The cyclical component of US asset returns*

David K. Backus,† Bryan R. Routledge,‡ and Stanley E. Zin§

(December 2008) April 6, 2010
Preliminary (yes, still)

Abstract

We show that equity returns, the term spread, and excess returns on a broad range of assets are positively correlated with future economic growth. The common tendency for excess returns to lead the business cycle suggests a macroeconomic factor in the cyclical behavior of asset returns and the equity premium in particular. We construct an exchange economy that illustrates how this might work. Its important ingredients are recursive preferences, stochastic volatility in consumption growth, and dynamic interaction between volatility and growth.

JEL Classification Codes: E32, E43, E44, G12.

Keywords: leading indicators, term spread, excess returns, business cycles, recursive preferences, stochastic volatility.

---

* Thanks to Stijn Van Nieuwerburgh, Tom Tallarini, Amir Yaron for helpful comments. We also appreciate the comments from seminar participants at University of Calgary, University of North Carolina at Chapel Hill, University of Western Ontario, Wharton, and feedback from conference attendees at the NBER and SED meetings in Istanbul. The latest version of the paper is available at:
`http://pages.stern.nyu.edu/~dbackus/GE_asset_pricing/ms/BRZ_returns_latest.pdf`.

† Stern School of Business, New York University, and NBER; dbackus@stern.nyu.edu.

‡ Tepper School of Business, Carnegie Mellon University; routledge@cmu.edu.

§ Stern School of Business, New York University, and NBER; stan.zin@nyu.edu.

1 Introduction

We look at asset prices from the perspective of macroeconomists and ask: What do they tell us about the structure of the economy that generated them? We document two sets of facts that we think are worth exploring further. The first is the well-known tendency for equity prices (or their growth rates) and term spreads (differences between long- and short-term interest rates) to lead the business cycle. In US data, and to some extent in data for other countries, fluctuations in these variables are positively correlated with economic growth up to 6 to 9 months in the future. Yet in virtually all business cycle models, everything — including equity prices and term spreads — moves up and down together. Indeed, this is often treated as the defining feature of the business cycle, and therefore an attractive property of a model. The second set of facts concerns excess returns. We show that excess returns on a broad range of equity and bond portfolios also lead the cycle: high excess returns are associated with high future growth. Moreover, they display similar cyclical patterns across a broad cross section. Portfolios with very different expected returns all display a strikingly similar lead of the business cycle. This property is less widely known and might even be new. It suggests that the cyclical component of excess returns is common across asset classes and has a macroeconomic origin.

If these are the facts, what do they tell us about the structure of the aggregate economy? In a general markov environment, prices and quantities are functions of the state of the economy. The facts suggest that asset prices reflect some feature of the state that is correlated with future economic growth. We construct an example that shows how this might work. To keep things simple, we consider a traditional exchange economy in which the only inputs are preferences and a stochastic process for consumption growth. The evidence points us to the consumption growth process. What we need, it seems, is a process in which predictable changes in future consumption growth show up now in equity prices and term spreads. Preferences play a role, too, since they affect the value given to expected future growth.

The evidence on excess returns poses more of a challenge. If we have constant variances (constant risk) in the consumption growth process and standard homothetic preferences (constant risk aversion), expected excess returns are constant, by construction. This point is made by Atkeson and Kehoe (2008) and we think it's a good one: you can't talk sensibly about the cyclical behavior of excess returns without cyclical variation in risk and/or risk aversion. In order for excess returns to vary, we need variation in either aggregate aversion to risk or in risk itself. Both paths have been taken in the literature. Time variation in aggregate risk aversion can have stem from risk preferences (e.g., generalized disappointment aversion Routledge and Zin (2010)) or, perhaps, heterogeneity across agents (Guvenen (2009)). Alternatively, the level of risk can be time varying through, for example, stochastic volatility (Bansal and Yaron (2004)) or markov switching (Constantinides and Ghosh (2010)). Without making any claim to superiority, we consider changes in risk generated by stochastic volatility in the consumption growth process. Crucially, recursive preferences allow us to assign volatility a nonzero price. Cyclical variation requires, in addition, some interaction between consumption growth (the cycle) and volatility (risk). The net result is a modest generalization of the environment of Bansal and Yaron (2004).

