

Consumption Growth Persistence and the Stock–Bond Correlation

Christopher S. Jones
University of Southern California Marshall School of Business
christoj@usc.edu

Sungjune Pyun 
Yonsei University School of Business
sjpyun@yonsei.ac.kr (corresponding author)

Abstract

We consider a model in which the correlation between shocks to consumption and to expected future consumption growth is nonzero and varies over time. We validate this assumption empirically using the model's implication that time variation in consumption growth persistence (CGP) drives the correlation between stock and bond returns. Our model implies that the stock–bond correlation is also related to the predictive relation between bond yields and future stock returns. Finally, we provide suggestive evidence that asset price fluctuations are the primary driver of changes in CGP.

I. Introduction


While the correlation between stock and bond returns has been the subject of research for some time, the abrupt change in the sign of this correlation, which Graph A of [Figure 1](#) shows turned from positive to negative in the late 1990s, has spurred renewed interest in its determinants. One explanation of this shift is an apparent regime change in the behavior of inflation, as demonstrated in [David and Veronesi \(2013\)](#), [Song \(2017\)](#), and [Campbell, Pflueger, and Viceira \(2020\)](#). However, this explanation is at best incomplete, because as pointed out by [Duffee \(2018a\)](#), expected inflation shocks are responsible for just 10–20% of the variation in nominal yields, implying that it is unlikely that inflation is the dominant driver of nominal bond returns or their correlation with stocks.

In this article, we propose a new explanation of this shift and of variation in the correlation between stock and bond returns (SB correlation) more generally. Our explanation relies on a shift in macroeconomic dynamics, but in a channel omitted by extant models. Specifically, we show that the stock–bond correlation

We thank Thierry Foucault (the editor), Allaudeen Hameed, Scott Joslin, Mete Kilic, Toomas Laarits (the referee), Lars Lochstoer, Miguel Palacios, Thomas Sargent, Johan Sulaeman, Selale Tuzel, participants at the 2021 Risk Management Conference, 2021 NFA Annual Meetings, Korea University, National University of Singapore, and University of Southern California. This research was supported by the Yonsei Business Research Institute and Yonsei University Research Fund 2023-22-0444. All errors are our own.

FIGURE 1
Pre-1999 Versus 1999+ Comparison

Figure 1 shows the time variation in the stock–bond return (SB) correlation in Graph A and the autocorrelations in consumption growth in Graph B for the pre-1999 (1962–1998) and 1999+ (1999–2019) sample periods. The SB correlation is estimated using first differences in daily 1-year or 10-year bond yields over a rolling window of 365 calendar days.



is related to variation in consumption growth persistence (CGP), which we define as the tendency of positive shocks to current consumption growth to raise expected future consumption growth. The logic is straightforward: Changes in current realized growth affect cash flows, while changes in expected growth drive real interest rates via intertemporal smoothing. When CGP increases, the correlation between real rates and cash flows rises, resulting in a lower (and likely negative) SB correlation. When CGP is negative, higher consumption growth forecasts lower growth in the future, and the SB correlation rises.

The persistence of consumption growth indeed appears to have changed over time. Graph B of Figure 1 shows that autocorrelations in consumption growth were moderate through 1998 but significantly higher in the period starting in 1999, which is around when the SB correlation changed sign. The goal of this study is to determine whether this suggestive evidence is indicative of a more systematic effect that CGP has on the SB correlation and other asset return moments.

We generalize the long-run risk (LRR) model of Bansal and Yaron (2004) by allowing a time-varying correlation between current and expected consumption growth shocks. As in the standard LRR model, a highly persistent expected consumption growth process induces modest but very long-run dependence in consumption growth. Adding variation in the correlation between current and expected consumption growth induces time variation in the degree of serial dependence while maintaining the long-run positive autocorrelation critical for matching the moments of asset returns.

We validate the model-implied relationship between CGP and the SB correlation by showing that the serial correlation in consumption growth is significantly higher when the SB correlation is low. This result holds at multiple horizons and is obtained whether we use returns on short- or long-term bonds and whether the bonds are nominal or inflation-indexed. Second, we show that the contemporaneous relationship between consumption growth and changes in survey forecasts of long-run consumption growth is more negative when the SB correlation is higher. This result is also consistent with a negative relation between CGP and the stock–bond correlation.

An additional feature of our model is the negative correlation between shocks to consumption growth and its volatility. This correlation is empirically motivated and is thought to arise from a precautionary savings motive.¹ With precautionary savings, a higher CGP implies that uncertainty shocks will not only lead to a lower current consumption, but also to a lower future expected consumption.

An implication of this additional feature is a conditional predictive relation between yields and future stock returns. The model implies that this negative relationship should be stronger when CGP is high or the SB correlation is low. As documented in prior work, we report a weak unconditional relationship between yields and future stock returns, but we observe a strong conditional relationship when the SB correlation is negative. Therefore, the insignificant unconditional relation, which is at odds with the predictions of many models, is a natural result of the SB correlation being positive over much of our sample.

While our analysis suggests that CGP drives the SB correlation, we further ask whether fluctuating asset values are a primary driving force behind consumption persistence. Using a sample that largely comes from the earlier period in which the SB correlation is positive, Lettau and Ludvigson (2004) find that the deviations from the cointegrating relationship between consumption and wealth do not predict future consumption growth, which implies that those deviations must predict returns. In contrast, in the later part of our sample, when the SB correlation was negative, we show that consumption growth tended to react with a delay in response to an asset growth shock, such that cointegrating residuals predict future consumption growth but not future returns.

There are a number of other explanations for why the SB correlation varies over time, and we believe it is unlikely that any single theory can explain all fluctuations. Aside from the other predictions that we confirm in our empirical analysis, we believe that CGP has certain merits that distinguish it from other

¹See, e.g., Carroll (1997) and Basu and Bundick (2017).

explanations. Baele, Bekaert, and Inghelbrecht (2010), for example, claim that “macroeconomic fundamentals contribute little to explaining stock and bond return correlations” and conclude that flights to quality/liquidity are the likely explanation for negative correlations. While Pástor and Stambaugh (2003), Connolly, Stivers, and Sun (2005), and others provide additional evidence for this channel, the mechanism seems inadequate in the period starting in 1999, during which the stock–bond correlation has remained negative even during periods of relative market stability.

Variation in the stock–bond correlation has also been attributed to changes in the dynamics of inflation. David and Veronesi (2013), Song (2017), and Campbell et al. (2020) present models in which the relation between inflation and real economic activity changes signs. Campbell, Pflueger, and Viceira (2020), for example, show that the correlation between inflation and the output gap was negative between 1979 and 2001 but positive in the following decade. If inflation shocks are the primary driver of bond returns, this result would appear to provide a clear explanation for the shift in correlation that occurred around that time.²

In contrast, the SB correlations shown in our model are entirely driven by variation in real interest rates and not in inflation. While changing properties of inflation are undoubtedly a reason for changes in SB correlations, a model based on real rates may be better positioned to explain interest rate behavior in environments such as that of the last 20 years, in which both inflation levels and inflation risk have mostly been low. From 2003 to 2019, the part of our sample for which reliable TIPS data are available, real and nominal bond yields have tracked each other closely, with a correlation above 90%. More importantly, the real stock–bond correlation and nominal stock–bond correlation are themselves closely related over the sample period in which real rates are available.³ While inflation is undoubtedly more important in other periods, it almost certainly does not tell the entire story.

Ours is not the only article to propose that variation in real yields is an important driver of changes in the SB correlation. For example, Duffee (2018a) argues that the inflation expectations that underlie long-term bond yields vary too little to explain much variation in such yields and that the stock–bond correlation is therefore primarily driven by changes in real yields. Kozak (2022) proposes a production model with 2 technologies that generates a time-varying correlation, attributing the shift in correlation in the late 1990s to a decline in high-risk capital.

In contemporaneous and complementary work, Chernov, Lochstoer, and Song (2021) also highlight the role of CGP in explaining the stock–bond correlation. Their model is substantially different from ours, relying on a regime-switching model to generate time variation in the relative importance of permanent and transient shocks. Similar to ours, it succeeds in explaining the stock–

²Ilmanen (2003), Hasseltoft (2012), Campbell, Sunderam, and Viceira (2017), and Swanson (2019) also advance inflation-based explanations of the stock–bond return correlation.

³When these 2 correlations are measured using non-overlapping monthly subsamples, the correlation between them is 86%.

bond correlation using a mechanism based on consumption persistence rather than inflation dynamics. Besides differences in modeling, our article focuses more on the stock–bond correlation and on the predictability of equity returns, while Chernov et al. (2021) concentrate on explaining the real and nominal term structures. The different empirical predictions derived in their article show that time-varying CGP can explain a broader range of phenomena than those we address here.

In Section II, we present and calibrate our model. Section III contains our empirical results. Section IV explores potential reasons for the recent shift in CGP, and Section V concludes.

II. A Model of the Real Stock–Bond Correlation

A. Model Dynamics

Our model is a generalization of the standard framework of Bansal and Yaron (2004). In our specification, the representative agent has Epstein and Zin (1991) preferences, and consumption growth (Δc_{t+1}) has a persistent time-varying component x_t and time-varying uncertainty σ_t^2 :

$$(1) \quad \begin{aligned} \Delta c_{t+1} &= \mu + x_t + \sigma_t \epsilon_{c,t+1}, \\ x_{t+1} &= \zeta_1 x_t + \phi_x \sigma_t \epsilon_{x,t+1}, \\ \sigma_{t+1}^2 &= s_0 + s_1 \sigma_t^2 + \sigma_v \sigma_t \epsilon_{v,t+1}, \end{aligned}$$

where $\epsilon_{c,t+1}$, $\epsilon_{x,t+1}$, and $\epsilon_{v,t+1}$ are $N(0, 1)$.

Our model deviates from Bansal and Yaron (2004) in several dimensions. Most importantly, we allow shocks to consumption growth ($\epsilon_{c,t+1}$) and expected long-run consumption growth ($\epsilon_{x,t+1}$) to be stochastically correlated. We refer to this correlation, which we denote ρ_t , as CGP, given that it determines whether a shock to current consumption growth is associated with more or less consumption growth in the future.

