Predictability in Consumption Growth and Equity Returns: An Empirical Investigation

Paseka, A.[1], and G. Theocharides[2]

[1]I.H. Asper School of Business, University of Manitoba, Winnipeg, MB Canada, R3T 5V4
[2]SKK GSB, Sungkyunkwan University, Seoul, Korea, 110-745
Email: georghio@skku.edu

Keywords*: Consumption-based asset pricing, consumption growth predictability, return predictability, equity volatility puzzle*

EXTENDED ABSTRACT

In the last couple of decades, researchers have discovered a number of asset pricing "puzzles" that cannot be understood in isolation. For example, Mehra and Prescott (1985) described the "equity premium puzzle" as "too high" expected excess asset returns. This phenomenon could be resolved by increasing the level of risk aversion, but that would create the "risk-free rate puzzle" documented by Weil (1989). It would be impossible to reconcile the high level of risk-aversion with the low level of interest rate. A third important puzzle is the fact that the unconditional volatility of real stock returns has been excessively high relative to the unconditional volatility of the real consumption growth. This is the "equity volatility puzzle" documented by Campbell (2003).

One of the leading responses to the above puzzles is to modify the representative investor's utility preferences to include habit formation, while modeling consumption growth as a random walk (Campbell and Cochrane (1999)). Alternative solution is to model a small, predictable component in consumption growth process so that its annualized volatility increases with the horizon to the levels of the volatility of the wealth growth (Bansal and Yaron (2004)). The conflict between the approach of the above two models points to an ongoing debate in the consumption-based asset pricing literature. Is it a gradual adjustment/high persistence of consumption or mean reversion in wealth (with consumption growth as random walk)? In this paper we seek to answer the above question and estimate a consumption-based model based on Bansal and Yaron (2004) using fully Bayesian Markov Chain Monte Carlo (MCMC) methodology. Our anlalysis proceeds as follows: First, we estimate the model using equity returns, consumption growth and dividend growth data, and excluding interest rate data. This procedure does not force the model to fit low volatility interest rates and high volatility equity returns. Also, the model implies that the latent process must be perfectly correlated with p/d ratios. Our results however show that the latent process proxies for a state variable other than the p/d ratio. Mean consumption growth contains a small predictable latent component, but the persistence is not large enough to make the long term volatility of consumption growth equal to the observed volatility of aggregate wealth. This is indirect evidence of mean-reverting wealth rather than an upward trend in long term consumption volatility.

Second, we estimate the model using the full data set, that is, incorporating the risk free rate of return. We find that the small consumption persistence is robust to the inclusion of interest rate data. We also argue that the inconsistency between the latent state and the observed dynamics of p/d ratios results in the equity volatility puzzle and zero predictability of excess equity returns by price/dividend ratios under the model. Our estimates suggest that the model accounts only for 52% of total variation in asset returns.

1. INTRODUCTION

Mehra and Prescott (1985) described the "equity premium puzzle" as "too high" expected excess asset returns. Researchers have since discovered a number of other related phenomena that cannot be understood in isolation. For example, a natural explanation of the equity premium puzzle would be to increase the level of risk aversion but that would create the "risk-free rate puzzle" documented by Weil (1989). Investors need to have a low or even negative rate of time preference (i.e. give the same or even more weight to future consumption than current consumption), to reconcile the high level of risk-aversion with the low level of interest rate. A third important puzzle is the fact that the unconditional volatility of real stock returns has been excessively high relative to the unconditional volatility of the real consumption growth in the metric of consumption based asset pricing models. Campbell (2003) called this the "equity volatility puzzle".

One of the responses to the above puzzles is to modify the representative investor's utility preferences to include habit formation, while modeling consumption growth as a random walk. An important model in this direction is by Campbell and Cochrane (1999) (hereafter, CC (1999)). The equity premium still remains a puzzle, however, since a relatively high value of risk aversion is needed to capture it.

