Consumption in Asset Returns

Svetlana Bryzgalova [1] Jiantao Huang[2] Christian Julliard[3]

[1]London Business School [2]University of Hong Kong [3]London School of Economics

The Big Picture

In most macro-financial models, risk premia are driven by consumption risk.

The models crucially depend on the parametrization of the consumption DGP.

Inherent fragility:

⇒ consumption process reflects the "dark matter" in asset pricing (Chen, Dou, and Kogan, 2024): something hard to independently test or detect, yet inevitable to rely on to make sense of the data given the currently available models.

This paper:

- use returns information to sharply identify the consumption process: conditional mean & volatilities
- intuition: if household is hit by a shock, Euler equation implies changes in both consumption and investment (returns)

⇒ new set of empirical facts that poses both challenges and solutions for macro-finance

This Paper: learning about consumption from returns

Methodology:

- Flexible parametric (Bayesian) approach to model the joint dynamics of stocks, bonds & consumption, and identify shocks to which consumption reacts (potentially) slowly.

⇒ time series *and* cross-sectional implications

Main findings:

1. the time-varying conditional mean of consumption growth:
   - drives a quarter of consumption variance and generates substantial (but not excessive) consumption predictability
   - reflects the IRF of consumption to asset innovations
2. the (stochastic) volatility of consumption:
   - very distinct from asset returns' one...
   - ... and neither of them drives time varying risk premia (sharp Bayesian selection)
3. consumption conditional mean shocks:
   - explain on average 79% of the time series variation in stocks, and
   - explain a significant (yet small) portion of bond returns time series
   - are priced in cross-section of returns and command a SR as large as the market
4. a standard recursive utility model based on our estimates explains both equity premium and risk-free rate puzzles

Recovering Consumption Dynamics

The consumption process recovery problem: a LRR example

Identifying the type of consumption processes we postulate in our models is hard.

Example: Bansal and Yaron (2004):

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

$$x_{t+1} = \rho x_t + \phi_e \sigma_t e_{t+1},$$

$$\sigma_t \sim SV, \quad \eta_t, e_t, \sim iid\, \mathcal{N}(0,1)$$

In Hansen, Heaton, Lee, and Roussanov (2007) calibration the conditional mean drives about 12% of the variance of consumption—the largest share in the literature.

Can canonical model selection recover this degree of consumption predictability in samples of the same size as the historical ones?

LRR example cont'd: the missing consumption predictability

Table 1: Frequencies of ARIMA(p,d,q) models selected by BIC / AIC in 1,000 simulations

(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															
p	d	q	freq	p	d	q	freq	p	d	q	freq	p	d	q	freq
0	0	0	537	0	0	0	251	0	1	0	500	0	1	1	215
0	1	1	248	0	1	1	146	0	1	1	156	0	1	0	137
1	0	1	59	1	0	1	138	1	0	1	122	1	0	1	116
1	1	1	51	0	0	1	63	2	0	0	101	2	0	0	76
1	0	0	39	1	1	1	53	1	1	0	51	1	1	0	75
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Panel B: Monthly frequency															
0	0	0	551	0	0	0	278	0	1	0	923	0	1	0	700
0	1	1	313	0	1	1	240	1	0	0	72	2	1	2	49
1	0	1	47	1	0	1	134	1	1	0	3	1	0	0	47
5	1	0	14	0	0	1	65	1	1	1	1	1	1	0	38
0	0	1	9	0	1	2	32	2	0	0	1	1	1	1	36
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Empirical frequencies of ARIMA(p,d,q) models selected by Bayesian information criterion (BIC) and Akaike information criterion (AIC) in 1,000 simulations. We only list the top five most frequent models.

Even if data were generated by the easiest to detect LRR process, a researcher would conclude that there is no (or very little) predictability in consumption growth.