To obtain closed-form solutions, we assume a “square root” process for consumption variance, and we parameterize the conditional covariance (as opposed to correlation) between c_{t+1} and x_{t+1} , which we label q_t . This covariance, which we call the CGP covariance, is related to CGP (ρ_t) by

$$q_t = \sigma_t^2 \rho_t.$$

Its dynamics are given as

$$(2) \quad q_{t+1} = \omega_0 + \omega_1 q_t + \sigma_q \sigma_t \epsilon_{q,t+1},$$

where $\epsilon_{q,t+1}$ is $i.i.d. N(0, 1)$. In addition to our full model, we also consider a baseline model, which is the special case with $q_t = \rho_t = 0$.

A stochastic correlation can be viewed as a reduced-form approach to modeling time variation in the relative importance of permanent and transitory shocks. For example, in the production economy described by Kaltenbrunner and

Lochstoer (2010), the assumption of permanent productivity shocks results in a positive CGP, while transitory shocks generate a negative CGP. This results from differences in how investment (and therefore consumption) responds to changing productivity and in how adjustment costs and mean reversion induce trends in future output. Given that both types of shocks are likely important, either effect could dominate depending on which type of shock is currently more volatile. Furthermore, this phenomenon is not limited to shocks to productivity. Permanent and transitory shocks to income generate similar responses, as discussed, for example, by Hall and Mishkin (1982) and Campbell and Deaton (1989).

We also allow consumption growth shocks to be correlated with consumption variance shocks. A negative correlation is the expected result of a precautionary savings motive, confirmed empirically in several studies, including Carroll and Samwick (1998) and Basu and Bundick (2017). For simplicity, we assume that this correlation, denoted ρ_{cv} , is constant.

Given the negative correlation between consumption growth and consumption volatility shocks, it is natural to expect a nonzero correlation between shocks to expected consumption growth and consumption volatility. For example, an increase in precautionary savings induced by greater uncertainty should reduce current consumption as households increase their savings, leading to a rise in expected long-run consumption growth as uncertainty wanes and consumption returns to normal. Empirically, a nonzero correlation between σ_t and x_t is found by Nakamura, Sergeyev, and Steinsson (2017), who show that it tends to be more negative during economic contractions. In another work, Parker and Preston (2005) find significant evidence from household survey data that the precautionary savings motive explains the predictable component of consumption growth.

In the interest of parsimony, we avoid introducing unnecessary additional parameters by assuming that this correlation between shocks to consumption volatility and expected consumption growth is equal to the product $\rho_t \rho_{cv}$.⁴

In addition to the consumption process, a dividend growth process is specified as

$$(4) \quad \Delta d_{t+1} = \mu_d + \phi_d x_t + \sigma_t \phi_{cd} \epsilon_{c,t+1} + \sigma_t \phi_d \epsilon_{d,t+1},$$

where $\epsilon_{d,t+1}$ is *IID* $N(0, 1)$. Thus, dividend growth shares similarities with consumption due to its dependence on the long-run growth process x_t and its sensitivity to the consumption growth shock $\epsilon_{c,t}$.

Similar to other models in the LRR framework, the wealth-to-consumption ratio z_t can be approximated as an affine function of long-run expected

⁴This correlation structure is consistent with the assumption that there are 3 orthonormal shocks, $[u_{c,t}, u_{x,t}, u_{v,t}]$, that drive the shocks to the 3 state variables via

$$(3) \quad \begin{aligned} \epsilon_{c,t} &= u_{c,t}, \\ \epsilon_{x,t} &= \rho_t u_{c,t} + \sqrt{1 - \rho_t^2} u_{x,t}, \\ \epsilon_{v,t} &= \rho_{cv} u_{c,t} + \sqrt{1 - \rho_{cv}^2} u_{v,t}. \end{aligned}$$

consumption growth (x_t), the variance of consumption growth (σ_t^2), and the CGP covariance (q_t). That is,

$$(5) \quad z_t = A_0 + A_x x_t + A_v \sigma_t^2 + A_q q_t,$$

where $A_x > 0$, $A_v < 0$, and $A_q < 0$ under conventional parameter assumptions ($\gamma > 1$ and $\psi > 1$), as shown in the Supplementary Material.

Bond yields of all maturities are also affine functions of the 3 state variables. The Supplementary Material derives an analytic formula for the one-period bond, which is increasing in x_t and decreasing in σ_t^2 and q_t , and provides a solution method for longer-term bonds.

Similar to the wealth–consumption ratio, we approximate the return on the market portfolio using the Campbell–Shiller decomposition and verify that the price–dividend ratio is also an affine function of the 3 state variables.

Given expressions for stock returns and for bond yields of any maturity, it is straightforward to solve for the stock and bond return variances and covariance. We show in the Supplementary Material that all 3 may be expressed as affine and increasing functions of σ_t^2 and q_t . Furthermore, the stock–bond return correlation is a univariate (though nonlinear) decreasing function of ρ_t .

B. Calibration

We calibrate the model to examine its quantitative implications. The consumption and consumption variance parameters mirror those of Bansal and Yaron (2004). We set a high bar for the exercise by choosing the parameters that govern the persistent component of consumption growth and the CGP covariance to match macrodata and survey forecasts, not asset returns. Coherence between these values requires that we deviate from Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2012), but the differences are minor. The parameter values we assume are summarized in Panel A of Table 1.

We proxy for the long-run growth process using the four-quarter-ahead forecast of real consumption from the Survey of Professional Forecasters (SPF), as detailed in the data section. The persistence (ξ_1) and volatility (ϕ_x) of the long-run growth process are set to match the persistence and volatility of the SPF forecast at the annual level, and the forecasts are only slightly less persistent than the values implied by Bansal and Yaron (2004) and Bansal et al. (2012). The unconditional correlation between the shocks to the survey forecasts and realized consumption is slightly negative (−0.12). For our primary specification, we set the correlation that determines the relationship between consumption growth and uncertainty shocks ρ_{cv} to −0.2 following the results of Basu and Bundick (2017), but we also show results for 2 alternative specifications by setting ρ_{cv} to either −0.1 or −0.4. We set the parameters of the q_t process such that the realized correlation between consumption growth and changes in the SPF forecast matches the correlation between Δc and x in the model in terms of the average, standard-deviation, and persistence.

Panel B of Table 1 compares the asset moments generated by the 2 specifications. These include the baseline model, in which the CGP covariance is set to 0, and

TABLE 1
Model Calibration

Table 1 summarizes the parameters that describe the representative investor's preferences and the consumption and dividend growth, volatility, and covariance processes used in the main specification, as well as asset pricing moments implied by these parameters. Panel A shows the values of the parameters, and Panel B shows the moments obtained by simulating the model dynamics. y denotes a bond yield, R_m is the return on the market portfolio, σ_m is the volatility of the market portfolio, and ρ_{SB} denotes the stock–bond return correlation. Values in Panel B are annualized. We also simulate stock returns, bond yields, and stock market variance using the consumption/dividend dynamics assumed in the model. Panel C displays model-implied simulated moments of 60-month rolling-window correlations of stock returns with 1- or 10-year bond yields. The simulated moments are compared with the data, in which we perform the same estimations.

Panel A. Parameters (in Monthly Unit)

Preferences		Consumption		Covariance	
γ	10	μ	0.0015	ω_0	-5.18×10^{-7}
ψ	1.5	ξ_1	0.95	ω_1	0.934
β	0.9985	ϕ_x	0.046	σ_q	6.3×10^{-4}
				ρ_{cv}	-0.2
Variance				Dividend	
s_0	8.52×10^{-7}			μ_d	0.0015
s_1	0.986			ϕ_d	2.5
σ_v	2.6×10^{-4}			ϕ_{cd}	3.50
				ϕ_d	4.50

Panel B. Unconditional Means (Annualized)

Variable	Maturity	Model		Data	
		Baseline	Full	Real	Nominal
R_m		5.80%	5.53%	6.64%	10.44%
y	1Y	2.09%	2.15%	1.36%	5.10%
	10Y	2.05%	2.14%	2.45%	6.20%
σ_m		15.88%	15.69%	15.26%	15.26%
$SD(y)$	1Y	2.31%	2.34%	2.37%	3.39%
	10Y	0.62%	0.63%	1.86%	2.89%
$SD(\Delta y)$	1Y	0.73%	0.74%	0.83%	1.65%
	10Y	0.17%	0.18%	0.95%	1.15%
ρ_{SB}	1Y	-0.241	-0.184	-0.183	-0.092
	10Y	-0.241	-0.184	-0.121	-0.100

Panel C. Rolling-Window Simulated Correlations

Variable	Statistic	Model		Data
		Baseline	Full	Nominal
ρ_{SB} (1Y)	Mean	-0.241	-0.187	-0.092
	AC(1)	0.000	0.971	0.978
	STD	0.000	0.133	0.179
ρ_{SB} (10Y)	Mean	-0.241	-0.187	-0.100
	AC(1)	0.000	0.971	0.971
	STD	0.000	0.128	0.173

the full model, in which the CGP covariance is time varying. For each specification, we generate 1 million monthly observations and evaluate the first 2 moments of stock and bond returns as well as several other relevant asset pricing moments. The table shows that the unconditional moments generated by the simulations are generally comparable to those of other standard LRR models, aside from the correlations between stock and bond returns and between returns and volatility changes, which our new model matches better.⁵

⁵The full model does not improve on the baseline model's inability to match the volatility of long-term yields. Thus, while the full model matches the SB correlation reasonably well, it does not match the SB covariance when measured using the long-term yield.

Finally, Panel C of Table 1 compares the dynamic behavior of estimated correlations implied by the models to those observed in the data. From both data and simulations, we compute various correlation measures using nonoverlapping 60-month windows. The panel displays the means and standard-deviations of these values. Monthly serial correlations are inferred by assuming a first-order autoregressive process and taking the 60th root of each autocorrelation. The results show that the level of persistence in the model matches the data quite well, though the correlations are somewhat less volatile in the full model compared to the data.