Here's the plan. In Sections 2, we document the cyclical behavior financial variables and economic growth. We do this with cross-correlation functions, which are useful visual representations of the dynamic relations between variables. In Sections 3 we describe an exchange economy that has the elements in place to deliver something like the facts documented earlier. Section 4 looks at some examples. The first strips down to a simplistic growth process that, hopefully, lets us see all the moving pieces of the model that contribute to the cross correlations. The second set of examples are numerical. There, we build on the "long-run risk" component of Bansal and Yaron (2004) and deliver a calibration that matches the cross correlations that cross correlation patterns we see in the data. The last two sections connect our work to the literature, clean up some loose ends, and point to issues that remain unresolved.

2 Financial indicators of business cycles

In the US and elsewhere, financial variables are commonly used as indicators of future economic growth. Two of the most popular are equity prices (typically the growth rate or return of a broad-based index) and the term spread (the difference between a long-maturity interest rate and a short rate). We describe the dynamic relations between these variables and aggregate economic growth with cross-correlation functions. We then go on to look at excess returns for a variety of assets with different cross-sectional properties.

Data

Our data cover the period from 1960 to the present. We use monthly series because the finer time interval allows clearer identification of leads and lags than the quarterly data commonly used in national income accounts and business cycle research. Definitions and sources are given in Appendix A, but here's a summary. Most of our financial variables come from CRSP and Ken French's web site. Returns, r_t, cover the whole month t; the return for October 2008, for example, refers to the period September 30 to October 31 and includes any interim cash-flows (e.g., dividends). In addition to bond yields, we look at one month holding period returns on portfolios of bonds of a fixed maturity range (CRSP Fama bond portfolios). Throughout, we calculate as the logarithms of gross returns because they line up more neatly with our theory. Bond yields refer to the end of the month — the last trading day — so the yield associated with October 2008 is that for October 31.

Most of our macroeconomic variables come from FRED, the online data repository of the Federal Reserve Bank of St. Louis. They are typically time averages (e.g., a quantity per month). Industrial production for October 2008 is an estimate of the average for that month. We also look at consumption and employment at similar frequencies. Aggregate profit and dividend information is from Robert Shiller's web site and represents monthly earnings and dividends at S&P 500 companies. For all economic quantity data, Y_t, we compute growth rates as log-differences over the previous month $y_t = \log Y_t - \log Y_{t-1}$, so that the October growth rate is the growth rate of October over September.

To filter the data to focus on business cycles, we use centered year-on-year growth rates $\hat{y}_t = \log Y_{t+6} - \log Y_{t-6}$. This smoothes out the high-frequency variation in these series without disturbing the timing. We use a similar filter with the financial data to calculate centered annual returns $\hat{r}_t = \log P_{t+6} - \log P_{t-6}$ where we construct P_t as the value of an index from the monthly returns for the asset (i.e., average annual returns centered around month t). We think of year-on-year growth rates and returns as crude approximations to the Hodrick-Prescott filter often used in business cycle analysis. We find the smoothed series more informative about the cyclical component we are interested in. However, we present both the raw data and the smoothed data below.

The dynamic interrelations between financial indicators and economic growth are conveniently summarized with cross-correlation functions. If r is a financial indicator and y is a measure of economic growth, their cross-correlation function is

$$ccf_{ry}(k) \quad = \quad corr(r_t, y_{t-k}),$$

a function of the lag k. For negative values of k, this is the correlation of the financial indicator with future economic growth. If these correlations are nonzero, we say the financial indicator leads the business cycle. Similarly, positive values of k correspond to correlations of the indicator with past economic growth; nonzero values suggest a lagging indicator. Our interest is in the former: financial variables that lead economic growth and thus serve as a source of information about the future.