A parametric framework for consumption and asset returns

Joint dynamics for stocks and bonds:

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + w_t^c; \qquad f_t \perp \{\mathbf{w}_t^r,\, w_t^c\}$$

$$\underset{N\times 1}{\mathbf{r}_t^e} = \underset{N\times 1}{\boldsymbol{\mu}_r} + \underset{N\times 1}{\boldsymbol{\rho}_f^r} f_t + \underset{N\times K}{\boldsymbol{\rho}_g^r} \mathbf{g}_t + \underset{N\times 1}{\mathbf{w}_t^r}; \qquad \mathbf{w}_t^r \perp w_t^c$$

where Δc and $\mathbf{r}^e$ denote, respectively, consumption growth and log excess returns.

A state-space representation with implications for:

- the conditional mean of consumption (can be filtered via a Bayesian approach)
- time series predictability and return patterns
- factor loadings and cross-sectional dispersion in returns
- variance decomposition

Note: IRF of consumption: $\frac{\partial \mathbb{E}\left[\Delta c_{t-1,t+S}\right]}{\partial f_t} = \sum_{j=0}^{S} \rho_j$

Also: allow for other factors not spanned by consumption $\mathbf{g}_t \in \mathbb{R}^K \perp \{\mathbf{w}_t^r,\, w_t^c,\, f_t\,,\, b_t\}$

Consumption: focus on the conditional mean

Joint dynamics for stocks and bonds. Can also allow for a bond-specific factor, b_t

$$\Delta c_{t-1,t} = \mu_c + \sum_{j=0}^{\bar{S}} \rho_j f_{t-j} + \sum_{j=0}^{\bar{S}} \theta_j b_{t-j} + w_t^c; \quad b_t \perp \{\mathbf{g}_t, \mathbf{w}_t^r, w_t^c, f_t, b_t\}$$

$$\underset{N\times 1}{\mathbf{r}_t^e} = \underset{N\times 1}{\boldsymbol{\mu}_r} + \underset{N\times 1}{\boldsymbol{\rho}_f^r} f_t + \left[\underset{N_b\times 1}{\boldsymbol{\theta}'^b}, \quad \mathbf{0}'_{N-N_b} \right]' b_t + \underset{N\times K}{\boldsymbol{\rho}_g^r} \mathbf{g}_t + \underset{N\times 1}{\mathbf{w}_t^r};$$

⇒ Use returns *and* consumption to figure out the DGP for consumption growth

- Important:
  1. allows for a potentially slow propagation of shocks through consumption
  2. no strict parametric structure for the consumption process and its memory
  3. no assumption on the structural origins of the shocks

⇒ embed a very high order MA component in the consumption process

SV: note that: i) misspecified mean process leads to spurious evidence of volatility clustering; ii) but, erroneously assuming constant vol affects only efficiency; iii) we allow for SVs in ALL shocks and obtain identical estimates as with constant cons vol.

Example: Would our framework recover the calibrated LRR predictability?

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.

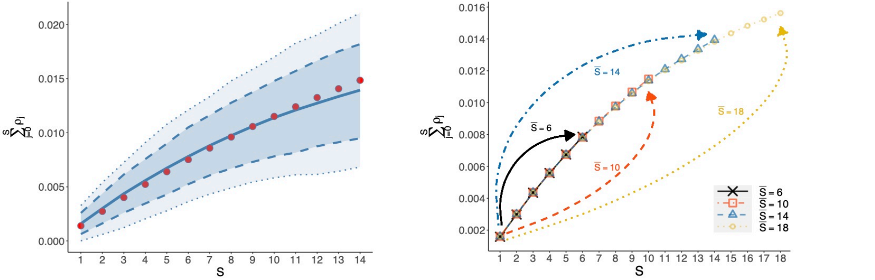

(a) Cumulative impulse response function of consumption growth ($\bar{S} = 14$)

(b) Cumulative impulse response function of consumption growth across different $\bar{S}$

Panel (a) mean, 5th, 16th, 84th, and 95th percentiles of estimated cumulative IRF and pseudo-true values (red dots). Panel (b) plots the average cumulative IRF across simulations for different $\bar{S}$. Consumption process calibrated as in Hansen et al. (2007). Parameters of the return process calibrated to their sample estimates. Simulated time-series size is 200 (quarters)

- conditional mean accurately captured
- robust to $\bar{S}$ choice (only affects max horizon and CIs)

Intuition: Wold representation theorem / state-space filtering with low signal to noise

Would our framework recover the conditional mean in general settings?