C. The Stock–Bond Return Correlation

This section shows the relationship between CGP and the SB correlation. Establishing the link between is essential because it provides a new explanation for why the correlation varies over time. Furthermore, since it can be computed easily as long as stock and bond returns are available, the SB correlation can be used as an empirical proxy for CGP (ρ_t).

To establish the relationship between CGP and the SB correlation, we first explain how bond yields, stock returns, and stock variances are determined in our model. For bonds, we examine both a short-term bond with a 1-year maturity and a long-term bond with a 10-year maturity. Table 2 summarizes the relationships between these quantities and the model’s state variables.

There are 2 channels that drive SB correlations under the baseline model, in which all shocks are assumed to be uncorrelated. The first channel occurs via shocks to expected consumption growth ($\epsilon_{x,t+1}$). If this shock is positive, the higher expected future cash flows result in higher stock prices. Bond yields will also increase, as the demand for money rises due to the intertemporal consumption smoothing motive. Since stock and bond returns will have opposite responses, this channel implies a negative SB correlation.

The second channel is the result of shocks to consumption growth uncertainty ($\epsilon_{v,t+1}$). Stock market variance will rise following a positive uncertainty shock, raising the risk premium and lowering equity valuations. At the same time, bond yields will drop due to a precautionary savings effect. Therefore, an increase in

TABLE 2
Correlations Between Simulated Values

Table 2 summarizes the correlations between macroeconomic and asset pricing variables based on the simulation of different models. The table shows the relationship between ΔC_{t+1} and the shocks to x_{t+1} , σ_{t+1} , and q_{t+1} and the first differences in 1-year ($y_{1,t+1}$) and 10-year ($y_{10,t+1}$) bond yields, the return on the market portfolio (R_m), and the first difference in the variance of the market portfolio ($\sigma_{m,t+1}$). We simulate 1,000,000 periods and drop the first 100,000 before calculating correlations.

	Model	$\Delta y_{1,t+1}$	$\Delta y_{10,t+1}$	$R_{m,t+1}$	$\Delta \sigma_{m,t+1}^2$
ΔC_{t+1}	Baseline	−0.023	−0.021	0.591	0.001
	Full	−0.148	−0.141	0.565	0.001
$x_{t+1} - E_t[x_{t+1}]$	Baseline	0.978	0.935	0.238	−0.001
	Full	0.970	0.918	0.194	0.002
$\sigma_{t+1}^2 - E_t[\sigma_{t+1}^2]$	Baseline	−0.098	−0.304	−0.048	0.992
	Full	−0.097	−0.308	−0.051	0.793
$q_{t+1} - E_t[q_{t+1}]$	Baseline	—	—	—	—
	Full	−0.160	−0.194	−0.045	0.598

uncertainty leads to stock and bond prices moving in opposite directions. The flight-to-quality phenomenon often refers to the negative SB correlation arising from this second channel.

While both channels in the baseline model imply a negative correlation between stock and bond returns, regardless of bond maturity, Panel A of Table 2 shows that the first channel is generally much stronger than the second. This is especially true for short-term yields, for which the correlation between yield changes and shocks to expected consumption growth is greater than 0.97 for both models. For long-term yields, the correlation between yield changes and shocks to the x_t process is approximately 0.94 for the baseline and 0.92 for the full model. In contrast, the correlation between yield changes and volatility shocks is -0.10 for short-term yields and -0.30 for long-term yields for the full model. These results establish the close connection between interest rates and the expected consumption growth process that underlies our empirical analyses. The table also shows that shocks to the market portfolio are highly correlated with realized consumption growth for both models.

While Table 1 shows that our generalized model exhibits a similar average SB correlation, CGP causes this correlation to vary over time. For example, the SB correlation should increase when CGP (ρ_t) decreases. To illustrate this point, suppose there is a positive shock to expected consumption growth ($\epsilon_{x,t+1} > 0$), which is likely to coincide with a decline in current consumption when ρ_t is negative. In this case, bond yields will increase as the economy expects higher levels of future growth, while the negative shock to current consumption will lower equity values. While the net effect may be that equity values rise due to higher expected long-run growth, the rise will be moderated by the negative shock to current consumption. Therefore, a negative ρ_t will lead to the SB correlation being less negative than usual and perhaps even positive.


Figure 2 shows how the SB correlation varies with CGP (ρ_t). Graph A shows the relationship for 2 different bond maturities. Graph B examines different values of the persistence of the expected growth process (ξ_1), and Graph C considers different values of intertemporal elasticity of substitution coefficient (ψ). Graphs B and C consider only the correlation based on the short-term bond but examine sensitivity to key parameters. For comparison, the horizontal solid lines depict the corresponding values under the baseline model, in which CGP is assumed to be constant at 0.

The average level of the SB correlation is higher when the expected growth process is less persistent. This outcome is natural, because the serial correlation of consumption growth is lower when ξ_1 is low. Moreover, the SB correlation is more negative with higher values of ψ , which is also expected since ψ measures the sensitivity of consumption growth to shocks to the real interest rate. Overall, these panels confirm the negative relation between ρ_t and the SB correlation is slightly convex in ρ_t .

In the Supplementary Material, we further show that the relationship between CGP and the SB correlation is almost unaffected by the risk-aversion coefficient, the correlation between consumption growth and consumption volatility, and the persistence of the CGP covariance process. Overall, the negative relation between CGP and the SB correlation appears very robust.

FIGURE 2
Consumption Persistence and Model-Based Correlations

Figure 2 shows the relationships between CGP and the stock-bond return correlation for different bond maturities and parameter assumptions. The relationships for the full model are shown in dashed lines. The solid horizontal line in Graph A shows the stock-bond correlation under the baseline model, which is constant and invariant to maturity. The graphs show the relationship for different maturities (Graph A), values of the persistence of the expected growth process (Graph B), and intertemporal elasticities of substitution (Graph C).



We also examine this relationship in our model using simulation by calculating the “correlation of correlations.”⁶ We find that ρ_t and the SB correlation are almost perfectly negatively related for both models, with correlations below -0.99 . Thus, our model suggests that the SB correlation is a very good proxy for the less easily observed ρ_t process.

D. Conditional Moments of Consumption Growth

The key assumption of our generalized model is the time-varying persistence of consumption growth shocks. In this section, we show how this assumption affects the conditional distribution of consumption growth for different values of ρ_t .


While higher CGP should clearly increase the serial correlation in consumption growth, it is difficult to assess the strength of this effect analytically. We therefore simulate 10 million months of data from our full model and compute approximate conditional moments by separating the simulated sample into narrow bins (e.g., $[-0.05, 0]$, $[0, 0.05]$, etc.) according to the value of ρ_t .⁷ We then compute moments of interest using all the observations in each bin. The simulation applied is at a monthly frequency, but we aggregate consumption growth to the quarterly level by taking the sum of 3 consecutive monthly growth rates.

⁶Using 10 million months of simulated data, we derive the value of SB correlation using its analytic formula and compute its correlation with ρ_t .

⁷In simulating the CGP process, the correlations are outside the boundary of -1 and 1 less than 0.05% of the time. If these boundaries are reached, we set the value to -1 or 1 .

FIGURE 3
Simulation-Based Regression Betas and Correlations Conditional on CGP

Figure 3 depicts the relationship between CGP and the serial correlation of consumption growth (Graph A), the conditional slope coefficient of shocks to expected consumption growth regressed on the contemporaneous consumption growth (Graph B), and the conditional predictive correlation between the monthly market risk premium and the lagged 1-month bond yield (Graph C). For the first 2 panels, quarterly values are in triangles, and the annual values implied by the model are shown in circles.



Graph A of Figure 3 shows how the first-order serial correlation of quarterly and annual consumption growth relates to CGP. Due to the presence of the LRR process, both quarterly and annual serial correlation are positive even for moderately negative values of ρ_t . For higher values of ρ_t the serial correlation is higher. The simulated ρ_t process has a mean of -0.128 and a standard-deviation of 0.23 , so it tends to fluctuate in the range from about -0.6 to 0.4 . The panel shows that when ρ_t increases from the bottom to the top of this range, the serial correlation of consumption growth increases from -0.01 to 0.09 at the quarterly frequency and from -0.01 to 0.18 at the annual frequency.

To further understand the driving forces behind the variation in serial correlation, we examine the conditional slope coefficient of the regression of shocks to expected consumption growth on contemporaneous realized consumption growth. These slope coefficients are plotted, as a function of CGP (ρ), in Graph B of Figure 3. The figure shows that a 1% increase in consumption growth leads to a -0.09% decrease in the expected growth when $\rho_t = -0.5$, but it leads to an increase of $+0.15\%$ when $\rho_t = 0.5$.⁸

E. Stock Return Predictability of Bond Yields

One implication of the model is that CGP should modulate the strength of the predictive relationship between bond yields and future stock returns. This result

⁸In the regressions, we match the horizons with which consumption growth and shocks to expected growth are measured.

enriches standard consumption-based asset pricing models, which imply that bond yields should negatively predict future stock returns. This occurs because the equity risk premium is increasing while bond yields are decreasing in consumption volatility, implying a negative relationship between the equity risk premia and bond yields.

While several studies, starting with Fama and Schwert (1977), find a negative relation between stock returns and lagged bond yields, the negative relationship appears to be sample dependent. Additionally, as evidenced by Welch and Goyal (2008) the significance level of this relation is well below that of other predictors, such as the aggregate dividend yield.

Our model implies that the strength of this form of stock market return predictability depends critically on the level of CGP. Bond yields are the inverse of the expected marginal utility of investors, which is closely related to the level of expected consumption growth. Meanwhile, the stock risk premium is increasing in volatility. Therefore, the negative relationship between bond yields and stock risk premia should be stronger when expected consumption growth is more negatively related to volatility.

As a result, given the negative correlation between uncertainty shocks and consumption, we expect the negative predictive relationship between future stock returns and bond yields to be stronger when CGP is high, or equivalently, when the SB correlation is low. Using the simulations described above, we examine the conditional correlation between stock risk premia and bond yields. This analysis uses the analytic formula for the market risk premium.