Equity

Consider equity returns. In Figure 2 we report the cross-correlation function for the (nominal) return on an aggregate portfolio of publicly-traded equity and the monthly growth rate of industrial production. Both series are inherently noisy — there's little persistence in either series — yet we see a modest but clear pattern. The correlations on the left show that equity returns are positively correlated with growth in industrial production up to one year in the future. The correlations are modest individually (the largest are between 0.12 and 0.25) but exhibit a clear pattern. The correlations on the right are generally smaller.

When we average growth over 12 months, the pattern emerges even more clearly: high equity returns are associated with high economic growth several months later. Figure 3a shows the correlation using year-on-year growth rates of industrial production with monthly equity returns and Figure 3b shows the cross correlations of year-on-year growth with year-on-year returns. This is a typical result: correlations are larger and smoother if we use year-on-year growth rates. These annualized series are closer to what is done in business cycle research.

We use industrial production industrial production as our measure of economic activity since it focusses focusses on easily measured quantities in industries like manufacturing, mining, and electric and gas utilities. These are the industries, along with construction, where the bulk of business cycle variation occurs.[1] However, the key feature of the data – equity returns lead the business cycle – is robust to changes in our measurement of economic growth. Figure 4 replaces industrial production with employment (nonfarm employment from the establishment survey) or consumption (total, per capita, real) and little changes. Similarly, we can measure economic activity with measures of corporate profits (S&P 500 earnings and dividend growth) with similar results. The employment figures are a bit sharper and consumption less so. The corporate profit measures are sharper still. Whether these differences reflects better measurement or something else is hard to say.

Equity returns leading the business cycle is also robust to a number of variations in measurement and methodology. For example, we can use nominal or real returns with little

[1] See: http://www.federalreserve.gov/releases/G17/About.htm.

difference (not reported). Similarly, we get similar results if we look at a subsample of the data, say, only from 1990 on. The pattern is similar, albeit a bit more choppy with a shorter sample. In particular, the recent recession exhibits a similar pattern. Figure 5 plots the year-on-year growth in industrial production and year-on-year equity returns. The market decline precedes the decline in industrial production. The difference, if any, is the shorter interval between the decline in equity and subsequent fall in economic activity.

Interest Rates

We turn next to interest rates. In Figure 6a, we show the cross-correlation function between the term spread (in this case the difference between continuously compounded nominal yields on 5-year and 1-month treasuries) and the monthly growth rate in industrial production. A large positive value for the spread indicates a steep yield curve, a small or negative value a flat or declining yield curve. Decades of research has found that steep yield curves (and large term spreads) are associated with above-average future economic growth. In Figure 6b shows the same result with year-on-year industrial production growth. Interesting in Figures 6c and 6d, the short rate (the 1-month yield) cross-correlation function is a mirror image of what we see with the term spread. This suggests that most of what we see in the cross-correlation function for the term spread comes from the short rate. Again, as you would expect, these results are also robust to changes in our measurement of economic activity and sample periods.

Most of these facts have been documented in earlier studies. Prominent examples include Ang, Piazzesi, and Wei (2006), Estrella and Hardouvelis (1991), Fama and French (1989), King and Watson (1996), Mueller (2008). Each contains an extensive set of references to related work.

Excess Returns

If these cyclical properties of equity returns and interest rates seem familiar, a little thought points to the excess return on equity and the equity premium. The equity premium is the difference of equity returns and the short-term interest rate (recall we work with log returns throughout). So roughly speaking, the evidence suggests that several months before an increase in economic growth, returns on equity rise and the return on the short bond falls. If we put the two together, we see that changes in economic growth are preceded (on average) by a increase in the expected excess return on equity; that is an increase in the equity premium.

We see precisely this cyclical variation in excess returns in Figure 7. The figure shows the cross-correlation function for the excess return on equity and the growth rate of industrial production (both monthly and year-on-year are shown). The correlations for excess returns are slightly larger than those for returns, but the pattern is similar. Evidently, most of the variation in excess returns comes from the return rather than the short rate.