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

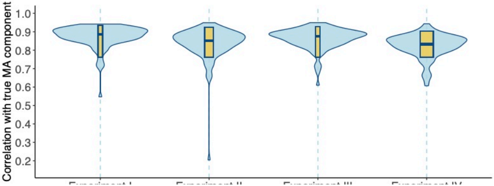

Experiment I: ρ^r loads on first two PCs of asset returns and $f_t = u_{1t}$.

Experiment II: ρ^r loads on first two PCs of asset returns and $f_t = u_{2t}$.

Experiment III: ρ^r loads on first two PCs of asset returns and $f_t = (u_{1t} + u_{2t})/\sqrt{2}$.

Experiment IV: ρ^r loads on first five PCs of asset returns and $f_t = (u_{1t} + u_{4t})/\sqrt{2}$.

Distribution of correlations between estimated MA components and their pseudo-true values in 1,000 simulations. Simulated time-series sample size is 200 (quarters). Loadings for simulation estimated from data as: $(\rho_1^r)'\rho_1^r = 16(\rho_2^r)'\rho_2^r = 25(\rho_3^r)'\rho_3^r = 36(\rho_4^r)'\rho_4^r = 144(\rho_5^r)'\rho_5^r$

$\Rightarrow$ **our method: 1) does not mechanically recover the 1st PC; 2) can identify f_t even if it is a small PC or a linear combination of several latent factors.**

Data

Data description

- Sample: 1963Q3 – 2019Q4
- Bonds: quarterly holding returns from the zero coupon yield data (constructed following Gurkaynak and Wright, 2007, from daily fitting of the Nelson-Siegel-Svensson curve).
- Maturities: 6 months, 1, 2, 3, 4, 5, 6, 7, and 10 years.
- Stock returns: 25 size and book-to-market Fama-French portfolios, 12 industry portfolios (cf. Lewellen, Nagel, and Shanken, 2010)
- Consumption flow: real (chain-weighted) consumption expenditure on non-durable goods per capita (end-of-period timing convention).

The Conditional Consumption Mean

Slow consumption adjustment to f shocks

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

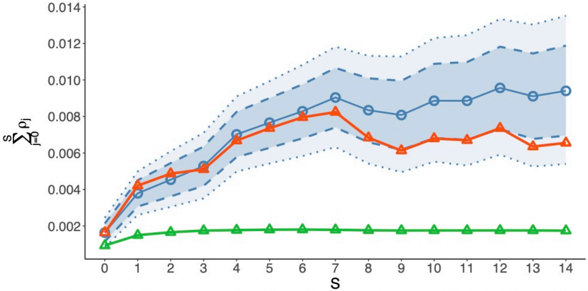

Posterior means of the cumulative response function of consumption growth (solid line with circles), with the centered posterior 90% (dotted lines) and 68% (dashed lines) coverage regions. Red line with triangles denotes the first principal component of $cov(r^{ex}_{i,t}, \Delta c_{t,t+1+S})$. Quarterly 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).

Large cumulative reaction after 2-3 years, flat thereafter. Robustness: 1) not driven by measurement error, benchmarking, time averaging; 2) same IRF with monthly data, mixed-frequency, and external predictors; 3) model implies 1-factor structure for $Cov\left(\Delta c_{t-1,t+S}; r^{e}_{i,t}\right)$ $\Rightarrow$ single factor drives over 99% of the uncentered variation and yields same IRF behaviour.

Shareholders vs non-shareholders consumption

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

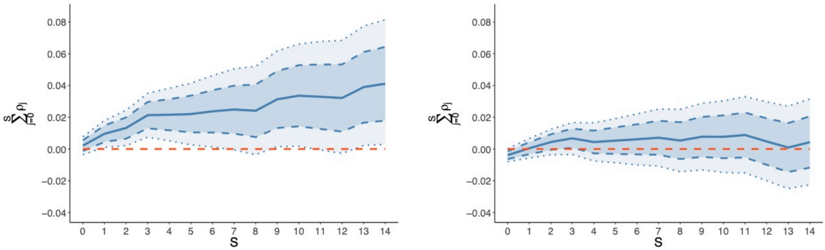

(a) Shareholders' consumption growth **(b)** Non-shareholders' consumption growth

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.