Alternative solution is to model a small, predictable component in consumption growth process so that its annualized volatility increases with the horizon to the levels of the volatility of the wealth growth. Bansal and Yaron (2004) (hereafter, BY(2004)) pursue this approach. They show that, if consumption growth is highly persistent, their model is able to explain the equity premium, the risk-free rate, and volatility puzzle, and achieve this with a relatively low level of risk aversion. The process by which this is achieved is as follows. As Lettau and Ludvigson (2001) showed, aggregate consumption, financial wealth, and labor income are cointegrated. Thus, this cointegration results in identical annual growth rates and volatilities between consumption and wealth in the long-run. Since consumption is very smooth in the short-run, this can be achieved in two ways; either the volatility of consumption increases with the horizon to reach the high volatility of wealth in the long run, or the volatility of wealth decreases with the horizon to reach the low volatility of consumption growth. The first scenario implies that consumption growth is not a random walk as suggested and modelled by CC, but rather that it has a small predictable component (i.e. positive serial correlation). This is the approach of BY(2004).

Although it is interesting that the assumed long run properties of consumption growth help to capture these empirical puzzles, it is not clear whether consumption growth is persistent enough in the data to explain the puzzles. The disparity between the above two modeling approaches highlights an ongoing debate in the consumption-based asset pricing literature. Is it a gradual adjustment/high persistence of consumption (similar to BY) or mean reversion in wealth (similar to CC)?

We estimate a consumption-based model based on BY(2004) using fully Bayesian Markov Chain Monte Carlo (MCMC) methodology. This technique has several advantages as opposed to classical methods. First, this approach allows us to infer the entire time series of the latent processes embedded in these models, in addition to parameter values. Second, we avoid parameter estimation risk, which could be a contributing factor to the above puzzles. Estimation risk can be especially problematic in problems involving latent state processes. However, using a Bayesian estimation the marginal posterior integrates out the parameters as opposed to classical methods, which obtain estimates conditional on parameters (see Johannes and Polson (2006)).

2. MODEL

In this section we provide a brief summary of a consumption-based asset pricing framework to learn about the equity premium, volatility of equity returns and the behavior of equilibrium interest rates. The asset pricing model we adopt here closely follows the model proposed by BY(2004).

Consumption smoothness and relatively low correlation between consumption growth and returns can be reconciled, at least formally, with high equity premia by assuming high relative risk aversion. Unfortunately, that may lead to high risk-free rates - the risk-free rate puzzle (see Weil (1989)). But even this can be managed by using a more general form of power utility developed by Epstein and Zin (1989, 1991) and Weil (1989). Under simple power utility, the elasticity of intertemporal substitution, ψ, is the reciprocal of the coefficient of relative risk aversion, γ. The more general form of utility does away with this tight link between the coefficients.

BY(2004) stress that a potential explanation of the puzzles may be provided through the flexible Epstein-Zin-Weil framework. Then, they use Campbell and Shiller's (1988) log-linear framework to derive the solution for the returns in their model. To avoid singularities in the likelihood function, we introduce measurement errors. Specifically, we assume that observed returns contain an i.i.d. mean zero measurement errors. Despite the beneficial theoretical properties of the model, the success of Epstein-Zin framework with a persistent latent factor relies heavily on the assumption that consumption growth has very persistent component (ρ close to 1). Yet empirical evidence on this is mixed at best (see Campbell (2003)). Compounding the problem is the estimation risk that is normally present in classical econometric approaches (e.g., maximum likelihood estimation with Kalman Filter) to state space models with latent processes.

We circumvent the estimation risk problem by using a Bayesian Markov Chain Monte Carlo approach to draw an inference on the model parameters and the latent process. The information in the data about the latent process is summarized in the marginal posterior distribution:

$$p(X \mid Data) = \int p(X, \Omega \mid Data) d\Omega \qquad (1)$$

where Ω is the parameter vector, X is the latent process and Y is the data. In classical estimation procedures we compute a latent variable estimate by conditioning on the parameters. This approach ignores sampling variation of the parameter estimates. In MCMC implementation all the remaining parameters are marginalized out of the likelihood.

