in financial markets to predict future consumption growth.

As shown in Panel A of Table IA.VIII, after controlling for our MA-implied conditional consumption mean, the news index of Liu and Matthies (2022) has no incremental predictive

¹³Specifically, we consider the Kalman filter estimate from an AR(1) latent expected growth model:

$$\Delta c_{t-1,t} = \mu_c + x_t + w_t^c, \quad x_t = \rho_x x_{t-1} + w_t^x, \quad w_t^c \sim \mathcal{N}(0, \sigma_{wc}^2), \quad w_t^x \sim \mathcal{N}(0, \sigma_{wx}^2), \quad (11)$$

which implies that the latent state x_t should predict future consumption growth.

¹⁴Two caveats are in order. First, the SPF forecasts are based on total real personal consumption (PCE). Second, Liu and Matthies (2022) show that their news index has forecasting power for nondurable *plus* service consumption. Our baseline consumption series is the nondurable one instead.

Table 2: Validating the predictability of the conditional consumption mean

$j =$	1	2	3	4	5	6	7	8	9	10	11	12
Panel A. Regress $\Delta c_{t+j-1,t+j}$ on $\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$												
$\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$	0.998	0.964	0.964	0.965	0.932	0.913	0.963	1.129	1.021	0.984	1.075	1.080
s.e. (OLS)	(0.146)	(0.196)	(0.202)	(0.204)	(0.271)	(0.287)	(0.313)	(0.353)	(0.409)	(0.423)	(0.561)	(0.561)
s.e. (NW, lag=12)	(0.196)	(0.199)	(0.195)	(0.203)	(0.190)	(0.210)	(0.182)	(0.204)	(0.330)	(0.374)	(0.461)	(0.461)
Predictive R^2	0.182	0.104	0.098	0.097	0.054	0.046	0.043	0.047	0.029	0.025	0.017	0.017
Panel B. Additional control: Lagged consumption growth												
$\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$	0.932	0.896	0.867	0.947	1.007	0.894	0.963	1.100	1.024	1.028	1.032	1.087
s.e. (OLS)	(0.151)	(0.200)	(0.210)	(0.208)	(0.278)	(0.295)	(0.314)	(0.350)	(0.411)	(0.435)	(0.572)	(0.563)
s.e. (NW, lag=12)	(0.192)	(0.190)	(0.213)	(0.222)	(0.241)	(0.224)	(0.185)	(0.226)	(0.325)	(0.376)	(0.453)	(0.466)
Δc_t	0.107	0.103	0.110	0.032	-0.079	0.019	-0.024	-0.150	-0.006	-0.030	0.026	-0.017
s.e. (OLS)	(0.064)	(0.066)	(0.067)	(0.066)	(0.068)	(0.068)	(0.067)	(0.065)	(0.064)	(0.066)	(0.066)	(0.065)
s.e. (NW, lag=12)	(0.067)	(0.074)	(0.074)	(0.066)	(0.111)	(0.075)	(0.058)	(0.087)	(0.057)	(0.059)	(0.065)	(0.061)
Predictive R^2	0.193	0.114	0.110	0.098	0.060	0.047	0.044	0.070	0.029	0.026	0.018	0.018
Panel C. Additional control: Lagged GDP growth												
$\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$	0.839	0.781	0.907	1.023	0.928	0.942	0.957	1.173	1.051	0.994	1.097	1.107
s.e. (OLS)	(0.149)	(0.211)	(0.210)	(0.208)	(0.276)	(0.287)	(0.313)	(0.353)	(0.415)	(0.436)	(0.563)	(0.563)
s.e. (NW, lag=12)	(0.207)	(0.223)	(0.202)	(0.224)	(0.192)	(0.211)	(0.186)	(0.198)	(0.331)	(0.342)	(0.436)	(0.491)
ΔGDP_t	0.199	0.136	0.061	-0.079	0.004	-0.072	-0.061	-0.094	-0.025	-0.005	-0.028	-0.045
s.e. (OLS)	(0.056)	(0.062)	(0.059)	(0.058)	(0.059)	(0.059)	(0.059)	(0.058)	(0.057)	(0.058)	(0.057)	(0.057)
s.e. (NW, lag=12)	(0.053)	(0.052)	(0.050)	(0.072)	(0.063)	(0.055)	(0.074)	(0.080)	(0.052)	(0.052)	(0.064)	(0.074)
Predictive R^2	0.229	0.124	0.103	0.105	0.054	0.053	0.048	0.058	0.030	0.025	0.018	0.020
Panel D. Additional control: Latent state (x_t) of AR(1) expected growth model												
$\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$	0.901	0.853	0.866	0.963	1.010	0.919	0.965	1.141	1.033	1.014	1.050	1.086
s.e. (OLS)	(0.152)	(0.204)	(0.213)	(0.210)	(0.281)	(0.293)	(0.313)	(0.351)	(0.415)	(0.436)	(0.571)	(0.563)
s.e. (NW, lag=12)	(0.197)	(0.195)	(0.227)	(0.234)	(0.253)	(0.222)	(0.187)	(0.218)	(0.322)	(0.374)	(0.460)	(0.468)
x_t	0.299	0.272	0.213	0.006	-0.155	-0.014	-0.141	-0.285	-0.027	-0.042	0.035	-0.034
s.e. (OLS)	(0.141)	(0.147)	(0.148)	(0.144)	(0.148)	(0.147)	(0.144)	(0.142)	(0.141)	(0.143)	(0.142)	(0.140)
s.e. (NW, lag=12)	(0.135)	(0.158)	(0.163)	(0.166)	(0.265)	(0.176)	(0.149)	(0.199)	(0.129)	(0.131)	(0.149)	(0.146)
Predictive R^2	0.200	0.118	0.107	0.097	0.059	0.046	0.048	0.065	0.029	0.026	0.018	0.018

The table summarizes the regressions in which future realized growth rates of nondurable consumption ($\Delta c_{t+j-1,t+j}$, $1 \leq j \leq 12$) are forecasted by several predictors. In Panel A, we regress $\Delta c_{t+j-1,t+j}$ on the conditional consumption mean ($\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$) implied by the MA model (see equation (9)). In Panels B–D, we include additional controls, such as time- t consumption growth, time- t GDP growth, and time- t latent state x_t from the AR(1) latent expected growth model in equation (11). We report (1) the point estimates of the slope coefficients, (2) the OLS and Newey-West (12 lags) standard errors within the parentheses, and (3) the predictive R^2 . Sample: 1963:Q3–2019:Q4.

power for our baseline nondurable consumption at any horizon (from one to 12 quarters). Nevertheless, in Panel B, we confirm that the news index does help predict nondurable *plus* service consumption.

III.2 Clustering and predictability of total consumption volatility

Note that the estimation of the conditional mean of consumption growth in equations (2) and (5) is generally consistent even in the presence of time-varying consumption volatility. Therefore, the presence of volatility clustering can be assessed by analyzing the serial correlation of the squared one-step-ahead forecast errors (see, e.g., Engle (1982)) of the consumption growth process. This proxy for volatility has been used extensively in the


empirical consumption-based asset pricing literature, for example, Bansal, Khatchatrian, and Yaron (2005), Bansal, Kiku, and Yaron (2012), Beeler and Campbell (2012), and Chen (2017). That is, we examine 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 $f_{\tau < t}$ and $b_{\tau < t}$.¹⁵ Note that this nonparametric volatility proxy does not distinguish the nature of the volatility being captured. For example, under the null of the consumption volatility process of Schorfheide, Song, and Yaron (2018), the square forecast errors would be a linear combination of the variances of innovations to immediate consumption growth (short-run news) and expected growth (long-run news). In Section III.6 we introduce a parametric model to be able to distinguish these potentially different volatility processes.

Figure 8 reports the autocorrelation function (left panel) as well as the p -values of the Ljung and Box (1978) and Box and Pierce (1970) tests (right panel) of joint significance of the autocorrelations of $\widehat{Var}_t(\Delta c_{t,t+1})$ and shows no evidence of volatility clustering in the consumption growth process. Nevertheless, conditional consumption volatility might still, in principle, be correlated with financial asset returns. We test this hypothesis by running linear predictive regressions of $\widehat{Var}_{t+h}(\Delta c_{t+h,t+h+1})$, at several horizons h , on the time- t first eight principal components of stock and bond returns – that is, we check whether asset returns can predict the nonparametric consumption volatility proxy. Note that, in line with the previous literature (see, e.g., Liew and Vassalou (2000)), this is the same test used to establish predictability of the first conditional moment of consumption growth, as shown in Appendix D. The p -values of the F -tests for these predictability regressions are depicted by the blue continuous line with circles in Figure 9. The p -values (which range from 0.2826 to 0.922) show that asset returns are not significant predictors of future consumption volatility.

We conduct two robustness checks of the above result. First, in Figures IA.9 and IA.10 of Internet Appendix G, we show that extremely similar results to those above are obtained in the mixed-frequency estimation setting, confirming that these are not due to a time-aggregation bias. Second, in Table IA.IX of the Internet Appendix, we run univariate regressions of $\widehat{Var}_t(\Delta c_{t,t+1})$ on persistent predictors that are possibly more informative about the level of volatility: the log price-dividend ratio of the market index as well as the measures of financial, macro, and real uncertainty of Jurado, Ludvigson, and Ng (2015) and

¹⁵ $\widehat{\mathbb{E}}_t[\Delta c_{t,t+1}]$ is the posterior mean of $\mu_c + \sum_{j=1}^S \rho_j f_{t+1-j} + \sum_{j=1}^S \theta_j b_{t+1-j}$ at each t .

Figure 8: 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) (blue circles) tests.

Ludvigson, Ma, and Ng (2021). The regression R^2 's are extremely small, ranging from 0% to 1.8%, and in only one case (that of the real uncertainty) we find a regressor to be statistically significant at canonical levels with one of the two types of standard errors that we consider.


Since the clustering in total consumption volatility is weak at best, 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 that a common latent factor drives both asset returns and consumption, if one were to erroneously model the conditional mean of consumption growth, one would be likely to find spurious evidence of volatility clustering. For instance, if one erroneously models 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 example, the k -th autocorrelation of $(\Delta c_{t,t+1} - \mu_c)^2$, for $k \leq \bar{S}$, is proportional to

$$Cov \left(\left(\sum_{j=k}^{\bar{S}} \rho_j f_{t-j} + \theta_j b_{t-j} \right)^2, \left(\sum_{j=k}^{\bar{S}} \rho_{j-k} f_{t-j} + \theta_{j-k} b_{t-j} \right)^2 \right) \neq 0. \quad (12)$$

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

First, we again run 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 used in the preliminary evidence presented in Appendix D) but, as a misspecification benchmark, we construct these measures assuming a constant conditional mean for consumption growth. Summary statistics for these regressions are depicted by the dashed red line with triangles in Figure 9. The figure shows that the misspecification of the mean process generates spurious predictability

Figure 9: 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 . We report the p -value of the F -test of joint significance of the covariates as well as the 5% and 1% significance thresholds (respectively, horizontal dot-dashed and dotted lines). The solid blue line with circles denotes statistics for the correctly specified conditional mean for the consumption growth process. The dashed red line with triangles corresponds to the assumption of a constant conditional mean.

of consumption volatility in three of the eight 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 (blue continuous line with circles of Figure 9).¹⁶

Second, in Table 3 we estimate ARCH(1), GARCH(1,1), and IGARCH(1,1) models¹⁷ for consumption volatility under different assumptions about the mean process. In Panel A, following the standard approach of Bansal, Khatchatrian, and Yaron (2005), Bansal, Kiku, and Yaron (2012), Beeler and Campbell (2012), Tédongap (2014), Chen (2017), and Zviadadze (2021), we identify the volatility process by postulating an AR(1) specification for the conditional mean of consumption growth. The AR(1)-GARCH(1,1) formulation in Panel A is the most often used in the literature to provide evidence of time-varying consumption volatility.¹⁸ In this case, there is statistically significant evidence of volatility clustering ($\beta > 0$ in Column (5)) with a half-life of volatility departures from the mean of about 5–6

¹⁶This is in line with the evidence of Dew-Becker and Giglio (2016, Appendix D) who, proxying consumption volatility with the realized volatility of the S&P500 index (as often assumed in the prior literature), find no predictability.

¹⁷See, respectively, Engle (1982), Bollerslev (1986), and Bollerslev and Engle (1986).

¹⁸In Section III.6 we provide an extensive analysis of SV processes in consumption and asset returns, of the type also commonly used in the literature.

Table 3: Estimates of ARCH(1), GARCH (1,1), and IGARCH(1,1) for different models

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

	ARCH(1)		GARCH(1,1)			IGARCH(1,1)		
	ω (1)	α (2)	ω (3)	α (4)	β (5)	ω (6)	α (7)	β (8)
Panel A: $\mu_t = \mu_0 + \mu_1^c \Delta c_t$								
Estimate	0.917	0.041	0.133	0.139	0.727	0.000	0.026	0.974
Std error	(0.200)	(0.109)	(0.076)	(0.125)	(0.108)	(0.000)	(0.034)	(0.034)
Akaike		2.824			2.799			2.820
Bayes		2.887			2.878			2.883
Shibata		2.823			2.798			2.819
Hannan-Quinn		2.850			2.831			2.845
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	0.902	0.087	0.000	0.000	0.999	0.000	0.000	1.000
Std error	(0.138)	(0.088)	(0.001)	(0.000)	(0.000)	(0.000)	(0.004)	(0.004)
Akaike		2.842			2.851			2.852
Bayes		2.874			2.899			2.884
Shibata		2.842			2.850			2.852
Hannan-Quinn		2.855			2.870			2.865
Panel C: $\mu_t = \mu_0 + \sum_{i=1}^S \rho_i f_{t+1-i}$								
Estimate	0.891	0.115	0.000	0.000	0.999	0.000	0.016	0.984
Std error	(0.150)	(0.102)	(0.000)	(0.000)	(0.000)	(0.000)	(0.060)	(0.060)
Akaike		2.843			2.850			2.856
Bayes		2.875			2.898			2.888
Shibata		2.843			2.850			2.856
Hannan-Quinn		2.856			2.869			2.869

ARCH(1), GARCH(1,1), and IGARCH(1,1) estimates for consumption growth volatility with different models for the conditional mean. Models estimated via QMLE. Robust standard errors constructed using Newey and West (1987). The table also reports Akaike, Bayesian, Shibata, and Hannan-Quinn information criteria.

quarters.¹⁹ These estimates are almost identical to those uncovered in the previous literature (see, e.g., Table 6 in Chen (2017)). However, the statistically insignificant α (Column (4)) indicates a potential identification failure for the model parameters, including β . This is why we also consider an IGARCH specification that once again yields (Columns (6)–(8)) a statistically significant β and an insignificant α , with point estimates compatible with constant consumption volatility.

However, as shown in both the previous literature (see, e.g., Liu and Matthies (2022) and Bansal, Dittmar, and Kiku (2007)) and in Appendix D, lagged consumption alone does not capture the full extent of consumption predictability. Therefore, in Panel B, we add to the AR(1), as drivers of the conditional mean, the same principal components of asset returns

¹⁹The half-life of a covariance-stationary GARCH (1,1) process is $1 + \log(1/2) / \log(\alpha + \beta)$.

that we found to predict consumption (see Appendix D). The resulting change is striking: When we better control for consumption predictability, the evidence in favor of time-varying volatility vanishes in all the model specifications, including the canonical GARCH(1,1). One cannot reject the hypothesis of constant consumption volatility in any specification.²⁰

Finally, in Panel C, we use the conditional mean of our moving average specification (without the contemporaneous shock) evaluated at its posterior mean. First, as in Panel B, there is no evidence of volatility clustering. Second, the sharply different results in Panel A relative to those in Panels B 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 would closely approximate it (as shown in Section II.1) and lead to similar implications for volatility. Empirically, however, using an $MA(\bar{S})$ parametrization for the conditional mean leads to drastically different estimations of the volatility models and shows that the canonical results are driven by the misspecification of the mean.

III.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 (2)–(3) and (5)–(6). The loadings of equity portfolios on the latent factor f_t are depicted in Figure 10.

The size- and book-to-market-sorted portfolios are ordered first (e.g., Portfolio 2 is the smallest quintile of size and the second smallest decile of book-to-market ratio), followed by the 12 industry portfolios (Portfolios 26–37 in the graph). All the portfolios have significant and positive exposures to the factor driving the conditional mean of consumption growth. Note, however, that this does not imply a single-factor model for returns, as our specification allows for multiple systematic factors in the return space orthogonal to f_t by the identification restriction. The findings remain unchanged when a bond-specific factor is added to the model as in equations (5)–(6) (see Figure IA.23 in the Internet Appendix).

The loadings are not only statistically but also quantitatively significant, as shown in Figure 11: The common (filtered) factor f explains on average 79% of the time-series variation of the 12 industry and 25 Fama and French (1992) size- and book-to-market-sorted


²⁰Note that a (I)GARCH(1,1) with $\alpha = \omega = 0$ and $\beta = 1$ is identical to a constant volatility model.

Figure 10: Common factor loadings (ρ^r) of the stock portfolios in the one-factor model.



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

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



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


equity portfolios, ranging from 36% to nearly 95% for individual portfolios.²¹

As shown in Figure IA.24 in the Internet Appendix, adding a bond-specific factor leaves the variance decomposition of stock returns basically unaffected.²²

²¹This observation is robust to a multi-factor model. In Panel A, Figure IA.46 of the Internet Appendix, we report the fraction of return variation explained by the f component when using a five-factor model to extract the conditional mean shocks in nondurable consumption growth. The results therein are extremely similar: the state-space model explains on average 73% of stock return variation.


²²The high explanatory power of the state-space model for asset returns is not a mechanical byproduct of our method. As the simulation studies in Figure 2 suggest, our method can capture the conditional dynamics

Figure 12: Bond loadings on common factor f_t .




The figure shows posterior means of the bond factor loadings on f_t (blue circles) and centered posterior 90% (dashed line) and 68% (dotted line) coverage regions in the one- and two-factor models.

Figure 13: Bond loadings (θ^b) on the bond-specific factor (b_t).



The figure shows posterior means (blue circles), the centered posterior 90% (dashed line) and 68% (dotted line) coverage regions, of bond loadings on the bond-specific factor b_t .

Figure 14: Variance decomposition box plots of bond returns.



The figure shows 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 b_t (right panel) shocks.

The loadings of the bond portfolios on the common consumption factor f_t are reported in Figures 12a and 12b, respectively, for 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 quite similar in the two specifications, with the main difference being that allowing for a bond-specific factor (b_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. Although 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 b_t are reported in Figure 13. These loadings are highly statistically significant and increase steeply and monotonically with maturity, revealing a rather pronounced term structure pattern.

Finally, Figure 14 reports the share of time-series variation of bond returns explained by the f_t shocks (left panel) and the f_t and b_t shocks (right panel). It 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%) but statistically significant proportion of the time-series variation in bond returns. The bond-specific factor, in turn, captures most of the residual time-series variation in returns. Although the model captures almost 55% of the variation in the six-month bond returns, its performance rapidly improves with maturity and results in a nearly perfect fit for time horizon of about five years.

Once again, our setting has an important testable implication that can be verified directly from the data without making all of the ancillary assumptions of the state-space model. In particular, equation (8) implies that, given our state-space results for the consumption loadings in Figure 3, 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 IA.35 of the Internet Appendix, 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

even when the conditional mean of consumption growth is determined by small PCs, in which case the f components should explain relatively small shares of return variation. To verify this property of the method in the data, we utilize the same state-space model to extract the f components in the growth rate of real GDP, durable consumption, consumer price index (CPI), and service and report the related results in Panels B, C, D, and E, Figure IA.46 of the Internet Appendix, respectively. On average, the state-space models of GDP, durable consumption, CPI, and service growth explain 69%, 35%, 34%, and 17% of the time-series variation of 37 stock returns.

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), who use $\Delta c_{t,t+1+S}$ to price time t asset returns.

III.4 External predictors of consumption and returns

The previous analysis assumes that asset returns are serially uncorrelated over time. However, our framework is flexible in incorporating other predictors in consumption growth and asset returns. We use two predictors to introduce serial correlations: the market’s price-dividend ratio and the Chicago Fed National Activity Index (CFNAI). The two-factor model is as follows:

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + \sum_{j=0}^{\bar{S}} \theta_j b_{t-j} + \gamma_1^c pd_{t-1} + \gamma_2^c cfnai_{t-1} + w_t^c, \quad (13)$$

$$\text{Stocks: } r_{st}^e = \mu_{sr} + \rho^{sr} f_t + \gamma^s pd_{t-1} + w_t^{sr}, \text{ and} \quad (14)$$

$$\text{Bonds: } r_{bt}^e = \mu_{br} + \rho^{br} f_t + \theta^b b_t + \gamma^b cfnai_{t-1} + w_t^{br}, \quad (15)$$

where the pd_{t-1} is the lagged price-dividend ratio of the market portfolio and $cfnai_{t-1}$ is the lagged CFNAI.²³ In addition, we allow these two predictors to drive the conditional dynamic of consumption growth, with the goal of testing whether they can crowd out the importance of the MA component of f_t .

We first test whether the P/D ratio and CFNAI can predict consumption growth and asset returns. Note that the CFNAI has been available since January 1967, so the sample used in this subsection differs slightly from that of other parts. Figure IA.25 in the Internet Appendix plots the distribution of γ^c , γ^s , and γ^b in equations (13)–(15). While the P/D ratio predicts neither consumption growth nor stock returns, CFNAI is essential in characterizing the conditional dynamics of longer-maturity bonds (3Y–10Y) and consumption growth. Because we standardize the predictors to have unit variances, the estimate implies that if CFNAI rises by one standard deviation, consumption growth will increase by 0.13% quarterly, with a 95% confidence interval of [0.01%, 0.25%].

Next, we study the identification of f_t shock and consumption’s impulse responses. The

²³CFNAI is a weighted average of 85 monthly indicators of national economic activity, delivering a rough measure of a common factor in these national economic data. In fact, this measure is similar to the canonical measure proposed by James Stock and Mark Watson, which has been used to predict inflation since the late-1990s.

empirical results, as shown in Figure IA.27, are reassuring: We still detect the slow adjustment of consumption growth to f_t shock, although the size slightly declines as CFNAI captures part of the predictability. Moreover, the MA component, loadings on f_t , and variance decompositions of consumption growth have almost identical patterns as previously.

III.5 Multifactor analysis

The simulation study in Section II.2 shows that our framework can identify the conditional dynamic of consumption growth even if the common component, f_t , is a small PC of asset returns or a linear combination of several PCs. To ensure that the single-factor analysis in the previous sections does not neglect important information in returns, we now consider a five-factor model for asset returns in equations (2)–(3).

Figure IA.29 in the Internet Appendix plots the related results. First, all the important estimates, including the MA component of consumption growth, impulse responses to f_t shock, stocks and bonds' loadings on f_t shock, and the variance decomposition of predicting consumption growth, are almost unchanged using the five-factor model. Hence, the single-factor model that we considered before can characterize the key patterns of the consumption mean process in the simplest form. Second, it is worth noting that the multifactor model introduces additional estimation noise. For instance, predictive R_{adj}^2 in Panels (e) and (f) are centered around the same estimates yet have wider posterior coverage than those in the single-factor case. As more latent factors do not identify any additional useful information but do increase estimation errors, we focus on the more parsimonious single-factor model in the following section and report the multifactor results as robustness checks.

III.6 Short- and long-run stochastic volatilities

Although our state-space model for consumption and returns in equations (2)–(3) 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 SV process for consumption and asset returns. This approach has been commonly used because it provides time-varying risk premia for equilibrium asset returns in representative agent models.

As shown in Section III.2, ad hoc specifications for the mean process (e.g., constant or AR(1)), which miss the true degree of predictability in the consumption process, mechanically

invalidate the identification of the consumption volatility process and its predictability. This is why we rely on a long MA representation that can correctly capture the conditional mean dynamics irrespective 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 manage inference on a class of structural models rather than a particular parametrization.

We generalize our state-space model in equations (2)–(3) to allow for stochastic volatilities in *all* innovations. Furthermore, we allow the volatilities to drive time-varying risk premia; that is, the process for log excess returns is now

$$\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}_r^r f_t + \boldsymbol{\beta}_c \sigma_{c,t-1}^2 + \boldsymbol{\beta}_f \sigma_{f,t-1}^2 + \boldsymbol{\beta}_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r, \quad (16)$$

where $\sigma_{c,t-1}$ is the stochastic volatility of w_t^c ; $\sigma_{f,t-1}$ is the stochastic volatility of f_t , the innovation to expected consumption growth; and $\sigma_{r,t-1}$ is a common market volatility process that affects the volatility of all excess return shocks (\mathbf{w}_t^r). In the language of Schorfheide, Song, and Yaron (2018), $\sigma_{c,t-1}$ and $\sigma_{f,t-1}$ are, respectively, the so-called short- and long-run volatilities of consumption, and the latter captures the time-varying volatility of expected consumption growth.

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 :

$$w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N} \left(0, \underbrace{\exp(h_{ct})}_{\sigma_{c,t-1}^2} \right), \quad f_t \stackrel{\text{iid}}{\sim} \mathcal{N} \left(0, \underbrace{\exp(h_{ft})}_{\sigma_{f,t-1}^2} \right), \text{ and} \quad (17)$$

$$\mathbf{w}_t^r \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{r,t-1}), \quad \boldsymbol{\Sigma}_{r,t-1} = \text{diag}\{\sigma_{r,1,t-1}^2, \dots, \sigma_{r,N,t-1}^2\}, \quad (18)$$

where $\sigma_{r,i,t-1}^2$ 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 ($\sigma_{r,t-1}^2 \equiv \exp(h_{rt})$); that is,

$$\sigma_{r,i,t-1}^2 = \exp\{\kappa_0^{(i)} + \kappa_1^{(i)} h_{rt} + e_{r,i,t}\}. \quad (19)$$

In what follows, for simplicity, we refer to $\sigma_{r,t-1}^2 \equiv \exp(h_{rt})$ as the “market” variance, albeit in our formulation of asset return dynamics, the total volatility of market returns is a combination of $\sigma_{r,t}$ and $\sigma_{f,t}$.

The formulation above is rather general and encompasses the standard models in the literature as particular cases. For instance, in the canonical long-run risk model of Bansal

and Yaron (2004), the short- and long-run volatilities of consumption are identical ($\sigma_{c,t}^2$ is proportional to $\sigma_{f,t}^2$), with the consumption dynamics summarized as follows:

$$\Delta c_{t-1,t} = \mu_c + x_{t-1} + w_t^c, \quad w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{t-1}^2), \quad \text{and} \quad x_t = \rho_x x_{t-1} + \varphi_e \sigma_{t-1} e_t, \quad e_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1),$$

which implies that $x_t = \varphi_e \sum_{j=0}^{\infty} \rho_x^j \sigma_{t-j-1} e_{t-j}$, where σ_t is the only time-varying economic uncertainty. This formulation can be mapped into our general framework for the consumption dynamics in equation (2), imposing the restrictions i) $\rho_0 = 0$ and $\rho_j = \varphi_e \rho_x^{j-1}$ and ii) $f_t = \sigma_{t-1} e_t$. Hence, in this case, w_t^c and f_t follow the same SV process.

Schorfheide, Song, and Yaron (2018) further separate the volatility of consumption into short- and long-run components, as follows:

$$\Delta c_{t-1,t} = \mu_c + x_{t-1} + \sigma_{c,t-1} \eta_{c,t} \quad \text{and} \quad x_t = \rho_x x_{t-1} + \sqrt{1 - \rho_x^2} \sigma_{x,t-1} \eta_{x,t},$$

where $\sigma_{i,t}$ ($i \in \{c, x\}$) denotes two distinct stochastic volatility processes, and $\eta_{c,t}$ and $\eta_{x,t}$ follow uncorrelated standard normal distributions. In this formulation, $\sigma_{c,t}$ ($\sigma_{f,t}$) is meant to capture the short-run (long-run) volatility in consumption. This setup also fits within our general MA formulation in equation (2) after imposing the following restrictions: i) $w_t^c = \sigma_{c,t-1} \eta_{c,t} \sim \mathcal{N}(0, \sigma_{c,t-1}^2)$, ii) $x_t = \sqrt{1 - \rho_x^2} \sum_{j=0}^{\infty} \rho_x^j \sigma_{x,t-j-1} \eta_{x,t-j}$, iii) $\rho_0 = 0$ and $\rho_j = \sqrt{1 - \rho_x^2} \rho_x^{j-1}$, and iv) $f_t = \sigma_{x,t-1} \eta_{x,t}$ and $\sigma_{f,t} = \sigma_{x,t}$.

In order to identify the model, without loss of generality, we normalize h_{ft} and the common (log) 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 our unrestricted formulation. See Internet Appendix B.2 for further details, including the MCMC algorithm needed to evaluate this model.

Using our general stochastic volatility formulation, the estimated parameters for the consumption ($\{\rho_j\}_{j=0}^S$) and asset returns ($\boldsymbol{\rho}^r$) loadings on the common shock f remain essentially unchanged (see Figure IA.30 in the Internet Appendix). Similarly, the variance decompositions for consumption and returns are almost identical to those presented in the previous sections. Although this is not surprising (because the formulations (2)–(3) are 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. These shocks drive more than a quarter of the consumption growth variance. Overall, all of our results are robust to the modeling choice for volatility.