$\Rightarrow$ **external validation of identification assumption**: 1) no returns Euler equation for non-shareholders; 2) shareholders' IRF (proportionally) larger than the aggregated one

The conditional mean of consumption growth

Figure 5: Moving average component of consumption growth.

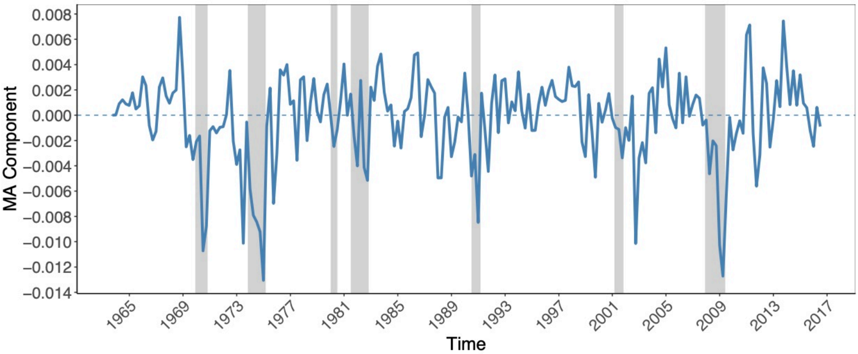

Posterior mean of the moving average f_t component of consumption growth. Grey areas denote NBER recessions. Estimate based on the single-factor model in equations (2)–(3), with $\bar{S} = 14$. The cross-section of base assets includes 25 size- and value-sorted portfolios, 12 industry portfolios, and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

⇒ clear business cycle pattern: conditional mean growth contracts right before/during recessions, and increases right before/during expansions

Consumption variance decomposition: $MA(f)$ share

Financial shocks drive over 25% of the consumption variance and generate substantial consumption predictability (more than double the largest calibration in the literature)

Figure 7: Share of consumption growth variance driven by its moving average component.

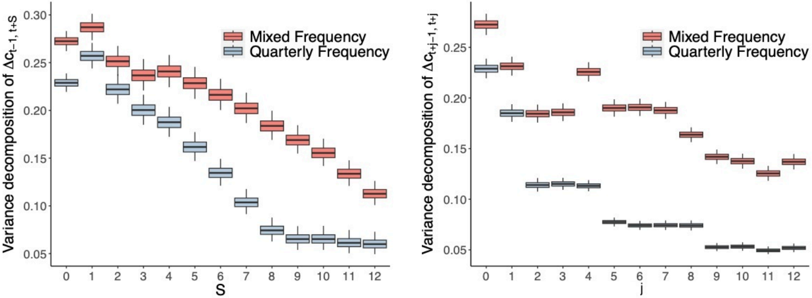

Box plots (posterior 95% coverage area) of the percentage of time-series variances of consumption growth explained by the MA component. The state-space models are estimated at both quarterly (blue bars) and mixed (red bars) frequencies. Plots report the unadjusted R-squared. Left panel: Cumulated consumption 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.

Are the consumption mean shocks priced?

Table 5: Cross-sectional pricing ability of the shock to conditional consumption mean

	Estimating $\Delta\mathbb{E}_t[m_t] = -b_f f_t - \boldsymbol{b}_g^\top \boldsymbol{g}_t$					$\mathbb{E}[r_t^{mkt}] =$	
	b_f	$\mathbb{E}[SR_m \mid \text{data}]$	$\mathbb{E}[SR_f \mid \text{data}]$	$\mathbb{E}[\frac{SR_f^2}{SR_m^2} \mid \text{data}]$	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 in a five-factor model							
Posterior median	0.236	0.806	0.474	0.371	0.772	0.069	0.070
90% CI	[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) anomaly portfolios in a five-factor model							
Posterior median	0.244	0.960	0.489	0.263	0.409	0.071	0.069
90% CI	[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]