We use a Gibbs sampler methodology to draw parameter as well as latent state values. Construction of the Gibbs sampler requires the specification of full conditional posteriors for all parameters and state variables. By Bayes rule a full conditional on a variable is proportional to the joint likelihood of the data given the parameters and latent process times the prior density on that variable, with all the remaining variables being treated as constants. To generate samples of parameters and state process and to draw inferences we use WinBUGS 1.4 package.

3. EMPIRICAL RESULTS

The data used in the estimation procedure are stock market index returns, personal consumption data on nondurables and services, a dividend series, as well as returns of the three-month T-Bill (as a proxy for the risk-free rate). These are all annual data for the period 1930-2005. Secondary market T-bill rates are available starting from 1934 only. We estimate model parameters based on both periods, 1930-2005 and 1934-2005. Although some sample autocorrelation and variance properties appear to be sample dependent, parameter estimates are fairly robust to the choice of the sample period. We report the estimates for 1934-2005 period, in which the data for T-bill rates are available. All nominal values are deflated using the Consumer Price Index - All Urban Consumers (CPI-U) taken from the Bureau of Labor Statistics.

For a stock market index, we use the value-weighted composite NYSE/AMEX/NASD index returns taken from CRSP. Personal consumption data on nondurables and services (in billions of dollars) comes from the Bureau of Economic Analysis (BEA). After deflating the data, real annual personal consumption growth is constructed. Dividends and the real dividend growth series is extracted from index returns with and without dividends. We obtain the returns on the three-month T-bill from the Federal Reserve web site.

Table 1 reports summary statistics for the index return and consumption and dividend growth. All three variables have non-zero excess kurtosis, evidence of fat tails. However, the magnitude is not large at the annual frequency we use. The skewness estimates are negative for returns and positive for consumption and dividend growth.

Table 1. Summary statistics for the value-weighted composite NYSE/AMEX/NASDAQ index returns, consumption growth, and dividend growth series for the period 1934-2005. The statistics are based on annual observations.

	Index Returns	Cons. Growth	Div. Growth
Mean	0.069	0.034	0.028
Median	0.109	0.035	0.021
Std. Dev.	0.176	0.016	0.137
Skewness	-0.726	0.511	0.593
Kurtosis	0.48	2.24	0.729

In Table 2 we present the parameter estimation results. The results are based on the entire sample of the data from 1934 to 2005. We monitored Gelman-Rubin statistic to track the end of the burn-in phase. This exercise resulted in running the Gibbs sampler for 5000 iterations. We discarded the first 1000 iterations of the burn-in phase and used the remaining 4000 for inference. We report all results in Table 2 with units corresponding to annual frequency. A casual comparison of our parameter estimates reveals certain non-trivial diffrences from those calibrated by BY. Most striking is the difference in the estimate of ρ, the parameter at the center of current model's ability to fit equity moments.

Table 2. Parameter estimates using annual observations for consumption and dividend growth and returns on NYSE/AMEX/NASD index over the period 1934-2005. All units correspond to annual frequency.

Parameter	Mean	Std. Dev.	Median
$A_{1,m}$	2.565	1.985	2.474
a_m	0.051	0.015	0.051
φ_m	9.513	1.479	9.397
φ_e	0.962	0.198	0.940
μ_m	0.020	0.015	0.020
μ	0.033	0.003	0.033
ϕ_m	0.478	1.676	0.455
$1/\psi$	-1.188	2.193	-1.163
ρ	0.327	0.194	0.337
σ	0.013	0.002	0.013
σ_m	0.1217	0.0121	0.1217