What do we learn about the volatility itself for both consumption and returns? Figure

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						
$cor(\sigma_{ct}^2, VXO_t^2)$	0.18	-0.03	-0.00	0.18	0.39	0.43
$cor(\sigma_{ft}^2, VXO_t^2)$	0.45	0.22	0.25	0.45	0.62	0.65
$cor(\sigma_{rt}^2, VXO_t^2)$	0.52	0.37	0.40	0.53	0.62	0.63
Panel B: Pairwise correlations of vol processes						
$cor(\sigma_{ct}^2, \sigma_{ft}^2)$	0.07	-0.09	-0.07	0.06	0.26	0.30
$cor(\sigma_{ct}^2, \sigma_{rt}^2)$	0.15	-0.02	0.00	0.14	0.32	0.36
$cor(\sigma_{ft}^2, \sigma_{rt}^2)$	0.22	0.06	0.08	0.21	0.38	0.41

The table summarizes 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.


15 reports the estimated volatility processes (posterior median and 95% credible intervals) under a diffuse prior for the autoregressive coefficients of the processes. Several observations are in order. First, there is clear evidence of time-varying volatility in stock returns (Panel C), whereas our findings are more nuanced for the consumption volatilities. Second, for the short-run consumption volatility ($\sigma_{c,t}$ in Panel A), there is an entire range of constant values that are within the posterior confidence bands. This finding holds even when we consider monthly consumption data as well as the mixed-frequency state-space formulation that accounts for time-aggregation bias, as we show in the Internet Appendix (see, respectively, Figures IA.6 and IA.12). Third, instead, while the estimated long-run consumption volatility ($\sigma_{f,t}$) in Panel B of Figure 15 contains an entire range of constant values within the posterior credible intervals, there is no such parameter region i) using monthly consumption data (see Figure IA.6), ii) estimating a mixed-frequency model (see Figure IA.12), and iii) directly incorporating a leverage effect into $\sigma_{f,t}$ (see Figure IA.16).²⁴ In other words, the long-run consumption volatility becomes more sharply identified in latter three cases. Note also that the confidence bands in both Panels A and B are sufficiently wide to allow for potential substantial time variation in both short-run and long-run consumption volatilities. Fourth, all asset volatilities have significant loadings (κ_1^i) on the common financial market volatility process (see Figure IA.45 of the Internet Appendix). Fifth, the common financial market volatility increases during recessions and market crashes, whereas short- and long-run consumption volatilities do not display such a clear pattern.

²⁴We extend our framework to allow for a leverage effect in the long-run consumption volatility following Omori, Chib, Shephard, and Nakajima (2007). Details are reported in Internet Appendix J.


Figure 15: Filtered stochastic volatilities of consumption and returns.



Panel A: Log volatility of the short-run consumption shock (w_t^c) of equation (2)




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 (19)).

Estimated stochastic volatilities of the model in equations (2), and (16)–(19), and Section B.2 under a diffuse prior for the autoregressive volatility coefficients. Solid blue lines depict the posterior median of the log volatility, whereas dotted red lines denote 2.5% and 97.5% credible intervals. Shaded (patterned) 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_c) on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$) in equation (16).



Panel B: 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 (16).



Panel C: Posterior distributions of excess return loadings (β_r) on the common financial return variance ($\sigma_{r,t-1}^2$) in equation (16).

This figure shows the box plots of the posterior distributions of the loadings of portfolio excess returns on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$), the variance of shocks to the conditional consumption growth ($\sigma_{f,t-1}^2$), and the common financial returns variance ($\sigma_{r,t-1}^2$). Portfolios are ordered with bonds first (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last.

We further explore the correlation between the estimated volatilities and (the square of) the VVO volatility index (a proxy for the underlying volatility process commonly used in the literature) in Table 4.²⁵ Both the financial market and long-run consumption volatilities

²⁵We use the VVO index, rather than VIX, due to the longer time-series available for the former. The VVO (as is the VIX), measured in discrete time, is adapted to both the physical underlying volatility process and the “true” latent stochastic discount factor. Hence, the correlations we measure are potentially generated by either or both of them.

have nontrivial correlations with VXO (.53 for the former and .45 for the latter, at the median, in Panel A), while being weakly correlated with each other (the median correlation is only .21, Panel B). 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 45% 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 (rather general) modeling choice for the joint dynamics of consumption and returns.

We provide additional external validation of our estimated volatility processes in Table IA.XII of the Internet Appendix, which reports the correlation between the filtered stochastic volatilities and the real/macro/financial uncertainty measures in Jurado, Ludvigson, and Ng (2015) and Ludvigson, Ma, and Ng (2021). We find that the long-run consumption and asset-market volatilities significantly load on all three uncertainty indices. On the contrary, the short-run consumption volatility weakly correlates with these uncertainty measures at best.

Overall, our findings suggest, as in Schorfheide, Song, and Yaron (2018), different dynamics for the short- and long-run consumption volatilities, but also a distinct process for the market common volatility.

But do the volatilities of returns and consumption drive excess return dynamics? Figure 16 reports the posterior distributions of the loadings in equation (16). Strikingly, none of the excess return series loads significantly on short- and long-run consumption volatilities (Panels A and B). This finding is robust even when we i) use monthly consumption data, ii) estimate a mixed-frequency model, and iii) directly incorporate a leverage effect into $\sigma_{f,t}$ (see, respectively, Figures IA.7, IA.13 and IA.17 of the Internet Appendix). This evidence poses a challenge to the literature that has modeled time-varying risk premia by assuming time variation in the volatility process of consumption. Panel C shows some support for the common financial market volatility being negatively associated with excess returns. Nevertheless, this latter finding is generally weak and fragile with respect to the various generalizations that we consider (see Figures IA.7, IA.13, and IA.17 of the Internet Appendix).

The generality of our setting also allows us to formally evaluate a multiplicity of alternative formulations considered in the previous literature. We do so in Section I of the Internet Appendix. The data favor a specification in which *i*) excess returns do not load on *any* of the stochastic volatilities and *ii*) short- and long-run consumption volatilities are different

from each other and distinct from the common market volatility – the posterior probability of such a formulation is almost 100%.

The finding that excess returns do not significantly load on the volatilities is perhaps not too surprising, given the mixed and inconclusive evidence on the volatility-return trade-off in the literature (see, e.g., French, Schwert, and Stambaugh (1987), Nelson (1991), and Campbell and Hentschel (1992)). However, Bandi and Perron (2008) show that this tradeoff becomes stronger at longer horizons. Motivated by this, we explore whether the connection between returns and volatilities strengthens at lower frequencies. In particular, we examine whether the common market and (short- and long-run) consumption volatilities predict multi-horizon excess returns.²⁶ Predictive regression coefficients for four- to 12-quarter cumulative returns are reported in Figures IA.42–IA.44 of the Internet Appendix. Overwhelmingly, we find insignificant coefficients at all horizons. That is, volatilities do not seem to predict excess returns even at these lower frequencies.

III.7 Further Robustness Checks

We conduct an extensive set of additional robustness checks and show that all our key results presented in the previous sections remain unchanged. In this section, we discuss the main exercises and report additional ones in Section K of the Internet Appendix.

First, one may be concerned that our empirical results strongly depend on a particular choice of the MA length of the consumption mean process, \bar{S} . The data indicate that this is not the case. When the MA component includes at least six quarters of data, our results remain almost unchanged. Figure A1 in the Appendix reproduces all the main findings for the state-space formulation estimated with $\bar{S} = 20$ quarters, that is, five years of data, and shows that they are qualitatively and quantitatively identical to those presented in the previous sections. We find that the similar size of the cumulative impulse response of consumption is approximately 1%. The reason for this similarity is that the incremental consumption impulse response to f_t is almost zero after two to three years. Note that this finding is consistent with the evidence in Bandi and Tamoni (2023), who show that most of the (priced) consumption risks occur within the business cycle frequency.

²⁶Given the AR(1) nature of our log stochastic volatility processes, this exercise is equivalent to investigating whether the conditional variance of the multi-period consumption, $\text{var}_t(\Delta c_{t,t+H})$ ($H > 1$), can forecast cumulative excess returns, $r_{t,t+H}^e$.

Second, because the state-space estimation results may depend on the choice of the cross-section of assets used for the analysis, we repeat our analysis using a multiplicity of alternative base assets, as discussed in Section K of the Internet Appendix. In particular, Figures IA.34, IA.40, and IA.41 of the Internet Appendix present all the key empirical findings of the paper, estimated on a wide cross-section of characteristic-based anomaly portfolios used in Kozak, Nagel, and Santosh (2020) (KNS). Our findings either remain virtually unchanged, or become even stronger, which again emphasizes their robustness to the choice of test assets.

To summarize, our findings are robust to a plethora of perturbations to the baseline specification, all supporting the same key takeaways: *i*) the conditional mean of consumption growth is spanned by returns; *ii*) consumption reacts to financial market innovations slowly, with the total cumulative effect peaking after two to three years; *iii*) short- and long-run consumption shocks and returns have distinct volatility processes, and neither seems to drive time-varying risk premia.

IV Asset Pricing Implications

The first-order importance of financial markets in driving the consumption process that we have uncovered raises the natural question of whether the conditional mean shocks command a significant risk premium. Furthermore, can our identified consumption process improve the performance of basic consumption-based representative-agent models with lower risk aversion? Can we also explain the level and volatility of the risk-free rate? That is, can we jointly provide an explanation of both the equity premium (Mehra and Prescott (1985)) and risk-free rate (Weil (1989)) puzzles? The answers to all of these questions, as illustrated in the next two subsections, are affirmative.

IV.1 The risk premium of consumption mean shocks

We now turn to assessing the cross-sectional pricing of the shocks to the consumption mean. We do not take a stand on the preferences determining risk prices, as we only want to assess whether the f_t shocks are part of the stochastic discount factor (SDF). Given the evidence in Bryzgalova, Huang, and Julliard (2023) that the SDF in the economy is *dense* in the space of observable factors, we need to account for the omitted variables. We do so following Giglio and Xiu (2021) and Bryzgalova, Huang, and Julliard (2024); that is, we maintain the

assumption about asset return dynamics in equation (3) and further postulate the following log linear SDF:

$$\Delta \mathbb{E}_t[m_t] := m_t - \mathbb{E}_{t-1}[m_t] = - \underbrace{b_f f_t}_{CCAPM} - \underbrace{\mathbf{b}_g^\top \mathbf{g}_t}_{omitted \ factors}, \quad (20)$$

where b_f is the risk price of the shock to conditional consumption mean, and the second component, $\mathbf{b}_g^\top \mathbf{g}_t$, is included to account for the omitted factor bias in estimating f_t 's risk price.²⁷ Also note that only the shocks to the SDF, $\Delta \mathbb{E}_t[m_t]$, matter for the pricing excess returns. Furthermore, since f_t is normalized to have unit variance and is (by construction) orthogonal to the omitted factors, b_f is both its risk premium, price of risk, and (tradable) Sharpe ratio.

Table 5 reports the estimates of f_t 's risk premium in several cross-sections of excess returns. First, f_t , the shock to both asset returns and conditional consumption mean, carries a stable and significant risk premium in all cross-sections. Hence, portfolios with higher consumption risk exposure, on average, earn higher returns. Furthermore, our estimates of the quarterly risk premium (b_f) are virtually identical to the ones delivered by the misspecification-robust estimator of Giglio and Xiu (2021). Note that both Bayesian and frequentist approaches recover the risk premium as defined by the negative covariance of the factor(s) with the tradable SDF (spanned by latent tradable factors). By doing so, both estimators are robust to weak identification and model misspecification.

Second, the Sharpe ratio of the consumption risk is about 0.44-0.55 per annum. Nevertheless, there are additional sources of risk driving the cross-section of asset returns: the Sharpe ratio of the entire SDF (which further includes the latent factors g_t) is about 0.8–1.0. The column reporting $\mathbb{E} \left[\frac{SR_f^2}{SR_m^2} \mid \text{data} \right]$ – the share of SDF variance explained by consumption mean shocks – further quantifies the importance of consumption risk in the SDF. Across panels, consumption risk explains 22-37% of the total SDF variance. In other words, consumption risk is a first-order component of the SDF, but there is also a significant amount of additional shocks, orthogonal to consumption risk, that drive the cross-section of asset returns.

Third, our SDF has high explanatory power for all the cross-sections, as indicated by both the high cross-sectional R^2 in the 41-86% range, and extremely small and insignificant average pricing errors.

²⁷We can interpret $\mathbf{b}_g^\top \mathbf{g}_t$ as the missing component in C-CAPM as in Ghosh, Julliard, and Taylor (2016).

Fourth, the estimated SDF can almost perfectly capture the 7% average annual market risk premium²⁸ (column $-\text{cov}(r_t^{mkt}, m_t)$). And crucially, it does so almost entirely thanks to the consumption risk component (column $-\text{cov}(r_t^{mkt}, -b_f f_t)$).

Finally, the results in Table 5 appear to be extremely robust. First, a comparison of the results in Panels A and D highlights that the estimates of consumption risk premia are not driven by the differential exposures of stocks and bonds to the f_t shocks—removing bonds does not affect risk premia estimates (see Panel D). Second, using a wide cross-section of stock anomalies (the ones of Kozak, Nagel, and Santosh (2020) in Panel B) yields almost identical risk premia estimates as with our baseline cross-section (in Panel A). Third, a comparison of Panel B vs. Panel C, and of Panel D vs. E, highlights that both the Bayesian (Bryzgalova, Huang, and Julliard (2024)) and frequentist (Giglio and Xiu (2021)) robust estimation methods are not affected by span-invariant repackaging of the test assets. Internet Appendix N presents estimates based on the frequentist Fama-MacBeth regressions. In contrast to the robust estimators of Table 5, risk premia estimates produced by Fama-MacBeth regressions are very sensitive (as theoretically expected) to model parametrization and repackaging of the test assets. This, however, is not a property of the shocks to consumption mean and the risk premia they carry, but an artifact of the fragility of Fama-MacBeth estimation (see the discussion in Internet Appendix N).

But what type of returns predict consumption? That is, what type of portfolios and which industries are strongly correlated with the conditional consumption mean shocks?

Table 6 reports the correlation of the filtered f_t shocks with common tradable risk factors and leading principal components of returns for two different cross-sections of base assets: our baseline one and the KNS anomaly portfolios. First and foremost, we observe a strong and stable correlation of the consumption innovations with the overall market index (81 – 92%) and the first principal component of the based assets (85 – 96%). That is, financial market-wide shocks are the main drivers of the consumption mean process. This explains why, in Table 5, the conditional mean shocks can, at least in reduced form, generate a risk premium that almost perfectly matches the historical mean excess returns on the market index. Second, the returns on small firms are particularly salient, with large loadings on f_t (see Figure 10) declining monotonically with the size characteristic of the portfolios. This, in turn, leads to a strong correlation in Table 6 with the SMB factor (about 60%

²⁸In the sample between 1963:Q3 and 2019:Q4, the time-series average of log market excess returns, $\mathbb{E}[r_t^m] + 0.5\text{var}(r_t^m)$, is around 6.8% per year.

Table 5: Cross-sectional estimates of risk premia of conditional consumption mean shocks

Posterior	Average α	Estimating $\Delta\mathbb{E}_t[m_t] = -b_f f_t - \mathbf{b}_g^\top \mathbf{g}_t$					$\mathbb{E}[r_t^{mkt}] =$	GX (2021)
		b_f	$\mathbb{E}[SR_m \mid \text{data}]$	$\mathbb{E}[SR_f \mid \text{data}]$	$\mathbb{E}\left[\frac{SR_f^2}{SR_m^2}\right] \mid \text{data}$	R^2	$-\text{cov}(r_t^{mkt}, m_t)$	
Panel A. 37 value-size-industry stock and nine bond portfolios in a five-factor model								
Median	0.000	0.236	0.806	0.474	0.371	0.772	0.069	0.070
90% CIs	[0.000, 0.000]	[0.079, 0.385]	[0.554, 1.128]	[0.180, 0.771]	[0.056, 0.756]	[0.477, 0.903]	[0.033, 0.109]	[0.025, 0.114]
Panel B. 74 Kozak, Nagel, and Santosh (2020) decile 1 and decile 10 anomaly portfolios in a five-factor model								
Median	0.000	0.244	0.960	0.489	0.263	0.409	0.071	0.069
90% CIs	[0.000, 0.000]	[0.109, 0.382]	[0.735, 1.200]	[0.220, 0.763]	[0.059, 0.563]	[0.213, 0.580]	[0.034, 0.111]	[0.030, 0.110]
Panel C. 37 Kozak, Nagel, and Santosh (2020) decile 10 portfolios plus 37 long-short portfolios in a five-factor model								
Median	-0.002	0.219	0.936	0.437	0.223	0.795	0.087	0.065
90% CIs	[-0.003, -0.002]	[0.069, 0.358]	[0.710, 1.170]	[0.139, 0.716]	[0.026, 0.550]	[0.657, 0.875]	[0.050, 0.125]	[0.020, 0.109]
Panel D. 37 value-size-industry stock portfolios in a five-factor model								
Median	0.000	0.219	0.715	0.438	0.399	0.449	0.068	0.063
90% CIs	[0.000, 0.000]	[0.096, 0.344]	[0.493, 0.938]	[0.192, 0.689]	[0.087, 0.770]	[0.050, 0.718]	[0.030, 0.106]	[0.027, 0.100]
Panel E. Transformed 37 value-size-industry stock portfolios in a five-factor model								
Median	0.001	0.276	0.944	0.553	0.354	0.863	0.079	0.078
90% CIs	[0.000, 0.001]	[0.134, 0.417]	[0.715, 1.170]	[0.269, 0.833]	[0.096, 0.670]	[0.666, 0.941]	[0.040, 0.116]	[0.038, 0.118]

Estimation results for five cross-sections of excess returns: 37 value-size-industry stock and nine bond portfolios (Panel A), Kozak, Nagel, and Santosh (2020) 74 decile 1 and 10 characteristic-sorted portfolios (Panel B), Kozak, Nagel, and Santosh (2020) 37 decile 10 portfolios plus 37 long-short (decile 1 minus decile 10) portfolios (Panel C), 37 value-size-industry stock portfolios (Panel D), and transformed 37 value-size-industry stock portfolios (Panel E). In the transformed cross-section of 37 value-size-industry stock portfolios (Panel E), we long portfolios 1–19; for the remaining 18 portfolios, we construct the long-short portfolios by buying portfolios 20–37 but short-selling portfolios 1–18. We report: (1) average cross-sectional mispricing (α); (2) risk price of the shock to the conditional consumption mean f_t (b_f); (3) annualized Sharpe ratio of the SDF in equation (20), defined as the annualized volatility of the SDF (SR_m); (4) annualized Sharpe ratio of $b_f f_t$ (SR_f); (5) ratio of SR_f^2 to SR_m^2 ; (6) cross-sectional R^2 ; (6) (annualized) market risk premium implied by the SDF, $-\text{cov}(r_t^{mkt}, m_t)$; (7) (annualized) market risk premium implied by the covariance between market excess return and $-b_f f_t$, $-\text{cov}(r_t^{mkt}, -b_f f_t)$; (8) risk premia estimates and their t -statistics based on the three-pass procedure in Giglio and Xiu (2021), with corresponding standard errors. We estimate the risk prices using the Bayesian approach of Bryzgalova, Huang, and Julliard (2024). Details are provided in Internet Appendix M.

in both cross-sections). Third, the correlation with the value anomaly (HML) is small and not statistically significant due to the U-shaped pattern of loadings in Figure 10 in the Growth/Value characteristic space. Fourth, the correlations with both the Robust-Minus-Weak (operating profitability) portfolio and also the Conservative-Minus-Aggressive (investment) portfolio are significantly negative and sizable (respectively, 22 – 30% and 20–30%). Fifth, there is more than solely the first PC of returns reflected in the consumption process, with higher-order PCs being often significant predictors of consumption growth (PC4 in Panel A and both PC2 and PC5 in Panel B).

Figure 10 also reports the loadings of the 12 industry portfolios on the consumption mean shocks. All industries have statistically significant and strong-to-moderate loadings on the conditional consumption mean innovations. In particular, Consumer Durables, Business Equipment, and Finance (Money) display the highest coefficients (11 – 12%) while Utilities, Energy, and Telephones and Television Transmission (Telecom) have the lowest (5 – 6%).

Table 6: What drives the consumption shocks spanned by financial markets?

Panel A: Baseline stock and bonds cross-section					
	MKT	SMB	HML	RMW	CMA
Correlation	0.917	0.605	-0.127	-0.221	-0.295
95% CI	[0.863, 0.959]	[0.515, 0.700]	[-0.292, 0.087]	[-0.267, -0.173]	[-0.425, -0.126]
	PC1	PC2	PC3	PC4	PC5
Correlation	0.964	0.072	-0.007	-0.132	-0.024
95% CI	[0.922, 0.990]	[-0.079, 0.258]	[-0.136, 0.120]	[-0.259, -0.048]	[-0.048, 0.000]
Panel B: Kozak, Nagel, and Santosh (2020) 74 anomaly portfolios					
	MKT	SMB	HML	RMW	CMA
Correlation	0.811	0.601	-0.0579	-0.299	-0.200
95% CI	[0.71 , 0.909]	[0.515 , 0.689]	[-0.201 , 0.0757]	[-0.405 , -0.176]	[-0.316 , -0.083]
	PC1	PC2	PC3	PC4	PC5
Correlation	0.851	0.291	0.0674	-0.0716	0.279
95% CI	[0.761 , 0.937]	[0.14 , 0.441]	[-0.117 , 0.244]	[-0.209 , 0.0624]	[0.126 , 0.406]

Correlation coefficients between f_t and Fama-French five factors and the first five principal components of asset returns. We also report their 95% posterior credible intervals under the coefficient estimates. In Panels A, the cross-section of base assets used to estimate f_t includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios, while in panel B we employ the Kozak, Nagel, and Santosh (2020) cross-section of 74 characteristic-sorted portfolios.

IV.2 Implications for Structural Models

The reduced-form estimate of f_t 's risk price obtained in the previous subsection (Table 5) has direct implications for structural models, in particular, for the coefficient of relative risk aversion (RRA) and the implied risk-free rate. We consider two types of representative agent preferences: (1) additively separable CRRA and (2) Epstein-Zin (EZ) recursive utility – that separates RRA from the elasticity of intertemporal substitution (IES). In Internet Appendix O, we show how the estimated market price of consumption mean risk and our estimated dynamics for the consumption process can be mapped into utility parameters and preference-implied risk-free rate level and volatility.

Note that the estimated market price of consumption risk in Table 5 already generates a market risk premium in line with its historical average. Hence, the questions to ask are whether, given a preference setting, too high RRA is required to yield the estimated price of risk and whether the implied risk-free rate is realistic. That is, we ask whether our uncovered dynamics for the consumption process can jointly match the equity premium and risk-free rate puzzles.

Panels A–D of Table 7 report the results based on the aggregate consumption data. For CRRA preferences, we report the implied RRA level in the canonical C-CAPM ($S = 0$) and also using the ultimate consumption risk approach of Parker and Julliard (2005) ($S = 2, \dots, 14$), which allows consumption to react slowly to return innovations and back out

Table 7: Implications for structural models

	$S = 0$	2	4	6	8	10	12	14	$\psi = 1.5$	$\psi = 0.5$
Panel A. Aggregate consumption in 37 stock and nine bond portfolios (constant volatility)										
RRA	144.3	52.1	33.7	28.5	28.3	26.7	24.7	25.1	26.0	27.0
90% CI	[47.9, 234.9]	[17.3, 84.9]	[11.2, 54.9]	[9.5, 46.5]	[9.4, 46.1]	[8.9, 43.4]	[8.2, 40.3]	[8.4, 40.9]	[9.0, 41.9]	[10.1, 43.0]
$\mathbb{E}[r_f]$	29.40%	51.60%	39.03%	34.51%	34.32%	32.77%	30.87%	31.28%	1.71%	3.72%
$\sigma(r_f)$	103.27%	37.31%	24.11%	20.43%	20.29%	19.10%	17.69%	18.00%	0.48%	1.43%
Panel B. Aggregate consumption in 37 stock and nine bond portfolios (stochastic volatility)										
RRA	180.2	76.5	41.4	32.6	30.2	25.0	20.8	20.1	21.1	22.3
90% CI	[59.8, 293.4]	[25.4, 124.6]	[13.8, 67.4]	[10.8, 53.1]	[10.0, 49.2]	[8.3, 40.6]	[6.9, 33.9]	[6.7, 32.8]	[7.4, 34.0]	[8.6, 35.1]
$\mathbb{E}[r_f]$	12.86%	65.02%	46.60%	39.13%	36.88%	31.68%	27.28%	26.56%	2.08%	3.72%
$\sigma(r_f)$	192.96%	65.77%	32.62%	25.11%	23.10%	18.82%	15.51%	15.00%	0.49%	1.43%
Panel C. Aggregate consumption in Kozak, Nagel, and Santosh (2020) 74 anomaly portfolios (constant volatility)										
RRA	150.7	42.6	27.7	21.7	19.9	17.7	16.0	15.7	16.5	17.7
90% CI	[67.3, 235.4]	[19.0, 66.5]	[12.4, 43.2]	[9.7, 33.9]	[8.9, 31.1]	[7.9, 27.7]	[7.1, 24.9]	[7.0, 24.5]	[7.7, 25.5]	[8.9, 26.7]
$\mathbb{E}[r_f]$	44.65%	47.59%	34.42%	28.28%	26.30%	23.87%	21.83%	21.55%	1.77%	4.28%
$\sigma(r_f)$	126.13%	35.66%	23.14%	18.17%	16.65%	14.84%	13.35%	13.15%	0.56%	1.67%
Panel D. Aggregate consumption in Kozak, Nagel, and Santosh (2020) 74 anomaly portfolios (stochastic volatility)										
RRA	159.7	46.4	30.2	22.1	19.5	16.8	14.6	14.4	15.2	16.4
90% CI	[71.3, 249.4]	[20.7, 72.4]	[13.5, 47.2]	[9.9, 34.6]	[8.7, 30.5]	[7.5, 26.2]	[6.5, 22.8]	[6.4, 22.4]	[7.1, 23.4]	[8.3, 24.6]
$\mathbb{E}[r_f]$	41.64%	50.78%	37.09%	28.83%	25.96%	22.86%	20.29%	20.00%	2.16%	3.86%
$\sigma(r_f)$	158.52%	40.89%	26.18%	19.00%	16.71%	14.33%	12.43%	12.22%	0.61%	1.71%
Panel E. Shareholders' consumption in 37 stock and nine bond portfolios (constant volatility)										
RRA	110.1	17.9	11.0	10.0	9.8	7.0	7.4	5.7	6.5	7.7
90% CI	[36.5, 179.2]	[5.9, 29.1]	[3.6, 17.8]	[3.3, 16.2]	[3.3, 16.0]	[2.3, 11.5]	[2.4, 12.0]	[1.9, 9.4]	[2.6, 10.2]	[3.8, 11.4]
$\mathbb{E}[r_f]$	49.52%	23.58%	15.67%	14.46%	14.27%	10.81%	11.22%	9.15%	1.47%	5.51%
$\sigma(r_f)$	78.77%	12.78%	7.84%	7.13%	7.02%	5.04%	5.27%	4.11%	0.48%	1.43%

Estimates of RRA, average implied annualized risk-free rate, and its annualized volatility, using two cross-sections of asset returns: (1) 37 stock and nine bond portfolios and (2) Kozak, Nagel, and Santosh (2020) 74 anomaly portfolios. The risk price of f_t (and its 90% CIs), denoted by b_f , are from Table 5. Cumulative impulse responses of *nondurable* consumption growth, $\sum_{j=0}^S \rho_j$, are estimated using our MA model, with $\bar{S} = 14$ quarters. We consider specifications with and without SV. In models with SV, we assume three separate SV processes, $\sigma_{c,t}^2$, $\sigma_{f,t}^2$, and $\sigma_{r,t}^2$, all of which are allowed to drive conditional mean returns. In the CRRA case, the RRA coefficients are given by $\gamma_S = b_f / \sum_{j=0}^S \rho_j$, where $S \in \{0, 2, \dots, 14\}$ (see equation (IA.44) of the Internet Appendix). In the Epstein-Zin case, we consider two different values of IES, i.e., $\psi \in \{0.5, 1.5\}$, and $\gamma = (b_f + \frac{1}{\psi} \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j) / (\rho_0 + \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j)$ (see equation (IA.54) of the Internet Appendix). Mean and volatility of the risk-free rate are obtained using equations (IA.42) and (IA.48) of the Internet Appendix, where $\delta = 0.9965$ (following the quarterly calibration of Bansal, Kiku, and Yaron (2016, Table 6)). In Panels A–D, we consider the aggregate nondurable consumption data, while Panel E uses shareholders' consumption data.

the coefficient of risk aversion.²⁹ For EZ preferences, we considered two possible calibrated values for the IES, $\psi = 1.5$ and $\psi = 0.5$. Not surprisingly, when we consider only the contemporaneous response of consumption growth to asset return shocks ($S = 0$), the implied RRA estimate in the CRRA preference is larger than 100. As we allow for a slow reaction of consumption to asset return shocks, for example, $S = 14$, the implied RRA is much lower (from 25 in Panel A to 15 in Panel D), and values below 10 are now within the 90% confidence bands. However, the implied real risk-free rate is too high (20–31%) and volatile (12–18%) with CRRA preferences.