There is, of course, much variety in the excess returns across assets. Surprisingly, at least to us, is that the same cyclical pattern appears in a wide range of equity portfolios. The basic theme is repeated in portfolios based on industry, portfolios based on firm market capitalization ("size"), and portfolios based on accounting measures ("book-to-market").

Consider industry portfolios. Cross correlations for four examples are pictured in Figure Figure 8. We report two industries we thought would have highly cyclical production and sales (automobiles and machinery) and two that would be less cyclical (food and drugs). Their cross-correlation functions are nevertheless quite similar. Indeed, we could say the same for virtually all 49 industry portfolios (based on SIC codes; not reported). At least to a first approximation, the cyclical behavior of these excess returns is the same.

The same is true of Fama-French (1992) portfolios. Four examples are given in Figure 9: the smallest and largest firms ranked by market market capitalization and the lowest and highest ranked by book-to-market ratio (accounting measured book value divided by market capitalization). All four have positive correlations of excess returns with growth 3-12 months in the future. Also striking (not reported) is that "difference portfolios" (the infamous "small minus big" or SMB and "high minus low" or HML) show little cyclical pattern: the cyclical behavior apparently cancels when you subtract one return from the other.

We find the common cyclical pattern in these portfolios surprising, because we know — from Fama and French (1992) and many others — that these portfolios have very different return characteristics. Their average returns, in particular, are wildly different and the difference portfolios are the standard factors used to fit the cross-section. The cross correlation functions suggest, however, that whatever these differences are, they are unrelated to the business cycle.

If equity portfolios exhibit similar cyclical behavior, what about bonds? Bond returns are noisier than the yields we looked at in the in Figure 6, but they display a clear cyclical pattern. The four panels of Figure 10 are based on one-month returns on portfolios of US treasuries with different maturities: 1-6 months, 19-24 months, 55-60 months, and 55-60 months, and 115-120 months. Curiously, their cyclical behavior is similar (the year-on-year growth and returns is plotted). This is, again, despite substantial differences in volatility and average returns across the bond portfolios. This pattern is similar to that of equities, but not identical: where equities have a contemporaneous correlation close to zero with current economic growth, bond returns have a noticeably negative correlation with growth 0-5 months in the future.

Figures 11 and 12 look at the cyclical behavior of commodity excess returns (long positions on a fully-collateralized futures contract). Oil, in Figure 11, mirrors the behavior bonds. Low excess returns in oil precede growth and there is a positive contemporaneous correlation. Figure 12 looks at other commodities. There is little cyclical behavior seen in the agricultural commodities of wheat and oil. Copper leads the cycle and also shows contemporaneous correlation. (Casassus, Liu and Tang (2010) take a different look at the correlation structure in commodity prices.)

So what do we make of all this? We have seen that excess returns on a variety of equity and bond portfolios lead the business cycle: they are positively correlated with future economic growth. Cyclical variation in excess returns suggests that risk premiums vary systematically over the business cycle. The common pattern across assets suggests that a single macroeconomic factor might be able to account for all of them.

3 A theoretical exchange economy

The rest of the paper is concerned with mimicking the observed cyclical behavior of excess returns in a theoretical environment. Before diving into the mechanics, it's worth thinking a little about what we need. Consider a stationary markov environment in which quantities (consumption, for example) and asset prices (equity and bonds) at any date t are functions of a finite state vector s_t. Returns and excess returns between dates t and $t+1$ are then functions of successive states, say $r(s_t, s_{t+1})$. Variation in each of these variables thus reflects variation in the underlying state vector. The evidence of the last two sections indicates that some of this variation is positively correlated with future economic growth.