From Table 3, Panel A, the value of the posterior standard deviation of consumption growth is about 1.9% per year, consistent both with sample statistics from the data (Table 3, Panel B) and with previous estimates (see, e.g., Campbell (2003)). With small calibrated values for σ and φ_e, BY are able to fit this consumption moment due to very high persistence $\rho = 0.979$. High ρ implies that most of the variation in consumption growth is due to variation in the latent factor - conditional mean of consumption growth. Our estimate of ρ is only moderate. Hence, both conditional variance and the variance of the conditional mean have similar contributions to the unconditional variance of consumption growth unlike in BY(2004). We fit consumption growth, dividend growth and equity returns simultaneously. To reconcile these processes within the model with the data, parameter estimates must correspond to the moments of these processes. Parameter σ^2 primarily determines consumption growth volatility as long as persistence parameter ρ is not very close to one. The mean estimate of ρ is 0.33 for 1934-2005 periods and $\rho = 0.575$ for 1930-2005 period (we report only the former estimate in Table 2). This runs contrary to the belief that there is a highly persistent component in consumption growth. At the same time the model can fit a much more volatile dividend growth. Table 3 reports the standard deviation of the dividend growth in the data of 13% per year (12% over 1930 - 2005). Despite small σ-estimate, the estimation procedure fits the dividend growth through a large φ_m-estimate. The resulting standard deviation of dividend growth implied by the model is 12.25% per year that compares to 13% in the data (Table 3).

Table 3. Model-implied consumption and dividend growth moment estimates using annual data on NYSE/AMEX/NASD index returns, consumption and dividend growth for the period 1934-2005. We use MCMC parameter estimates to build posterior summaries of autocorrelations, γ, at different lags as well as unconditional standard deviations.

Panel A: Posterior Model Estimates (MCMC output)

Autocorr.(Lag)	Mean	Std. Dev.	Median
Consumption Growth Autocorrelations			
$\gamma(1)$	0.1703	0.1139	0.1636
$\gamma(2)$	0.0755	0.0741	0.0553
$\gamma(5)$	0.0110	0.0269	0.0020
$\gamma(7)$	0.0041	0.0170	0.0002
$\gamma(10)$	0.0013	0.0108	0.0000
Dividend Growth Autocorrelations			
$\gamma(1)$	0.0106	0.0192	0.0037
$\gamma(2)$	0.0048	0.0104	0.0011
$\gamma(5)$	0.0007	0.0032	0.0000
$\gamma(7)$	0.0003	0.0020	0.0000
$\gamma(10)$	0.0001	0.0013	0.0000
Unconditional Correlation Between Consumption and Dividend Growth			
	0.0360	0.1244	0.0348
Unconditional Standard Deviations			
Consumption Growth	0.0186	0.0017	0.0184
Dividend Growth	0.1225	0.0108	0.1218

Panel B: Estimates from the Data

Autocorr.(Lag)	Estimate	Std. Error	Estimate	Std. Error
	1934-2005		**1930-2005**	
Consumption Growth Autocorrelations				
$\gamma(1)$	0.224	0.118	0.416	0.115
$\gamma(2)$	-0.034	0.124	0.132	0.133
$\gamma(5)$	0.159	0.125	0.005	0.142
$\gamma(7)$	0.063	0.128	0.010	0.142
$\gamma(10)$	0.232	0.135	0.060	0.143
Dividend Growth Autocorrelations				
$\gamma(1)$	-0.177	0.118	-0.145	0.115
$\gamma(2)$	-0.148	0.121	-0.161	0.117
$\gamma(5)$	-0.025	0.124	-0.123	0.120
$\gamma(7)$	-0.119	0.125	-0.080	0.122
$\gamma(10)$	0.003	0.129	-0.005	0.126
Unconditional Corr. Between Consumption and Dividend Growth				
	-0.005	0.125	0.120	0.125
Unconditional Standard Deviations				
Cons. Growth	0.015		0.023	
Div. Growth	0.130		0.123	

Another way to assess the evidence of moderate persistence in consumption growth is to directly compare its sample autocorrelations against those implied by the model, specifically, according to how fast they decay over time. Model-implied autocovariances die out exponentially at the rate $\rho = 0.33$ (at the posterior mean). It follows then that the autocorrelations also must decrease exponentially at the same speed.