Estimation results for two cross-sections of excess returns: 37 stock and nine bond portfolios (Panel A) and Kozak, Nagel, and Santosh (2020) 74 characteristic-sorted portfolios (Panel B). We report: (1) risk price of the shock to the conditional consumption mean f_t (b_f); (2) annualized Sharpe ratio of the SDF in equation (20), defined as the annualized volatility of the SDF (SR_m); (3) annualized Sharpe ratio of $b_f f_t$ (SR_f); (4) ratio of SR_f^2 to SR_m^2; (5) cross-sectional R^2; (6) (annualized) market risk premium implied by the SDF, $-\text{cov}(r_t^{mkt}, m_t)$; and (7) (annualized) market risk premium implied by the covariance between market excess return and $-b_f f_t$, $-\text{cov}(r_t^{mkt}, -b_f f_t)$. We estimate the risk prices using the Bayesian approach of Bryzgalova, Huang, and Julliard (2024). Details are provided in Internet Appendix M. We consider five-factor models of asset returns. Both the posterior median and the 90% Bayesian credible intervals are reported.

***f* shocks**: **1) price assets; 2) command a large market price of risk and a SR as large as the market one; 3) are sufficient to match the equity premium of the market; 4) but there are also other risks in the true, latent, SDF**

A simple explanation of the Equity Premium and Risk Free Rate Puzzles?

Table 7: Implications for structural models

	CRRA								Epstein-Zin	
	$S=0$	2	4	6	8	10	12	14	$\psi=1.5$	$\psi=0.5$
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%

A standard recursive utility model based on our estimated process can explain both equity premium and risk-free rate puzzles with low risk aversion.

Consumption Volatility

What about Consumption SV? 1. Volatility clustering & forecast errors

Autocorrelated consumption vol implies serial correlation of the squared one-step ahead forecast errors of consumption (see, e.g., Engle, 1982)

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

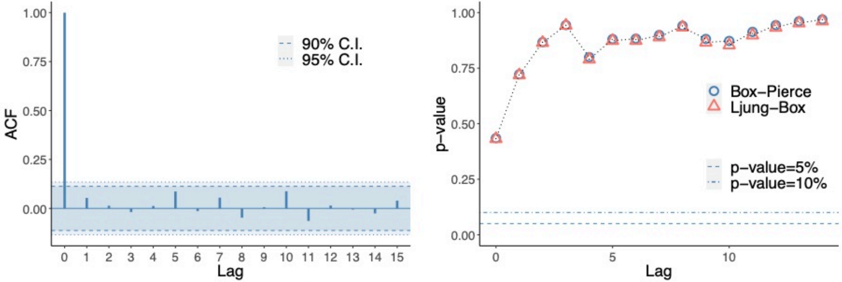

Left panel: Autocorrelation function of $\widehat{Var}_t(\Delta c_{t,t+1})$ with 95% and 99% confidence bands. Right panel: p-values of Ljung and Box (1978) (red triangles) and Box and Pierce (1970) (blue circles) tests.

... but **no evidence in the data**

What about Consumption SV? 2. GARCH

Misspecified consumption mean process leads to evidence of time-varying volatility:

$$\Delta c_{t+1} = \mu_t + \epsilon_{t+1}$$
$$\sigma^2_{t+1} = \omega + \alpha\epsilon^2_t + \beta\sigma^2_t$$

	ARCH(1)		GARCH(1,1)			IGARCH(1,1)		
	Panel A: $\mu_t = \mu_0 + \mu^c_1 \Delta c_t$							
	ω	α	ω	α	β	ω	α	β
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)
	Panel B: $\mu_t = \mu_0 + \mu^c_1 \Delta c_t + \sum_{i=1}^{8} \mu^r_i r^{ex}_{PC_i,t}$							
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)
	Panel C: $\mu_t = \mu_0 + \sum_{i=1}^{\bar{S}} \rho_i f_{t+1-i}$							
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)

The table presents ARCH(1), GARCH(1,1), and IGARCH(1,1) estimates for consumption growth volatility with different models for the conditional mean. Models are estimated using QMLE, and robust standard errors are constructed using Newey-West (1987).