Does our method improve the performance of the CRRA model compared to Parker

²⁹See equations (IA.43)–(IA.44) of the Internet Appendix.

(2001)? Figure IA.36 of the Internet Appendix plots the RRA coefficients implied by the VAR approach of Parker (2001). Similar to our state-space formulation, the VAR-implied RRA declines from more than 100 to 28 at the three-year horizon. However, the RRA coefficients based on our method can be significantly lower, particularly in the cross-section of KNS anomalies. What drives the difference between VAR and our MA model? The former assumes that shocks to the market index are the *only* spanned innovations in consumption. Instead, our MA model, as we show in simulations, and which seems to be the case in the data (see Table 6), can capture an arbitrary linear combination of systematic latent factors spanned by financial markets. Hence, when consumption growth reacts to other factors beyond only the market return (as in the cross-section of KNS anomalies), we detect a significant difference between the VAR- and MA-implied estimates of RRA.

EZ preferences in Panels A–D of Table 7 yield moderate levels of risk aversion (15 – 27), with a substantial likelihood of values below 10, irrespective of the calibrated IES coefficient. Instead, the risk-free rate level and volatility are quite sensitive to the choice of ψ . For $\psi = 0.5$, the implied level (3.7 – 4.3%) overshoots the historical average (1.68%), and the implied volatility (1.43 – 1.71%) roughly matches its historical, *ex-post*, volatility (2.44%).³⁰ For $\psi = 1.5$, EZ preferences almost perfectly match the risk-free rate level (1.71 – 2.16%) and provide reasonable values for its implied, *ex-ante*, volatility (0.48 – 0.61%).

Finally, in Panel E of Table 7, we compute the implied RRA coefficient and risk-free rate based on the impulse responses of shareholders' consumption growth (in Figure 4). In this case, due to the much larger responses of shareholders' consumption to return shocks, the RRA coefficients are about 6 – 8 for $S = 14$ and EZ preference.³¹ Nevertheless, only EZ preferences with $\psi = 1.5$ can also match well the observed risk-free rate.

V 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

³⁰Note that this historical volatility is an upper bound for the *ex-ante* real risk-free rate volatility that the models should match.

³¹A caveat is that the risk price estimate, b_f , is based on a larger sample (1963:Q3–2019:Q4), whereas the IRFs of shareholders' consumption growth are based on a limited sample (1982:Q1–2004:Q3).

function: Most shocks affecting the household force it to adjust both investment and consumption plans. We confirm this mechanism by using the micro-data on shareholders' and non-shareholders' consumption from Malloy, Moskowitz, and Vissing-Jorgensen (2009).

We show that a flexible parametric model with common factors driving returns and consumption identifies the slow-varying conditional mean of consumption growth. This component is persistent at business cycle frequency and is economically substantial, capturing more than a quarter of the variance of consumption growth. Furthermore, the predictability that we uncover does not appear to be caused by measurement issues of consumption such as time averaging, measurement error, and benchmarking. Hence, our findings indicate that the shocks spanned by financial markets are first-order drivers of consumption risk.

The conditional consumption mean shocks carry a significant risk premium in the cross-section of asset returns, and demand an annualized Sharpe ratio of about 0.5. Furthermore, we embed the estimated consumption dynamics into an otherwise standard Epstein and Zin (1989) preference setting with elasticity of intertemporal substitution above one (as in Bansal and Yaron (2004)). Such a calibration can match both the unconditional equity premium and the risk-free rate level and volatility, with low relative risk aversion.

We also show that detecting time variation in consumption volatility requires properly accounting for both the conditional mean process that we uncover and the measurement issues in consumption data. Nevertheless, we find no empirical evidence for the stochastic volatility of consumption driving time-varying risk premia.

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. In addition, we provide a sizable set of complementary evidence that does not rely on our parametric formulation but rather supports the main findings of our state-space analysis.

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

AÏT-SAHALIA, Y. (2004): "Disentangling diffusion from jumps," *Journal of Financial Economics*, 74(3), 487–528.
ANG, A., M. PIAZZESI, AND W. MIN (2006): "What Does the Yield Curve Tell Us about GDP Growth?," *Journal of Econometrics*, 131, 359–403.

BANDI, F. M., AND A. TAMONI (2023): "Business-Cycle Consumption Risk and Asset Prices," *Journal of Econometrics*, forthcoming.

BANSAL, R., R. DITTMAR, AND D. KIKU (2007): "Cointegration and Consumption Risks in Asset Returns," *The Review of Financial Studies*, 22(3), 1343–1375.

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.

——— (2016): "Risks for the Long Run: Estimation with Time Aggregation," *Journal of Monetary Economics*, (82), 52–69.

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. (2005): "Rare Events and the Equity Premium," NBER Working Paper 11310, National Bureau of Economic Research, Inc.

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

BAUWENS, L., M. LUBRANO, AND J.-F. RICHARD (1999): *Bayesian Inference in Dynamic Econometric Models*. Oxford University Press, Oxford, England.

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.

BLACK, F. (1976): "Studies of stock market volatility changes," *Proceedings of the American Statistical Association, Business & Economic Statistics Section*.

BLANCHARD, O. J., AND G. M. CONSTANTINIDES (2008): "Macroeconomic Crises since 1870. Comments and Discussion," *Brookings Papers on Economic Activity*, 2008(1), 336–350.

BLANCHARD, O. J., AND D. QUAH (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79, 655–673.

BOLLERSLEV, T. (1986): "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307–327.

BOLLERSLEV, T., AND R. F. ENGLE (1986): "Modelling the Persistence of Conditional Variances," *Econometric Reviews*, 5, 1–50.

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), 1509–1526.

BREEDEN, D. T., M. R. GIBBONS, AND R. H. LITZENBERGER (1989): "Empirical Test of the Consumption-Oriented CAPM," *Journal of Finance*, 44(2), 231–262.

BRYZGALOVA, S., J. HUANG, AND C. JULLIARD (2023): "Bayesian Solutions for the Factor Zoo: We Just Ran Two Quadrillion Models," *Journal of Finance*, 78(1), 487–557.

——— (2024): "Macro Strikes Back: Term Structure of Risk Premia and Market Segmentation," *Available at SSRN 4752696*.

BRYZGALOVA, S., A. QUAINI, F. TROJANI, AND Y. MING (2024): "Tradable Factor Risk Premia and Oracle Tests of Asset Pricing Models," *Available at SSRN 4752696*.

BURNSIDE, C. (2011): "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Comment," *American Economic Review*, 101(7), 3456–3476.

——— (2015): "Identification and Inference in Linear Stochastic Discount Factor Models with Excess Returns," *Journal of Financial Econometrics*, 14(2), 295–330.

CABALLERO, R., AND A. SIMSEK (2023): "A Monetary Policy Asset Pricing Model," NBER Working Paper 30132.

CAMPBELL, J. Y. (1993): "Intertemporal Asset Pricing without Consumption Data.," *American Economic Review*, 83(3).

——— (1999): "Asset Prices, Consumption, and the Business Cycle," *Handbook of Macroeconomics*, 1, 1231–1303.

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

Press, Princeton, New Jersey.

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.

CAMPBELL, J. Y., AND L. HENTSCHEL (1992): “No News Is Good News: An Asymmetric Model of Changing Volatility in Stock Returns,” *Journal of Financial Economics*, 31(3), 281–318.

CAMPBELL, J. Y., AND S. B. THOMPSON (2008): “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?,” *Review of Financial Studies*, 21(4), 1509–1531.

CHEN, A. Y. (2017): “External Habit in a Production Economy: A Model of Asset Prices and Consumption Volatility Risk,” *The Review of Financial Studies*, 30(8), 2890–2932.

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.

CHRISTIANO, L., M. EICHENBAUM, AND C. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113(1), 1–45.

COCHRANE, J. H. (2005): “Financial Markets and the Real Economy,” *Foundations and Trends® in Finance*, 1(1), 1–101.

COLACITO, R., AND M. M. CROCE (2011): “Risks for the Long Run and the Real Exchange Rate,” *Journal of Political Economy*, 119(1), 153–181.

COLLIN-DUFRESNE, P., M. JOHANNES, AND L. A. LOCHSTOER (2016): “Parameter Learning in General Equilibrium: The Asset Pricing Implications,” *American Economic Review*, 106(3), 664–698.

CONSTANTINIDES, G. M., AND A. GHOSH (2017): “Asset Pricing with Countercyclical Household Consumption Risk,” *The Journal of Finance*, 72(1), 415–460.

DEW-BECKER, I. (2017): “How Risky Is Consumption in the Long-Run? Benchmark Estimates from a Robust Estimator,” *Review of Financial Studies*, 30(2), 631–666.

DEW-BECKER, I., AND S. GIGLIO (2016): “Asset Pricing in the Frequency Domain: Theory and Empirics,” *Review of Financial Studies*, 29, 2029–2068.

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.

FRENCH, K. R., G. W. SCHWERT, AND R. F. STAMBAUGH (1987): “Expected Stock Returns and Volatility,” *Journal of Financial Economics*, 19(1), 3–29.

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.

GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2021): “Economic Predictions with Big Data: The Illusion of Sparsity,” *Econometrica*, 89(5), 2409–2437.

GIGLIO, S., AND D. XIU (2021): “Asset Pricing with Omitted Factors,” *Journal of Political Economy*, 129(7), 1947–1990.

GIGLIO, S., D. XIU, AND D. ZHANG (2021): “Test Assets and Weak Factors,” working paper, available at SSRN: <https://ssrn.com/abstract=3768081>.

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), 971–987.

HANSEN, L., AND T. J. SARGENT (2001): “Robust Control and Model Uncertainty,” *American Economic*

Review, 91(2), 60–66.

HANSEN, L. P., J. HEATON, J. LEE, AND N. ROUSSANOV (2007): “Intertemporal Substitution and Risk Aversion,” *Handbook of Econometrics*, 6, 3967–4056.

HANSEN, L. P., AND K. J. SINGLETON (1983): “Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns,” *Journal of Political Economy*, 91(2), 249–265.

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

HEATON, J. (1995): “An empirical investigation of asset pricing with temporally dependent preference specifications,” *Econometrica: Journal of the Econometric Society*, pp. 681–717.

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.

JOHANNES, M., L. A. LOCHSTOER, AND Y. MOU (2016): “Learning about Consumption Dynamics,” *The Journal of Finance*, 71(2), 551–600.

JULLIARD, C., AND A. GHOSH (2012): “Can Rare Events Explain the Equity Premium Puzzle?,” *Review of Financial Studies*, 25(10), 3037–3076.

JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105(3), 1177–1216.

KAN, R., AND C. ZHANG (1999): “Two-Pass Tests of Asset Pricing Models with Useless Factors,” *The Journal of Finance*, 54(1), 203–235.

KANDEL, S., AND R. F. STAMBAUGH (1990): “Expectations and Volatility of Consumption and Asset Returns,” *The Review of Financial Studies*, 3(2), 207–232.

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.

KLEIBERGEN, F., AND Z. ZHAN (2020): “Robust Inference for Consumption-Based Asset Pricing,” *The Journal of Finance*, 75(1), 507–550.

KOOP, G., M. PESARAN, AND S. M. POTTER (1996): “Impulse Response Analysis in Nonlinear Multivariate Models,” *Journal of Econometrics*, 74(1), 119–147.

KOZAK, S., S. NAGEL, AND S. SANTOSH (2018): “Interpreting Factor Models,” *Journal of Finance*, 73(3), 1183–1223.

——— (2020): “Shrinking the Cross-Section,” *Journal of Financial Economics*, 135(2), 271–292.

KROENCKE, T. A. (2017): “Asset Pricing without Garbage,” *Journal of Finance*, 72(2), 47–98.

LETTAU, M., AND S. LUDVIGSON (2001): “Consumption, Aggregate Wealth, and Expected Stock Returns,” *Journal of Finance*, 56(3), 815–849.

LI EW, 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 (2022): “Long-Run Risk: Is It There?,” *Journal of Finance*, 77(3), 1587–1633.

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, England.

LUDVIGSON, S. C. (2013): “Advances in Consumption-Based Asset Pricing: Empirical Tests,” *Handbook of the Economics of Finance*, 2, 799–906.

LUDVIGSON, S. C., S. MA, AND S. NG (2021): “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?,” *American Economic Journal: Macroeconomics*, 13(4), 369–410.

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.

MARIANO, R. S., AND Y. MURASAWA (2003): “A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series,” *Journal of Applied Econometrics*, 18(4), 427–443.

MEHRA, R., AND E. C. PRESCOTT (1985): “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 15(1), 225–250.

Economics, 15(2), 145–161.

NELSON, D. B. (1991): “Conditional Heteroskedasticity in Asset Returns: A New Approach,” *Econometrica: Journal of the Econometric Society*, pp. 347–370.

NEWHEY, W. K., AND K. D. WEST (1987): “A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.

OMORI, Y., S. CHIB, N. SHEPHARD, AND J. NAKAJIMA (2007): “Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference,” *Journal of Econometrics*, 140(2), 425–449.

ORTU, F., A. TAMONI, AND C. TEBALDI (2013): “Long-Run Risk and the Persistence of Consumption Shocks,” *The Review of Financial Studies*, 26(11), 2876–2915.

PARKER, J. A. (2001): “The Consumption Risk of the Stock Market,” *Brookings Papers on Economic Activity*, 2001(2), 279–333.

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.

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

SCHERVISH, M. J. (1996): *Theory of Statistics*, Springer Series in Statistics. Springer, New York.

SCHORFHEIDE, F., D. SONG, AND A. YARON (2018): “Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach,” *Econometrica*, 86, 617–654.

SHANKEN, J. (1992): “On the Estimation of Beta-Pricing Models,” *The Review of Financial Studies*, 5, 1–33.

SIMS, C. A., AND T. ZHA (2006): “Were There Regime Switches in U.S. Monetary Policy?,” *American Economic Review*, 96(1), 54–81.

TÉDONGAP, R. (2014): “Consumption Volatility and the Cross-Section of Stock Returns,” *Review of Finance*, 19(1), 367–405.

TRIPPLETT, J. E. (1997): “Measuring Consumption: The Post-1973 Slowdown and The Research Issues,” *Review, Federal Reserve Bank of Saint Louis*, 79, 9–42.

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.

VASICEK, O. (1977): “An Equilibrium Characterization of the Term Structure,” *Journal of Financial Economics*, 5, 177–188.

WACHTER, J. (2013): “Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?,” *Journal of Finance*, 180, 987–1035.

WEIL, P. (1989): “The Equity Premium Puzzle and the Risk-Free Rate Puzzle,” *Journal of Monetary Economics*, 24, 401–421.


WELCH, I., AND A. GOYAL (2008): “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction,” *Review of Financial Studies*, 21(4), 1455–1508.

WILCOX, D. W. (1992): “The Construction of US Consumption Data: Some Facts and Their Implications for Empirical Work,” *American Economic Review*, pp. 922–941.

YU, J. (2005): “On Leverage in a Stochastic Volatility Model,” *Journal of Econometrics*, 127(2), 165–178.

ZVIADADZE, I. (2021): “Term Structure of Risk in Expected Returns,” *Review of Financial Studies*, 34, 6030–6086.

Figure A1: Robustness check: $\bar{S} = 20$.



The figure plots (a) the cumulative IRF of consumption growth ($\bar{S} = 20$), (b) the MA component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$ (adjusted R^2), and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$ (adjusted R^2). We consider a single-factor model in equations (2)–(3), with $\bar{S} = 20$. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Red triangles in Panel (a) denote the cumulated IRF of the long-run risk model (Bansal and Yaron (2004)) in the monthly calibration of Hansen, Heaton, Lee, and Roussanov (2007) aggregated to quarterly frequency.

Internet Appendix for Consumption in Asset Returns

Svetlana Bryzgalova,^a Jiantao Huang,^b and Christian Julliard^{c,*}

^a*London Business School*

^b*University of Hong Kong*

^c*London School of Economics, FMG, SRC, and CEPR*

Abstract

The Internet Appendix provides additional tables, figures, and results supporting the main text.

*Email addresses: sbryzgalova@london.edu (S. Bryzgalova), huangjt@hku.hk (J. Huang), and c.julliard@lse.ac.uk (C. Julliard).

A Simulation design

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

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

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

We use this particular long-run risk (LRR) calibration because it ensures the nonnegativity of the volatility process and, most importantly, it 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% at the quarterly frequency, 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

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

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 selected to be their sample estimates in the main text.² After simulating monthly sequences of $(\Delta c_{t,t+1}, \mathbf{r}_{t+1}^e)$, we aggregate them into quarterly sequences by summing all three monthly observations within a quarter. We simulate 1,000 sample paths with 214 quarterly observations. Because we approximate the log consumption growth using an MA(14) process, there are 200 effective quarterly observations.

²More precisely, we obtain the estimates of $\boldsymbol{\mu}_r$, $\boldsymbol{\rho}^r$, and $\boldsymbol{\Sigma}_r$ using quarterly data. In the simulation of monthly returns, we divide all three parameters above by 3.

B Additional methodological details

B.1 State-space estimation with constant volatilities

We can rewrite the dynamic model in equations (2) and (3) in state-space form, assuming Gaussian innovations with constant volatility (the general SV version is presented in Appendix B.2), as

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

$$\mathbf{y}_t = \boldsymbol{\mu} + \mathbf{H}\mathbf{z}_t + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}_{N+1}, \Sigma), \quad (\text{IA.2})$$

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

$$\Psi := \underbrace{\begin{bmatrix} 1 & \mathbf{0}'_{\bar{S}} & \mathbf{0}'_{K-1} \\ \mathbf{0}_S & \mathbf{0}_{S \times \bar{S}} & \mathbf{0}_{S \times (K-1)} \\ \mathbf{0}_{K-1} & \mathbf{0}_{(K-1) \times \bar{S}} & \mathbf{I}_{K-1} \end{bmatrix}}_{(\bar{S}+K) \times (\bar{S}+K)}, \quad \mathbf{F} := \underbrace{\begin{bmatrix} \mathbf{0}'_{\bar{S}} & \mathbf{0}'_K \\ \mathbf{I}_{\bar{S}} & \mathbf{0}_{\bar{S} \times K} \\ \mathbf{0}_{(K-1) \times \bar{S}} & \mathbf{0}_{(K-1) \times K} \end{bmatrix}}_{(\bar{S}+K) \times (\bar{S}+K)}, \quad \Sigma := \underbrace{\begin{bmatrix} \sigma_c^2 & \mathbf{0}'_N \\ \mathbf{0}_N & \Sigma_r \end{bmatrix}}_{(N+1) \times (N+1)}, \quad \mathbf{H} := \underbrace{\begin{bmatrix} \rho_0 & \rho_1 & \dots & \rho_S & \mathbf{0}'_{K-1} \\ \rho_f^r & \mathbf{0}_N & \dots & \mathbf{0}_N & \rho_g^r \end{bmatrix}}_{(N+1) \times (\bar{S}+K)}, \quad (\text{IA.3})$$

and $\mathbf{I}_{\bar{S}}$ 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 (5)–(6) can be represented in the state-space form (IA.1)–(IA.2) with $\mathbf{z}_t := [f_t, \dots, f_{t-\bar{S}}, b_t, \dots, b_{t-\bar{S}}]', \mathbf{v}_t := [f_t, \mathbf{0}_{\bar{S}}', b_t, \mathbf{0}_{\bar{S}}']' \sim \mathcal{N}(\mathbf{0}; \Psi); \Psi$ and \mathbf{F} being block diagonal with blocks repeated twice and given by the first two matrices in equation (IA.3) (when $K = 1$); and space equation matrix of coefficients given by

$$\mathbf{H} := \underbrace{\begin{bmatrix} \rho_0 & \dots & \dots & \rho_{\bar{S}} & \theta_0 & \dots & \dots & \theta_{\bar{S}} \\ \rho_1^r & 0 & \dots & 0 & \theta_1^b & 0 & \dots & 0 \\ \dots & \dots \\ \rho_{N_b}^r & 0 & \dots & 0 & \theta_{N_b}^b & 0 & \dots & 0 \\ \dots & 0 & \dots & 0 & 0 & 0 & \dots & 0 \\ \rho_N^r & 0 & \dots & 0 & 0 & 0 & \dots & 0 \end{bmatrix}}_{(N+1) \times 2(\bar{S}+1)}. \quad (\text{IA.4})$$

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

$$\mathbf{y}_t | \mathcal{I}_{t-1}, \boldsymbol{\mu}, \mathbf{H}, \Psi, \Sigma, \mathbf{z}_t \sim \mathcal{N}(\boldsymbol{\mu} + \mathbf{H}\mathbf{z}_t; \Sigma), \quad (\text{IA.5})$$

where \mathcal{I}_{t-1} denotes the history of the state variables until time $t - 1$. Hence, under a diffuse (Jeffreys') prior, 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

of the unobservable \mathbf{z}_t , conditional on the data and the parameters, is constructed using a standard Kalman filter and smoother approach (see, e.g., Primiceri (2005)).

Let $\boldsymbol{\Pi}' := [\boldsymbol{\mu}, \mathbf{H}]$ and $\mathbf{x}'_t := [1, \mathbf{z}'_t]$. Under a diffuse prior, the likelihood of the data in equation (IA.5) implies the posterior distribution

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

where \mathbf{x} contains the stacked regressors, and the posterior distribution of each element on the main diagonal of $\boldsymbol{\Sigma}$ is given by³

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

where m_j is the number of estimated coefficients in the j -th equation. That is, the conditional posterior has a Normal-Inverse-Gamma structure. Moreover, \mathbf{F} and $\boldsymbol{\Psi}$ have a Dirac posterior distribution at the points defined in equation (IA.3). Therefore, the missing part necessary for taking draws via MCMC using a Gibbs sampler is the conditional distributions of \mathbf{z}_t . Because

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

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

B.1.1 Kalman filter algorithm

Let $\mathbf{z}_{t|\tau} := E[\mathbf{z}_t | \mathbf{y}^\tau, \mathbf{H}, \boldsymbol{\Psi}, \boldsymbol{\Sigma}]$ and $\mathbf{V}_{t|\tau} := \text{Var}(\mathbf{z}_t | \mathbf{y}^\tau, \mathbf{H}, \boldsymbol{\Psi}, \boldsymbol{\Sigma})$, 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: $\mathbf{z}_{t|t-1} = \mathbf{F} \mathbf{z}'_{t-1|t-1}$; $\mathbf{V}_{t|t-1} = \mathbf{F} \mathbf{V}_{t-1|t-1} \mathbf{F}' + \boldsymbol{\Psi}$; $\mathbf{K}_t = \mathbf{V}_{t|t-1} \mathbf{H}' (\mathbf{H} \mathbf{V}_{t|t-1} \mathbf{H}' + \boldsymbol{\Sigma})^{-1}$; $\mathbf{z}_{t|t} = \mathbf{z}_{t|t-1} + \mathbf{K}_t (\mathbf{y}_t - \boldsymbol{\mu} - \mathbf{H} \mathbf{z}_{t|t-1})$; and $\mathbf{V}_{t|t} = \mathbf{V}_{t|t-1} - \mathbf{K}_t \mathbf{H} \mathbf{V}_{t|t-1}$. 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 formulas for the following generic time- t recursion: $\mathbf{z}_{t|t+1} = \mathbf{z}_{t|t} + \mathbf{V}_{t|t} \mathbf{F}' \mathbf{V}_{t+1|t}^{-1} (\mathbf{z}_{t+1} - \mathbf{F} \mathbf{z}_{t|t})$ and $\mathbf{V}_{t|t+1} = \mathbf{V}_{t|t} - \mathbf{V}_{t|t} \mathbf{F}' \mathbf{V}_{t+1|t}^{-1} \mathbf{F} \mathbf{V}_{t|t}$.

Hence, parameters and states can be drawn via the Gibbs sampler, as follows:

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

³By relaxing the diagonality assumption, the posterior distribution of $\boldsymbol{\Sigma}^{-1}$ becomes a Wishart centered at the OLS estimates.

and use it to construct initial draws for $\boldsymbol{\mu}$ and \mathbf{H} . Using \mathbf{F} and $\boldsymbol{\Psi}$, 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 in the previous step) and $\{\mathbf{y}_t\}_{t=1}^T$, run *OLS* imposing the zero restrictions and get $\hat{\boldsymbol{\Pi}}'_{OLS}$ and $\hat{\boldsymbol{\Sigma}}_{OLS}$, and draw $\tilde{\boldsymbol{\Pi}}'$ and $\tilde{\boldsymbol{\Sigma}}^{-1}$ from the Normal-Inverse-Gamma (N-i- Γ step). Use these draws as the initial guess for the previous point of the algorithm, and repeat.

B.2 A generalized model with stochastic volatilities

To complete the specification in equations (2) and (16)–(19), 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.

B.2.1 Stochastic volatility of w_t^c

Let $y_{ct}^* = \log((w_t^c)^2) = h_{ct} + \log(\epsilon_{ct}^2)$ and $h_{ct} = \psi_c + \delta_c(h_{c,t-1} - \psi_c) + \sigma_{\eta c} \eta_{ct}$, 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, $e_{ct} = \sum_{j=1}^7 q_j \mathcal{N}(m_j - 1.2704, v_j^2)$, or, equivalently, $e_{ct} | S_{ct} = j \sim \mathcal{N}(m_j - 1.2704, v_j^2)$, where $\sum_{j=1}^7 q_j = 1$, and values of $\{m_j\}_{j=1}^7$ and $\{v_j\}_{j=1}^7$ can be found in Table 4 of Kim, Shephard, and Chib (1998).

Following is the algorithm used to estimate the unknown parameters $\{h_{ct}, S_{ct}, \psi_c, \delta_c, \sigma_{\eta c}\}$:

Step 1: Initialize $\{S_{ct}\}_{t=1}^T, \psi_c, \delta_c, \sigma_{\eta c}$.

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

Step 3: Sample S_{ct} from $p(S_{ct} | y_{ct}^*, h_{ct})$, where

$$p(S_{ct} = j | y_{ct}^*, h_{ct}) \propto q_j \times p(y_{ct}^* | h_{ct}, S_{ct} = j)$$

$$\propto q_j \times f_N(y_{ct}^* | h_{ct} + m_j - 1.2704, v_j^2).$$

Step 4: Update $\psi_c, \delta_c, \sigma_{\eta c}$. An inverse-gamma prior for $\sigma_{\eta c}^2$, that is, $\sigma_{\eta c}^2 \sim \Gamma^{-1}(\frac{\sigma_0}{2}, \frac{s_\sigma}{2})$, yields

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

Therefore, the conditional posterior is

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

Let $\delta_c = 2\phi - 1$, and $\phi \sim Beta(\phi^{(1)}, \phi^{(2)})$. 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

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

The function above 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

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

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 , as follows:

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

Therefore, we obtain the posterior distribution of ψ_c : $\psi_c | h_c, \delta_c, \sigma_{\eta c}^2 \sim \mathcal{N}(\hat{\psi}_c, \sigma_\psi^2)$, 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: Return to Step 2 until convergence.

B.2.2 Stochastic volatility of f_t

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

$$y_{ft}^* = \log(f_t^2) = h_{ft} + \log(\epsilon_{ft}^2), \quad h_{ft} = \delta_f h_{f,t-1} + \sqrt{1 - \delta_f^2} \eta_{ft},$$

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 $e_{ft} = \log(\epsilon_{ft}^2) \sim \log(\chi^2(1))$. The process above for h_{ft} is simply an AR(1) normalized to have unconditional zero mean and unit variance. As previously, we approximate e_{ft} using a mixture of Gaussians; 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 $p(S_{ft} \mid y_f^*, h_{ft})$.

Step 4: Update δ_f using the Metropolis algorithm. Similar to the prior distribution of δ_c , we assume that $\delta_f = 2\phi - 1$, 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 the Metropolis algorithm: (M1) initialize δ_f^0 ; (M2) draw δ_f^* from a normal distribution $\mathcal{N}(\delta_f^{(i-1)}, c_{mh}^2)$ ⁴; (M3) calculate $\rho(\delta_f^*, \delta_f^{(i-1)}) = \min\{1, \frac{p(\delta_f^* \mid h_f)}{p(\delta_f^{(i-1)} \mid h_f)}\}$; (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: Return to Step 2 until convergence.

B.2.3 Stochastic volatility of asset returns

Let $\mathbf{y}_{rt}^* = \boldsymbol{\kappa}_0 + \boldsymbol{\kappa}_1 h_{rt} + \mathbf{e}_{rt}$ and $h_{rt} = \delta_r h_{r,t-1} + \sqrt{1 - \delta_r^2} \eta_{rt}$, where

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

and the ϵ_{it} shocks are independent across different assets. In the model above, 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 h_{rt} to have zero mean and unit variance. In

⁴ c_{mh} determines the step size in the Metropolis algorithm. We aim to choose c_{mh} such that the frequency of accepting a new δ_f is about 50%.