Graph C of Figure 3 shows that for values of ρ_t greater than -0.5 , there is a negative relationship between bond yields and the market risk premium, which is the case in most asset pricing models. But whereas other models imply that the degree of predictability is constant, our model suggests that it is highly time-varying. The figure implies that bond yields will be poor predictors of market returns when CGP is low or the SB correlation is high, but that predictability will increase as CGP rises and the SB correlation drops.

III. Empirical Results

The primary prediction of our model is the negative relationship between CGP and the stock–bond correlation. In this section, we test this prediction using quarterly consumption growth, several annual consumption measures, and an expected consumption growth series from the Survey of Professional Forecasters.

We estimate the stock–bond return correlation as the negative of the correlation between first differences in yields and stock returns, computed on a rolling basis using daily observations over the last 365 calendar days. This estimate approximates the true SB correlation, as it ignores the effect of convexity, but it nevertheless should be highly accurate.

Since the SB correlation can be measured using different bond maturities, we compute several such correlation series. Specifically, we report results based on 1-year and 10-year constant maturity bonds, although using other maturities produces very similar results. Given that inflation likely contaminates our measures of the SB correlation, we also calculate the SB correlation using real yields inferred

from TIPS prices. Real yields are the difference between the 10-year nominal yields and the 10-year break-even inflation rate. We refer to this correlation, also estimated using a 1-year rolling window, as the real SB correlation.

Establishing the connection between CGP, which is difficult to measure, and the SB correlation, which can be measured accurately using high-frequency asset price data, allows us to test additional implications of our model that would otherwise prove difficult. The later parts of this section test these implications using the SB correlation as a proxy for CGP.

A. Serial Correlation of Consumption Growth

A direct implication of our model is that the persistence of consumption growth shocks should be reflected in the level of the SB correlation. We test the relationship by examining autocorrelations in consumption growth at different horizons, where our model predicts that serial correlation will be larger in periods when the SB correlation is more negative. In interpreting these results, we note that the first-order autocorrelations obtained from consumption growth data are likely high due to time-aggregation effects and measurement issues that are absent from our theoretical model. As shown by Breeden, Gibbons, and Litzenberger (1989) and Heaton (1993), if investors make consumption decisions more frequently than the interval over which consumption is measured, first-order autocorrelation in growth rates may be as high as 0.25 in quarterly data even if higher frequency changes are unpredictable. However, serial correlations at longer lags should be immune to this effect.

We first estimate a predictive regression of quarterly consumption growth on its own lag. We then test whether this relationship is stronger during high or low SB correlation periods by adding an interaction term. The regression we estimate is

$$(6) \quad \Delta c_{t+k} = \alpha_0 + \alpha_1 \Delta c_t + \alpha_2 \hat{\rho}_{SB,t} \times \Delta c_t + \alpha_3 \hat{\rho}_{SB,t} + \epsilon_{t+k},$$

for $k = 1, 2, 3, 4$, where Δc_t is quarterly consumption growth and $\hat{\rho}_{SB,t}$ is one of the SB correlation series described previously. If, as implied by our model, the serial correlation is stronger for periods in which the SB correlation is negative, we should see a negative slope on the interaction term ($\alpha_2 < 0$).

Table 3 summarizes the results of these regressions. Panel A of the table first shows the simple regressions in which the only explanatory variable is a single lag of consumption growth. We observe significant positive autocorrelations at up to a four-quarter horizon, and the first-order autocorrelation is much larger than the value implied by time aggregation. These results are consistent with the base assumption of LRR in consumption growth, and the results are comparable to values reported by previous studies (e.g., Savov (2011)).

However, our primary interest is in testing the sign and the significance of the interactive coefficient, α_2 . The results of this regression are summarized in Panel B of Table 3 for multiple horizons. The regression results in the first 2 columns show that the first-order serial correlation increases when the SB correlation is more negative. The interaction coefficients are negative and statistically significant for both the 1-year and 10-year SB correlations. Quantitatively, a 0.1 increase in the SB

TABLE 3
CGP and the SB Correlation

Table 3 summarizes the slopes and Newey–West-adjusted (12 lags) *t*-statistics from quarterly regressions of the form

$$\Delta C_{t+k} = \alpha_0 + \alpha_1 \Delta C_t + \alpha_2 \hat{\rho}_{S,t} \times \Delta C_t + \alpha_3 \hat{\rho}_{S,t} + \varepsilon_{t+k}$$

for $k = 1, 2, 3, 4$, where $\hat{\rho}_{S,t}$ is the stock–bond return correlation (or stock–inflation correlation) estimated in one of several ways. Panel A summarizes the results under the restriction that $\alpha_2 = \alpha_3 = 0$, while Panel B uses the 1-year or 10-year nominal Treasury yield to compute the stock–bond correlation. The real stock–bond correlation and the stock–inflation correlation are used in Panel C after decomposing 10-year nominal yields into real and inflation components based on TIPS yields. The sample is quarterly from 1963 to 2019 (228 observations) in Panels A and B and from 2003 to 2019 (68 observations) in Panel C.

Panel A. Serial Correlation in Consumption Growth

Dependent Variable:	ΔC_{t+1}	ΔC_{t+2}	ΔC_{t+3}	ΔC_{t+4}
ΔC_t	0.494 (7.39)	0.370 (4.30)	0.420 (6.54)	0.186 (2.22)
Adj. R^2	0.219	0.134	0.173	0.030

Panel B. Persistence in Consumption Growth and the SB Correlation Based on Nominal Yields

Dependent Variable:	ΔC_{t+1}		ΔC_{t+2}		ΔC_{t+3}		ΔC_{t+4}	
	1Y	10Y	1Y	10Y	1Y	10Y	1Y	10Y
ΔC_t	0.480 (10.80)	0.482 (10.48)	0.367 (7.34)	0.371 (6.78)	0.397 (8.03)	0.406 (7.62)	0.150 (2.47)	0.164 (2.79)
$\Delta C_t \times \hat{\rho}_{SB,t}$	−0.400 (−2.79)	−0.312 (−2.17)	−0.690 (−3.76)	−0.520 (−2.87)	−0.372 (−1.92)	−0.294 (−1.94)	−0.475 (−2.39)	−0.164 (−3.38)
$\hat{\rho}_{SB,t}$	0.004 (3.38)	0.003 (3.46)	0.005 (4.39)	0.004 (4.00)	0.005 (3.19)	0.003 (2.59)	0.005 (2.96)	0.004 (2.88)
Adj. R^2	0.242	0.243	0.173	0.171	0.202	0.194	0.069	0.068

Panel C. Persistence in Consumption Growth and Real SB or Stock–Inflation Correlation

Dependent Variable:	ΔC_{t+1}		ΔC_{t+2}		ΔC_{t+3}		ΔC_{t+4}	
	$\hat{\rho}_{SR,t}$	$-\hat{\rho}_{SR,t}$	$\hat{\rho}_{SR,t}$	$-\hat{\rho}_{SR,t}$	$\hat{\rho}_{SR,t}$	$-\hat{\rho}_{SR,t}$	$\hat{\rho}_{SR,t}$	$-\hat{\rho}_{SR,t}$
ΔC_t	0.451 (2.85)	0.684 (3.40)	0.240 (1.69)	0.742 (4.20)	0.218 (3.55)	0.804 (3.12)	0.125 (1.68)	0.738 (4.29)
$\Delta C_t \times \hat{\rho}_{S,t}$	−0.451 (−0.87)	0.412 (0.92)	−1.138 (−2.42)	0.739 (1.73)	−1.394 (−4.46)	0.721 (0.91)	−0.895 (−2.29)	1.193 (1.85)
$\hat{\rho}_{S,t}$	0.005 (3.49)	0.002 (0.53)	0.007 (5.14)	−0.002 (−0.88)	0.008 (3.08)	−0.004 (−1.00)	0.007 (2.99)	−0.006 (−1.89)
Adj. R^2	0.232	0.231	0.346	0.275	0.424	0.338	0.186	0.153

correlation leads to a 0.03–0.04 decrease in the first-order serial correlation of consumption growth, which is slightly higher than what is implied by our model. We then increase the forecast horizon by replacing the one-quarter-ahead dependent variable with 1 that is between 2 and 4 quarters ahead. Overall, the results in this panel are all consistent with our model predictions.

A potential concern is that a significant fraction of the variation in nominal yields may be driven by inflation, which falls outside the scope of our model. Whether nominal yields change more due to expected inflation or real yields is controversial.⁹

⁹For example, Fama (1975) finds a strong relationship between the nominal interest rate and future inflation for the pre-1970 sample. More recently, Ang, Bekaert, and Wei (2008) also find that most of the monthly variation in nominal interest rates results from fluctuations in expected inflation. However, Mishkin (1992), Barr and Campbell (1997), and Duffee (2018a) suggest that most of the variation in the nominal interest rate is due to the real interest rate.

To study whether the above results originate from the correlation of stock returns with changes in the inflation rate or real bond yields, we repeat the previous analysis using the real SB correlation, which is based on real yields from 10-year Treasury Inflation-Protected Securities (TIPS). We start this analysis in 2003 to avoid well-known illiquidity problems in the early years of the TIPS market (e.g., Dudley, Roush, and Ezer (2009), Gürkaynak, Sack, and Wright (2010), and D'Amico, Kim, and Wei (2018)). We also estimate the correlation between stock returns and changes in the breakeven inflation rate. We multiply this correlation by negative one to be consistent with the signs in the first 2 panels.¹⁰

Panel C of Table 3 shows the results of estimating regression (6) after replacing the SB correlation with one of these 2 alternative measures. While no time variation in consumption persistence is found at the one-quarter horizon, higher-order correlations do appear to be significantly lower when the real SB correlation ($\hat{\rho}_{SR,t}$) is higher. This finding is again consistent with a negative relation between the SB correlation and CGP.

The panel also shows that the correlation between stock returns and changes in breakeven inflation ($\hat{\rho}_{S\pi,t}$) is generally uninformative about future consumption and is not significantly related to consumption persistence. If anything, a lower stock–inflation correlation is associated with more persistence in consumption growth, which again confirms the hypothesis that inflation effects are not responsible for the results given in Table 3. These results support our earlier conclusion that nominal SB correlations are informative in this setting because they are related to the corresponding correlation based on real yields.