Table 3 reports the standard deviations of consumption and dividend growth, their correlation and autocorrelations from the data (Panel B) and from the model (Panel A). Parameter estimates are consistent with autocorrelations in the data. Looking at the posterior means of autocorrelation estimates, there is some evidence that they decrease exponentially roughly as ρ^k, where ρ is the persistence parameter in the conditional mean of consumption growth and k is the lag. This is especially evident for a 1930 - 2005 sample period. In this period consumption growth autocorrelations tend to decrease even at a faster rate than the estimated 0.575 for that period. However, they still remain within one standard deviation of the model-implied values. The marginal posterior densities of the model-implied consumption growth autocorrelations are right skewed. This is mainly the result of the model's inability to generate negative autocovariances as long as the estimate of ρ is non-negative.

With respect to dividend growth autocorrelations, the model is not rich enough to reproduce negative autocorrelations observed in the data. As long as the estimate of ρ is positive with little probability mass on negative values, the model implies positive dividend growth autocovariances. Nevertheless, the fit is reasonable as the estimates from the data are not statistically different from zero at conventional significance levels. Similarly, the model-implied results are all very close to zero within two standard deviations margin. Covariance between consumption and dividend growth of -0.0046 in the data is well within one standard deviation of the model posterior mean of 0.036.

A model with high persistence in the mean consumption growth, represented by parameter ρ here, would predict similar high persistence in both dividend and consumption growth autocorrelations. Specifically, if ρ is close to one, we should not see much change in the autocorrelations with horizon.

3.1. Equity Volatility

In step one we estimate the parameters listed in Table 2 by forcing the model to fit the moments of the consumption and dividend growth as well as the equity returns. Breaking the singularity in the likelihood requires a measurement error in the return dynamics. The quality of the model will then depend on the volatility of the measurement error, σ_m. The model is good, if most of the variation in returns is explained by the model and only a small portion of the total variation is due to the volatility of the measurement error. In this case the R^2 of regressing net return (return net of dividend growth) on latent variable should be high. R^2 is high only if the latent process has similar characteristics to those of the p/d ratio. It is the p/d ratio that serves as a sufficient statistic for net returns in the Campbell-Shiller identity. Otherwise, low R^2 will reflect a conflict between the properties of the two processes. The conflict will result in misrepresentation of equity volatility among other things. Below we build the R^2 measure to assess the severity of the equity volatility puzzle under the model.

Researchers have shown that in the data most of the variation in returns comes from shocks to future returns rather than dividend growth. Thus, Campbell-Shiller identity (2) implies that most of return variation should come from p/d ratios. In the model, the p/d ratio is perfectly correlated with the latent variable. The model, however, prescribes two conflicting roles to the latent process, x_t. On one hand, it inherits the properties of consumption growth including low volatility. On the other hand, it has to explain both the high equity return volatility and relatively low volatility of the risk free rate.

In terms of numbers, the problem the model faces can be described as follows. Unconditional equity volatility consists of two components. One is the conditional equity volatility which is around 12.3% in annual data (see Table 4). The other is the volatility of the latent process scaled down by ψ^2. The sample standard deviation of equity return is 17.4%. It implies that the scaled latent factor volatility must explain the remaining 4.7%. The estimated unconditional standard deviation of the latent process, x, is 1.33%. Given these estimates matching sample return volatility of 17.4% requires ψ of around 0.11, or, alternatively, 1/ ψ of 9.26. However, the posterior mean of 1/ψ based on full data is only 0.31 with 97.5% confidence limit of 1.32. The reason is that low IES is incompatible with low volatility of the risk free rate. To sum up, the estimate of ψ is too large to generate reasonable equity volatility. Even when the model is not constrained by the interest rate data (Table 2), the posterior mean of 1/ ψ is -1.188 with posterior bands (-5.526, 3.122), not enough to match equity volatility.

Table 4. Conditional and unconditional volatility of the equity return that the model implies based on the parameter estimates reported in Table 2 for the period 1934-2005.