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

$$\kappa_0 = \mathbf{y}_{rt}^* - \bar{\mathbf{y}}_{rt}^*, \quad \bar{\mathbf{y}}_{rt}^* = \kappa_1 h_{rt} + \mathbf{e}_{rt}, \quad h_{rt} = \delta_r h_{r,t-1} + \sqrt{1 - \delta_r^2} \eta_{rt},$$

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, $\bar{\mathbf{y}}_{i,rt}^* \mid \kappa_1^{(i)}, h_{rt}, S_{i,rt} = j \stackrel{\text{iid}}{\sim} \mathcal{N}(\kappa_1^{(i)} h_{rt} + m_j - 1.2704, v_j^2)$. 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 diffuse prior for $\kappa_1^{(i)}$, we have

$$\begin{aligned} p(\kappa_1^{(i)} \mid \bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)}, \mathbf{H}_{\mathbf{r}}, \{S_{i,rt}\}_{t=1}^T) &\propto p(\bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)} \mid \kappa_1^{(i)}, \mathbf{H}_{\mathbf{r}}, \{S_{i,rt}\}_{t=1}^T) \\ &\propto \exp\left\{-\frac{1}{2}(\bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)} - \mathbf{H}_{\mathbf{r}} \kappa_1^{(i)})^\top (\mathbf{V}_{\mathbf{r}}^{(i)})^{-1} (\bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)} - \mathbf{H}_{\mathbf{r}} \kappa_1^{(i)})\right\}, \text{ and} \\ \kappa_1^{(i)} \mid \bar{\mathbf{Y}}_{\mathbf{r}}^{*,(i)}, \mathbf{H}_{\mathbf{r}}, \{S_{i,rt}\}_{t=1}^T &\sim \mathcal{N}(\hat{\kappa}_1^{(i)}, [\mathbf{H}_{\mathbf{r}}^\top (\mathbf{V}_{\mathbf{r}}^{(i)})^{-1} \mathbf{H}_{\mathbf{r}}]^{-1}), \end{aligned}$$

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 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 previously.

B.2.4 The consumption growth excess return mean-equation coefficients

Because shocks in equations (2) and (16) are uncorrelated, we can sample the unknown parameters equation by equation. We introduce the following notations:

$$\begin{aligned} \boldsymbol{\rho}^c &= \begin{pmatrix} \mu_c \\ \rho_0 \\ \vdots \\ \rho_{\bar{S}} \end{pmatrix}, \quad \Delta \mathbf{C} = \begin{pmatrix} \Delta C_{0,1} \\ \Delta C_{1,2} \\ \vdots \\ \Delta C_{T-1,T} \end{pmatrix}, \quad \boldsymbol{\Sigma}_w^c = \begin{pmatrix} \sigma_{c,0}^2 & 0 & \dots & 0 \\ 0 & \sigma_{c,1}^2 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & \sigma_{c,T-1}^2 \end{pmatrix}, \quad \mathbf{X}^c = \begin{pmatrix} 1 & f_1 & \dots & f_{1-\bar{S}} \\ 1 & f_2 & \dots & f_{2-\bar{S}} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & f_T & \dots & f_{T-\bar{S}} \end{pmatrix}, \\ \boldsymbol{\rho}_i^r &= \begin{pmatrix} \mu_{ri} \\ \rho_i^r \\ \beta_{ci} \\ \beta_{fi} \\ \beta_{ri} \end{pmatrix}, \quad \mathbf{X}_i^r = \begin{pmatrix} 1 & f_1 & \sigma_{c,0}^2 & \sigma_{f,0}^2 & \sigma_{r,0}^2 \\ 1 & f_2 & \sigma_{c,1}^2 & \sigma_{f,1}^2 & \sigma_{r,1}^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & f_T & \sigma_{c,T-1}^2 & \sigma_{f,T-1}^2 & \sigma_{r,T-1}^2 \end{pmatrix}, \text{ and} \quad \boldsymbol{\Sigma}_{wr,i} = \begin{pmatrix} \sigma_{r,i,0}^2 & 0 & \dots & 0 \\ 0 & \sigma_{r,i,1}^2 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & \sigma_{r,i,T-1}^2 \end{pmatrix}. \end{aligned}$$

Then the posterior distribution of ρ^c under a flat prior is as follows:⁵

$$\rho^c | \Delta C, \mathbf{X}^c, \Sigma_w^c \sim \mathcal{N}(\hat{\rho}_{gls}, [(\mathbf{X}^c)^\top (\Sigma_w^c)^{-1} \mathbf{X}^c]^{-1}), \quad \hat{\rho}_{gls} = [(\mathbf{X}^c)^\top (\Sigma_w^c)^{-1} \mathbf{X}^c]^{-1} (\mathbf{X}^c)^\top (\Sigma_w^c)^{-1} \Delta C.$$

The conditional posterior for the return equation coefficients under a flat prior is then

$$\rho_i^r | \mathbf{R}_i, \mathbf{X}_i^r, \{\sigma_{ft}^2\}_{t=0}^{T-1}, \{\sigma_{r,i,t}^2\}_{t=0}^{T-1} \sim \mathcal{N}(\hat{\rho}_{i,gls}^r, [(\mathbf{X}_i^r)^\top \Sigma_{wr,i}^{-1} \mathbf{X}_i^r]^{-1}), \text{ and}$$

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

B.2.5 Drawing the conditional consumption mean shocks

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

$$\mathbf{y}_t = \begin{pmatrix} \Delta c_{t-1,t} \\ \mathbf{r}_t^e \end{pmatrix}, \quad \tilde{\mu}_t = \begin{pmatrix} \mu_c \\ \mu_r + \beta_c \sigma_{c,t-1}^2 + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 \end{pmatrix}, \quad \mathbf{H} = \begin{pmatrix} \rho_0 & \rho_1 & \dots & \rho_{\bar{S}} \\ \mathbf{0}_N & \mathbf{0}_N & \dots & \mathbf{0}_N \end{pmatrix},$$

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

Hence, the joint distribution of observables and f shocks is

$$\begin{pmatrix} \mathbf{y}_t \\ \mathbf{z}_t \end{pmatrix} | \mathcal{I}_{t-1}, \tilde{\mu}_t, \mathbf{H}, \Sigma_t, \sigma_{f,t-1}^2 \sim \mathcal{N} \left(\begin{pmatrix} \tilde{\mu}_t \\ \mathbf{F} \mathbf{z}_{t-1} \end{pmatrix}, \begin{pmatrix} \mathbf{\Omega}_t & \mathbf{H} \\ \mathbf{H}^\top & \mathbf{\Psi}_t \end{pmatrix} \right), \quad (\text{IA.6})$$

where $\mathbf{\Psi}_t = \begin{pmatrix} \sigma_{f,t-1}^2 & \mathbf{0}_{\bar{S}}^\top \\ \mathbf{0}_{\bar{S}} & \mathbf{0}_{\bar{S} \times \bar{S}} \end{pmatrix}$, $\mathbf{F} := \begin{bmatrix} \mathbf{0}'_{\bar{S}} & 0 \\ I_{\bar{S}} & \mathbf{0}_{\bar{S}} \end{bmatrix}$, and $\mathbf{\Omega}_t = \mathbf{H} \mathbf{\Psi}_t \mathbf{H}^\top + \Sigma_t$. We then use the Kalman smoother to draw f_t , following the same procedure as in Internet Appendix B.1.1.

B.3 Model comparison

The model comparison for the restricted and unrestricted specification in Table IA.I of Section I is performed using Bayes factors (i.e., the the marginal likelihoods of the various models) and posterior probabilities. Based on equation (IA.6), the likelihood function of the

⁵The likelihood function of the data is $\Delta C | \rho^c, \mathbf{X}^c, \Sigma_w^c \sim \mathcal{N}(\mathbf{X}^c \rho^c, \Sigma_w^c)$. By applying the diffuse prior for ρ^c , the posterior distribution of ρ^c becomes

$$p(\rho^c | \Delta C, \mathbf{X}^c, \Sigma_w^c) \propto p(\Delta C | \rho^c, \mathbf{X}^c, \Sigma_w^c) \propto \exp\left\{-\frac{1}{2}(\rho^c - \hat{\rho}_{gls})^\top (\mathbf{X}^c)^\top (\Sigma_w^c)^{-1} \mathbf{X}^c (\rho^c - \hat{\rho}_{gls})\right\}.$$

observed data $\{\mathbf{y}_t\}_{t=1}^T$ (after integrating out f_t) is

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

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 $(\boldsymbol{\mu}, \mathbf{H}, \boldsymbol{\beta}_f, \boldsymbol{\beta}_r, \boldsymbol{\beta}_c)$ 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 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

$$\log(BF_{1,2}) \approx \log p(\mathbf{Y} \mid \hat{\theta}(\mathbb{M}_1)) - \log p(\mathbf{Y} \mid \hat{\theta}(\mathbb{M}_2)) - \frac{d_1 - d_2}{2} \log(T), \quad (\text{IA.7})$$

where \mathbf{Y} is the observed data, T is the sample size, \mathbb{M}_1 and \mathbb{M}_2 represent Models 1 and 2, $\hat{\theta}(\mathbb{M}_1)$ and $\hat{\theta}(\mathbb{M}_2)$ are the posterior mean of parameter θ under Models 1 and 2, and d_1 and d_2 are model dimensions.⁶ The Bayes factor in equation (IA.7) 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 (IA.7)) or the BIC-based model selection.

Posterior model probabilities are then computed using the approximation above 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.

⁶Note that the vector of parameters encompasses both the frequentist parameters $(\mu_c, \boldsymbol{\mu}_r, \boldsymbol{\beta}_f, \boldsymbol{\beta}_r, \boldsymbol{\beta}_c, \boldsymbol{\rho}^r, \psi_c, \delta_c, \sigma_{\eta c}, \delta_f, \boldsymbol{\kappa}_0, \boldsymbol{\kappa}_0, \delta_r)$ and the latent states $(\{f_t, \sigma_{ct}^2, \sigma_{ft}^2, \sigma_{rt}^2\}_{t=1}^T)$.

C Data description

Bond holding returns are calculated on a quarterly basis using the zero coupon yield data constructed by Gurkaynak, Sack, and Wright (2007),⁷ which fit the Nelson-Siegel-Svensson curves daily since June 1961, and excess returns are computed by subtracting the return on a three-month Treasury bill. We consider the following maturities: six months, one, two, three, four, five, six, seven, and 10 years, which gives us a set of nine bond portfolios.

We consider several portfolios of stock returns. In addition to the bond portfolios, the baseline specification relies on the 25 size and book-to-market Fama-French portfolios (Fama and French (1992)) and 12 industry portfolios. In robustness checks, we consider four additional sets of characteristic-sorted portfolios, including (1) 32 size-profitability-investment-sorted portfolios, (2) 32 size-value-investment-sorted portfolios, (3) 32 size-value-profitability-sorted portfolios, and (4) 74 decile 1 and 10 portfolios from Kozak, Nagel, and Santosh (2018). These portfolios are sorted by 37 firm characteristics that have data since July 1963.⁸ The stock portfolio data of the first three cross-sections are available from the Kenneth French data library, while we obtain the data of 74 decile 1 and 10 portfolios from Serhiy Kozak’s website. We consider monthly returns from July 1963 to December 2019 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 the real (chain-weighted) expenditure on nondurable goods per capita available from the National Income and Product Accounts (NIPA). As in, for example, Parker and Julliard (2005), we do not include services in our baseline definition of consumption because these are likely to mechanically bias both the persistence of the consumption proxy (due to, e.g., utilities and health care) and the comovement with market returns (due to, e.g., financial services and insurance). Furthermore, nondurable consumption seems less affected by interpolation and other measurement issues (see, e.g., Savov (2011) and Kroencke (2017)). Our results are robust to the usage of alternative measures and refinements of the consumption proxy (e.g., the exclusion of shoes and clothing,

⁷The data are regularly updated and available at <http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html>.

⁸The author provides the data of portfolios sorted by more than 50 firm characteristics. However, some characteristic-sorted portfolios do not have available data before the 1970s; hence, we use 37 characteristics to ensure that the time-series sample is consistent with the baseline analysis.

as in Lettau and Ludvigson (2001, 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 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.

In Section III.1.2, we conduct horse races with the canonical predictors to explore whether the conditional consumption mean process uncovered by our state-space formulation indeed captures the consumption predictability. In particular, the real GDP growth comes from FRED St. Louis. The SPF forecasts are downloaded from the Federal Reserve Bank of Philadelphia. The news indices in Liu and Matthies (2022) are obtained from the authors' websites. Shareholders' and nonshareholders' consumption growth are downloaded from Tobias Moskowitz's website.

In some analyses, we include the market's price-dividend (P/D) ratio and the Chicago Fed National Activity Index (CFNAI) to predict stock and bond portfolio returns, respectively. We download the CFNAI data from the Federal Reserve Bank of Chicago, and the sample begins in January 1967.

We estimate the P/D ratio by comparing the gross return with the ex-dividends return. Suppose we have a cross-section of test assets, $i = 1, 2, \dots, N$, and D_{it} and P_{it} are the dividends and market prices of portfolio i between time $t - 1$ and t . The total and ex-dividend portfolio returns are defined as follows:

$$R_{it} = \frac{P_{it} + D_{it} - P_{i,t-1}}{P_{i,t-1}}, \quad R_{it,-d} = \frac{P_{it} - P_{i,t-1}}{P_{i,t-1}}, \implies \frac{D_{it}}{P_{i,t-1}} = R_{it} - R_{it,-d}, \text{ and } P_{i,t-1} = \frac{P_{it}}{1 + R_{it,-d}}.$$

We define the market's P/D ratio as follows:

$$PD_t = \frac{\sum_{i=1}^N P_{it}}{\sum_{s=0}^3 \sum_{i=1}^N D_{i,t-s}},$$

which is a smooth version P/D ratio used in past research, such as Welch and Goyal (2008) and Campbell and Thompson (2008).

Finally, we use several uncertainty measures commonly studied in the previous literature. Specifically, the VIX volatility index is downloaded from CRSP (and has been publicly available since 1986:Q1). The financial/real/macro uncertainty measures from Jurado, Ludvigson, and Ng (2015) and Ludvigson, Ma, and Ng (2021) are downloaded from the authors'


websites.

D 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 and *b*) not only asset returns predict future levels of consumption growth, but they also do it better than the past values of consumption itself. A detailed description of the data is reported in Appendix C.

First, Figure IA.1 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 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 14 quarters (and significant at lower confidence levels at even longer horizons). That is, there is substantial persistence in the time-series of consumption growth.⁹

Figure IA.1: Autocorrelation structure of consumption growth.



Left panel: Autocorrelation function of consumption growth ($\Delta c_{t,t+1}$) with 90% and 95% confidence bands.
 Right panel: *p*-values of Ljung and Box (1978) (triangles) and Box and Pierce (1970) (circles) tests.


Second, we run multivariate linear predictive regressions of cumulated log consumption growth $\Delta c_{t,t+1+S}$ (for several values of S) on the first eight principal components of time t asset returns.¹⁰ Figure IA.2 depicts summary statistics for these predictive regressions at

⁹Note 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.

¹⁰We use the first eight principal components of the 25 size and book-to-market Fama-French portfolios, 12

different horizons (S). In particular, the left panel plots the time-series adjusted R^2 of these regressions, and the right panel plots the p -value of the F -test of joint significance of the regressors for this as well as other additional specifications.

Figure IA.2: Predictive regressions of $\Delta c_{t,t+1+S}$ on time t asset returns and consumption.



The figure shows the 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 (blue line with circles) and theoretical adjusted R^2 (yellow dashed line with triangles) if all the predictability was driven by the first period. The red dashed line with pluses stands for the adjusted R^2 when only $t - 1$ consumption growth is used as a predictor. The dotted line with rhombi 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%, 5%, and 1% significance thresholds.

Several observations are to be mentioned. 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 $S > 0$, because $\Delta c_{t,t+1+S} \equiv \Delta c_{t,t+1} + \Delta c_{t+1,t+1+S}$, 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 yellow dashed lines with triangles in the left panel of Figure IA.2. Instead, for $S > 0$, the figure shows no such decrease in the data (blue dashed 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.

Could one achieve the same level of predictability by using only 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? This is unlikely. The red crosses in Figure

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 those reported.

IA.2 depict the degree of predictability obtained using only lagged consumption growth, $\Delta c_{t-1,t}$. Although highly significant at horizons 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. Although non-classical errors could possibly contribute to a wide range of statistical artefacts, most of their impact should either disappear within a horizon of about one year (if it is 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)¹¹ (purple diamonds in Figure IA.2). If the predictability result is 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, because only yearly data are 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 is there substantial predictability in consumption growth, but it is also best captured by asset returns.

E Dew-Becker (2017) bounds

Dew-Becker (2017) provides bounds for the long-run standard deviation (LRSD) of consumption growth. Within our framework, this can be expressed as

$$LRSD \equiv \text{std}(\Delta \mathbb{E}_t \sum_{j=0}^{\infty} \Delta c_{t+j}) = \text{std}\left(\left(\sum_{j=0}^{\bar{S}} \rho_j\right) f_t + w_t^c\right) = \sqrt{\left(\sum_{j=0}^{\bar{S}} \rho_j\right)^2 + \sigma_c^2}.$$

In our baseline analysis using nondurable consumption growth, we find that the posterior mean of (annualized) LRSD is about 2.3%, with [1.6%, 3.1%] as 95% CI. Using instead nondurable plus service consumption, Dew-Becker (2017) estimates that the LRSD of con-

¹¹We are grateful to Tim Kroenke for making the data available on his website.

sumption growth in the postwar sample is 2.5% per year, with an upper bound being the 95% confidence interval of 4.9%. Using the same consumption measure, we estimate a posterior mean of 2.2% for the LRSD, with [1.6%, 2.9%] as 95% CI. Hence, the estimates of the LRSD implied by our state-space method are close to the one in Dew-Becker (2017) and within his bounds.

F Estimation with monthly consumption data

In this section, we consider the consumption dynamics at the monthly frequency. We download the real nondurable consumption data from Table 2.8.3 of BEA and the monthly US population from FRED St. Louis to calculate the real monthly nondurable consumption growth per capita.¹²


We first consider the estimation without stochastic volatility. Figure IA.3 plots the cumulative response function of monthly consumption growth to a one-standard-deviation f_t shock. Even at the monthly frequency without the concern of time-aggregation bias, we find that consumption growth slowly adjusts to the asset return shocks. The cumulative effects are substantial: They keep increasing even after 15 months, and they jointly explain about 11% of the time-series variation of (one-period) monthly consumption growth, as can be seen in Figure IA.4. Moreover, monthly consumption growth shows significantly negative autocorrelations (see the left panel of Figure IA.5), which may potentially originate from the mean-reverting measurement error in the monthly consumption data. Consequently, the null hypothesis of uncorrelated forecast errors in consumption growth is rejected in the data (see the right panel of Figure IA.5).

F.1 Stochastic Volatility at the monthly frequency

To study the stochastic volatilities of consumption at the monthly frequency, we need to explicitly model the mean-reverting measurement error; otherwise, we would obtain spurious volatility clustering (see Figure IA.5). We consider a simple IID measurement error in


¹²The monthly consumption data are used by previous literature, such as Hansen and Singleton (1983) and Heaton (1995). However, Breeden, Gibbons, and Litzenberger (1989) show several reasons for using quarterly consumption rather than monthly data, such as sampling error and infrequent reporting of consumption.

Figure IA.3: Cumulative response function of consumption growth to a one-standard-deviation shock spanned by asset returns: Monthly frequency.




The graph presents posterior means of the cumulative impulse response function of consumption growth (solid line with circles), along with the centered posterior 90% (dotted lines) and 68% (dashed lines) coverage regions. These shocks account for 11% of monthly consumption growth time-series variation (monthly data, 07/1963–12/2019).

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



The figure presents box plots (posterior 95% coverage area) of the percentage of time-series variances of consumption growth explained by the MA component. These plots report the unadjusted R-squared. Left panel: Cumulated consumption growth $\Delta c_{t-1,t+S}$. Right panel: One-period consumption growth $\Delta c_{t-1+j,t+j}$. We study a single-factor model of asset returns, with $\bar{S} = 36$. The cross-section of test assets includes 25 size- and value-sorted portfolios, 12 industry portfolios, and nine bond portfolios (monthly data, 07/1963–12/2019).

Figure IA.5: Autocorrelation structure of consumption growth & its forecast errors.



Autocorrelation function of $\Delta c_{t,t+1}$ (left panel) and its forecast error (right panel) with 95% and 99% confidence bands (monthly data, 07/1963–12/2019).

monthly log consumption level (as in Schorfheide, Song, and Yaron (2018)),

$$c_t = c_t^* + \sigma_{\epsilon c} \cdot \epsilon_{c,t}, \quad \text{where } \epsilon_{c,t} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1),$$

where c_t^* is the latent true log consumption level without measurement errors, and this implies the following consumption growth dynamics,

$$\Delta c_t = \mu_c + \underbrace{\sum_{j=0}^{\bar{S}} \rho_j f_{t-j}}_{MA(\bar{S})} + \underbrace{\sigma_{\epsilon c} \cdot (\epsilon_{c,t} - \epsilon_{c,t-1})}_{\text{measurement error}} + w_t^c, \quad w_t^c \sim \mathcal{N} \left(0, \exp(h_{ct}) \frac{\sigma_{\epsilon c}^2}{\sigma_{c,t-1}^2} \right), \quad (\text{IA.8})$$

where we assume that the measurement error is uncorrelated with the innovation to expected consumption growth, f_t , and the short-run consumption shock, w_t^c . The measurement error in monthly consumption growth, $\sigma_{\epsilon c} \cdot (\epsilon_{c,t} - \epsilon_{c,t-1})$, can explain the significantly negative one-lag autocorrelation coefficient in Figure IA.5. The return dynamics is the same as in equation (16).

To estimate the system with an additional measurement error, we need to expand the space of latent states. Specifically, we write down the following state-space model:

$$\mathbf{z}_t = \mathbf{F} \mathbf{z}_{t-1} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}_{\bar{S}+3}, \boldsymbol{\Psi}_t), \quad \text{and} \quad (\text{IA.9})$$

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{H} \mathbf{z}_t + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}_{N+1}, \boldsymbol{\Sigma}_t), \quad (\text{IA.10})$$

where $\mathbf{y}_t := [\Delta c_{t-1,t}, \mathbf{r}_t^{er}]$, $\mathbf{v}_t := [f_t, \mathbf{0}'_{\bar{S}}, \epsilon_{c,t}, 0]'$, $\mathbf{w}_t := [\tilde{w}_t^c, \mathbf{w}_t^{er}]'$,

$$\mathbf{z}_t = \begin{bmatrix} f_t \\ \vdots \\ f_{t-\bar{S}} \\ \epsilon_{c,t} \\ \epsilon_{c,t-1} \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \mathbf{0}_{\bar{S}}^\top & 0 & 0 & 0 \\ \mathbf{I}_{\bar{S}} & \mathbf{0}_{\bar{S}} & \mathbf{0}_{\bar{S}} & \mathbf{0}_{\bar{S}} \\ \mathbf{0}_{\bar{S}}^\top & 0 & 0 & 0 \\ \mathbf{0}_{\bar{S}}^\top & 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{\Psi}_t = \begin{bmatrix} \exp(h_{ft}) & \mathbf{0}_{\bar{S}}^\top & 0 & 0 \\ 0_{\bar{S}} & \mathbf{0}_{\bar{S} \times \bar{S}} & \mathbf{0}_{\bar{S}} & \mathbf{0}_{\bar{S}} \\ 0 & \mathbf{0}_{\bar{S}}^\top & 1 & 0 \\ 0 & \mathbf{0}_{\bar{S}}^\top & 0 & 0 \end{bmatrix},$$

$$\boldsymbol{\mu}_t = \begin{bmatrix} \mu_c \\ \mu_r + \beta_c \sigma_{c,t-1}^2 + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 \end{bmatrix}, \quad \mathbf{H} = \begin{bmatrix} \rho_0 & \dots & \rho_{\bar{S}} & \sigma_{ec} & -\sigma_{ec} \\ \boldsymbol{\rho}^r & \dots & \mathbf{0}_N & \mathbf{0}_N & \mathbf{0}_N \end{bmatrix}, \quad \boldsymbol{\Sigma}_t = \begin{bmatrix} \exp(h_{ct}) & & & \\ & \boldsymbol{\Sigma}_{r,t-1} & & \end{bmatrix}.$$

In the monthly data, the measurement error, quantified by σ_{ec} , is sizable. The posterior mean of σ_{ec} is 0.0036. Note that the volatility of the monthly consumption growth is 0.0072; hence, ignoring measurement errors would induce mechanical volatility clustering.


Results are reported in Figures IA.6 and IA.7. First, the short-run consumption volatility is still hard to detect. As Panel A of Figure IA.6 shows, there is an entire range of constant volatility levels for $\sigma_{c,t}^2$. Second, the long-run consumption volatility, $\sigma_{f,t}^2$, is sharply identified using the monthly data. Third, the common market volatility, $\sigma_{r,t}^2$, becomes much spikier at the monthly than quarterly frequency, especially in the early 2000s. Finally, neither short- nor long-run consumption volatility ($\sigma_{c,t}^2$ or $\sigma_{f,t}^2$) significantly predicts excess returns (see Panels A and B of Figure IA.7). In contrast, the common market volatility negatively predicts the next-month excess returns, although most of the coefficients are not significant.

G Mixed-frequency state-space estimation

In the main analysis of the quarterly consumption growth, we ignore the fact that consumption is measured over a time interval, whereas asset prices are observed at a higher frequency (e.g., monthly). In this section, we derive a mixed-frequency state-space model to handle the time-aggregation bias. Specifically, we assume that we observe only the quarterly consumption growth in the final month of a quarter but monthly portfolio returns. However, the estimation in this section does not rely on monthly consumption growth. The estimation based on monthly consumption data is reported in Section F.


Throughout this section, we use t to denote the month t . For instance, \mathbf{r}_t^e denote the excess returns in month t . Consumption data, C_t , are observed at the quarterly frequency, which equals the sum of the monthly consumption within the quarter: $C_t = C_t^* + C_{t-1}^* + C_{t-2}^*$, where C_t^* denotes the unobserved monthly consumption. Using the log-linearization (see,

Figure IA.6: Filtered stochastic volatilities of *monthly* consumption and returns.




The figure shows the estimated stochastic volatilities of the model in equations (16) and (IA.8). Solid blue lines depict the posterior median of the log volatility, whereas dotted red lines denote 2.5% and 97.5% credible intervals. Shaded (pattered) areas reflect constant volatility levels that would not be rejected given the credible intervals. Both consumption and asset returns are observed at the monthly frequency.


Figure IA.7: Loadings of excess returns on *monthly* consumption and returns volatilities.



Panel A: Posterior distribution of excess return loadings (β_c) on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$) in equation (16).



Panel B: 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 (16).



Panel C: Posterior distributions of excess return loadings (β_r) on the common financial return variance ($\sigma_{r,t-1}^2$) in equation (16).

The figure shows the 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 (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last. Both consumption and asset returns are observed at the monthly frequency.

for example, Schorfheide, Song, and Yaron (2018)), we obtain $\log C_t = \frac{1}{3}(\log C_t^* + \log C_{t-1}^* + \log C_{t-2}^*)$. Hence, the quarterly consumption growth, for all t , is

$$\Delta c_{t-3,t} = \log C_t - \log C_{t-3} = \frac{1}{3}\Delta c_t^* + \frac{2}{3}\Delta c_{t-1}^* + \Delta c_{t-2}^* + \frac{2}{3}\Delta c_{t-3}^* + \frac{1}{3}\Delta c_{t-4}^*, \quad (\text{IA.11})$$

where $\Delta c_t^* = \log C_t^* - \log C_{t-1}^*$ denotes the unobserved monthly consumption growth.

We assume the same data-generating process of asset returns as before,

$$\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}_f^r f_t + \mathbf{w}_t^r, \quad (\text{IA.12})$$

where the only difference is that we now observe \mathbf{r}_t^e at the monthly frequency. Moreover, monthly consumption growth can react slowly to the f_t shocks, as follows:

$$\Delta c_t^* = \mu_c^* + \underbrace{\sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + w_t^{c*}}_{MA(\bar{S})}, \quad (\text{IA.13})$$

which implies the following MA representation of quarterly consumption growth $\Delta c_{t-3,t}$:

$$\Delta c_{t-3,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j \left(\frac{1}{3} f_{t-j} + \frac{2}{3} f_{t-j-1} + f_{t-j-2} + \frac{2}{3} f_{t-j-3} + \frac{1}{3} f_{t-j-4} \right) + w_t^c, \text{ and} \quad (\text{IA.14})$$

$$\mu_c = 3\mu_c^*, \quad w_t^c = \frac{1}{3}w_t^{c*} + \frac{2}{3}w_{t-1}^{c*} + w_{t-2}^{c*} + \frac{2}{3}w_{t-3}^{c*} + \frac{1}{3}w_{t-4}^{c*}.$$

We take an agnostic view on the dynamics of w_t^{c*} ; hence, we do not impose a MA(4) process for w_t^c . It is worth noting that the distribution of w_t^c does not affect the consistency of μ_c and ρ^c , as long as it is a zero-mean process and uncorrelated with f_t .