While the regressions shown in Table 3 examine nonoverlapping growth rates at different horizons, Figure A2 in the Supplementary Material examines consumption growth autocorrelations using overlapping longer-horizon growth rates. Results for these longer horizons are consistent with those in Table 3.

Although NIPA consumption measures are widely used in both macroeconomics and macrofinance literatures, several problems with these data have been identified. The aggregation bias mentioned above implies that the measured consumption growth at the quarterly level may have a serial correlation of 0.25 or more even in the absence of any true persistence in growth rates. Furthermore, Breeden et al. (1989) suggests that the effect of time-aggregation bias is worsened by the use of interpolation in the construction of NIPA consumption data. Finally, consumption decisions made at different fixed intervals across agents, as suggested by Grossman and Laroque (1990) and Lynch (1996), generate serial correlation even in the absence of any correlation between current and expectations of consumption growth.

While some of these concerns are alleviated by the evidence showing longer-term persistence, we perform additional analysis using alternative measures of consumption for added robustness. First, we measure consumption only at the fourth quarter, following Jagannathan and Wang (2007). Given that the tax year

¹⁰In computing the stock–bond correlation, we approximate the value by taking the negative of the correlation between stock returns and first differences in bond yields. Since changes in nominal bond yields can be decomposed into the sum of changes in real bond yields and changes in expected inflation, we also take the negative value of the correlation between stock returns and inflation shocks.

TABLE 4
CGP and the SB Correlation Using Alternative Consumption Measures

Table 4 summarizes the slopes and Newey–West-adjusted (3 lags) *t*-statistics of the regression considered in Table 3 but estimated with annual consumption data. “NIPA Annual” is the annual consumption on nondurables and services, “Q4” uses the Q4-to-Q4 consumption growth series of Jagannathan and Wang (2007), while “Unfiltered” is the unfiltered consumption series of Kroencke (2017). Consumption data used in this table are from Tim Kroencke’s website. The sample period used for Panel A is from 1962 to 2018 (56 observations). The sample in Panel B begins in 2003 (16 observations).

Dependent Variable: ΔC_{t+1}									
Consumption Series	NIPA Annual			Q4			Unfiltered		
<i>Panel A. Persistence in Consumption Growth and the SB Correlation Based on Nominal Yields</i>									
Maturity:		1Y	10Y		1Y	10Y		1Y	10Y
ΔC_t	0.497 (4.25)	0.480 (7.05)	0.476 (6.63)	0.393 (3.21)	0.390 (4.88)	0.386 (5.10)	0.030 (0.25)	0.070 (1.16)	0.239 (3.89)
$\Delta C_t \times \hat{\rho}_{SB,t}$		−0.844 (−2.60)	−0.647 (−3.06)		−0.844 (−2.74)	−0.674 (−2.58)		−1.069 (−3.71)	−0.691 (−2.88)
$\hat{\rho}_{SB,t}$		0.023 (2.41)	0.016 (3.28)		0.022 (2.23)	0.017 (2.90)		0.014 (1.15)	0.007 (1.41)
Adj. R^2	0.234	0.262	0.266	0.139	0.160	0.172	−0.018	0.006	0.033
<i>Panel B. Persistence in Consumption Growth and Real SB or Stock–Inflation Correlation</i>									
$\hat{\rho}_{S,t}$ Used:		$\hat{\rho}_{SR,t}$	$-\hat{\rho}_{SR,t}$		$\hat{\rho}_{SR,t}$	$-\hat{\rho}_{SR,t}$		$\hat{\rho}_{SR,t}$	$-\hat{\rho}_{SR,t}$
ΔC_t	0.583 (8.51)	−0.021 (−0.43)	0.844 (1.59)	0.393 (3.21)	−0.083 (−1.52)	1.852 (13.92)	0.030 (0.27)	−0.963 (−6.78)	1.968 (6.07)
$\Delta C_t \times \hat{\rho}_{S,t}$		−5.143 (−4.23)	0.630 (0.44)		−3.568 (−4.70)	4.581 (7.48)		−5.854 (−5.20)	5.087 (5.66)
$\hat{\rho}_{S,t}$		0.075 (4.14)	−0.011 (−0.37)		0.055 (10.72)	−0.094 (−6.94)		0.112 (4.78)	−0.117 (−4.44)
Adj. R^2	0.338	0.683	0.250	0.139	0.551	0.400	−0.018	0.472	0.337

ends in December, they argue that investors are more likely to make joint consumption/investment decisions in the fourth quarter (Q4) of each year. In addition, as noted by Savov (2011) and Kroencke (2017), using fourth-quarter consumption also mitigates time-aggregation bias. Second, we use the unfiltered consumption measure of Kroencke (2017), who applies a backward recursion to the filtered NIPA consumption process. The resulting unfiltered series also mitigates time-aggregation bias by adding a correction factor to the recursion.

Table 4 summarizes the results of regression (6) using these alternative consumption series. Since these measures are constructed at the annual frequency, we compare them with NIPA consumption measured at the annual frequency. Panel A shows results based on the nominal SB correlation, while Panel B uses the real SB correlation as well as the correlation between stock returns and the breakeven inflation rate.

Overall, the results using these alternative consumption data are qualitatively similar to those of Table 3. Focusing on Panel A, there is a positive serial correlation of consumption growth measured either using NIPA annual consumption or NIPA Q4 consumption. The positive first-order serial correlation disappears, however, if the unfiltered data of Kroencke (2017) is used to compute consumption growth.

The main focus of this article is whether this serial correlation varies with the SB correlation. Similar to previous results, we test this relation using the interactive

regression in equation (6). For all 3 annual consumption measures considered, the coefficient on the interaction term is negative and statistically significant. These results confirm the hypothesis that consumption growth is more persistent when the SB correlation is negative.

Panel B of Table 4 provides results based instead on the real SB correlation or the inflation component of the SB correlation. We find comparable results to Panel C of Table 3, in that inflation does not appear to be the primary driver of our main findings. The difference in this table is that the inflation component of the SB correlation is now statistically significant, though with a positive sign, which is again inconsistent with an inflation driving our results based on nominal yields.

B. Realized and Expected Consumption Growth

The previous section confirms that the serial correlation of consumption growth is related to the SB correlation. In this section, we examine whether the relationship between shocks to the long-run consumption growth forecast ($\Delta\hat{x}_t$) and to current consumption growth (Δc_t) also depends on the level of the SB correlation.

Table 5 presents the results of this arguably more direct test of the model. Since the timing of SPF survey forecasts during the quarter is vague, we use annual data to test this relationship. Each panel differs by how consumption is measured. Focusing on the first regression, when changes in expected growth are regressed on contemporaneous consumption growth, we find a weak negative relationship between the 2 shocks.

We next test whether the growth forecast is more positively correlated for the period beginning 1999, during which time the SB correlation was usually negative. If there was a regime shift around 1999 that increased the overall level of CGP, we expect a more positive relationship between consumption growth and its forecast in the later period. We therefore estimate the regression

$$(7) \quad \Delta\hat{x}_t = \alpha_0 + \alpha_1 \Delta c_t + \alpha_2 1_{99+,t} \times \Delta c_t + \alpha_3 1_{99+,t} + \varepsilon_{1,t},$$

where 1_{99+} is a dummy variable that takes a value of 1 starting in 1999 and 0 before, and where $\hat{x}_t$ is the long-run SPF forecast of consumption growth.

The second column of each panel of Table 5 shows the results of this regression, where our primary focus is on whether $\alpha_2 > 0$, as our hypothesis suggests. All 3 regressions show that the relationship between long-run SPF forecasts and current consumption growth is more negative in the earlier sample.

Given the negative relation between CGP and the SB correlation, the model predicts that the α'_2 coefficient in the regression

$$(8) \quad \Delta\hat{x}_t = \alpha'_0 + \alpha'_1 \Delta c_t + \alpha'_2 \hat{\rho}_{SB,t} \times \Delta c_t + \alpha'_3 \hat{\rho}_{SB,t} + \varepsilon_{2,t}$$

to be negative. We report these regression results in the last 3 columns of each panel of Table 5. The results show that the relationship between current consumption growth and long-run expectations is more positive when the SB correlation is lower, confirming our theoretical prediction in Graph B of Figure 3. Similar results are

TABLE 5
Consumption Growth and Survey Forecasts

Table 5 summarizes the slopes and Newey–West-adjusted (3 lags) t -statistics from regressing the first difference in the SPF long-run consumption growth forecast on contemporaneous consumption growth. Some regressions also include interactions with a 19+ year dummy or with the SB correlation, which is estimated using 1-year or 10-year nominal yields or with the 10-year real yield, as well as main effects for these variables. “NIPA Annual” is the annual consumption on nondurables and services, “Q4” uses the Q4-to-Q4 consumption growth series of Jagannathan and Wang (2007), while “Unfiltered” is the unfiltered consumption series of Kroencke (2017). There are 38 observations (1981–2018), except for when the 10-year real SB correlation is used (2003–2018, 16 observations).