What about Consumption SV? 3. Allowing for SVs in all shocks

Figure 15: Filtered stochastic volatilities of consumption and returns.

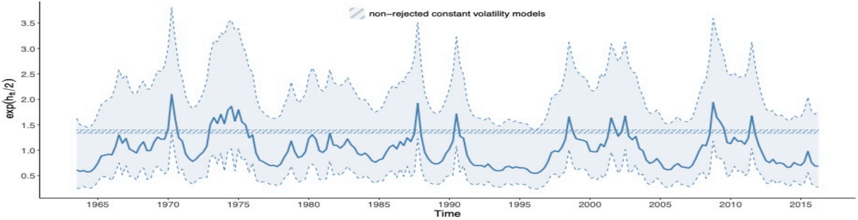

Panel B: Log volatility of the shock to the conditional mean of consumption growth (f_t).

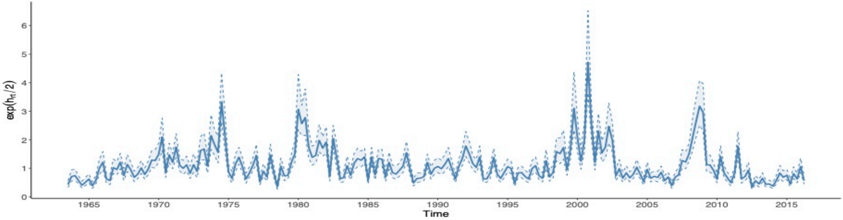

Panel C: Common log volatility of asset return (h_{rt} of equation (19)).

⇒ **null of common SV for consumption and asset returns <u>strongly</u> rejected**

What about Consumption SV? 4. Driving time varying risk premia?

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

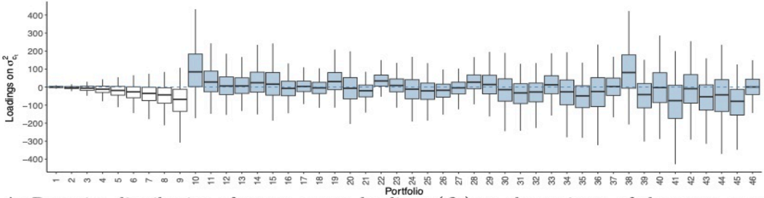

Panel A: Posterior distribution of excess return loadings ($\boldsymbol{\beta_c}$) on the variance of short-run consumption shocks ($\sigma^2_{c,t-1}$) in equation (16).

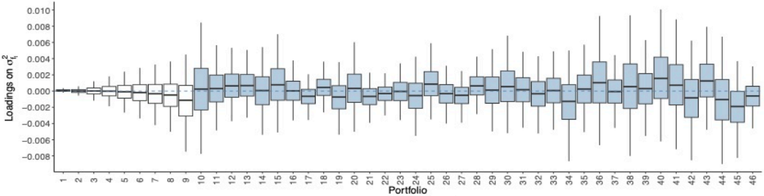

Panel B: Posterior distribution of excess return loadings ($\boldsymbol{\beta_f}$) on the variance of shocks to the conditional consumption growth mean ($\sigma^2_{f,t-1}$) in equation (16).

Sharply rejected by the data. Nevertheless, $cor(\sigma^2_{f,t}, VXO^2_t) = 45\%$

Consumption and Returns

Stocks and bonds reaction to common financial shocks (f)

Stocks:

- On average, common shocks (f_t) explain 79% of the time series variation in stocks
- Stock loadings are sharply estimated and have a pronounced pattern across Size and B/M portfolios

$\Rightarrow$ f is a strong factor

Bonds:

- Bond loadings increase with the maturity of the bond ("slope" component)
- Adding a bond-specific shock, on which consumption has insignificant (MA) loading, does not change the loadings on f, but greatly sharpens identification
- Most of time-series variation explained by the bond-specific factor, while f drives a small, yet significant, share at horizons below 4 years.