We introduce the following matrix notations:

$$\mathbf{z}_t = \begin{bmatrix} f_t \\ \vdots \\ f_{t-\bar{S}} \\ f_{t-\bar{S}-1} \\ f_{t-\bar{S}-2} \\ f_{t-\bar{S}-3} \\ f_{t-\bar{S}-4} \end{bmatrix}, \quad \boldsymbol{\rho}_0^c = \begin{bmatrix} \rho_0 \\ \vdots \\ \rho_{\bar{S}} \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \boldsymbol{\rho}_1^c = \begin{bmatrix} 0 \\ \vdots \\ \rho_{\bar{S}-1} \\ \rho_{\bar{S}} \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \boldsymbol{\rho}_2^c = \begin{bmatrix} 0 \\ \vdots \\ \rho_{\bar{S}-2} \\ \rho_{\bar{S}-1} \\ \rho_{\bar{S}} \\ 0 \\ 0 \end{bmatrix}, \quad \boldsymbol{\rho}_3^c = \begin{bmatrix} 0 \\ \vdots \\ \rho_{\bar{S}-3} \\ \rho_{\bar{S}-2} \\ \rho_{\bar{S}-1} \\ \rho_{\bar{S}} \\ 0 \end{bmatrix}, \quad \boldsymbol{\rho}_4^c = \begin{bmatrix} 0 \\ \vdots \\ \rho_{\bar{S}-4} \\ \rho_{\bar{S}-3} \\ \rho_{\bar{S}-2} \\ \rho_{\bar{S}-1} \\ \rho_{\bar{S}} \end{bmatrix},$$

which imply the following dynamics for quarterly consumption growth:

$$\Delta c_{t-3,t} = \mu_c + (\boldsymbol{\rho}^c)^\top \mathbf{z}_t + w_t^c, \quad \text{where: } \boldsymbol{\rho}^c = \frac{1}{3}\boldsymbol{\rho}_0^c + \frac{2}{3}\boldsymbol{\rho}_1^c + \boldsymbol{\rho}_2^c + \frac{2}{3}\boldsymbol{\rho}_3^c + \frac{1}{3}\boldsymbol{\rho}_4^c. \quad (\text{IA.15})$$

Note that we observe $\Delta c_{t-3,t}$ only once per quarter. To handle the missing data issue,

we follow Mariano and Murasawa (2003) and define $\Delta c_{t-3,t}^+$ as follows:

$$\Delta c_{t-3,t}^+ = \begin{cases} \Delta c_{t-3,t}, & t \text{ is the last month of the quarter,} \\ d_t, & \text{otherwise,} \end{cases}$$

where d_t is a random draw from a distribution that does not depend on the model parameters (denoted by $\boldsymbol{\theta}$). As in Mariano and Murasawa (2003), we assume $d_t \sim \mathcal{N}(0, 1)$, yielding the likelihood

$$p(\{\Delta c_{t-3,t}^+, \mathbf{r}_t^e\}_{t=1}^T \mid \boldsymbol{\theta}) = p(\{\Delta c_{t-3,t}, \mathbf{r}_t^e\}_{t=1}^T \mid \boldsymbol{\theta}) \prod_{t \in A} p(d_t),$$

where $\Delta c_{t-3,t}$ is missing for all $t \in A$. Therefore, the inference for $\boldsymbol{\theta}$ does not depend on d_t . For simplicity, we follow Mariano and Murasawa (2003) and set $d_t = 0$ for all t .

We can now write down a mixed-frequency state-space system,

$$\mathbf{z}_t = \mathbf{F}\mathbf{z}_{t-1} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}_{\bar{S}+5}, \boldsymbol{\Psi}), \quad \text{and} \quad (\text{IA.16})$$

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{H}_t \mathbf{z}_t + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}_{N+1}, \boldsymbol{\Sigma}_t) \quad (\text{IA.17})$$

where $\mathbf{y}_t := [\Delta c_{t-3,t}^+, \mathbf{r}_t^{e'}]', \boldsymbol{\mu}_t := [\mu_{c,t}, \boldsymbol{\mu}_r']', \mathbf{v}_t := [f_t, \mathbf{0}_{\bar{S}+4}']', \mathbf{w}_t := [\tilde{w}_t^c, \mathbf{w}_t^{r'}]',$

$$\boldsymbol{\Psi} := \underbrace{\begin{bmatrix} 1 & \mathbf{0}_{\bar{S}+4}' \\ \mathbf{0}_{\bar{S}+4} & \mathbf{0}_{(\bar{S}+4) \times (\bar{S}+4)} \end{bmatrix}}_{(\bar{S}+5) \times (\bar{S}+5)}, \quad \mathbf{F} := \underbrace{\begin{bmatrix} \mathbf{0}_{\bar{S}+4}' & 0 \\ \mathbf{I}_{\bar{S}+4} & \mathbf{0}_{\bar{S}+4} \end{bmatrix}}_{(\bar{S}+5) \times (\bar{S}+5)}, \quad \boldsymbol{\Sigma}_t := \underbrace{\begin{bmatrix} \sigma_{c,t}^2 & \mathbf{0}_N' \\ \mathbf{0}_N & \boldsymbol{\Sigma}_r \end{bmatrix}}_{(N+1) \times (N+1)}, \quad \mathbf{H}_t := \underbrace{\begin{bmatrix} & \mathbf{H}_{1,t} \\ \rho_f^r & \mathbf{0}_N & \dots & \mathbf{0}_N \end{bmatrix}}_{(N+1) \times (\bar{S}+5)},$$

$$\mu_{c,t} = \begin{cases} \mu_c, & t \text{ is the last month of the quarter} \\ 0, & \text{otherwise} \end{cases} \quad \mathbf{H}_{1,t} = \begin{cases} (\boldsymbol{\rho}^c)^\top, & t \text{ is the last month of the quarter} \\ \mathbf{0}_{\bar{S}+5}^\top, & \text{otherwise} \end{cases}$$

$$\tilde{w}_{c,t} = \begin{cases} w_{c,t}, & t \text{ is the last month of the quarter} \\ d_t, & \text{otherwise} \end{cases} \quad \sigma_{c,t}^2 = \begin{cases} \sigma_c^2, & t \text{ is the last month of the quarter} \\ 1, & \text{otherwise} \end{cases}.$$


There are two major differences from the benchmark analysis based on the time-aggregated quarterly data. First, there is a measurement error, denoted by d_t , in the consumption growth equation when $\Delta c_{t-3,t}$ is missing (that is, when t is not the last month of the quarter). Second, we infer the common factor f_t using monthly asset returns and quarterly observed consumption growth. However, when $\Delta c_{t-3,t}$ is missing, it does not play a role in estimating

f_t . To see the last point, let's consider the standard Kalman filter,

$$\begin{aligned}\mathbf{z}_{t|t-1} &= \mathbf{F} z'_{t-1|t-1}; \quad \mathbf{V}_{t|t-1} = \mathbf{F} \mathbf{V}_{t-1|t-1} \mathbf{F}' + \Psi; \quad \mathbf{K}_t = \mathbf{V}_{t|t-1} \mathbf{H}_t' \left(\mathbf{H}_t \mathbf{V}_{t|t-1} \mathbf{H}_t' + \Sigma_t \right)^{-1} \\ \mathbf{z}_{t|t} &= \mathbf{z}_{t|t-1} + \mathbf{K}_t (\mathbf{y}_t - \boldsymbol{\mu}_t - \mathbf{H}_t \mathbf{z}_{t|t-1}); \quad \mathbf{V}_{t|t} = \mathbf{V}_{t|t-1} - \mathbf{K}_t \mathbf{H}_t \mathbf{V}_{t|t-1}; \\ \mathbf{z}_{t|\tau} &:= \mathbb{E} [\mathbf{z}_t | \mathbf{y}^\tau, \mathbf{H}_t, \boldsymbol{\Psi}, \Sigma_t]; \quad \text{and } \mathbf{V}_{t|\tau} := \text{Cov} (\mathbf{z}_t | \mathbf{y}^\tau, \mathbf{H}_t, \boldsymbol{\Psi}, \Sigma_t).\end{aligned}$$

When t is not the final month of a quarter, the first row of \mathbf{H}_t consists of zeros, so the Kalman gain \mathbf{K}_t depends only on the factor loadings of asset returns ($\boldsymbol{\rho}_f^r$). Moreover, the first element of $\mathbf{y}_t - \boldsymbol{\mu}_t - \mathbf{H}_t \mathbf{z}_{t|t-1}$ is zero when $\Delta c_{t-3,t}$ is not observed. This observation echoes our previous argument that the assumption of d_t 's data-generating process is not essential as long as it does not rely on $\boldsymbol{\theta}$.

Figure IA.8: MA component of consumption growth: Mixed-frequency estimation.




The figure shows the posterior mean of the moving average f_t component of quarterly consumption growth. Grey areas denote NBER recessions. We study a mixed-frequency single-factor model in equations (3) and (IA.14), with $\bar{S} = 36$ (months). The cross-section of test assets includes 25 size- and value-sorted portfolios, 12 industry portfolios, and nine bond portfolios.

G.1 Incorporating stochastic volatility at mixed frequencies

We consider the following return dynamics at the monthly frequency:


$$\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r f_t + \boldsymbol{\beta}_f \sigma_{f,t-1}^2 + \boldsymbol{\beta}_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r, \quad f_t \sim \mathcal{N} \left(0, \frac{\exp(h_{ft})}{\sigma_{f,t-1}^2} \right), \quad (\text{IA.18})$$

Figure IA.9: Autocorrelation structure of consumption growth squared forecast errors: Quarterly frequency.



Left panel: Autocorrelation function of $\widehat{Var}_t(\Delta c_{t+1})$ with 95% and 99% confidence bands. Right panel: p -values of Ljung and Box (1978) (red triangles) and Box and Pierce (1970) (blue circles) tests (quarterly data, 1963:Q3–2019:Q4). The model is based on the mixed-frequency state-space estimation.


Figure IA.10: 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 quarterly asset returns at several horizons h . We report the p -value of the F -test of joint significance of the covariates as well as the 10% and 5% significance thresholds (respectively, horizontal dot-dashed and dotted lines). The solid blue line with circles denotes statistics for the correctly specified conditional mean for the consumption growth process, while the dashed red line with triangles corresponds to the assumption of a constant conditional mean. In the correct mean process, we estimate a mixed-frequency state-space model of quarterly consumption growth and monthly asset returns.

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 long-run consumption volatility (in equation (IA.15)). The key distinction from the analysis in the main text is that we consider monthly asset returns in equation (IA.18) and, hence, monthly stochastic volatility of f_t and \mathbf{w}_t^r ; meanwhile, we still estimate the stochastic volatility process of the quarterly consumption growth shock (w_t^c). At the mixed frequency, we do not include $\sigma_{c,t-1}^2$ into equation (IA.18). In Section F.1, we consider monthly consumption data and further incorporate the short-run stochastic volatility of consumption, $\sigma_{c,t-1}^2$, into the return dynamics.

Figure IA.11: Cumulative response function of consumption growth to a one-standard-deviation shock spanned by asset returns: Mixed-frequency estimation with $\bar{S} = 0$.



The figure presents posterior means of the cumulative response function of consumption growth (solid line with circles), along with the centered posterior 90% (dotted lines) and 68% (dashed lines) coverage regions. We estimate a mixed-frequency state-space model of quarterly consumption growth and monthly asset returns, with $\bar{S} = 0$. These monthly f_t shocks account for **6.4%** of quarterly consumption growth time-series variation (quarterly data, 1963:Q3–2019:Q4).

We next consider the consumption dynamics, as follows:

$$\Delta c_{t-3,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j \left(\frac{1}{3} f_{t-j} + \frac{2}{3} f_{t-j-1} + f_{t-j-2} + \frac{2}{3} f_{t-j-3} + \frac{1}{3} f_{t-j-4} \right) + w_t^c, \quad w_t^c \sim \mathcal{N} \left(0, \exp(h_{ct}) \frac{\sigma_{c,t-1}^2}{\sigma_{c,t-1}^2} \right). \quad (\text{IA.19})$$

We can now write a mixed-frequency state-space system with three separate stochastic volatility processes,

$$\mathbf{z}_t = \mathbf{F} \mathbf{z}_{t-1} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}_{\bar{S}+5}, \mathbf{\Psi}_t), \quad \text{and} \quad (\text{IA.20})$$

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{H}_t \mathbf{z}_t + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}_{N+1}, \mathbf{\Sigma}_t) \quad (\text{IA.21})$$

$$\text{where } \mathbf{y}_t := [\Delta c_{t-3,t}^+, \mathbf{r}_t^{e'}]', \quad \boldsymbol{\mu}_t := [\mu_{c,t}, \boldsymbol{\mu}_{r,t}']', \quad \mathbf{v}_t := [f_t, \mathbf{0}'_{\bar{S}+4}]', \quad \mathbf{w}_t := [\tilde{w}_t^c, \mathbf{w}_t^r]',$$

$$\mathbf{\Psi}_t := \underbrace{\begin{bmatrix} \exp(h_{ft}) & \mathbf{0}'_{\bar{S}+4} \\ \mathbf{0}_{\bar{S}+4} & \mathbf{0}_{(\bar{S}+4) \times (\bar{S}+4)} \end{bmatrix}}_{(\bar{S}+5) \times (\bar{S}+5)}, \quad \mathbf{F} := \underbrace{\begin{bmatrix} \mathbf{0}'_{\bar{S}+4} & 0 \\ \mathbf{I}_{\bar{S}+4} & \mathbf{0}_{\bar{S}+4} \end{bmatrix}}_{(\bar{S}+5) \times (\bar{S}+5)}, \quad \mathbf{\Sigma}_t := \underbrace{\begin{bmatrix} \tilde{\sigma}_{c,t}^2 & \mathbf{0}'_N \\ \mathbf{0}_N & \mathbf{\Sigma}_{r,t-1} \end{bmatrix}}_{(N+1) \times (N+1)}, \quad \mathbf{H}_t := \underbrace{\begin{bmatrix} \mathbf{H}_{1,t} & \mathbf{0}_N & \dots & \mathbf{0}_N \\ \rho_f^r & \mathbf{0}_N & \dots & \mathbf{0}_N \end{bmatrix}}_{(N+1) \times (\bar{S}+5)},$$

$$\mu_{c,t} = \begin{cases} \mu_c, & t \text{ is the last month of the quarter} \\ 0, & \text{otherwise} \end{cases}, \quad \mathbf{H}_{1,t} = \begin{cases} (\boldsymbol{\rho}^c)^\top, & t \text{ is the last month of the quarter} \\ \mathbf{0}_{\bar{S}+5}^\top, & \text{otherwise} \end{cases},$$

$$\tilde{w}_{c,t} = \begin{cases} w_{c,t}, & t \text{ is the last month of the quarter} \\ d_t, & \text{otherwise} \end{cases}, \quad \tilde{\sigma}_{c,t}^2 = \begin{cases} \exp(h_{ct}), & t \text{ is the last month of the quarter} \\ 1, & \text{otherwise} \end{cases}.$$

Finally, we apply the Kalman smoother to extract the latent state f_t . The estimation method

of other parameters is the same as in the previous analysis.

H Benchmarking

The covariance between consumption growth and asset returns can be biased by the benchmarking process used by the US Bureau of Economic Analysis (BEA). The benchmarking of NIPA data have been discussed in the past literature (see, e.g., Wilcox (1992) and Triplett (1997)). Simply speaking, the Monthly Retail Trade Survey (MRTS), which is used to construct monthly and quarterly consumption estimates, is of lower quality than the Annual Retail Trade Survey (ARTS). Hence, the four quarters of consumption will never precisely equal the corresponding annual measure; consequently, the quarterly estimates are ex post revised to benchmark the annual ones. As explained in Chapter 4, estimating methods, of the NIPA handbook,¹³ “*for the periods for which annual estimates are available and the quarterly estimates must be forced to average to these annual totals*”.

Without observing the unfiltered data, we cannot directly incorporate the benchmarking equation into our state-space model. Nevertheless, we can explore the mechanical effect originating from benchmarking through simulation. However, the exact benchmarking procedure (e.g., how BEA smoothes the consumption data) is unknown. Therefore, in this section, we consider several simulation settings and aim to explore when the benchmarking process will (and will not) distort the impulse responses of quarterly consumption growth to asset return shocks.

We simulate monthly consumption in log units; that is, $c_t^* = c_{t-1}^* + \mu_c + \rho_0 f_t + w_t^c$, $w_t^c \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_c^2)$, and $f_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. Note that the true monthly consumption growth is IID and correlates with only the contemporaneous f_t in simulations. However, we do not observe c_t^* but only its noisy proxy based on the MRTS,


$$c_t = \beta c_t^* + (1 - \beta) \mu_c t + \epsilon_{c,t}, \quad \epsilon_{c,t} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{\epsilon c}^2), \quad (\text{IA.22})$$

where $\beta \leq 1$, and $\epsilon_{c,t}$ is the measurement error. Equation (IA.22) implies the following dynamics of monthly consumption growth:


$$\Delta c_t = \beta \Delta c_t^* + (1 - \beta) \mu_c + (\epsilon_{c,t} - \epsilon_{c,t-1}). \quad (\text{IA.23})$$

¹³See <https://www.bea.gov/resources/methodologies/nipa-handbook/pdf/chapter-04.pdf>.


Figure IA.12: Filtered stochastic volatilities of consumption and returns: Mixed-frequency state-space model.



Panel A: Log volatility of the short-run consumption shock (w_t^c) of equation (IA.19)




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}).

The figure shows the estimated stochastic volatilities of the mixed-frequency model in equations (IA.18)–(IA.19). Solid blue lines depict the posterior median of the log volatility, whereas dotted red lines denote 2.5% and 97.5% credible intervals. Shaded (patterned) areas reflect constant volatility levels that would not be rejected given the credible intervals.

Figure IA.13: Loadings of excess returns on consumption and returns volatilities: Mixed-frequency state-space model.



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 (IA.18).



Panel B: Posterior distributions of excess return loadings (β_r) on the common financial return variance ($\sigma_{r,t-1}^2$) in equation (IA.18).

This figure shows the 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 (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last.

That is, monthly consumption growth from MRTS (Δc_t) has the same unconditional mean μ_c as the pseudo-true process (Δc_t^*) that we do not observe.

We further observe the annual consumption data from the ARTS, which is considerably more comprehensive than the MRTS. Suppose that t is the final month of the year. The annual consumption flow is denoted by $C_{t,t+12}^a$, which equals $\sum_{j=1}^{12} \exp(c_{t+j}^*)$. Since $C_{t,t+12}^a$ never matches the sum of monthly consumption flows obtained from the MRTS ($\sum_{j=1}^{12} \exp(c_{t+j})$), benchmarking implies the following revision of the monthly estimates:

$$\exp(\tilde{c}_{t+j}) = \exp(c_{t+j}) \times \frac{\sum_{j=1}^{12} \exp(c_{t+j}^*)}{\sum_{j=1}^{12} \exp(c_{t+j})}, \quad (\text{IA.24})$$

where \tilde{c}_{t+j} denotes the revised estimates of monthly consumption in log units. Using \tilde{c}_{t+j} , we compute the quarterly consumption growth (denoted by $\tilde{c}_{t-3,t}^q$) and estimate its impulse responses to quarterly asset return shocks, defined as $f_{t-3,t}^q = (f_t + f_{t-1} + f_{t-2})/\sqrt{3}$.

There are two representative simulation settings to consider: (1) $\beta = 1$ and (2) $\beta < 1$ in equation (IA.22). First, when $\beta = 1$, the benchmarking process in equation (IA.24) does not distort (in population) the impulse responses of quarterly consumption growth to f_t shocks. To see this intuitively, we can rewrite the adjustment factor in equation (IA.24) as follows:

$$\frac{\sum_{j=1}^{12} \exp(c_{t+j}^*)}{\sum_{j=1}^{12} \exp(c_{t+j})} = \frac{\sum_{j=1}^{12} \exp(c_{t+j}^*)}{\sum_{j=1}^{12} \exp(c_{t+j}^* + \epsilon_{c,t+j})}.$$


Smoothing the monthly estimates based on the MRTS to match the annual ones, on average, reallocates only measurement errors over different periods and, hence, does not change the impulse responses to f_t . Note that assuming $\beta = 1$ is a common way of modeling consumption measurement error in the literature (see, e.g., Schorfheide, Song, and Yaron (2018)).

Second, instead, when $\beta < 1$, benchmarking introduces a mechanical autocorrelation between f_t shocks and observed consumption growth. The intuition is that, in this case, benchmarking smooths both measurement errors and f_t shocks over 12 months, leading to the slow adjustment of consumption growth to f_t shocks.

To illustrate the above points, we simulate consumption data following equations (IA.22)–(IA.23). In the simulations, the contemporaneous correlation between the true monthly consumption growth (Δc_t^*) and f_t is 0.40, whereas the measurement error is calibrated such that the observed monthly consumption growth (Δc_t) has a correlation of 0.30 with f_t . After obtaining the monthly growth rates, we aggregate the simulated data to quarterly consumption growth, which is then normalized to have the same standard deviation (0.0072) as in the true data.

Panel (a) of Figure IA.14 shows the simulated impulse responses based on $\beta = 1$ at the quarterly frequency. Due to the time-aggregation bias (the blue line with circles), the cumulative impulse responses increase from 0.002 to about 0.0035 at $S = 1$ and stay flat after one quarter. The benchmarking process does not introduce additional autocorrelations between consumption growth and f_t (see orange line with triangles). Panel (b) displays the case with $\beta = 0.5$ at the quarterly frequency. Unlike Panel (a), quarterly consumption growth responds to both the one-quarter lagged f_t shocks and the lagged f_t of the preceding four quarters. The cumulative impulse responses become flat after four quarters. Panels (c) and (d) of Figure IA.14 repeat similar simulations with mixed-frequency estimation; that is, it displays how quarterly consumption growth responds to a one-standard-deviation *monthly*

Figure IA.14: Simulated cumulative response function of quarterly consumption growth.




The figure plots the cumulative impulse response function (CIRF) of quarterly consumption growth. The blue lines with circles show the impulses responses with the time aggregation bias, whereas the orange lines with triangualrs display those with both the time aggregation bias and benchmarking. Panels (a) and (b), which assume $\beta = 1$ and $\beta = 0.5$ in equation (IA.23) respectively, consider the impulse responses at the quarterly frequency, i.e., how quarterly consumption growth responds to quarterly f_t shocks. Panels (c) and (d), instead, study the mixed-frequency CIRFs of quarterly consumption growth to a one-standard-deviation monthly f_t shock. We simulate 100,000 years of data. Quarterly consumption growth is standardized to have the same standard deviation (0.0072) as in the data.

f_t shock. With only the time-aggregation bias, quarterly consumption growth reacts to both contemporaneous and four lagged monthly f_t shocks. However, as we further incorporate benchmarking of consumption data, the CIRFs continue to increase until 12 months after the impulse and stay flat afterwards.¹⁴

We further explore the impulse responses with time-aggregation bias and benchmarking

¹⁴Note that the quarterly CIRFs roughly equal the mixed-frequency estimates divided by $\sqrt{3}$.

Figure IA.15: Simulated cumulative response function of quarterly consumption growth: Different coefficients of $\text{corr}(\Delta c_t^*, f_t)$.



The figure plots the cumulative impulse response function (CIRF) of quarterly consumption growth. We consider both time-aggregation bias and benchmarking ($\beta = 0.5$ in equation (IA.23)). The contemporaneous correlation between true monthly consumption growth and f_t , $\text{corr}(\Delta c_t^*, f_t)$, ranges from 0.10 to 0.40. Panel (a) considers the impulse responses at the quarterly frequency, i.e., how quarterly consumption growth responds to quarterly f_t shocks. Panel (b), instead, presents the mixed-frequency CIRFs of quarterly consumption growth to a one-standard-deviation monthly f_t shock. We simulate 100,000 years of data. Quarterly consumption growth is standardized to have the same standard deviation (0.0072) as in the data.

of consumption data, assuming different contemporaneous correlations between true monthly consumption growth and f_t . Figure IA.15 plots these CIRFs for correlations ranging from 0.10 to 0.40. In the monthly data, this contemporaneous correlation coefficient is estimated to be about 0.17; hence, the assumption of a correlation of 0.20 (the solid green line) represents the most realistic case.

A crucial observation is that the benchmarking process cannot fully explain the patterns that we detected in the real data. First, in Figures 3 and 6, the CIRFs continue to increase until the seventh quarter (or 21 months), whereas the benchmarking process in our simulations leads to slowly increasing CIRFs within only the first four quarters. Second, the CIRFs of the IID consumption data with benchmarking are less than 0.002 at the quarterly frequency (0.003 at the mixed frequency) under the assumption of 0.2 correlation (in the most realistic calibration), as shown in green solid lines. Instead, we estimate much more sizable CIRFs of about 0.01 at the quarterly frequency in Figure 3 (0.015 in the mixed-frequency estimation in Figure 6).

I Model selection: Consumption mean and volatility risks in asset returns

Commonly, time-varying risk premia in consumption-based asset pricing models are induced by the time variation in the volatility of consumption shocks; that is, the models (a) assume the existence of time-varying volatility and (b) postulate that returns are linear in this volatility. Our state-space formulation allows us to formally test these common theoretical assumptions because it nests, as particular cases, the most popular structural models.

Table IA.I, Panel A, reports Bayes factors and posterior probabilities for several restricted and unrestricted versions of the specification in equations (2) and (16)–(19), and Section B.2. In particular, we test both commonalities in the various volatility processes (for returns and consumption mean shocks) and 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 solely the model with the highest likelihood). See, for example, Schervish (1996).

The data favor a specification in which *i*) volatilities do *not* affect excess returns (Models I–V) and *ii*) returns and consumption have *distinct* volatility processes—the posterior probability of such a formulation (Model I) is almost 100%.¹⁵

Panel B of Table IA.I shows that the share of variance of asset returns explained by the shock spanned by consumption (f_t) is stable across specification, even when we include the stochastic volatilities in the return equation (Models VI–X). This stresses the robustness of our identification approach to recover the conditional mean shocks to consumption. Furthermore, Panel C emphasizes that even when allowing for stochastic volatility in all the latent shocks, the correlation between the variances of consumption and asset returns remains considerably small.

One may wonder whether our model selection findings are largely driven by the likelihood function for the joint dynamics of returns and consumption being dominated by the former due to their much higher volatility and, hence, reflect only a model for returns, rather than their joint dynamics. Consequently, Table IA.X repeats the same model selection exercise after scaling returns to have unit variance and consumption to have variance equal to the

¹⁵As shown in Table IA.II, these results are stable if we exclude bond returns from the state-space model.

Table IA.I: Model comparison using log Bayes factors: Baseline stock and bond portfolios

	Models									
	I	II	III	IV	V	VI	VII	VIII	IX	X
Panel A: Model probabilities and Bayes factors										
Log of Bayes factor:	0	-27	-37	-51	-85	-105	-161	-223	-236	-332
Posterior probability.:	100%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Panel B: Variance Decomposition of Returns										
f_t	62.3%	62.4%	62.2%	62.3%	62.2%	57.6%	57.8%	57.7%	58.4%	57.8%
$\sigma_{c,t-1}^2$										0.2%
$\sigma_{f,t-1}^2$								0.3%	0.3%	0.3%
$\sigma_{r,t-1}^2$						5.1%	4.8%	4.9%	4.5%	4.9%
Panel C: Correlations among stochastic volatility processes										
$cor(\sigma_{ct}^2, \sigma_{ft}^2)$	0.06	0.12			0.14		0.08			0.07
	[-0.09, 0.29]	[-0.03, 0.32]			[-0.02, 0.35]		[-0.08, 0.31]			[-0.09, 0.30]
$cor(\sigma_{ct}^2, \sigma_{rt}^2)$	0.12	0.12	0.12		0.14		0.15			0.15
	[-0.03, 0.32]	[-0.03, 0.32]	[-0.03, 0.32]		[-0.02, 0.35]		[-0.02, 0.36]			[-0.02, 0.36]
$cor(\sigma_{ft}^2, \sigma_{rt}^2)$	0.23	1		0.23		1	0.22	0.22	0.22	
	[0.07, 0.42]	[1, 1]		[0.07, 0.43]		[1, 1]	[0.06, 0.41]	[0.06, 0.41]	[0.06, 0.41]	

Model I: w_t^c , f_t and w_t^r follow SV processes, and $r_t^e = \mu_r + \rho^r f_t + w_t^r$.

Model II: w_t^c , f_t and w_t^r follow SV processes, $r_t^e = \mu_r + \rho^r f_t + w_t^r$, and $h_{ft} = h_{rt}$.

Model III: w_t^c and w_t^r follow SV processes, $f_t \stackrel{iid}{\sim} \mathcal{N}(0, 1)$, and $r_t^e = \mu_r + \rho^r f_t + w_t^r$.

Model IV: f_t and w_t^r follow SV processes, $w_t^c \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_c^2)$, and $r_t^e = \mu_r + \rho^r f_t + w_t^r$.

Model V: w_t^r follows SV process, $w_t^c \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_c^2)$, $f_t \stackrel{iid}{\sim} \mathcal{N}(0, 1)$, and $r_t^e = \mu_r + \rho^r f_t + w_t^r$.

Model VI: w_t^c , f_t and w_t^r follow SV processes, $r_t^e = \mu_r + \rho^r f_t + \beta_f \sigma_{f,t-1}^2 + w_t^r$, and $h_{ft} = h_{rt}$.