		Dependent Variable: $\Delta \hat{x}_t$			
Maturity:			1Y	10Y	10Y Real
<i>Panel A. NIPA Annual Consumption</i>					
Δc_t	−0.161 (−2.12)	−0.450 (−3.46)	−0.217 (−3.81)	−0.204 (−2.51)	−0.265 (−7.57)
$\Delta c_t \times 1_{99+}$		0.460 (3.47)			
$\Delta c_t \times \hat{\rho}_{SB,t}$			−1.013 (−4.56)	−0.650 (−3.13)	−1.101 (−3.61)
1_{99+}		−0.010 (−3.03)			
$\hat{\rho}_{SB,t}$			0.020 (5.33)	0.012 (2.70)	0.013 (3.22)
Adj. R^2	0.103	0.325	0.371	0.324	0.247
<i>Panel B. Q4 Consumption</i>					
Δc_t	−0.121 (−1.86)	−0.343 (−2.75)	−0.016 (−2.86)	−0.145 (−2.09)	−0.276 (−5.71)
$\Delta c_t \times 1_{99+}$		0.371 (2.98)			
$\Delta c_t \times \hat{\rho}_{SB,t}$			−0.736 (−3.47)	−0.473 (−2.52)	−0.940 (−4.81)
1_{99+}		−0.008 (−2.45)			
$\hat{\rho}_{SB,t}$			0.015 (3.58)	0.009 (2.22)	0.010 (4.87)
Adj. R^2	0.057	0.212	0.241	0.193	0.201
<i>Panel C. Unfiltered Consumption</i>					
Δc_t	−0.073 (−1.76)	−0.177 (−2.89)	−0.101 (−4.27)	−0.085 (−2.52)	−0.067 (−2.61)
$\Delta c_t \times 1_{99+}$		0.192 (3.13)			
$\Delta c_t \times \hat{\rho}_{SB,t}$			−0.470 (−8.18)	−0.315 (−3.63)	−0.408 (−2.54)
1_{99+}		−0.005 (−2.43)			
$\hat{\rho}_{SB,t}$			0.010 (3.83)	0.006 (2.56)	0.004 (1.48)
Adj. R^2	0.062	0.175	0.271	0.229	−0.023

obtained for SB correlations based on the 1-year and 10-year nominal yields and the 10-year real yield.

C. The Equity Risk Premium and Consumption Growth Persistence

The model presented suggests that the equity risk premium should be higher when CGP is higher. This would be the case assuming other state variables remain constant. In this section, we examine the evidence on whether the equity risk premium was lower in the pre-1999 period compared to the 1999+ period. If

TABLE 6
Moments of Consumption Growth and Stock Returns

Table 6 shows the standard-deviations, the first-order autocovariances and autocorrelations of consumption growth, and mean excess stock market returns for the pre-1999 and 1999+ samples. "NIPA Annual" is the annual consumption on nondurables and services, "Q4" uses the Q4-to-Q4 consumption growth series of Jagannathan and Wang (2007), while "Unfiltered" is the unfiltered consumption series of Kroencke (2017). The standard-deviations are in monthly terms, and the autocovariances are in monthly units multiplied by 10,000.

Panel A. Moments of Consumption Growth

		1962–1998	1999–2019
NIPA quarterly	Standard-deviation	0.0026	0.0021
	Autocovariance	0.0264	0.0245
	Autocorrelation	0.3939	0.5912
NIPA annual	Standard-deviation	0.0035	0.0034
	Autocovariance	0.0400	0.0652
	Autocorrelation	0.3236	0.6121
Q4 consumption	Standard-deviation	0.0041	0.0035
	Autocovariance	0.0383	0.0608
	Autocorrelation	0.2265	0.5356
Unfiltered	Standard-deviation	0.0077	0.0058
	Autocovariance	−0.0704	0.0645
	Autocorrelation	−0.1207	0.1946

Panel B. Average Market of Consumption Growth

		1962–1998	1999–2019
Market returns	Average	0.0052	0.0054

CGP was higher for the latter period, we expect the equity risk premium to be higher as long as other state variables, such as consumption volatility, do not have offsetting effects.

Table 6 describes the standard-deviation, autocovariance, and autocorrelation of the consumption process as well as the average stock market excess return for each the 2 subsamples. Panel A summarizes the moments of consumption growth, measured in several ways, whereas Panel B shows mean excess stock returns.

Table 6 shows that consumption volatility has decreased somewhat in the 1999+ sample, particularly when measured using the less noisy unfiltered consumption series or the fourth-quarter consumption measure. While the autocorrelation of consumption growth is much higher in the second sample, as reported earlier, the lower volatility in the later sample should be expected to have an offsetting effect on equity risk premia. In fact, Panel B of Table 6 shows similar mean excess stock returns for the two sample periods.

These moments suggest why we do not observe a higher equity risk premium in a sample with higher CGP. Our results are also consistent with Kozak (2022), who shows no clear evidence of any predictive relation between the SB correlation and future stock market returns, arguing that the negative relation between changes in the SB correlation and stock return nevertheless represents indirect evidence.

D. Stock Return Predictability Using Bond Yields

While we find no direct relation between CGP and future stock returns, our model also implies that CGP should drive the conditional relationship between

bond yields and future stock returns. The results we show here provide direct evidence of stock market predictability, but it is driven by the interaction between the SB correlation and bond yields rather than the SB correlation alone.

A number of articles have studied Treasury bond yields as stock market return predictors. Fama and Schwert (1977) estimate a simple predictive regression of future stock returns on lagged bond yields and find a negative slope, which they interpret as the result of stocks being inflation hedges. Breen, Glosten, and Jagannathan (1989) further confirm the economic significance of this predictability. More recently, Ang and Bekaert (2007) find that short-term Treasury yields and dividend yields jointly predict stock returns in many international markets. They argue that the yields represent a component of the discount rate used by investors to value equities. Campbell and Thompson (2008) also document statistically significant in-sample predictability, but Welch and Goyal (2008) report weak in- and out-of-sample performance.

Our model suggests that the extent to which the bond yield predicts stock returns depends on CGP. Specifically, a higher CGP is associated with a more negative predictive slope coefficient between bond yields and future returns. Given the relation between CGP and the stock–bond correlation, we test this hypothesis in monthly regressions of the form

$$(9) \quad R_{S,t,t+\tau}^e = b_0 + b_1 y_t + b_2 y_t \times \hat{\rho}_{SB,t} + \epsilon_{t+1},$$

where $R_{S,t,t+\tau}^e$ is the τ -month excess market return, y_t is a bond yield, and $\hat{\rho}_{SB,t}$ is a SB correlation estimated from a 1-year rolling window of daily returns. Panels A and B of Table 7 consider regressions based on nominal yields, while Panel C examines real yields over the post-2002 subsample.

Panel A of Table 7 first considers the simple predictive regression in which leading stock returns are regressed on bond yields alone, separately for samples before and after 1999, using either the 1-year or 10-year Treasury yield. Comparing the 2 subsamples, we see that the regression slope coefficient is uniformly negative and significant in the latter sample, with high R^2 s, but generally insignificant in the earlier period. These results are consistent with consumption growth being more persistent for the sample beginning in 1999 and the sign change in the SB correlation that occurred around that time.

The novel implication of our model is that the slope should be more negative when the SB correlation is lower, implying $b_2 > 0$ in equation (9). The results, summarized in Panel B of Table 7, show strong support for this prediction, with b_2 coefficient estimates that are positive and significant across all horizons for both yields. To understand the strength of the relation, consider the relation between 1-month excess market returns and 1-year Treasury yields that would hold if the SB correlation were equal to 0.4. In this case, the conditional slope coefficient would be a paltry -0.045 ($-0.267 + 0.554 \times 0.4$), implying that yields have essentially no predictive power for future returns. Similar conclusions hold for longer investment horizons as well. However, were the SB correlation instead equal to -0.5 , the conditional slope coefficient would increase in magnitude to around -0.544 . A 1 percentage point increase in the 1-year Treasury yield would then be associated with a 0.5% decline in monthly stock returns and, following the same logic, a 1.7%

TABLE 7
Market Return Predictability Using Bond Yields

Table 7 summarizes the results of the regression

$$R_{S,t+\tau}^e = \beta_0 + \beta_1 y_t + \beta_2 y_t \times \hat{\rho}_{SB,t} + \epsilon_{t+1},$$

where $R_{S,t+\tau}^e$ is the value-weighted market excess return over a 1-month, 3-month, 6-month, or 12-month horizon (τ), and y_t is either the 1-year nominal ($y_{1,t}$), 10-year nominal ($y_{10,t}$), or 10-year real yield ($y_{r,10,t}$). $\hat{\rho}_{SB,t}$ is the estimated correlation between stock and bond returns, which is constructed using the same yield used for y_t . Panel A shows the results in which future stock returns are regressed on nominal bond yields separately for 2 periods before and after the end of 1998 (with $\beta_2 = 0$). Panel B estimates the full model using the entire sample, again using nominal yields, for the sample period of 1963–2019. Panel C shows regression results for simple and interactive regressions based on 10-year real yields for the sample period of 2004–2019. T -statistics, in parentheses, are adjusted for heteroscedasticity and autocorrelation using Newey–West standard errors with 12 lags.

Panel A. Simple Predictive Regressions Using Nominal Yields

Horizon:	1962–1998				1999–2019			
	1-Month	3-Month	6-Month	12-Month	1-Month	3-Month	6-Month	12-Month
$y_{1,t}$	−0.154 (−1.94)	−0.296 (−1.71)	−0.368 (−0.97)	−0.212 (−0.28)	−0.315 (−2.20)	−0.972 (−3.07)	−2.081 (−3.48)	−4.553 (−4.05)
Adj. R^2	0.007	0.008	0.006	−0.001	0.015	0.052	0.111	0.254
$y_{10,t}$	−0.083 (−0.98)	−0.126 (−0.59)	−0.095 (−0.21)	0.310 (0.36)	−0.584 (−2.90)	−1.709 (−3.79)	−3.462 (−3.98)	−6.503 (−4.28)
Adj. R^2	0.000	−0.001	−0.002	0.000	0.027	0.079	0.150	0.250

Panel B. Interactive Predictive Regressions Using Nominal Yields (684 Observations)

Horizon:	1-Month		3-Month		6-Month		12-Month	
	$k = 1$	$k = 10$	$k = 1$	$k = 10$	$k = 1$	$k = 10$	$k = 1$	$k = 10$
$y_{k,t}$	−0.267 (−2.86)	−0.436 (−3.74)	−0.770 (−3.72)	−1.176 (−4.11)	−1.440 (−3.70)	−2.120 (−3.81)	−2.142 (−2.86)	−2.759 (−2.78)
$y_{k,t} \times \hat{\rho}_{SB,t}$	0.554 (2.35)	0.581 (3.40)	1.762 (2.94)	1.616 (3.52)	3.465 (3.13)	2.980 (3.40)	5.138 (2.43)	4.006 (2.70)
Adj. R^2	0.016	0.020	0.044	0.052	0.079	0.085	0.088	0.075

Panel C. Predictive Regressions Using 10-Year Real Yields (192 Observations)

Horizon:	1-Month		3-Month		6-Month		12-Month	
$y_{r,10,t}$	−0.788 (−2.54)	−0.207 (−0.75)	−2.119 (−2.52)	−0.308 (−0.69)	−3.306 (−2.34)	−0.063 (−0.04)	−5.458 (−2.01)	−1.741 (−0.61)
$y_{r,10,t} \times \hat{\rho}_{SR,t}$		4.834 (2.81)		15.079 (3.63)		27.005 (3.13)		30.953 (2.18)
Adj. R^2	0.012	0.075	0.060	0.256	0.065	0.343	0.095	0.284

decline in 3-month returns, a 3.2% decline in 6-month returns, and a 4.7% decline in 12-month returns. Economic magnitudes are similarly large when based on 10-year yields.