Parameter	Mean	St. Dev.	2.5% Post. B.	Median	97.5% Post. B.
$\sqrt{\text{var}_t(r_{t+1})+\sigma_m^2}$	0.1738	0.01031	0.1554	0.1733	0.1957
$\sqrt{\text{var}_t(r_{t+1})}$	0.1232	0.01151	0.1034	0.1223	0.1483
Ξ_m	0.1217	0.01205	0.09758	0.1217	0.1459
$\sqrt{\text{var}(r_{t+1})}$	0.1272	0.01361	0.1049	0.1256	0.1583
σ_x^2	1.78E-04	5.73E-05	9.76E-05	1.70E-04	3.05E-04
R^2	0.5187	0.07783	0.3804	0.5138	0.6848

The estimates reported in Table 4 (interest rate excluded) suggest that the model accounts only for R^2 = 52% of the total return variance. The result is robust to interest rate data. Once we use the full data, the new R^2 = 50% is essentially the same as the one in the first case with interest rate excluded. The result demonstrates that Epstein-Zin framework with persistent consumption growth is not sufficient to reconcile the relatively smooth consumption growth and much more volatile return processes.

3.2. Equity Premium

To estimate the equity premium and the Sharpe ratio we augment the empirical model by adding the risk-free rate of return. The model can explain equity premium. However, it requires high risk aversion parameter. To have a reasonably low risk aversion parameter estimate requires high persistence of the latent process that we do not observe in consumption data.

4. CONCLUSIONS

To answer a question of whether consumption growth is i.i.d. or not, we estimate a consumption-based model of Epstein-Zin utility class. We use Bayesian Markov Chain Monte Carlo (MCMC) methodology that allows joint estimation of the entire time series of the latent process and parameter values. This in turn eliminates parameter estimation risk, which could be a confounding factor in inferences about the latent process. This is in contrast to classical approaches where latent state estimates are obtained conditional on parameter estimates.

Using equity returns, consumption growth and dividend growth data we find that mean consumption growth contains small predictable component. The model imposes certain restrictions on the latent process. In particular, it must have autocorrelation properties of consumption growth. Our estimation results suggest that like consumption growth, the latent state variable is only moderately persistent with corresponding half life of consumption growth shocks of around 1 year. Yet another restriction of the model is that the latent process must be perfectly correlated with p/d ratios. It implies that for the model to be internally consistent, p/d ratio autocorrelations must have the same decay rates as those of consumption growth. This contradicts empirical evidence that p/d ratios are highly persistent as opposed to the estimated latent process. When recast in terms of p/d ratios, the estimates of the model suggest the half life of price/dividend shocks of approximately 13 years. It appears that the latent process proxies for something other than p/d ratios contrary to model restrictions.

We argue that this disparity between the properties of the latent state and the observable dynamics of p/d ratios is also responsible for the equity volatility puzzle. The model explains only 52% of the total return volatility.

5. REFERENCES

Bansal, R. and A. Yaron (2004), Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance*, 59, 1481-1509.

Campbell, J.Y. (2003), Consumption-based asset pricing, Chapter 13 in George Constantinides, Milton Harris, and Rene Stulz, eds. *Handbook of the Economics of Finance*, Vol. IB, North-Holland, Amsterdam, 803-887.

Campbell, J.Y. and J.H. Cochrane (1999), By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy*, 107, 205-251.

Campbell, J.Y. and R.J. Shiller (1988), The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies,* 1, 195-228.

Epstein, L. and S. Zin (1989), Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica*, 57, 937-968.

Epstein, L., and S. Zin (1991), Substitution, risk aversion, and the temporal behaviour of consumption and asset returns: An empirical investigation, *Journal of Political Economy*, 99, 263-286.

Johannes, M. and N. Polson (2006), MCMC Methods for Financial Econometrics, *Handbook of Financial Econometrics*, edited by Yacine Ait Sahalia and Lars Hansen.

Lettau, M. and S. Ludvigson (2001), Consumption, aggregate wealth and expected stock returns, *Journal of Finance*, 56, 815-849.

Mehra, R. and E. Prescott (1985), The equity premium: A puzzle, *Journal of Monetary Economics*, 15, 145-161.

Weil, P. (1989), The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics*, 24, 401-421.