Model VII: w_t^r follows SV process, $w_t^c \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_c^2)$, $f_t \stackrel{iid}{\sim} \mathcal{N}(0, 1)$, and $r_t^e = \mu_r + \rho^r f_t + \beta_f \sigma_{r,t-1}^2 + w_t^r$.

Model VIII: w_t^c , f_t and w_t^r follow SV processes, and $r_t^e = \mu_r + \rho^r f_t + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 + w_t^r$.

Model IX: f_t and w_t^r follow SV processes, $w_t^c \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_c^2)$, and $r_t^e = \mu_r + \rho^r f_t + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 + w_t^r$.

Model X: w_t^c , f_t and w_t^r follow SV processes, and $r_t^e = \mu_r + \rho^r f_t + \beta_c \sigma_{c,t-1}^2 + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 + w_t^r$.

The table summarizes the model comparison for restricted and unrestricted versions of the specification in equations (2) and (16)–(19), and Section B.2. Panel A reports log Bayes factors and posterior probabilities. We approximate the Bayes factor using the Schwartz criterion. We use Model I as a 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. Panel B reports the variance decomposition of asset returns for the model-specific sources of time variation. Panel C reports the correlations among the stochastic volatility processes.

number of assets. Again, we find that there is no support for stochastic volatilities driving time-varying risk premia or for the volatility of consumption and returns being proportional to each other.

J Stochastic volatility with leverage

In Section III.6, we consider stochastic volatility models without the leverage effect; that is, shocks to asset returns are independent of the shocks to their stochastic volatility processes.

Table IA.II: Model comparison using log Bayes factors with only equity portfolios

	Models									
	I	II	III	IV	V	VI	VII	VIII	IX	X
Panel A: Model probabilities and Bayes factors										
Log of Bayes factor:	0	-26	-39	-48	-87	-103	-146	-185	-201	-278
Posterior probability.:	100%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Panel B: Variance Decomposition of Returns										
f_t	77.3%	77.3%	77.0%	77.2%	77.0%	73.2%	71.4%	71.5%	71.8%	71.2%
$\sigma_{c,t-1}^2$										0.2%
$\sigma_{f,t-1}^2$									0.2%	0.2%
$\sigma_{r,t-1}^2$						3.4%	5.2%	5.4%	4.9%	5.2%
Panel C: Correlations among stochastic volatility processes										
$cor(\sigma_{ct}^2, \sigma_{ft}^2)$	0.07	0.12			0.12		0.07		0.07	
	[-0.09, 0.29]	[-0.03, 0.34]				[-0.03, 0.34]		[-0.09, 0.30]		[-0.09, 0.30]
$cor(\sigma_{ct}^2, \sigma_{rt}^2)$	0.12	0.12	0.12			0.12	0.13		0.14	
	[-0.03, 0.34]	[-0.03, 0.34]	[-0.03, 0.33]				[-0.03, 0.34]	[-0.02, 0.36]		[-0.02, 0.38]
$cor(\sigma_{ft}^2, \sigma_{rt}^2)$	0.25	1		0.25		1	0.22	0.21	0.23	
	[0.07, 0.45]	[1, 1]		[0.08, 0.46]		[1, 1]		[0.05, 0.45]	[0.05, 0.44]	[0.06, 0.45]

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

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

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

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

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

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

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

Model VIII: w_t^c , f_t and \mathbf{w}_t^r follow SV processes, and $\mathbf{r}_t^e = \mu_r + \rho^r f_t + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r$.

Model IX: f_t and \mathbf{w}_t^r follow SV processes, $w_t^c \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_c^2)$, and $\mathbf{r}_t^e = \mu_r + \rho^r f_t + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r$.

Model X: w_t^c , f_t and \mathbf{w}_t^r follow SV processes, and $\mathbf{r}_t^e = \mu_r + \rho^r f_t + \beta_c \sigma_{c,t-1}^2 + \beta_f \sigma_{f,t-1}^2 + \beta_r \sigma_{r,t-1}^2 + \mathbf{w}_t^r$.

The table summarizes the model comparison for restricted and unrestricted versions of the specification in equations (2) and (16)–(19), and Section B.2. Panel A reports log Bayes factors and posterior probabilities. We approximate the Bayes factor using the Schwartz criterion. We use Model I as a 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. Panel B reports the variance decomposition of asset returns for the model-specific sources of time variation. Panel C reports the correlations among the stochastic volatility processes.

In this Appendix, we further consider the leverage effect in the stochastic volatility of shocks to the conditional consumption mean, f_t . Specifically, f_t follows a normal distribution with the stochastic volatility process given by

$$f_t = \exp\left(\frac{h_{ft}}{2}\right) \epsilon_{ft}, \quad y_{ft}^* = \log(f_t^2) = h_{ft} + \log(\epsilon_{ft}^2), \quad h_{ft} = \delta_f h_{f,t-1} + \sqrt{1 - \delta_f^2} \eta_{f,t-1}, \quad (\text{IA.25})$$

where $\epsilon_{ft} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $\eta_{ft} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. Although macro-finance literature normally assumes that ϵ_{ft} and η_{ft} are independent, many earlier papers (e.g., Black (1976), Nelson (1991), and Yu (2005)) document strong leverage effects in asset returns: An increase in volatility often follows a drop in equity returns; that is, ϵ_{ft} and η_{ft} tend to be negatively correlated.

To incorporate the leverage effect, we consider the following distributional assumption:

$$\begin{pmatrix} \epsilon_{ft} \\ \eta_{ft} \end{pmatrix} \stackrel{\text{iid}}{\sim} \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} 1 & \zeta \\ \zeta & 1 \end{bmatrix} \right), \quad (\text{IA.26})$$

where $\zeta < 0$ if the leverage effect exists: A negative shock to y_{ft}^* (negative ϵ_{ft}) leads to a spike in η_{ft} ; therefore, the stochastic volatility of f_{t+1} , $\exp(h_{f,t+1}/2)$, tends to increase.

We follow the solution proposed by Omori, Chib, Shephard, and Nakajima (2007). In particular, they approximate the joint distribution of (e_{ft}, η_{ft}) using a mixture of bivariate Gaussian densities:

$$p(e_{ft}, \eta_{ft} \mid d_t) \approx \sum_{i=1}^K q_i \cdot \mathcal{N}(e_{ft} \mid m_i, v_i^2) \cdot \mathcal{N}(\eta_{ft} \mid d_t \zeta \exp(m_i/2)[a_i + b_i(e_{ft} - m_i)], 1 - \zeta^2), \quad (\text{IA.27})$$

where $\mathcal{N}(x \mid m, v^2)$ means that $x \sim \mathcal{N}(m, v^2)$, $d_t = \text{sign}(f_t)$, $\{(q_i, m_i, v_i^2, a_i, b_i)\}_{i=1}^K$ are fixed real numbers given by Table 1 of Omori, Chib, Shephard, and Nakajima (2007). As they do, we select $K = 10$, so ten different normal distributions are used to approximate the joint distribution of (e_{ft}, η_{ft}) .

The SV model can be expressed as

$$\begin{pmatrix} y_{ft}^* \\ h_{f,t+1} \end{pmatrix} = \begin{pmatrix} h_{f,t} \\ \delta_f h_{f,t} \end{pmatrix} + \begin{pmatrix} e_{f,t} \\ \sqrt{1 - \delta_f^2} \eta_{f,t} \end{pmatrix}.$$

Using the mixture approximation (IA.27) and introducing the mixture component indicator $S_{ft} \in \{1, \dots, K\}$, we have that

$$\begin{pmatrix} e_{f,t} \\ \eta_{f,t} \end{pmatrix} \mid S_{ft} = i, d_t = \begin{pmatrix} m_i + v_i z_t \\ d_t \zeta \exp(m_i/2)[a_i + b_i v_i z_t] + \sqrt{1 - \zeta^2} z_t^* \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} z_t \\ z_t^* \end{pmatrix} \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_2).$$

Since (e_{ft}, η_{ft}) are not independent of each other, we augment the space of latent states as

follows:

$$y_{ft}^* - m_i = h_{f,t} + v_i z_t, \quad (\text{observation equation}), \quad (\text{IA.28})$$

and we treat both $h_{f,t}$ and z_t as latent states, with the following state equation:

$$\begin{pmatrix} h_{f,t+1} \\ z_{t+1} \\ c_{t+1} \end{pmatrix} = \begin{pmatrix} \delta_f & d_t \zeta \exp(m_i/2) b_i v_i \sqrt{1 - \delta_f^2} & d_t \zeta \exp(m_i/2) a_i \sqrt{1 - \delta_f^2} \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} h_{f,t} \\ z_t \\ c_t \end{pmatrix} + \begin{pmatrix} \sqrt{(1 - \zeta^2)(1 - \delta_f^2)} z_t^* \\ z_{t+1} \\ 0 \end{pmatrix}, \quad (\text{IA.29})$$

where c_t is introduced to remove the mechanical drift term: $c_{t+1} = c_t = \dots = c_0 = 1$. Using the system (IA.28)–(IA.29), we can use the Kalman smoother to estimate $(h_{f,t}, z_t)$.

The inference for S_{ft} and δ_f are the same as before. The posterior distribution of S_{ft} is

$$\begin{aligned} p(S_{ft} = i \mid e_{f,t}, \eta_{f,t}, \zeta, d_t) \\ \propto \Pr(S_{ft} = i) \cdot v_i^{-1} \exp \left\{ -\frac{(e_{ft} - m_i)^2}{2v_i^2} - \frac{[\eta_{f,t} - d_t \zeta \exp(m_i/2)[a_i + b_i(e_{ft} - m_i)]]^2}{2(1 - \zeta^2)} \right\}. \end{aligned}$$

We update δ_f using a Metropolis algorithm as before.

Finally, we discuss the posterior inference for ζ , the key parameter governing the leverage effect. We rely on the distributional assumption in equation (IA.26),

$$\begin{pmatrix} f_t \cdot \exp(-\frac{h_{ft}}{2}) \\ \frac{h_{f,t+1} - \delta_f h_{f,t}}{\sqrt{1 - \delta_f^2}} \end{pmatrix} \stackrel{\text{iid}}{\sim} \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} 1 & \zeta \\ \zeta & 1 \end{bmatrix} \right),$$

which implies the following posterior distribution of ζ under a flat prior $\pi(\zeta) \propto 1$:

$$\begin{aligned} p(\zeta \mid \mathbf{h}_f, \mathbf{f}, \delta_f) &\propto (1 - \zeta^2)^{-\frac{T-1}{2}} \exp \left\{ -\frac{\text{SSE}(\mathbf{h}_f, \mathbf{f}, \delta_f, \zeta)}{2} \right\}, \text{ and} \\ \text{SSE}(\mathbf{h}_f, \mathbf{f}, \delta_f, \zeta) &= \sum_{t=1}^{T-1} \begin{bmatrix} f_t \cdot \exp(-\frac{h_{ft}}{2}) & \frac{h_{f,t+1} - \delta_f h_{f,t}}{\sqrt{1 - \delta_f^2}} \end{bmatrix} \begin{bmatrix} 1 & \zeta \\ \zeta & 1 \end{bmatrix}^{-1} \begin{pmatrix} f_t \cdot \exp(-\frac{h_{ft}}{2}) \\ \frac{h_{f,t+1} - \delta_f h_{f,t}}{\sqrt{1 - \delta_f^2}} \end{pmatrix}. \end{aligned}$$

We draw ζ using a Metropolis algorithm, as follows:

- a. initialize ζ^0 ;
- b. draw ζ from a normal distribution $\zeta^* \sim \mathcal{N}(\zeta^{(j-1)}, c_{mh}^2)$;
- c. calculate $\rho(\zeta^*, \zeta^{(j-1)}) = \min\{1, \frac{p(\zeta^* \mid h_f, -)}{p(\zeta^{(j-1)} \mid h_f, -)}\}$;

d. set $\zeta^{(j)} = \zeta^{(j-1)}$ with probability $1 - \rho(\zeta^*, \zeta^{(j-1)})$ and $\zeta^{(j)} = \zeta^*$ with probability $\rho(\zeta^*, \zeta^{(j-1)})$.

c_{mh} determines the step size in the Metropolis algorithm. We choose c_{mh} such that the frequency of accepting a new ζ is about 50%.


We repeat the analysis in Section III.6 but allow for the leverage effect in f_t . First, we detect significantly negative ζ : The posterior median of ζ is -0.38 , with 90% CIs $[-0.59, -0.08]$. This finding is consistent with the leverage effect. Second, as we show in Figure IA.16, the estimates of long-run consumption volatility, σ_{ft}^2 , become sharper after introducing the leverage effect. Third, Figure IA.17 suggests that the returns' loadings on three stochastic volatility processes are almost unchanged after we introduce the leverage effect to f_t .

K Additional robustness checks


Because the state-space estimation results may depend on the choice of the cross-section of assets used for the analysis, we repeat our analysis using alternative test assets and report the results in Figures IA.31 –IA.33. They present all the key empirical results for a five-factor consumption-returns model estimated on 12 industry portfolios, nine bond portfolios, and different cross-sections of stock returns (32 size-profitability-investment-sorted portfolios, 32 size-value-investment-sorted portfolios, and 32 size-value-profitability-sorted portfolios, respectively). Figures IA.31–IA.33 indicate that all our results remain almost unchanged. We also observe the same fanning-out pattern in the term structure of consumption exposure for different cross-sections of test assets, as shown in Figure IA.37.

Throughout all the empirical analyses in the main text, we have used only nondurable consumption growth per capita. However, another popular choice for empirical work in macro-finance is a combination of nondurable consumption growth and services. We reestimate the state-space formulation of our model using this proxy for consumption growth. Because adding services might introduce additional latent dynamics, we present our findings in Figures IA.38 and IA.39 that allow for up to five latent factors. Our results are similar both qualitatively and quantitatively. In particular, the five-factor model, as shown in Figure IA.39, displays the cumulative impulse response function of consumption of about 1% after three years (Panel (a)), and the variance decompositions are rather similar to those previously reported (but with larger confidence bands due to a higher number of parame-


Figure IA.16: Filtered stochastic volatilities of consumption and returns with leverage effect in the stochastic volatility of f_t .



Panel A: Log volatility of the idiosyncratic shock to consumption (w_t^c) of equation (2).




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 (19)).

The figure shows the estimated stochastic volatilities of the model in equations (2) and (16)–(19), and Section B.2 under a diffuse prior for the autoregressive volatility coefficients. Solid blue lines depict the posterior median of the log volatility, whereas dotted red lines denote 2.5% and 97.5% credible intervals. Shaded (patterned) areas reflect constant volatility levels that would not be rejected given the credible intervals. f_t follows the stochastic volatility process with the leverage effect (see equations (IA.25)–(IA.26)).


Figure IA.17: Loadings of excess returns on consumption and returns volatilities with leverage effect in the stochastic volatility of f_t .



Panel A: Posterior distribution of excess return loadings (β_c) on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$) in equation (16).



Panel B: 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 (16).



Panel C: Posterior distributions of excess return loadings (β_r) on the common financial return variance ($\sigma_{r,t-1}^2$) in equation (16).

The figure shows the box plots of the posterior distributions of the loadings of portfolio excess returns on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$), the variance of shocks to the conditional consumption growth ($\sigma_{f,t-1}^2$), and the common financial returns variance ($\sigma_{r,t-1}^2$). Portfolios are ordered with bonds first (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last. f_t follows the stochastic volatility process with the leverage effect (see equations (IA.25)–(IA.26)).

ters). One caveat is that the estimation is more demanding in the five-factor model than the single-factor setup due to a much larger parameter space, hence, leading to more disperse distributions of, for instance, R_{adj}^2 .

L VAR estimation of the conditional consumption mean

As a benchmark for comparison of our method, in this Appendix, we employ the Parker (2001) VAR-based estimation of the conditional mean of consumption. Specifically, we estimate a three-variable VAR in market excess return ($r_{r-1,t}^m$), the log consumption level ($\log(C_{t-1,t})$), and the dividend-to-price ratio of the market portfolio (D_t/P_{t-1}) at the quarterly frequency. The number of lags is four, the same as in Parker (2001). Let \mathbf{y}_t denote the three-variable vector in quarter t , yielding the following VAR system:

$$\mathbf{y}_t = \boldsymbol{\mu}_y + \sum_{j=1}^4 \mathbf{A}_j \mathbf{y}_{t-j} + \boldsymbol{\epsilon}_{y,t}. \quad (\text{IA.30})$$

According to Parker (2001, equation (15) therein), the impulse responses of log consumption *level* to innovations in excess return measures the covariance between future cumulative consumption growth and excess returns, as follows:


$$\text{cov} \left[\log \left(\frac{C_{1+t+s}}{C_t} \right), r_{t,t+1}^m \right] = \sigma(r_{t,t+1}^m) \cdot \text{IRF}_s(\log(C)), \quad (\text{IA.31})$$

where $\text{IRF}_s(\log(C))$ represents the impulse response of the log consumption at horizon s to one-standard-deviation asset return shock. We can compare $\text{IRF}_s(\log(C))$ with the cumulative impulse responses of log consumption growth implied by our MA model. The estimation of equation (IA.30) attempts to fit the unit root in the log consumption level. We use the Bootstrap to estimate the 90% confidence intervals of $\text{IRF}_s(\log(C))$.

Figure IA.18 displays the impulse responses to a one-standard-deviation shock to excess returns. The left two panels report the IRF of excess returns and the dividend-price ratio, while Panels (b) and (d) report (twice) the IRF of the consumption level with overlay the one estimated in our state-space framework using our two cross-sections: our baseline stock and bond portfolios (Panel (b)) and the Kozak, Nagel, and Santosh (2020) anomaly portfolios.

First, we confirm that the patterns and magnitudes of impulse responses in Figure IA.18 are similar to those in Parker (2001). In particular, quarterly consumption slowly adjusts to the market excess return shock, as can be seen in Panel (b). Nevertheless, the VAR approach seems to underestimate the consumption responses in the medium- to long-run compared to the state-space framework. The long-horizon covariance at $S = 14$ quarters implied by our MA model is 28.7% larger than that implied by the VAR(4) in Panel (b), and almost twice

Figure IA.18: VAR impulse responses to a one-standard-deviation shock to excess returns.



The blue lines present the impulse response functions of quarterly market excess return (Panel (a)), log nondurable consumption (Panels (b) and (d)), and dividend-price ratio (Panel (c)) to a one-standard-deviation market excess return shock, along with the centered posterior 90% coverage regions. In Panels (b) and (d), we include the cumulative impulse responses of log nondurable consumption growth estimated based on our MA model (see the pink lines).

as large as the VAR one in Panel (d).

What drives the difference between Parker (2001) VAR and our MA model? The VAR model assumes that the shock in market excess return is the only source of asset return driver in consumption growth. Instead, our MA model, as we show in simulations, can capture an arbitrary linear combination of systematic latent factors, which are identified by a large cross-section of portfolio returns. Hence, when consumption growth reacts to other factors beyond solely the market excess return (as shown in Table 6), we can detect a significant difference between the VAR- and MA-implied IRFs.

But which model—the VAR one or our state-space formulation—does a better job at capturing the conditional consumption mean? To answer this question, we first rewrite the

Table IA.III: Predicting nondurable consumption growth using VAR(4) in Parker (2001)

$s =$	1	2	3	4	5	6	7	8	9	10	11	12
$\mathbb{E}_t^{VAR(4)}[\Delta c_{t+s-1,t+s}]$	0.984	1.028	1.024	0.958	0.853	0.468	0.389	0.307	0.516	0.692	0.531	0.462
s.e. (OLS)	(0.137)	(0.212)	(0.252)	(0.285)	(0.469)	(0.481)	(0.508)	(0.535)	(0.546)	(0.551)	(0.559)	(0.566)
s.e. (NW, lag=12)	(0.172)	(0.228)	(0.256)	(0.298)	(0.654)	(0.636)	(0.656)	(0.760)	(0.773)	(0.792)	(0.785)	(0.780)
Predictive R^2	0.200	0.103	0.074	0.052	0.016	0.005	0.003	0.002	0.004	0.008	0.004	0.003

We regress the future realized growth rates in nondurable consumption on the conditional consumption mean ($\mathbb{E}_t[\Delta c_{t+s-1,t+s}^{VAR(4)}]$) implied by the VAR(4) model in Parker (2001). s ranges from one to 12 quarters. We report (1) the point estimates of the slope coefficients, (2) the OLS and Newey-West (12 lags) standard errors within the parentheses, and (3) the predictive R^2 . Sample: 1963:Q3–2019:Q4.

VAR(4) in companion form; that is,

$$\mathbf{z}_t = \boldsymbol{\mu}_z + \mathbf{B}\mathbf{z}_{t-1} + \boldsymbol{\epsilon}_{z,t}, \quad (\text{IA.32})$$

$$\mathbf{z}_t = \begin{bmatrix} \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \mathbf{y}_{t-2} \\ \mathbf{y}_{t-3} \end{bmatrix}, \quad \boldsymbol{\mu}_z = \begin{bmatrix} \boldsymbol{\mu}_y \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \mathbf{A}_3 & \mathbf{A}_4 \\ \mathbf{I}_3 & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_3 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_3 & \mathbf{0} \end{bmatrix}, \quad \text{and } \boldsymbol{\epsilon}_{z,t} = \begin{bmatrix} \boldsymbol{\epsilon}_{y,t} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}.$$

Using equation (IA.32), we can then predict the future log consumption:

$$\mathbb{E}_t[\mathbf{z}_{t+s}] = \boldsymbol{\mu}_z + \mathbf{B}\mathbb{E}_t[\mathbf{z}_{t+s-1}], \quad s \geq 2, \quad \text{and } \mathbb{E}_t[\mathbf{z}_{t+1}] = \boldsymbol{\mu}_z + \mathbf{B}\mathbf{z}_t. \quad (\text{IA.33})$$

It then follows that the VAR-implied expected consumption growth between horizon $t+s-1$ and $t+s$ is $\mathbb{E}_t[\log(C_{t+s})] - \mathbb{E}_t[\log(C_{t+s-1})]$.

Then, to verify the accuracy of the VAR conditional mean approximation, we can regress $\Delta c_{t+s-1,t+s}$ on $\mathbb{E}_t[\log(C_{t+s})] - \mathbb{E}_t[\log(C_{t+s-1})]$ and test whether the slope coefficients are close to one (as we do for the state-space forecasts in Table 2). The results of the predictive regressions are reported in Table IA.III. Compared with the MA model (see Panel A of Table 2 in the main text), the predictions for one- and two-period ahead consumption growth are similar in VAR(4). However, as the forecast horizon increases, the performance of the VAR(4) prediction tends to deteriorate. Starting from $s \geq 5$, the predictive R^2 is close to zero in the VAR(4), with slope coefficients much smaller than one, although we cannot reject the null hypothesis of the slope coefficients being equal to one.

Finally, under the assumption of additively separable CRRA preferences, we can compute

the coefficient of relative risk aversion implied by the VAR specification,

$$\gamma_s = \frac{\mathbb{E}[r_{t,t+1}^m] + \frac{1}{2}\text{var}(r_{t,t+1}^m)}{\text{cov}\left[\log\left(\frac{C_{1+t+s}}{C_t}\right), r_{t,t+1}^m\right]} = \frac{\mathbb{E}[r_{t,t+1}^m] + \frac{1}{2}\text{var}(r_{t,t+1}^m)}{\sigma(r_{t,t+1}^m) \cdot IRF_s(\log(C))}. \quad (\text{IA.34})$$

M Cross-sectional pricing: Estimation details

We now describe the estimation of f_t 's risk price in Section IV.1 of the main text. The Bayesian framework comes from Bryzgalova, Huang, and Julliard (2024), and details can be found therein.

As in our generalized return dynamics in equation (7), we postulate an approximate factor structure for the cross-section of excess returns, as follows:

$$\mathbf{r}_t^e = \boldsymbol{\mu}_r + \boldsymbol{\rho}^r \mathbf{u}_t + \mathbf{w}_t^r, \quad \mathbf{u}_t \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_K, \mathbf{I}_K), \quad (\boldsymbol{\rho}^r)' \boldsymbol{\rho}^r = \mathbf{I}_K, \quad \mathbf{w}_t^r \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_N, \boldsymbol{\Sigma}_r), \quad (\text{IA.35})$$

where \mathbf{u}_t are K largest latent factors of asset returns. As in equation (2), consumption growth slowly adjusts to asset return shocks \mathbf{u}_t , as follows:

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j \underbrace{\boldsymbol{\eta}^\top \mathbf{u}_{t-j}}_{f_{t-j}} + w_t^c, \quad \boldsymbol{\eta}^\top \boldsymbol{\eta} = 1, \quad \text{var}(f_t) = 1, \quad (\text{IA.36})$$

so f_t is the common driver of consumption growth and asset returns. \mathbf{g}_t , the additional $(K - 1)$ latent factors in equations (3) and (20) of the main text, are linear combinations of \mathbf{u}_t orthogonal to f_t . It is worth noting that since asset return drivers, \mathbf{u}_t , are latent, we cannot identify the exact rotation of \mathbf{g}_t . However, f_t is uniquely identified (up to the sign restriction) due to equation (IA.36).

Moreover, to model the cross-sectional dimension, we consider a log linear SDF as follows:

$$m_t = 1 - \boldsymbol{\lambda}_u^\top \mathbf{u}_t = 1 - b_f f_t - \mathbf{b}_g^\top \mathbf{g}_t, \quad (\text{IA.37})$$

where $\boldsymbol{\lambda}_u$ are the risk prices associated with \mathbf{u}_t (also their risk premia since $\text{cov}(\mathbf{u}_t) = \mathbf{I}_K$). The last equality is due to the fact that we can always find a non-singular linear rotation of \mathbf{u}_t such that $(f_t, \mathbf{g}_t^\top)^\top = \mathbf{H} \mathbf{u}_t$ and $(b_f, \mathbf{b}_g^\top) = \boldsymbol{\lambda}_u^\top \mathbf{H}^{-1}$, where \mathbf{H} is a $K \times K$ non-singular matrix.

As standard, we postulate that the exposures to latent factors, $\boldsymbol{\rho}^r$, can partially explain the cross-sectional variation of expected returns:

$$\tilde{\boldsymbol{\mu}}_r = \boldsymbol{\mu}_r + \frac{1}{2} \mathbf{V}_r = \boldsymbol{\rho}^r \boldsymbol{\lambda}_u + \boldsymbol{\alpha}, \quad (\text{IA.38})$$

where we allow for pricing errors $\boldsymbol{\alpha}$ but require them to be orthogonal to the factor loadings $\boldsymbol{\rho}^r$. The latter requirement is standard in the literature (e.g., Giglio and Xiu (2021)) and acts as an identification assumption to recover risk premia estimates. We also include the Jensen correction term, $\frac{1}{2} \mathbf{V}_r$ (a vector of asset returns' variances), into equation (IA.38). Finally, after identifying $\boldsymbol{\lambda}_u$, the risk price of f_t is

$$b_f = -\text{cov}(m_t, f_t) = -\text{cov}(1 - \boldsymbol{\lambda}_u^\top \mathbf{u}_t, \boldsymbol{\eta}^\top \mathbf{u}_t) = \boldsymbol{\eta}^\top \boldsymbol{\lambda}_u, \quad (\text{IA.39})$$

where the first equality comes from the fact that f_t is orthogonal to g_t and has a unit variance.

In the data, asset return drivers, \mathbf{u}_t , are unidentified. That is, we can identify only a linear rotation of \mathbf{u}_t , denoted by $\tilde{\mathbf{u}}_t = \tilde{\mathbf{H}} \mathbf{u}_t$, and $\boldsymbol{\Sigma}_{\tilde{u}} = \text{cov}(\tilde{\mathbf{u}}_t) = \tilde{\mathbf{H}} \tilde{\mathbf{H}}^\top$. However, the risk price of f_t is point-identified due to the rotation invariance property emphasized in Giglio and Xiu (2021). To show the rotation invariance of b_f , we rewrite the equation system as

$$\begin{aligned} \mathbf{r}_t^e &= \boldsymbol{\alpha} + \underbrace{\boldsymbol{\rho}^r \tilde{\mathbf{H}}^{-1} \tilde{\mathbf{H}} \boldsymbol{\lambda}_u}_{\tilde{\boldsymbol{\rho}}^r \tilde{\boldsymbol{\lambda}}_{\tilde{u}}} - \frac{1}{2} \mathbf{V}_r + \underbrace{\boldsymbol{\rho}^r \tilde{\mathbf{H}}^{-1} \tilde{\mathbf{H}} \tilde{\mathbf{u}}_t}_{\tilde{\boldsymbol{\rho}}^r \tilde{\mathbf{u}}_t} + \mathbf{w}_t^r, \\ \Delta c_{t-1,t} &= \mu_c + \sum_{j=0}^{\bar{S}} \rho_j \underbrace{\boldsymbol{\eta}^\top \tilde{\mathbf{H}}^{-1} \tilde{\mathbf{H}} \tilde{\mathbf{u}}_{t-s}}_{\tilde{\boldsymbol{\eta}}^\top \tilde{\mathbf{u}}_{t-s}} + w_t^c, \text{ and} \\ m_t &= 1 - \boldsymbol{\lambda}_{\tilde{u}}^\top (\tilde{\mathbf{H}}^{-1})^\top \tilde{\mathbf{H}}^{-1} \tilde{\mathbf{u}}_t = 1 - \boldsymbol{\lambda}_{\tilde{u}}^\top \boldsymbol{\Sigma}_{\tilde{u}}^{-1} \tilde{\mathbf{u}}_t, \quad b_f = \underbrace{\boldsymbol{\eta}^\top \tilde{\mathbf{H}}^{-1} \tilde{\mathbf{H}} \boldsymbol{\lambda}_u}_{\tilde{\boldsymbol{\eta}}^\top \tilde{\boldsymbol{\lambda}}_{\tilde{u}}}. \end{aligned} \quad (\text{IA.40})$$

The posterior sampler for this estimation framework, as well as its finite-sample performance based on simulation studies, is provided in Bryzgalova, Huang, and Julliard (2024). Note that we estimate the means of the latent factors in our Bayesian approach and fully account for estimation uncertainty in the second (centred) moments.