Panel C of Table 7 repeats these regressions using real yields and the real SB correlation rather than nominal values. Overall, we see similar results, albeit with noticeably higher interaction coefficients. The coefficients are all statistically significant, echoing previous panels. A potential reason for the large interaction effects is the shorter sample period considered, in which the real SB correlation does not vary as much as it does over the full sample.

Many asset pricing models imply a negative relationship between bond yields and the stock risk premium, as high levels of uncertainty mean lower bond yields and a higher risk premium. Therefore, it is puzzling that the empirical relationship is so weak. Our results show that the predictive relationship is

stronger than it may appear, but only during periods when proxies indicate that CGP is high.

IV. The Recent Shift in Consumption Growth Persistence

The results presented so far suggest that an increase in the persistence of consumption growth was largely responsible for the shift in signs of the SB correlation that happened around 1999. But what was the reason for the significant increase in CGP observed around this time?

Production-based models, such as Kaltenbrunner and Lochstoer (2010), suggest that an increase in the magnitude of permanent productivity shocks will cause CGP to rise when consumers face adjustment costs. More volatile persistent shocks will also increase the variability in valuation ratios to a much greater extent than more volatile transient shocks. Consistent with this, the period starting in 1999 is notable for its inclusion of several major asset market “bubbles” and crashes (e.g., the “dot-com” crash, the real estate boom, and the Great Financial Crisis).

Regardless of the reason for such fluctuations, any permanent shock to asset values may be expected to produce some level of consumption persistence. While this claim cannot be demonstrated within our model, which features an exogenous consumption process, it is intuitive. According to the permanent income hypothesis, a positive wealth shock will raise the consumption level in perpetuity. However, if agents face adjustment costs, this higher level will not be reached immediately. Instead, they will experience a sequence of consecutive positive growth rates as consumption rises to its new steady state level. If wealth falls, a sequence of negative consumption growth rates will result. In a period with multiple booms and busts, consumption is constantly trending toward some target value, but that target is moving, sometimes higher and sometimes lower than current consumption, raising CGP.

In this section, we present some evidence, which we view as suggestive, that is consistent with the greater consumption persistence of this period being driven by fluctuating asset values. In particular, we show that consumption growth rates have become more responsive to past asset returns since 1999. These results are in line with recent articles by Laibson and Mollerstrom (2010), Mian and Sufi (2011), and Chen, Michaux, and Roussanov (2020), which show a stronger tendency for consumption to be driven by fluctuating asset valuations over this period. Our new finding is that asset returns affect consumption growth at longer horizons than documented previously, particularly in the latter period.


We first demonstrate the changing relation between asset returns and consumption growth by computing the predictive correlation ($\text{Corr}(R_t, \Delta c_{t+k})$) for different horizons k and using returns R_t on different wealth proxies, as defined in the Supplementary Material. We analyze horizons of 1 to 12 quarters and compute the correlations separately for the period before 1999 and starting in 1999.

The results, shown in Figure 4, suggest a major shift in the predictive relationship, particularly at longer horizons. Prior to 1999, the predictive

FIGURE 4

Consumption Growth Response to Past Asset Returns: Pre-1999 Versus 1999+

Figure 4 computes the predictive correlation $\text{Corr}(\Delta c_{t+k}, R_{t,t})$ for $k = 1, \dots, 12$, where $R_{t,t}$ is the market return (Graph A), asset growth (Graph B), the housing price index return (Graph C), or net worth growth (Graph D) for the pre-1999 and 1999+ sample. Solid lines denote estimates, while dashed lines represent critical values for a 5% significance level. The lower critical values are omitted for Graphs A and B.



relationship between asset returns and future consumption growth was relatively weak. Statistical significance, which is indicated by the estimated correlation exceeding the corresponding dashed line, is observed in some cases, but mainly at short horizons. For the sample starting in 1999, correlations are, in many cases, twice as large, if not more, and highly significant even at multiyear lags. The largest shift occurs in response to housing returns, to which consumption responded little before 1999 (and perhaps even negatively at long horizons) and very significantly thereafter.

The relationship between wealth and consumption growth also features critically in the model of Lettau and Ludvigson (2004). They find that fluctuations in the wealth-to-consumption ratio only weakly predict consumption growth but are strongly negatively related to future changes in aggregate wealth. These findings explain the success of the *cay* variable of Lettau and Ludvigson (2001) in predicting returns on the stock market. Because Lettau and Ludvigson's sample ends in 2003, their analysis focuses mainly on a period in which CGP is low, meaning that shocks tend to be transient. The effect of wealth shocks on consumption is therefore immediate, occurring mostly during the current period. Since the expected growth rate (x) responds little to wealth shocks, deviations in the wealth-to-consumption ratio must revert by a decline in the value of aggregate wealth.

TABLE 8
Consumption, Wealth, and Human Capital Dynamics Pre-1999 Versus 1999+

Panel A of Table 8 summarizes the results of estimating the vector error correction model

$$\Delta \mathbf{X}_{t+1} = \mathbf{u} + \zeta \text{cay}_t + \Gamma \Delta \mathbf{X}_t + \mathbf{e}_{t+1},$$

where $\mathbf{X}_t = [c_t, a_t, y_t]$ includes the log consumption, wealth, and human capital series from Sydney Ludvigson's website and cay_t is the deviation in their common trend, estimated separately for the pre-1999 (1962–1998) and 1999+ (1999–2019) sample. Panel B shows results for the predictive regression in which quarterly excess market returns are regressed on the lagged value of cay_t . In the first row of the panel, cay_t is taken directly from Sydney Ludvigson's website and is estimated using a single extended sample. In the second row, cay_t is estimated separately for each subsample, as in Panel A. The t -statistics reported in parenthesis uses the Newey–West standard errors with 12 lags.

Panel A. Vector Error Correction Model

	Pre-1999			1999+		
	Δc_{t+1}	Δa_{t+1}	Δy_{t+1}	Δc_{t+1}	Δa_{t+1}	Δy_{t+1}
Δc_t	0.099 (1.44)	0.204 (1.20)	0.248 (2.11)	0.790 (3.92)	0.700 (1.77)	0.378 (1.76)
Δa_t	0.165 (3.53)	0.175 (1.62)	0.083 (1.08)	0.040 (1.63)	0.149 (0.97)	0.062 (1.50)
Δy_t	0.156 (1.32)	−0.196 (−1.16)	−0.072 (−0.34)	0.034 (0.61)	0.142 (0.49)	−0.199 (−2.06)
cay_t	−0.026 (−0.81)	0.181 (1.98)	0.027 (0.62)	−0.213 (−3.11)	0.200 (0.34)	0.423 (3.31)

Panel B. Predicting Excess Stock Market Returns

	Pre-1999	1999+
cay_t estimated over extended sample	0.708 (2.81)	−0.831 (−1.42)
Adj. R^2	0.026	0.012
cay_t estimated separately for each subsample	0.946 (2.15)	−0.823 (−0.58)
Adj. R^2	0.024	−0.008

In contrast, during the high-CGP period starting around 1999, our previous analysis shows that growth in aggregate wealth has a substantial positive effect on the long-term consumption growth rate. Since shocks are persistent during this period, any growth in wealth relative to current consumption is also likely to increase future consumption. When transient deviations in the wealth–consumption ratio resolve by changing future consumption, such deviations no longer need forecast wealth changes or asset returns. This shift in the predictive information contained in the consumption–wealth ratio may underlie the poor postsample performance of cay_t that has been documented by Goyal, Welch, and Zafirov (2021) and others.

We estimate the cointegrated vector error correction model examined by Lettau and Ludvigson (2004). We use the data from Sydney Ludvigson's website and estimate the following model separately for samples before and after 1999:

(10)
$$\Delta \mathbf{X}_{t+1} = \mathbf{u} + \zeta \text{cay}_t + \Gamma \Delta \mathbf{X}_t + \mathbf{e}_{t+1}.$$

Here, $\mathbf{X}_t = [c_t, a_t, y_t]$ is a vector of log consumption, wealth, and human capital, and cay_t is the deviation from common trend, which we estimate separately for each subsample.¹¹

¹¹In order to follow standard notation, this section uses y_t to denote human capital, which is different from the rest of this article.

Panel A of Table 8 shows the estimation results. For the pre-1999 sample, we find similar results to Lettau and Ludvigson (2004). *cay* positively predicts asset growth, and there is no predictability of consumption growth. However, focusing on the later sample, there is no predictability of future asset growth. Instead, *cay* negatively predicts future consumption growth.

In Panel B of Table 8, we re-examine the predictive regression, where quarterly excess market returns are regressed on *cay*. We construct different *cay* series from each of the split samples analyzed in Panel A. For completeness, we also report the results from using the *cay* series directly taken from Ludvigson, which is estimated over a single extended sample period.

The results show that *cay* only predicts market excess returns for the pre-1999 sample, which approximately coincides with the sample studied by Lettau and Ludvigson (2001). For the sample beginning in 1999, we find statistically insignificant relationship between the consumption–wealth ratio and the leading excess market returns. These results support our hypothesis that the degree of CGP determines the shocks' influence on future asset growth.

These results are not meant to be definitive, and a more careful analysis would likely require a further enhancement of the theoretical framework that we adopt here. Nevertheless, they suggest a mechanism that may underlie the changes in CGP and related changes in the stock–bond correlation that we have documented.

V. Conclusion

While the consumption process examined by Bansal and Yaron (2004) is highly successful in replicating key moments of asset returns, its assumption of independent shocks is inconsistent both with macroeconomic theory and with consumption data. In particular, the model does not account for the relationship between shocks to current consumption growth and expected future consumption growth, which we term CGP.

Because of these assumptions, the model cannot match several well-documented features of financial markets. Most significantly, the correlation between stocks and bonds is highly time-varying in the data and varies with the level of stock market volatility. These effects are absent in the model of Bansal and Yaron, which features a constant stock–bond correlation.

We propose a model that allows for a significantly more realistic dependence structure. Shocks to current and expected future consumption growth are stochastically correlated, which we view as a reduced form approach to modeling the relative importance of transitory and permanent shocks. Shocks to current consumption and consumption growth are negatively correlated at a fixed value, which maintains parsimony and reflects the likely importance of the precautionary savings motive.

The model implies that the correlation between stock and bond returns is decreasing in CGP. Empirically, we see that consumption growth tends to become more serially correlated in periods of more negative stock–bond correlations. This result provides evidence of time variation in CGP and also links CGP to correlations that are readily estimable from high-frequency asset price data.

Our model also predicts the negative relation between stock market volatility and the SB correlation observed in prior studies, such as Connolly et al. (2005) and Baele et al. (2010). This is the case because high consumption persistence makes cash flows and discount rates negatively correlated, which amplifies the effects of these shocks. Empirically, we find strong evidence for this relation.

Finally, we show in our model that bond yields should negatively predict future market excess returns, but only if CGP is sufficiently high. Again using the SB correlation to proxy for the unobserved CGP process, we find strong confirmation of this prediction in the significance of the coefficient on the interaction of bond yields and the SB correlation. We also show that the source of this predictability is the real SB correlation rather than the component related to inflation.

Combined with the observation that CGP has increased markedly since 1999, our model provides a new explanation for the dramatic downward shift in the stock–bond correlation occurring around this time. Additional evidence suggests that the increase in persistence may be the result of a greater role of asset valuations in driving long-run consumption growth rates. The result may be a marked change in the cointegrating relationship, studied by Lettau and Ludvigson (2004), between consumption, labor income, and aggregate wealth.

Overall, time-varying CGP accounts for a number of stylized facts that are typically not linked together and whose explanations are still not fully understood. Furthermore, it uses an intuitive and relatively simple generalization of the standard LRR framework. As researchers examine the conditional implications of LRR more closely, it seems natural that time-varying correlations should play an important role.

Supplementary Material

To view supplementary material for this article, please visit <http://doi.org/10.1017/S002210902400019X>.

References

- Ang, A., and G. Bekaert. “Stock Return Predictability: Is it There?” *Review of Financial Studies*, 20 (2007), 651–707.
- Ang, A.; G. Bekaert; and M. Wei. “The Term Structure of Real Rates and Expected Inflation.” *Journal of Finance*, 63 (2008), 797–849.
- Baele, L.; G. Bekaert; and K. Inghelbrecht. “The Determinants of Stock and Bond Return Comovements.” *Review of Financial Studies*, 23 (2010), 2374–2428.
- Bansal, R.; D. Kiku; and A. Yaron. “An Empirical Evaluation of the Long-Run Risks Model for Asset Prices.” *Critical Finance Review*, 1 (2012), 183–221.
- Bansal, R., and A. Yaron. “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles.” *Journal of Finance*, 59 (2004), 1481–1509.
- Barr, D. G., and J. Y. Campbell. “Inflation, Real Interest Rates, and the Bond Market: A Study of UK Nominal and Index-Linked Government Bond Prices.” *Journal of Monetary Economics*, 39 (1997), 361–383.
- Basu, S., and B. Bundick. “Uncertainty Shocks in a Model of Effective Demand.” *Econometrica*, 85 (2017), 937–958.
- Breeden, D. T.; M. R. Gibbons; and R. H. Litzenberger. “Empirical Test of the Consumption-Oriented CAPM.” *Journal of Finance*, 44 (1989), 231–262.

- Breen, W.; L. R. Glosten; and R. Jagannathan. "Economic Significance of Predictable Variations in Stock Index Returns." *Journal of Finance*, 44 (1989), 1177.
- Campbell, J., and A. Deaton. "Why is Consumption so Smooth?" *Review of Economic Studies*, 56 (1989), 357–373.
- Campbell, J. Y.; C. Pflueger; and L. M. Viceira. "Macroeconomic Drivers of Bond and Equity Risks." *Journal of Political Economy*, 128 (2020), 3148–3185.
- Campbell, J. Y.; A. Sunderam; and L. M. Viceira. "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds." *Critical Finance Review*, 6 (2017), 263–301.
- Campbell, J. Y., and S. B. Thompson. "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21 (2008), 1509–1531.
- Carroll, C. D. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, 112 (1997), 1–55.
- Carroll, C. D., and A. A. Samwick. "How Important is Precautionary Saving?" *Review of Economics and Statistics*, 80 (1998), 410–419.
- Chen, H.; M. Michaux; and N. Roussanov. "Houses as ATMs: Mortgage Refinancing and Macroeconomic Uncertainty." *Journal of Finance*, 75 (2020), 323–375.
- Chernov, M.; L. A. Lochstoer; and D. Song. "The Real Explanation of Nominal Bond-Stock Puzzles." UCLA and Johns Hopkins (2021).
- Connolly, R.; C. Stivers; and L. Sun. "Stock Market Uncertainty and the Stock-Bond Return Relation." *Journal of Financial and Quantitative Analysis*, 40 (2005), 161–194.
- D'Amico, S.; D. H. Kim; and M. Wei. "Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices." *Journal of Financial and Quantitative Analysis*, 53 (2018), 395–436.
- David, A., and P. Veronesi. "What Ties Return Volatilities to Price Valuations and Fundamentals?" *Journal of Political Economy*, 121 (2013), 682–746.
- Dudley, W.; J. Roush; and M. S. Ezer. "The Case for TIPS: An Examination of the Costs and Benefits." Federal Reserve Economic Policy Review, New York, NY (2009).
- Duffee, G. R. "Expected Inflation and Other Determinants of Treasury Yields." *Journal of Finance*, 73 (2018a), 2139–2180.
- Epstein, L. G., and S. E. Zin. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy*, 99 (1991), 263–286.
- Fama, E. F. "Short-Term Interest Rates as Predictors of Inflation." *American Economic Review*, 65 (1975), 269–282.
- Fama, E. F., and G. Schwert. "Asset Returns and Inflation." *Journal of Financial Economics*, 5 (1977), 115–146.
- Goyal, A.; I. Welch; and A. Zafirov. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction II." UCLA and University of Lausanne (2021).
- Grossman, S. J., and G. Laroque. "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica*, 58 (1990), 25–51.
- Gürkaynak, R. S.; B. Sack; and J. H. Wright. "The TIPS Yield Curve and Inflation Compensation." *American Economic Journal: Macroeconomics*, 2 (2010), 70–92.
- Hall, R. E., and F. S. Mishkin. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica*, 50 (1982), 461–481.
- Hasseltoft, H. "Stocks, Bonds, and Long-Run Consumption Risks." *Journal of Financial and Quantitative Analysis*, 47 (2012), 309–332.
- Heaton, J. "The Interaction Between Time-Nonseparable Preferences and Time Aggregation." *Econometrica*, 61 (1993), 353–385.
- Ilmanen, A. "Stock-Bond Correlations." *Journal of Fixed Income*, 13 (2003), 55–66.
- Jagannathan, R., and Y. Wang. "Lazy Investors, Discretionary Consumption, and the Cross-Section of Stock Returns." *Journal of Finance*, 62 (2007), 1623–1661.
- Kaltenbrunner, G., and L. A. Lochstoer. "Long-Run Risk Through Consumption Smoothing." *Review of Financial Studies*, 23 (2010), 3190–3224.
- Kozak, S. "Dynamics of Bond and Stock Returns." *Journal of Monetary Economics*, 126 (2022), 188–209.
- Kroencke, T. A. "Asset Pricing Without Garbage." *Journal of Finance*, 72 (2017), 47–98.
- Laibson, D., and J. Mollerstrom. "Capital Flows, Consumption Booms and Asset Bubbles: A Behavioural Alternative to the Savings Glut Hypothesis." *Economic Journal*, 120 (2010), 354–374.
- Lettau, M., and S. Ludvigson. "Consumption, Aggregate Wealth, and Expected Stock Returns." *Journal of Finance*, 56 (2001), 815–849.
- Lettau, M., and S. C. Ludvigson. "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption." *American Economic Review*, 94 (2004), 276–299.

- Lynch, A. W. "Decision Frequency and Synchronization Across Agents: Implications for Aggregate Consumption and Equity Return." *Journal of Finance*, 51 (1996), 1479–1497.
- Mian, A., and A. Sufi. "House Prices, Home Equity–Based Borrowing, and the US Household Leverage Crisis." *American Economic Review*, 101 (2011), 2132–2156.
- Mishkin, F. S. "Is the Fisher Effect for Real?" *Journal of Monetary Economics*, 30 (1992), 195–215.
- Nakamura, E.; D. Sergeyev; and J. Steinsson. "Growth-Rate and Uncertainty Shocks in Consumption: Cross-Country Evidence." *American Economic Journal: Macroeconomics*, 9 (2017), 1–39.
- Parker, J. A., and B. Preston. "Precautionary Saving and Consumption Fluctuations." *American Economic Review*, 95 (2005), 1119–1143.
- Pástor, L., and R. F. Stambaugh. "Liquidity Risk and Expected Stock Returns." *Journal of Political Economy*, 111 (2003), 642–685.
- Savov, A. "Asset Pricing with Garbage." *Journal of Finance*, 66 (2011), 177–201.
- Song, D. "Bond Market Exposures to Macroeconomic and Monetary Policy Risks." *Review of Financial Studies*, 30 (2017), 2761–2817.
- Swanson, E. "A Macroeconomic Model of Equities and Real, Nominal, and Defaultable Debt." Federal Reserve Bank of San Francisco (2019).
- Welch, I., and A. Goyal. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies*, 21 (2008), 1455–1508.