N Frequentist Fama-MacBeth Estimation

To make our findings on the cross-sectional fit comparable to the existing literature, we repeat the same analysis using the traditional Fama-MacBeth (FMB) regressions. We first extract the posterior mean of the consumption mean shock, denoted as \hat{f}_t . Next, we regress portfolio returns on \hat{f}_t to obtain the asset return residuals and extract from them four principal components, denoted $\hat{\mathbf{g}}_t$, which are used as a control for omitted factors.¹⁶ Note that both \hat{f}_t and $\hat{\mathbf{g}}_t$ are normalized to have unit variances, and they are uncorrelated by construction. Therefore, the risk premia estimates based on the Fama-MacBeth regressions are also their risk prices in the SDF. Unlike the Bayesian method of Bryzgalova, Huang, and Julliard (2024) used in the main body of the paper, the frequentist approach used in this section does not account for the uncertainty in estimating f_t and g_t .

Table IA.IV reports the estimation results using frequentist Fama-MacBeth regressions with and without the common pricing error (i.e., with or without an intercept) or the control for omitted factors (i.e., with or without $\hat{\mathbf{g}}_t$). The first thing to notice is that, when an intercept is not included, the FMB estimates of consumption risk premia in Table IA.IV are statistically significant and virtually identical to the (Bayesian and frequentist) robust estimates in Table 5.

Note that the standard FMB regressions are not robust to neither omitted variable bias nor simple transformations of the test assets that preserve their span. Both types of fragility are on display in Table IA.IV. First, controlling for omitted factors ($\hat{\mathbf{g}}_t$) can drastically change the cross-sectional fit of the model. Second, both point estimates and statistical significance of the risk premia can radically change when a common intercept is included in the model and when test assets are repackaged without changing their span (see Panels D and E). This is due to the fact that a truly priced factor often appears insignificant in FMB regressions if it does not have sufficient cross-sectional dispersion in betas—but this is a property of the test assets, not of the candidate source of risk. Furthermore, this property of FMB regressions is not invariant to a simple transformation of the test assets that leaves their spanning (Sharpe ratios achievable and payoffs) unchanged (see Panel B vs. C and Panel D vs. E). That is, the popular empirical distinction between a “level” and a “slope” risk factor is an artifact of FMB regressions and not an inherent property of the risk factor. That is exactly why

¹⁶To avoid collinearity in the time-series step of the Fama-MacBeth regressions, we add a small amount of noise to $\hat{\mathbf{g}}_t$.

modern methods of risk premia estimation have largely abandoned the usage of canonical Fama-MacBeth regressions in favor of more robust approaches—for detailed discussions of these issues, see Giglio, Xiu, and Zhang (2021), Giglio and Xiu (2021), and Bryzgalova, Quaini, Trojani, and Ming (2024).

Conversely, Table 5 shows (see Panel B vs. C and Panel D vs. E therein) that both Bayesian (Bryzgalova, Huang, and Julliard (2024)) and frequentist (Giglio and Xiu (2021)) robust inference methods deliver virtually identical estimates of risk premia regardless of the asset repackaging.

O Mapping to structural models

In this Appendix, we show how our state-space formulation can be used to deliver the calibration in Section IV.2 of the main text. In particular, we show how to map the reduced-form estimate of f_t 's risk price to the structural parameters governing investors' preferences. We consider two types of investor preferences.

First, in the additively separable power utility setting, the log SDF is¹⁷:

$$m_t = \log(\delta) - \gamma \Delta c_{t-1,t}, \text{ and } \Delta \mathbb{E}_t[m_t] := m_t - \mathbb{E}_{t-1}[m_t] = -\gamma \cdot \Delta \mathbb{E}_t[\Delta c_{t-1,t}], \quad (\text{IA.41})$$

where $-\log(\delta)$ is the real risk-free rate in a hypothetical economy without growth or uncertainty, and γ is the coefficient of relative risk aversion (RRA). Assuming the log normality of the consumption shock w_t^c , we obtain a closed-form solution for the real interest rate:

$$r_{f,t} = -\log(\delta) + \gamma \cdot \mathbb{E}_{t-1}[\Delta c_{t-1,t}] - \frac{\gamma^2 \text{var}_{t-1}(\Delta c_{t-1,t})}{2} = -\log(\delta) + \gamma \cdot (\mu_c + \sum_{j=1}^S \rho_j f_{t-j}) - \frac{\gamma^2 \text{var}_{t-1}(\Delta c_{t-1,t})}{2}, \quad (\text{IA.42})$$

where $\text{var}_{t-1}(\Delta c_{t-1,t}) = \sigma_{c,t-1}^2 + \rho_0^2 \sigma_{f,t-1}^2$.

Our MA representation of consumption growth implies that $\Delta \mathbb{E}_t[m_t] = -\gamma \rho_0 f_t - \gamma w_t^c$. As seen in the previous analysis, the contemporaneous covariance between consumption growth and f_t , ρ_0 , is small, which implies an unreasonably high RRA – the canonical equity premium puzzle. Parker and Julliard (2005) show that, iterating forward the intertemporal Euler equation, a valid log SDF can be constructed as

$$m_t^S = r_{f,t,t+S} - \gamma_S \Delta c_{t-1,t+S}. \quad (\text{IA.43})$$

¹⁷Note that only the shocks to the time- t SDF matter for the pricing of time- t excess returns.

Table IA.IV: Cross-sectional pricing ability of the shock to conditional consumption mean: Frequentist Fama-MacBeth regressions

		No intercept		With intercept	
		b_f	R^2	intercept	b_f
Panel A. 37 value-size-industry stock and nine bond portfolios					
No control for omitted factors:	coefficient	0.226	0.555	0.006	0.168
	<i>t</i> -stat	(3.355)		(2.679)	(2.314)
Control for omitted factors:	coefficient	0.251	0.794	0.002	0.226
	<i>t</i> -stat	(3.751)		(1.659)	(3.319)
Panel B. 74 Kozak, Nagel, and Santosh (2020) decile 1 and decile 10 anomaly portfolios					
No control for omitted factors:	coefficient	0.202	-0.316	0.020	-0.010
	<i>t</i> -stat	(2.746)		(3.647)	(-0.115)
Control for omitted factors:	coefficient	0.247	0.377	0.012	0.127
	<i>t</i> -stat	(3.702)		(1.995)	(1.400)
Panel C. 37 Kozak, Nagel, and Santosh (2020) decile 10 portfolios plus 37 long-short portfolios					
No control for omitted factors:	coefficient	0.231	0.542	-0.008	0.316
	<i>t</i> -stat	(3.265)		(-3.436)	(4.213)
Control for omitted factors:	coefficient	0.233	0.803	-0.008	0.312
	<i>t</i> -stat	(3.480)		(-6.638)	(4.567)
Panel D. 37 value-size-industry stock portfolios					
No control for omitted factors:	coefficient	0.238	-0.218	0.013	0.092
	<i>t</i> -stat	(3.340)		(2.372)	(0.993)
Control for omitted factors:	coefficient	0.227	0.359	0.027	-0.025
	<i>t</i> -stat	(3.359)		(4.300)	(-0.258)
Panel E. Transformed 37 value-size-industry stock portfolios					
No control for omitted factors:	coefficient	0.214	0.515	0.006	0.178
	<i>t</i> -stat	(2.883)		(3.248)	(2.290)
Control for omitted factors:	coefficient	0.283	0.882	0.004	0.217
	<i>t</i> -stat	(4.136)		(2.346)	(2.900)

Frequentist Fama-MacBeth estimation results for five cross-sections of excess returns: 37 value-size-industry stock and nine bond portfolios (Panel A), Kozak, Nagel, and Santosh (2020) 74 decile 1 and 10 characteristic-sorted portfolios (Panel B), Kozak, Nagel, and Santosh (2020) 37 decile 10 portfolios plus 37 long-short (decile 1 minus decile 10) portfolios (Panel C), 37 value-size-industry stock portfolios (Panel D), and transformed 37 value-size-industry stock portfolios (Panel E). In the transformed cross-section of 37 value-size-industry stock portfolios (Panel E), we long portfolios 1–19; for the remaining 18 portfolios, we construct the long-short portfolios by buying portfolios 20–37 but short-selling portfolios 1–18. We report: (1) the risk price of the shock to the conditional consumption mean f_t (b_f); (2) the cross-sectional R^2 ; and (3) the common pricing error (see column intercept) if it is included in the cross-sectional regression. The *t*-statistics are based on the Shanken (1992) corrected standard errors. In each cross-section of test assets, we report the results for the specification (1) with and without a common intercept in the cross-sectional step and (2) with and without the control for omitted factors.

Table IA.V: Individual pricing errors: 37 stock and nine bond portfolios

	No control for omitted factors		Control for omitted factors	
	No intercept	With intercept	No intercept	With intercept
0.5Y	0.001 (2.744)	-0.005 (-2.433)	0.001 (2.223)	-0.001 (-1.228)
1Y	0.002 (3.143)	-0.004 (-2.213)	0.001 (2.141)	-0.001 (-1.022)
2Y	0.002 (2.385)	-0.003 (-2.187)	0.001 (0.919)	-0.001 (-1.562)
3Y	0.003 (2.212)	-0.002 (-1.908)	0.000 (0.496)	-0.001 (-2.128)
4Y	0.004 (2.142)	-0.001 (-1.328)	0.000 (0.245)	-0.001 (-2.010)
5Y	0.005 (2.082)	-0.001 (-0.651)	0.000 (-0.001)	-0.001 (-1.418)
6Y	0.005 (2.009)	0.000 (-0.122)	0.000 (-0.285)	-0.001 (-1.072)
7Y	0.006 (1.923)	0.000 (0.209)	-0.001 (-0.593)	-0.001 (-0.929)
10Y	0.006 (1.644)	0.001 (0.552)	-0.002 (-1.281)	-0.001 (-0.849)
SizeValue11	-0.020 (-4.502)	-0.017 (-4.443)	-0.011 (-5.519)	-0.011 (-5.815)
SizeValue12	-0.002 (-0.658)	0.000 (-0.175)	0.002 (1.443)	0.002 (1.314)
SizeValue13	-0.001 (-0.207)	0.000 (0.150)	0.000 (0.230)	0.000 (0.082)
SizeValue14	0.007 (2.704)	0.008 (2.973)	0.007 (4.525)	0.006 (4.446)
SizeValue15	0.008 (2.558)	0.009 (2.974)	0.006 (3.837)	0.006 (3.810)
SizeValue21	-0.011 (-3.381)	-0.009 (-3.038)	-0.004 (-2.634)	-0.004 (-2.557)
SizeValue22	-0.001 (-0.373)	0.000 (0.116)	0.001 (0.473)	0.001 (0.566)
SizeValue23	0.004 (1.989)	0.004 (2.244)	0.002 (1.216)	0.002 (1.364)
SizeValue24	0.006 (2.604)	0.006 (2.635)	0.002 (1.522)	0.002 (1.593)
SizeValue25	0.006 (2.246)	0.007 (2.516)	0.003 (1.647)	0.003 (1.618)
SizeValue31	-0.008 (-2.656)	-0.007 (-2.372)	-0.002 (-1.236)	-0.002 (-1.132)
SizeValue32	0.001 (0.799)	0.002 (1.033)	0.001 (0.635)	0.001 (0.847)
SizeValue33	0.002 (1.393)	0.002 (1.291)	0.000 (0.009)	0.000 (0.101)
SizeValue34	0.006 (2.829)	0.006 (2.773)	0.002 (1.132)	0.002 (1.333)
SizeValue35	0.008 (2.689)	0.008 (2.741)	0.003 (1.346)	0.003 (1.466)
SizeValue41	-0.003 (-0.857)	-0.002 (-0.750)	0.003 (1.643)	0.003 (1.779)
SizeValue42	-0.002 (-0.884)	-0.002 (-0.942)	-0.003 (-1.585)	-0.002 (-1.464)
SizeValue43	0.002 (0.890)	0.001 (0.697)	-0.002 (-1.060)	-0.001 (-0.908)
SizeValue44	0.005 (2.502)	0.005 (2.379)	0.001 (0.732)	0.002 (0.985)
SizeValue45	0.005 (1.588)	0.005 (1.648)	0.000 (-0.053)	0.000 (0.035)
SizeValue51	0.000 (-0.058)	-0.001 (-0.464)	0.003 (2.628)	0.003 (2.569)
SizeValue52	0.000 (-0.061)	-0.002 (-0.699)	-0.001 (-0.596)	-0.001 (-0.619)
SizeValue53	0.003 (1.048)	0.001 (0.425)	0.000 (0.118)	0.000 (0.075)
SizeValue54	-0.001 (-0.338)	-0.002 (-0.763)	-0.005 (-2.631)	-0.005 (-2.714)
SizeValue55	0.002 (0.558)	0.001 (0.291)	0.000 (0.108)	0.000 (-0.063)
NonDurables	0.005 (1.597)	0.003 (1.224)	0.002 (0.901)	0.002 (1.081)
Durables	-0.008 (-1.940)	-0.008 (-1.893)	-0.008 (-2.527)	-0.009 (-2.889)
Manufact	-0.001 (-0.562)	-0.002 (-0.693)	-0.001 (-0.584)	-0.001 (-0.646)
Energy	0.005 (0.973)	0.002 (0.517)	0.004 (1.644)	0.003 (1.462)
Chemicals	0.001 (0.232)	-0.001 (-0.239)	-0.001 (-0.207)	-0.001 (-0.272)
BusEq	-0.001 (-0.184)	-0.001 (-0.184)	0.008 (2.555)	0.008 (2.508)
Telecom	0.002 (0.479)	0.000 (-0.062)	0.002 (0.503)	0.001 (0.413)
Utilities	0.003 (0.834)	0.001 (0.155)	-0.004 (-1.455)	-0.003 (-1.278)
Shops	0.001 (0.339)	0.001 (0.204)	0.001 (0.544)	0.001 (0.602)
Health	0.007 (1.649)	0.005 (1.276)	0.009 (2.860)	0.009 (2.896)
Money	-0.002 (-0.631)	-0.002 (-0.669)	-0.006 (-2.245)	-0.005 (-2.084)
Other	-0.007 (-3.791)	-0.007 (-3.726)	-0.007 (-3.810)	-0.007 (-3.758)
# of significance at 5% level	19	15	13	13
# of significance at 10% level	22	18	14	14

This table displays the individual pricing errors (with the t -statistics given in the parentheses) of 37 stock and nine bond portfolios based on the frequentist Fama-MacBeth regressions. The final two rows report the number of significant pricing errors at the 5% or 10% significance level. We report the results for the specification (1) with and without a common intercept in the cross-sectional step and (2) with and without the control for omitted factors.

Table IA.VI: Individual pricing errors: 74 KNS anomaly portfolios

	No control for omitted factors		Control for omitted factors	
	No intercept	With intercept	No intercept	With intercept
accruals_p1	-0.009 (-3.462)	-0.007 (-2.832)	-0.006 (-2.487)	-0.006 (-2.246)
accruals_p10	0.005 (1.766)	0.003 (1.270)	0.008 (3.088)	0.006 (2.688)
aturnover_p1	-0.003 (-1.031)	-0.009 (-3.099)	0.000 (0.000)	-0.003 (-1.232)
aturnover_p10	0.005 (1.845)	0.003 (1.186)	-0.001 (-0.207)	0.000 (0.028)
betaarb_p1	-0.009 (-2.858)	-0.003 (-0.897)	-0.004 (-1.310)	-0.001 (-0.567)
betaarb_p10	0.007 (2.408)	-0.002 (-0.556)	0.001 (0.347)	-0.002 (-1.059)
cfp_p1	-0.002 (-0.650)	-0.003 (-1.333)	-0.003 (-2.032)	-0.002 (-1.512)
cfp_p10	0.008 (2.295)	0.006 (1.921)	0.001 (0.537)	0.001 (0.590)
ciss_p1	-0.009 (-4.150)	-0.009 (-4.240)	-0.009 (-3.997)	-0.008 (-4.025)
ciss_p10	0.011 (4.294)	0.005 (2.131)	0.006 (2.772)	0.005 (2.502)
nissa_p1	-0.010 (-4.960)	-0.011 (-5.729)	-0.009 (-4.901)	-0.008 (-4.711)
nissa_p10	0.013 (5.186)	0.008 (3.616)	0.006 (2.837)	0.005 (2.739)
divp_p1	-0.001 (-0.565)	-0.003 (-1.209)	-0.004 (-1.771)	-0.002 (-1.418)
divp_p10	0.005 (1.025)	0.001 (0.242)	-0.004 (-1.081)	-0.004 (-1.096)
dur_p1	-0.001 (-0.374)	-0.004 (-1.509)	-0.001 (-1.255)	-0.001 (-1.127)
dur_p10	0.007 (1.698)	0.007 (1.609)	0.003 (0.845)	0.002 (0.552)
ep_p1	-0.011 (-3.165)	-0.006 (-1.832)	-0.007 (-2.204)	-0.006 (-2.118)
ep_p10	0.010 (2.852)	0.008 (2.361)	0.000 (0.060)	0.001 (0.506)
gmargins_p1	0.000 (-0.222)	-0.002 (-0.930)	-0.003 (-1.406)	-0.003 (-1.489)
gmargins_p10	0.004 (1.484)	-0.001 (-0.359)	0.003 (1.915)	0.001 (1.128)
growth_p1	-0.004 (-1.574)	-0.004 (-1.686)	0.000 (-0.039)	0.001 (0.280)
growth_p10	0.003 (1.249)	0.003 (1.141)	0.000 (-0.003)	0.000 (-0.035)
igrowth_p1	-0.007 (-2.632)	-0.006 (-2.426)	-0.003 (-1.411)	-0.003 (-1.424)
igrowth_p10	0.002 (1.093)	0.003 (1.665)	0.001 (0.379)	0.001 (0.426)
indmom_p1	-0.008 (-1.931)	-0.005 (-1.308)	0.003 (0.868)	0.003 (0.918)
indmom_p10	0.013 (2.963)	0.006 (1.690)	-0.001 (-0.211)	-0.001 (-0.240)
indrrev_p1	-0.013 (-4.879)	-0.015 (-5.836)	-0.016 (-6.553)	-0.016 (-7.006)
indrrev_p10	0.009 (3.560)	0.014 (6.200)	0.015 (6.165)	0.016 (6.926)
inv_p1	-0.004 (-1.485)	-0.005 (-1.937)	-0.003 (-1.090)	-0.002 (-0.666)
inv_p10	0.009 (3.837)	0.007 (3.065)	0.005 (2.189)	0.005 (2.079)
invcap_p1	-0.001 (-0.174)	0.001 (0.297)	0.005 (2.118)	0.004 (1.907)
invcap_p10	0.007 (2.324)	0.002 (0.571)	0.003 (1.475)	0.003 (1.244)
lev_p1	0.001 (0.188)	-0.003 (-1.092)	0.000 (-0.650)	-0.001 (-1.226)
lev_p10	0.003 (0.687)	0.004 (0.922)	-0.007 (-1.824)	-0.003 (-1.025)
lrrev_p1	-0.001 (-0.193)	-0.002 (-0.701)	0.000 (0.071)	0.002 (0.780)
lrrev_p10	0.002 (0.469)	0.008 (2.000)	-0.003 (-0.929)	-0.003 (-0.947)
mom_p1	-0.005 (-1.218)	0.000 (0.145)	0.008 (3.064)	0.008 (3.383)
mom_p10	0.015 (3.796)	0.011 (3.189)	0.006 (3.521)	0.007 (3.842)
mom12_p1	-0.029 (-6.420)	-0.021 (-6.021)	-0.014 (-4.571)	-0.013 (-4.603)
mom12_p10	0.018 (4.940)	0.014 (4.414)	0.012 (5.738)	0.012 (5.710)
momrev_p1	-0.002 (-0.721)	-0.005 (-1.436)	0.000 (0.144)	0.000 (0.017)
momrev_p10	0.004 (1.091)	0.008 (2.368)	0.001 (0.270)	0.002 (0.604)
noa_p1	-0.007 (-2.431)	-0.009 (-3.440)	-0.002 (-0.832)	-0.002 (-1.097)
noa_p10	0.004 (1.582)	0.003 (1.185)	0.004 (1.863)	0.004 (1.732)
price_p1	-0.013 (-2.368)	-0.001 (-0.256)	-0.003 (-1.096)	-0.004 (-1.666)
price_p10	0.003 (1.452)	-0.003 (-2.332)	0.000 (-0.044)	-0.001 (-1.131)

Table IA.VI: Individual pricing errors: 74 KNS anomaly portfolios

	No control for omitted factors		Control for omitted factors	
	No intercept	With intercept	No intercept	With intercept
prof_p1	-0.001 (-0.396)	-0.007 (-2.629)	-0.001 (-0.531)	-0.004 (-1.758)
prof_p10	0.008 (2.790)	0.005 (1.885)	0.002 (1.476)	0.003 (1.720)
roaa_p1	-0.009 (-2.565)	-0.005 (-1.484)	-0.004 (-1.393)	-0.005 (-1.814)
roaa_p10	0.004 (1.606)	0.000 (-0.035)	0.001 (0.839)	0.000 (0.200)
roea_p1	-0.006 (-1.718)	0.000 (-0.159)	-0.001 (-0.357)	-0.002 (-0.646)
roea_p10	0.004 (1.727)	0.001 (0.368)	0.000 (0.086)	0.001 (0.419)
season_p1	-0.011 (-3.763)	-0.011 (-3.801)	-0.015 (-4.990)	-0.014 (-5.051)
season_p10	0.013 (4.970)	0.012 (4.824)	0.011 (4.588)	0.012 (5.017)
sgrowth_p1	0.002 (0.745)	0.001 (0.458)	0.006 (3.270)	0.006 (3.312)
sgrowth_p10	0.000 (0.032)	-0.002 (-0.686)	-0.001 (-0.412)	-0.002 (-0.980)
shvol_p1	-0.006 (-1.644)	-0.001 (-0.297)	0.001 (0.615)	0.002 (1.079)
shvol_p10	0.004 (1.327)	-0.004 (-1.595)	-0.001 (-0.688)	-0.004 (-1.690)
size_p1	0.002 (0.869)	-0.004 (-3.096)	-0.001 (-1.089)	-0.002 (-2.221)
size_p10	0.001 (0.193)	0.005 (1.359)	-0.002 (-0.840)	-0.003 (-1.023)
sp_p1	-0.001 (-0.427)	-0.004 (-2.053)	-0.001 (-1.179)	-0.002 (-1.938)
sp_p10	0.007 (1.762)	0.010 (2.605)	-0.001 (-0.254)	0.001 (0.473)
strev_p1	-0.007 (-2.364)	-0.010 (-3.915)	-0.011 (-4.064)	-0.012 (-4.685)
strev_p10	-0.002 (-0.656)	0.003 (0.849)	0.006 (1.998)	0.007 (2.383)
valmom_p1	-0.005 (-1.622)	-0.005 (-1.667)	0.002 (1.050)	0.002 (1.106)
valmom_p10	0.012 (3.671)	0.008 (2.468)	0.002 (0.819)	0.001 (0.605)
valmomprof_p1	-0.007 (-2.413)	-0.007 (-2.487)	0.002 (0.972)	0.002 (0.922)
valmomprof_p10	0.018 (5.901)	0.016 (5.275)	0.007 (3.397)	0.007 (3.327)
valprof_p1	-0.006 (-2.137)	-0.009 (-3.855)	-0.002 (-1.144)	-0.002 (-1.321)
valprof_p10	0.011 (3.019)	0.012 (3.377)	0.003 (1.360)	0.004 (1.541)
value_p1	0.000 (-0.005)	-0.003 (-1.291)	-0.001 (-0.770)	-0.001 (-0.782)
value_p10	0.008 (2.040)	0.007 (1.959)	0.002 (0.679)	0.001 (0.535)
valuem_p1	0.002 (0.763)	-0.001 (-0.579)	0.000 (-0.340)	0.000 (-0.202)
valuem_p10	0.004 (0.828)	0.009 (1.956)	0.013 (4.010)	0.011 (3.918)
# of significance at 5% level	32	30	23	22
# of significance at 10% level	39	40	27	30

This table displays the individual pricing errors (with the t -statistics given in the parentheses) of 74 Kozak, Nagel, and Santosh (2020) anomaly portfolios based on the frequentist Fama-MacBeth regressions. The final two rows report the number of significant pricing errors at the 5% or 10% significance level. We report the results for the specification (1) with and without a common intercept in the cross-sectional step and (2) with and without the control for omitted factors.

When pricing the time- t *excess* returns, only the time- t shocks to the SDF matters, as follows:

$$\begin{aligned}\Delta \mathbb{E}_t[m_t^S] &= (r_{f,t,t+S} - \mathbb{E}_{t-1}[r_{f,t,t+S}]) - \gamma_S (\mathbb{E}_t[\Delta c_{t-1,t+S}] - \mathbb{E}_{t-1}[\Delta c_{t-1,t+S}]) \\ &= (r_{f,t,t+S} - \mathbb{E}_{t-1}[r_{f,t,t+S}]) - \gamma_S \left(\sum_{j=0}^S \rho_j \right) f_t - \gamma_S w_t^c.\end{aligned}$$

The first term in the above equation is the shock to the long-term risk-free rate, which is relatively small compared to the f_t shock. w_t^c is orthogonal to excess returns and, therefore, unpriced. Hence, the SDF representation in equation (20) implies that we can elicit the

RRA coefficient as follows:

$$\gamma_S \approx \frac{b_f}{\sum_{j=0}^S \rho_j}, \quad (\text{IA.44})$$

where b_f is the risk price of the f shock. An extremely high risk price, or a low cumulative impulse response in consumption growth, lead to a extremely high RRA. Moreover, a high RRA implies high mean and volatility of the risk-free rate in equation (IA.42).

The second setting that we consider is the Epstein-Zin recursive utility,

$$V_t = \left[(1 - \delta) C_t^{\frac{1-\gamma}{\theta}} + \delta \left(\mathbb{E}_t [V_{t+1}^{1-\gamma}] \right)^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \quad (\text{IA.45})$$

where $0 < \delta < 1$ is the time preference parameter, ψ is the elasticity of intertemporal substitution (IES), and $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$. The intertemporal budget constraint of the representative agent is

$$W_{t+1} = (1 + R_{w,t+1})(W_t - C_t), \quad (\text{IA.46})$$

where W_{t+1} denotes the total wealth and $R_{w,t+1}$ is the gross return on the wealth portfolio. The log SDF is then

$$m_t = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t-1,t} - (1 - \theta) r_{w,t}, \text{ where } r_{w,t} = \log(1 + R_{w,t}). \quad (\text{IA.47})$$

Assuming jointly lognormal $r_{w,t}$ and $\Delta c_{t-1,t}$, we obtain the risk-free rate as

$$r_{f,t} = \begin{cases} -\log \delta + \frac{1}{\psi} \mathbb{E}_{t-1} [\Delta c_{t-1,t}] - \frac{\theta}{2\psi^2} \text{var}_{t-1} (\Delta c_{t-1,t}) + \frac{\theta-1}{2} \text{var}_{t-1} (r_{w,t}), & \text{if } \psi \neq 1, \\ -\log \delta + \mathbb{E}_{t-1} [\Delta c_{t-1,t}] - (\gamma - \frac{1}{2}) \text{var}_{t-1} (\Delta c_{t-1,t}), & \text{if } \psi = 1, \end{cases} \quad (\text{IA.48})$$

where $\text{var}_{t-1} (r_{w,t})$ is the variance of the shock to $r_{w,t}$. Note that $\text{var}_{t-1} (r_{w,t})$ and $\text{var}_{t-1} (\Delta c_{t-1,t})$ are closely related because of the budget constraint in equation (IA.46). Using the log-linear return approximation of Campbell (1993) and the state-space formulation, we have

$$\Delta \mathbb{E}_t [r_{w,t}] = \Delta \mathbb{E}_t [\Delta c_{t-1,t}] + \left(1 - \frac{1}{\psi}\right) \Delta \mathbb{E}_t \sum_{j=1}^{\infty} \kappa_z^j \Delta c_{t+j-1,t+j} \quad (\text{IA.49})$$

$$= \left[\rho_0 + \left(1 - \frac{1}{\psi}\right) \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j \right] f_t + u_t^c, \quad (\text{IA.50})$$

which implies

$$\text{var}_{t-1}(r_{w,t}) = \left[\rho_0 + \left(1 - \frac{1}{\psi}\right) \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j \right]^2 \sigma_{f,t-1}^2 + \sigma_{c,t-1}^2, \quad (\text{IA.51})$$

where κ_z is a constant determined by the steady-state consumption-to-wealth ratio. In Bansal and Yaron (2004), $\kappa_z \approx 0.997$.

Using the same logic, we can express the shock to the log SDF as follows:

$$\Delta \mathbb{E}_t[m_t] = m_t - \mathbb{E}_{t-1}[m_t] = -\gamma \cdot \Delta \mathbb{E}_t[\Delta c_{t-1,t}] - \left(\gamma - \frac{1}{\psi}\right) \cdot \Delta \mathbb{E}_t \sum_{j=1}^{\infty} \kappa_z^j \Delta c_{t+j-1,t+j}. \quad (\text{IA.52})$$

When the inverse IES equals RRA, equation (IA.52) can be simplified to the CRRA preference in equation (IA.41). Replacing the MA representation of $\Delta c_{t-1,t}$ in equation (IA.52), and ignoring the unpriced component $-\gamma w_t^c$, yields

$$\Delta \mathbb{E}_t[m_t] = -\gamma \rho_0 f_t - \left(\gamma - \frac{1}{\psi}\right) \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j f_t. \quad (\text{IA.53})$$


Finally, fixing the value of IES, and matching the above to the risk price of f_t in equation (20), we can reverse engineer γ as

$$\gamma \rho_0 + \left(\gamma - \frac{1}{\psi}\right) \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j = b_f \implies \gamma = \frac{b_f + \frac{1}{\psi} \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j}{\rho_0 + \sum_{j=1}^{\bar{S}} \kappa_z^j \rho_j}. \quad (\text{IA.54})$$

Using the estimated impulse responses $\{\rho_j\}_{j=0}^{\bar{S}}$, we can infer the risk price of the conditional consumption mean shock given the values of γ and ψ . We can further use these calibrated risk prices to infer the equity risk premium, which equals $\text{cov}(b_f f_t, r_t^{mkt})$. For example, assuming that $\gamma = 10$, $\psi = 1.5$, and there is no stochastic volatility in consumption, the implied annualized equity risk premium is about 2.6% and 4.1%, respectively, in the baseline cross-section of 46 assets and KNS anomalies. If we consider the estimates with stochastic volatilities, the implied annualized equity risk premium is 3.2% and 4.4% in those two cross-sections.


P Additional figures

Figure IA.19: Variance decomposition of asset returns (average of 1,000 simulations).




The figure shows the 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 hypothetical true calibrated values.

Figure IA.20: Share of consumption growth variance driven by its moving average component.




The figure shows the box plots (posterior 95% coverage area) of the percentage of time-series variances of consumption growth explained by the MA component. These plots report the adjusted R-squared, $R_{adj}^2 = 1 - \frac{T-1}{T-1-df} (1 - R^2)$, where df denotes the degree of freedom, which equals 15 in this model. Left panel: Cumulated consumption growth $\Delta c_{t-1, t+S}$. Right panel: One-period consumption growth $\Delta c_{t-1+j, t+j}$. We study a single-factor model in equations (2) and (3), with $\bar{S} = 14$. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios.

Figure IA.21: Consumption growth response to the latent factors f_t and b_t shocks.




The figure shows the posterior means (continuous solid line) and centered posterior 90% (dashed line) and 68% (dotted line) coverage regions. The estimation is based on the two-factor model in equations (5) and (6). Left panel: Cumulated consumption response to common factor (f_t) shocks. Right panel: Cumulated consumption response to bond factor (b_t) shocks.

Figure IA.22: Variance of consumption growth explained by the MA components f and b .




The figure shows the box plots (posterior 95% coverage area) of the percentage of time-series variances of consumption growth explained by the MA components f and b . These plots report the adjusted R-squared, $R_{adj}^2 = 1 - \frac{T-1}{T-1-df} (1 - R^2)$, where df denotes the degree of freedom, which equals 30 in this model. Left panel: Cumulated consumption growth $\Delta c_{t-1,t+S}$. Right panel: One-period consumption growth $\Delta c_{t+j-1,t+j}$.

Figure IA.23: 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 (circles) and centered posterior 90% (dashed line) and 68% (dotted line) coverage regions 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 IA.24: Share of stock portfolios' return variance explained by the f component in the two-factor model.




The figure shows the 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 IA.25: Box plots of γ^c , γ^s , and γ^b in equations (13)–(15).



This figure plots the distribution of the predictors' coefficient estimates in equations (13)–(15).

Figure IA.26: Box plots of cay .



The two-factor model is as follows:


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

Stocks: $r_{st}^e = \mu_{sr} + \rho^{sr} f_t + \gamma^s cay_{t-1} + w_t^{sr}$, and

Bonds: $r_{bt}^e = \mu_{br} + \rho^{br} f_t + \theta^b b_t + \gamma^b cay_{t-1} + w_t^{br}$,


where cay_{t-1} comes from Lettau and Ludvigson (2001).

Figure IA.27: Plots of the two-factor model with P/D ratio and CFNAI in Section III.4.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We consider a two-factor model for asset returns, including the market's P/D ratio and the CFNAI index as the predictors for stock and bond portfolios, respectively. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Details can be found in Section III.4.

Figure IA.28: Plots of the two-factor model with *cay* in Section III.4.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We consider a two-factor model for asset returns, including *cay* in Lettau and Ludvigson (2001) as the predictors for stock and bond portfolios, respectively. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Details can be found in Section III.4.

Figure IA.29: Plots of a five-factor model in Section III.5.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We consider a five-factor model for asset returns. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Details can be found in Section III.5.

Figure IA.30: Plots under the stochastic volatility assumption in Section III.6.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We assume different stochastic volatility processes for f_t , w_t^c , and w_t^r . The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Details can be found in Section III.6.

Figure IA.31: Robustness check: 32 size-profitability-investment-sorted portfolios, 12 industry portfolios, and nine bond portfolios in a five-factor model.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We study the cross-section of 32 size-profitability-investment-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Moreover, asset returns are modeled using a five-factor model, with $\bar{S} = 14$.

Figure IA.32: Robustness check: 32 size-value-investment-sorted portfolios, 12 industry portfolios, and nine bond portfolios in a five-factor model.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We study the cross-section of 32 size-value-investment-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Moreover, asset returns are modeled using a five-factor model, with $\bar{S} = 14$.

Figure IA.33: Robustness check: 32 size-value-profitability-sorted portfolios, 12 industry portfolios, and nine bond portfolios in a five-factor model.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We study the cross-section of 32 size-value-profitability-sorted portfolios, 12 industry portfolios, and nine bond portfolios. Moreover, asset returns are modeled using a five-factor model, with $\bar{S} = 14$.

Figure IA.34: Robustness check: Kozak, Nagel, and Santosh (2020) 74 portfolios in a five-factor model.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) variance decomposition of $\Delta C_{t-1,t+S}$, and (e) variance decomposition of $\Delta C_{t+j-1,t+j}$. We study the cross-section of 74 characteristic-sorted portfolios used in Kozak, Nagel, and Santosh (2020). Specifically, they create value-weighted decile portfolios sorted by firm characteristics, and we use the long and short legs (those in deciles 1 and 10) that have data since July 1963. Hence, we end up with 74 portfolios (both long and short legs of 37 characteristics). Moreover, asset returns are modeled using a five-factor model, with $\bar{S} = 14$.

Figure IA.35: Term structure of consumption exposure.




The figure shows the spreads of consumption betas measured as $\text{Cov}(\Delta c_{t,t+1+S}, r_{j,t+1}^e)$ for different horizon S (0-14) and asset j : nine bonds (circle), 25 Fama-French size and book-to-market (triangle), and 12 industry (cross) portfolios. The asset exposures are in basis point units.

Figure IA.36: Implied coefficients of relative risk aversion.




The figure reports the point estimates of the coefficients of relative risk aversion (RRA) implied by our MA(14) state-space formulation and the VAR(4) in Parker (2001). We consider the additively separable CRRA preferences. Details can be found in the footnote of Table 7 and Internet Appendices L and O.


Figure IA.37: Cross-sectional spread of exposure to slow consumption adjustment risk.



(a) 32 size-profitability-investment-sorted portfolios




(b) 32 size-value-investment-sorted portfolios



(c) 32 size-value-profitability-sorted portfolios


The figure presents the spread of consumption betas measured as $Cov(\Delta c_{t,t+1+S}, r_{j,t+1}^e)$ for different horizon S (0-14) and asset j : (a) 32 size-profitability-investment-sorted portfolios, (b) 32 size-value-investment-sorted portfolios, and (c) 32 size-value-profitability-sorted portfolios. The asset exposures are in basis point units.

Figure IA.38: Robustness check: Modeling the growth rate in nondurable consumption goods plus service in a single-factor model of asset returns.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We consider a single-factor model in equations (2) and (3), but $\Delta C_{t-1,t}$ is the quarterly growth rate in nondurable consumption goods and service. Moreover, asset returns are modeled using a single-factor model, with $\bar{S} = 14$. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios.

Figure IA.39: Robustness check: Modeling the growth rate in nondurable consumption goods plus service in a five-factor model of asset returns.




The figure plots (a) cumulative impulse response function of consumption growth ($\bar{S} = 14$), (b) moving average component of consumption growth, (c) stock loadings on f_t , (d) bond loadings on f_t , (e) variance decomposition of $\Delta C_{t-1,t+S}$, and (f) variance decomposition of $\Delta C_{t+j-1,t+j}$. We consider a single-factor model in equations (2) and (3), but $\Delta C_{t-1,t}$ is the quarterly growth rate in nondurable consumption goods and service. Moreover, asset returns are modeled using a five-factor model, with $\bar{S} = 14$. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios.

Figure IA.40: Filtered SVs with Kozak, Nagel, and Santosh (2020) 74 portfolios.




The figure shows the estimated stochastic volatilities of the model in equations (2) and (16)–(19), and Section B.2 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 (patterned) areas reflect constant volatility levels that would not be rejected given the credible intervals. We study the cross-section of 74 characteristic-sorted portfolios used in Kozak, Nagel, and Santosh (2020). Specifically, they create value-weighted decile portfolios sorted by firm characteristics, and we use the long and short legs (those in deciles 1 and 10) that have data since July 1963. Hence, we end up with 74 portfolios (both long and short legs of 37 characteristics).


Figure IA.41: Loadings of excess returns on consumption and returns volatilities in Kozak, Nagel, and Santosh (2020) 74 portfolios.



Panel A: Posterior distribution of excess return loadings (β_c) on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$) in equation (16).




Panel B: 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 (16).




Panel C: Posterior distributions of excess return loadings (β_r) on the common financial return variance ($\sigma_{r,t-1}^2$) in equation (16).

The figure shows the box plots of the posterior distributions of the loadings of portfolio excess returns on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$), the variance of shocks to the conditional consumption growth ($\sigma_{f,t-1}^2$), and the common financial returns variance ($\sigma_{r,t-1}^2$). We study the cross-section of 74 characteristic-sorted portfolios used in Kozak, Nagel, and Santosh (2020). Specifically, they create value-weighted decile portfolios sorted by firm characteristics, and we use the long and short legs (those in deciles 1 and 10) that have data since July 1963. Hence, we end up with 74 portfolios (both long and short legs of 37 characteristics).


Figure IA.42: Loadings of excess returns on consumption and returns volatilitie: Predicting cumulative excess returns over the next *four quarters*.



Panel A: Posterior distribution of excess return loadings on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$).




Panel B: Posterior distribution of excess return loadings on the variance of shocks to the conditional consumption growth mean ($\sigma_{f,t-1}^2$).




Panel C: Posterior distributions of excess return loadings on the common financial return variance ($\sigma_{r,t-1}^2$).

The figure shows the box plots of the posterior distributions of the loadings of portfolio excess returns on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$), the variance of shocks to the conditional consumption growth ($\sigma_{f,t-1}^2$), and the common financial returns variance ($\sigma_{r,t-1}^2$). Portfolios are ordered with bonds first (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last. f_t follows the stochastic volatility process with the leverage effect (see equations (IA.25)–(IA.26)).


Figure IA.43: Loadings of excess returns on consumption and returns volatilitie: Predicting cumulative excess returns over the next *eight quarters*.



Panel A: Posterior distribution of excess return loadings on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$).




Panel B: Posterior distribution of excess return loadings on the variance of shocks to the conditional consumption growth mean ($\sigma_{f,t-1}^2$).




Panel C: Posterior distributions of excess return loadings on the common financial return variance ($\sigma_{r,t-1}^2$).

This figure shows the box plots of the posterior distributions of the loadings of portfolio excess returns on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$), the variance of shocks to the conditional consumption growth ($\sigma_{f,t-1}^2$), and the common financial returns variance ($\sigma_{r,t-1}^2$). Portfolios are ordered with bonds first (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last. f_t follows the stochastic volatility process with the leverage effect (see equations (IA.25)–(IA.26)).


Figure IA.44: Loadings of excess returns on consumption and returns volatilitie: Predicting cumulative excess returns over the next 12 quarters.



Panel A: Posterior distribution of excess return loadings on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$).




Panel B: Posterior distribution of excess return loadings on the variance of shocks to the conditional consumption growth mean ($\sigma_{f,t-1}^2$).



Panel C: Posterior distributions of excess return loadings on the common financial return variance ($\sigma_{r,t-1}^2$).


The figure shows the box plots of the posterior distributions of the loadings of portfolio excess returns on the variance of short-run consumption shocks ($\sigma_{c,t-1}^2$), the variance of shocks to the conditional consumption growth ($\sigma_{f,t-1}^2$), and the common financial returns variance ($\sigma_{r,t-1}^2$). Portfolios are ordered with bonds first (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last. f_t follows the stochastic volatility process with the leverage effect (see equations (IA.25)–(IA.26)).

Figure IA.45: Loadings on the common financial market volatility process.



This figure shows the box plots of the posterior distributions of the loadings (κ_1^i) of asset volatilities on the common financial market volatility process. Portfolios are ordered with bonds first (1–9), Fama-French 25 size and book-to-market second (10–34), and industry portfolios last.

Figure IA.46: Share of stock portfolios' return variance explained by the f component in five-factor models.



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 five-factor model for alternative macro variables.

Q Additional tables

Table IA.VII: Predicting nondurable consumption growth using SPF forecasts

$j =$	1	2	3	4
Panel A. Only SPF Forecast				
$\mathbb{E}_t^{SPF}[\Delta c_{t+j-1,t+j}]$	0.957	1.305	0.465	1.029
s.e. (OLS)	(0.224)	(0.315)	(0.373)	(0.425)
s.e. (NW, lag=12)	(0.221)	(0.294)	(0.472)	(0.397)
Predictive R^2	0.117	0.111	0.011	0.041
Panel B. MA Model vs. SPF Forecast				
$\mathbb{E}_t[\Delta c_{t+j-1,t+j}]$	0.580	0.622	0.856	0.829
s.e. (OLS)	(0.165)	(0.213)	(0.218)	(0.214)
s.e. (NW, lag=12)	(0.140)	(0.192)	(0.194)	(0.209)
$\mathbb{E}_t^{SPF}[\Delta c_{t+j-1,t+j}]$	0.718	0.980	0.297	0.935
s.e. (OLS)	(0.225)	(0.326)	(0.357)	(0.406)
s.e. (NW, lag=12)	(0.157)	(0.277)	(0.403)	(0.366)
Predictive R^2	0.190	0.163	0.112	0.135

The table summarizes the regressions in which future realized growth rates of nondurable consumption are forecasted by several predictors. In Panel A, we regress the future one-period nondurable consumption growth on the mean SPF expected consumption growth $\mathbb{E}_t^{SPF}[\Delta c_{t+j-1,t+j}]$, where j ranges from one to four quarters. In Panel B, we further include the conditional consumption mean implied by our MA model (see equation (9)). We report (1) the point estimates of the slope coefficients, (2) the OLS and Newey-West (12 lags) standard errors within the parentheses, and (3) the predictive R^2 . Sample: 1981:Q3–2019:Q4.

Table IA.VIII: Validating the predictability of the conditional consumption mean

$j =$	1	2	3	4	5	6	7	8	9	10	11	12
Panel A. Predicting nondurable consumption growth												
$\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$	0.997	0.964	0.970	0.973	0.946	0.921	0.974	1.119	1.027	0.971	1.039	1.025
s.e. (OLS)	(0.149)	(0.200)	(0.206)	(0.206)	(0.276)	(0.291)	(0.318)	(0.358)	(0.416)	(0.431)	(0.571)	(0.572)
s.e. (NW, lag=12)	(0.213)	(0.212)	(0.211)	(0.217)	(0.203)	(0.219)	(0.185)	(0.202)	(0.314)	(0.343)	(0.446)	(0.446)
$HNIt$	0.167	0.161	0.153	0.139	0.149	0.155	0.158	0.168	0.145	0.137	0.133	0.127
s.e. (OLS)	(0.070)	(0.073)	(0.072)	(0.072)	(0.074)	(0.074)	(0.074)	(0.074)	(0.071)	(0.072)	(0.072)	(0.072)
s.e. (NW, lag=12)	(0.081)	(0.094)	(0.095)	(0.101)	(0.116)	(0.117)	(0.118)	(0.119)	(0.112)	(0.108)	(0.105)	(0.103)
Predictive R^2	0.201	0.123	0.117	0.115	0.073	0.068	0.065	0.070	0.048	0.042	0.033	0.031
Panel B. Predicting nondurable plus service consumption growth												
$\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$	1.002	0.893	0.831	0.820	0.678	0.635	0.604	0.517	0.542	0.550	0.469	0.376
s.e. (OLS)	(0.124)	(0.144)	(0.147)	(0.147)	(0.183)	(0.186)	(0.193)	(0.228)	(0.231)	(0.252)	(0.364)	(0.515)
s.e. (NW, lag=12)	(0.132)	(0.123)	(0.140)	(0.160)	(0.226)	(0.251)	(0.237)	(0.274)	(0.296)	(0.301)	(0.406)	(0.578)
$HNIt$	0.216	0.208	0.202	0.197	0.210	0.219	0.225	0.231	0.220	0.219	0.221	0.223
s.e. (OLS)	(0.039)	(0.041)	(0.042)	(0.042)	(0.043)	(0.043)	(0.043)	(0.043)	(0.042)	(0.042)	(0.042)	(0.043)
s.e. (NW, lag=12)	(0.049)	(0.048)	(0.047)	(0.049)	(0.054)	(0.053)	(0.053)	(0.054)	(0.054)	(0.055)	(0.056)	(0.056)
Predictive R^2	0.341	0.256	0.230	0.223	0.165	0.165	0.163	0.149	0.147	0.142	0.130	0.125

The table summarizes the regressions in which future realized growth rates of consumption ($\Delta c_{t+j-1,t+j}$, $1 \leq j \leq 12$) are forecasted by several predictors. We regress $\Delta c_{t+j-1,t+j}$ on the conditional consumption mean ($\widehat{\mathbb{E}}_t[\Delta c_{t+j-1,t+j}]$) implied by the MA model and the HN-index in Liu and Matthies (2022) (one-sided HP filter of the news index). Panels A and B report the results for nondurable consumption and nondurable plus service consumption growth. We report (1) the point estimates of the slope coefficients, (2) the OLS and Newey-West (12 lags) standard errors within the parentheses, and (3) the predictive R^2 . Sample: 1963:Q3–2013:Q4. The sample ends at 2013:Q4 due to the availability of the HN-index data.

Table IA.IX: Predicting squared forecast errors of consumption growth

	Market P/D ratio	Financial uncertainty	Macro uncertainty	Real uncertainty
Coefficient	-0.007	0.123	0.069	0.134
s.e. (OLS)	(0.069)	(0.068)	(0.069)	(0.068)
s.e. (NW, lag=12)	(0.068)	(0.080)	(0.079)	(0.051)
R^2	0.000	0.015	0.005	0.018

We regress the squared forecast error, $(\Delta c_{t,t+1} - \hat{\mathbb{E}}_t[\Delta c_{t,t+1}])^2$, on several persistent economic variables measured at time t , including the market price-dividend ratio and financial/real/macro uncertainty measures from Jurado, Ludvigson, and Ng (2015) and Ludvigson, Ma, and Ng (2021). All variables are standardize to have unit variances in the regressions. The uncertainty measures have horizons of three months, consistent with the quarterly frequency of consumption growth.

Table IA.X: Model comparison using log Bayes factors with scaled consumption growth and asset returns: $\text{var}(r_{it}^e) = 1$ for $i = 1, \dots, N$, and $\text{var}(\Delta c_{t-1,t}) = N$

	Models									
	I	II	III	IV	V	VI	VII	VIII	IX	X
Log of Bayes Factor:	0	-156	-38	-49	-87	-233	-165	-439	-234	-551
Posterior Probability.:	100%	0%	0%	0%	0%	0%	0%	0%	0%	0%

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

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

The table summarizes the model comparison for restricted and unrestricted versions of the specification in equations (2) and (16)–(19), and Section B.2. We approximate the Bayes factor using the Schwartz criterion. We use Model I as a 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. Different from Table IA.X, we rescale the variance of consumption growth and asset returns: $\text{var}(r_{it}^e) = 1$ for $i = 1, \dots, N$, and $\text{var}(\Delta c_{t-1,t}) = N$. Under such a normalization, we ensure that we put equal weights to fit consumption growth and asset returns equations.

Table IA.XI: Correlations of stochastic volatility processes and VXO: Kozak, Nagel, and Santosh (2020) 74 portfolios

	Mean	2.5%	5%	50%	95%	97.5%
Panel A: correlations of vol processes with VXO ² index						
$cor(\sigma_{ct}^2, VXO_t^2)$	0.10	-0.12	-0.09	0.09	0.31	0.35
$cor(\sigma_{ft}^2, VXO_t^2)$	0.43	0.22	0.26	0.44	0.58	0.60
$cor(\sigma_{rt}^2, VXO_t^2)$	0.53	0.41	0.44	0.53	0.63	0.65
Panel B: Pairwise correlations of vol processes						
$cor(\sigma_{ct}^2, \sigma_{ft}^2)$	0.07	-0.11	-0.08	0.05	0.27	0.31
$cor(\sigma_{ct}^2, \sigma_{rt}^2)$	0.08	-0.09	-0.07	0.06	0.26	0.31
$cor(\sigma_{ft}^2, \sigma_{rt}^2)$	0.30	0.12	0.14	0.30	0.48	0.52

The table summarises the 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 volatility processes. We study the cross-section of 74 characteristic-sorted portfolios in Kozak, Nagel, and Santosh (2020). Specifically, they create value-weighted decile portfolios sorted by firm characteristics, and we use the long and short legs (those in deciles 1 and 10) that have data since July 1963. Hence, we end up with 74 portfolios (both long and short legs of 37 characteristics).

Table IA.XII: Correlations of stochastic volatility processes and real/macro/financial uncertainty

	Posterior Mean	2.5%	5%	50%	95%	97.5%
Panel A: correlations of vol processes with real uncertainty						
$cor(\sigma_{ct}^2, real_t^2)$	0.22	0.02	0.04	0.22	0.42	0.45
$cor(\sigma_{ft}^2, real_t^2)$	0.26	0.05	0.08	0.26	0.44	0.47
$cor(\sigma_{rt}^2, real_t^2)$	0.36	0.25	0.27	0.37	0.44	0.46
Panel B: correlations of vol processes with macro uncertainty						
$cor(\sigma_{ct}^2, macro_t^2)$	0.16	-0.03	-0.01	0.15	0.36	0.41
$cor(\sigma_{ft}^2, macro_t^2)$	0.28	0.07	0.10	0.28	0.47	0.51
$cor(\sigma_{rt}^2, macro_t^2)$	0.50	0.36	0.39	0.51	0.59	0.61
Panel C: correlations of vol processes with financial uncertainty						
$cor(\sigma_{ct}^2, financial_t^2)$	0.16	-0.04	-0.01	0.15	0.34	0.38
$cor(\sigma_{ft}^2, financial_t^2)$	0.45	0.21	0.25	0.46	0.62	0.64
$cor(\sigma_{rt}^2, financial_t^2)$	0.54	0.43	0.46	0.55	0.61	0.62

The table summarizes 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 *squared* financial/real/macro uncertainty measures from Jurado, Ludvigson, and Ng (2015) and Ludvigson, Ma, and Ng (2021). The uncertainty measures have horizons of three months, consistent with the quarterly frequency of consumption growth.

Table IA.XIII: Cross-sectional estimates of risk premia in six-factor models

Estimating $\Delta\mathbb{E}_t[m_t] = b_f f_t - \mathbf{b}_g^\top \mathbf{g}_t$								$\mathbb{E}[r_t^{mkt}] =$
b_f	$\mathbb{E}[SR_m \mid \text{data}]$	$\mathbb{E}[SR_f \mid \text{data}]$	$\mathbb{E}\left[\frac{SR_f^2}{SR_m^2} \mid \text{data}\right]$	R^2	$-\text{cov}(r_t^{mkt}, m_t)$	$-\text{cov}(r_t^{mkt}, -b_f f_t)$		
Panel A. 37 stock and nine bond portfolios								
Posterior median	0.264	0.969	0.529	0.325	0.784	0.068	0.073	
90% CIs	[0.062, 0.460]	[0.652, 1.496]	[0.156, 0.919]	[0.022, 0.688]	[0.497, 0.905]	[0.031, 0.107]	[0.018, 0.123]	
Panel B. Kozak, Nagel, and Santosh (2020) 74 anomaly portfolios								
Posterior median	0.247	0.975	0.495	0.260	0.422	0.072	0.068	
90% CIs	[0.106, 0.374]	[0.754, 1.212]	[0.217, 0.749]	[0.056, 0.548]	[0.216, 0.592]	[0.034, 0.108]	[0.029, 0.108]	

The table reports estimation results for two cross-sections of excess returns: (1) 37 stock and nine bond portfolios and (2) Kozak, Nagel, and Santosh (2020) 74 characteristic-sorted portfolios. We report seven statistics: (1) the risk price of the shock to the conditional consumption mean f_t (b_f), (2) the annualized Sharpe ratio of the SDF in equation (20), defined as the annualized volatility of the SDF (SR_m), (3) the annualized Sharpe ratio of $b_f f_t$ (SR_f), (4) the ratio of SR_f^2 to SR_m^2 , (5) the cross-sectional R-squared (R^2), (6) the (annualized) market risk premium implied by the SDF, $-\text{cov}(r_t^{mkt}, m_t)$, and (7) the (annualized) market risk premium implied by the covariance between market excess return and $-b_f f_t$, $-\text{cov}(r_t^{mkt}, -b_f f_t)$. We estimate the risk prices using the Bayesian approach developed in Bryzgalova, Huang, and Julliard (2024). Details are provided in Internet Appendix M. We consider six-factor models of asset returns. Both the posterior median and the 90% Bayesian credible intervals are reported.

Table IA.XIV: What drives the consumption shocks spanned by financial markets?

Panel A: Single-factor model in nondurable consumption growth					
	MKT	SMB	HML	RMW	CMA
Correlation	0.938	0.634	-0.123	-0.224	-0.297
95% CI	[0.930, 0.946]	[0.613, 0.654]	[-0.150, -0.097]	[-0.238, -0.210]	[-0.318, -0.276]
	PC1	PC2	PC3	PC4	PC5
Correlation	0.994	0.052	-0.043	-0.027	-0.001
95% CI	[0.991, 0.996]	[0.024, 0.080]	[-0.070, -0.016]	[-0.051, -0.003]	[-0.018, 0.015]
Panel B: Two-factor model in nondurable consumption growth (P/D, CFNAI as predictors)					
	MKT	SMB	HML	RMW	CMA
Correlation	0.937	0.626	-0.132	-0.231	-0.288
95% CI	[0.924, 0.946]	[0.605, 0.647]	[-0.160, -0.105]	[-0.247, -0.214]	[-0.310, -0.266]
	PC1	PC2	PC3	PC4	PC5
Correlation	0.991	0.060	-0.037	0.030	0.001
95% CI	[0.982, 0.995]	[0.031, 0.089]	[-0.064, -0.009]	[0.012, 0.047]	[-0.016, 0.017]
Panel C: Two-factor model in nondurable consumption growth (<i>cay</i> as predictor)					
	MKT	SMB	HML	RMW	CMA
Correlation	0.938	0.619	-0.136	-0.225	-0.29
95% CI	[0.929, 0.947]	[0.596, 0.641]	[-0.163, -0.109]	[-0.243, -0.207]	[-0.311, -0.269]
	PC1	PC2	PC3	PC4	PC5
Correlation	0.990	0.058	-0.030	0.032	-0.003
95% CI	[0.981, 0.995]	[0.030, 0.086]	[-0.058, -0.002]	[0.016, 0.049]	[-0.019, 0.014]

The table reports the correlation coefficients between f_t and Fama-French five factors and the first five principal components of asset returns. We also report their 95% posterior credible intervals under the coefficient estimates. Moreover, asset returns are modeled using a single-factor model in Panel A and two-factor models in Panels B and C (as in Section III.4), with $\bar{S} = 14$. The cross-section of test assets includes 25 size-and-value-sorted portfolios, 12 industry portfolios, and nine bond portfolios.



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■



Economic
and Social
Research Council



The London School of Economics
and Political Science
Houghton Street
London WC2A 2AE
United Kingdom

tel: +44 (0)20 7405 7686
systemicrisk.ac.uk
src@lse.ac.uk