We illustrate this in a simple macroeconomic setting: an exchange economy with a representative agent. As usual in this sort of model, we specify preferences of a representative agent and the stationary markov growth dynamics for the endowment that the agent consumes. Our version has a few key ingredients. First, our agent has recursive preferences. Second, the endowment process has predictable variation in both consumption growth and its conditional variance. Variation in the conditional mean is the "long run risk" component. Variation in the conditional variance (stochastic volatility) along with the recursive preferences induce time variation in the equity premium. Finally, correlation with future consumption growth is produced directly, by specifying consumption growth as a process that depends on past volatility. All of these features are necessary to account for the cyclical behavior of the excess returns, and in particular, the equity premium. The resulting model closely resembles the structure in Bansal and Yaron's (2004). We think of it as a one-parameter extension. To make all the calculations we use a loglinear approximation method adapted from Hansen, Heaton, and Li (2008).

Environment

Preferences have the now-familiar recursive structure described by Kreps and Porteus (1978), Epstein and Zin (1989), and Weil (1989). If U_t is "utility from date t on," preferences follow from the time aggregator V,

$$U_t = V[c_t, \mu_t(U_{t+1})] \tag{1}$$

$$= [(1-\beta)c_t^{\rho} + \beta\mu_t(U_{t+1})^{\rho}]^{1/\rho}, \tag{2}$$

and (expected utility) certainty equivalent function μ,

$$\mu_t(U_{t+1}) = [E_t(U_{t+1}^{\alpha})]^{1/\alpha}. \tag{3}$$

The conventional interpretation is that $\rho < 1$ captures time preference (the intertemporal elasticity of substitution is $1/(1-\rho)$) and $\alpha < 1$ captures risk aversion (the coefficient of relative risk aversion is $1-\alpha$). Additive utility is a special case with $\alpha = \rho$.

Both the time aggregator and the certainty equivalent function are homogeneous of degree one, which allows us to scale everything by current consumption and convert our problem to one in growth rates. If we define scaled utility $u_t = U_t/c_t$, equation (2) can be expressed

$$u_t = [(1-\beta) + \beta\mu_t(g_{t+1}u_{t+1})^{\rho}]^{1/\rho}, \tag{4}$$

where $g_{t+1} = c_{t+1}/c_t$ is the growth rate of the endowment (consumption).

With these preferences, the pricing kernel (marginal rate of substitution) is

$$\begin{aligned} m_{t+1} &= \beta (c_{t+1}/c_t)^{\rho-1} \left[U_{t+1}/\mu_t(U_{t+1}) \right]^{\alpha-\rho} \\ &= \beta g_{t+1}^{\rho-1} \left[g_{t+1} u_{t+1}/\mu_t(g_{t+1} u_{t+1}) \right]^{\alpha-\rho} . \end{aligned} \tag{5}$$

See Appendix B. The pricing kernel is the heart of any asset pricing model, so (5) is central to the properties of asset prices and returns. We can see the role of recursive preferences in the pricing kernel. The first part of the kernel is risk-neutral discount factor, β, adjusted for "short-run risk" consumption growth risk $(c_{t+1}/c_t)^{\rho-1}$ The next term, $(U_{t+1}/\mu_t(U_{t+1}))^{\alpha-\rho}$. is like an innovation in future utility (note the presence of the certainty equivalent operator rather than expectation operator). This term determines how predictable changes in consumption growth and its volatility are priced. Note it drops out if preferences are time additive, $\alpha = \rho$.

We specify a general linear process for the logarithm of consumption growth. Let state variables x_t (a vector of arbitrary dimension) and v_t (volatility, a scalar) follow

$$x_{t+1} = A x_t + a(v_t - v) + v_t^{1/2} B w_{t+1} \tag{6}$$

$$v_{t+1} = (1 - \varphi_v) v + \varphi_v v_t + b w_{t+1}, \tag{7}$$

where v is the unconditional mean of v_t, $\{w_t\} \sim \text{NID}(0, I)$, and $Bb^\top = 0$ (innovations in x_t and v_t are uncorrelated).[2] The aggregate state is therefore $s_t = (x_t, v_t)$. Consumption growth is tied to x_t: $\log g_t = g + e^\top x_t$ for some constant vector e. This gives us flexible dynamics for x_t, and therefore consumption growth. The conditional variance of consumption growth is proportional to v_t:

$$\text{Var}_t (\log g_{t+1}) = e^\top B B^\top e \, v_t.$$

This structure also allows some interaction between the dynamics of x_t and v_t through the vector a. If $a = 0$, consumption growth and volatility are uncorrelated. In the examples that follow we a will be a positive (and scalar). We'll see later that all of these features — a predictable component in consumption growth, stochastic volatility, and interaction between the two — are needed to account for the cyclical behavior of asset returns.

Loglinear approximation of the pricing kernel

Asset prices in this setting are functions of the state variables and returns are functions of prices. We derive loglinear approximations to equilibrium asset prices with the goal of

[2] Formally, w_t is approximately normal since it is truncated to ensure $v_t \geq 0$.

having something that is both easy to compute and relatively transparent. We break the solution process into steps to show how it works.

Step 1. Approximate time aggregator. The starting point is equation (4), which is not loglinear unless $\rho = 0$. A first-order approximation of $\log u_t$ in $\log \mu_t$ around the point $\log \mu_t = \log \mu$ is

$$\begin{aligned} \log u_t &= \rho^{-1} \log\left[(1-\beta) + \beta \mu_t (g_{t+1} u_{t+1})^{\rho}\right] \\ &= \rho^{-1} \log\left[(1-\beta) + \beta e^{\rho \log \mu_t (g_{t+1} u_{t+1})}\right] \\ &\approx \kappa_0 + \kappa_1 \log \mu_t (g_{t+1} u_{t+1}), \end{aligned} \quad (8)$$

where

$$\begin{aligned} \kappa_1 &= \beta e^{\rho \log \mu} / [(1-\beta) + \beta e^{\rho \log \mu}] \\ \kappa_0 &= \rho^{-1} \log[(1-\beta) + \beta e^{\rho \log \mu}] - \kappa_1 \log \mu. \end{aligned}$$

The approximation is exact when $\rho = 0$, in which case $\kappa_0 = 0$ and $\kappa_1 = \beta$. See Hansen, Heaton, and Li (2008, Section III).

The rest of the solution follows those of many approximate solutions to dynamic programs: we guess a value function of a specific form with unknown parameters, substitute optimal decisions into the Bellman equation, and solve for the unknown parameters. In this case the decision is trivial (consume the endowment), but the rest of the solution is the same. Equation (8) serves as the (approximate) Bellman equation with κ_1 in the role of discount factor.

Step 2. Guess value function. We conjecture an approximate scaled value function

$$\log u_t = u + p_x^\top x_t + p_v v_t$$

with coefficients (u, p_x, p_v) to be determined.

Step 3. Compute certainty equivalent. The novel ingredient of (8) is the certainty equivalent $\mu_t(g_{t+1} u_{t+1})$. Note that

$$\begin{aligned} \log(g_{t+1} u_{t+1}) &= u + g + (e + p_x)^\top x_{t+1} + p_v v_{t+1} \\ &= u + g + (e + p_x)^\top [A x_t + a(v_t - v) + v_t^{1/2} B w_{t+1}] + p_v[(1-\varphi_v) v + \varphi_v v_t + b w_{t+1}]. \end{aligned}$$

The certainty equivalent is

$$\begin{aligned} \log \mu_t (g_{t+1} u_{t+1}) &= u + g - (e + p_x)^\top a v + (e + p_x)^\top [A x_t + a(v_t - v)] + p_v [(1-\varphi_v) v + \varphi_v v_t] \\ &\quad + (\alpha/2) [(e + p_x)^\top B B^\top (e + p_x) v_t + p_v^2 b b^\top]. \end{aligned}$$

This follows from common properties of lognormal random variables: if an arbitrary random variable $\log x \sim \mathrm{N}(a, b)$, then $\log E(x) = a + b/2$ and $\log \mu(x) = a + \alpha b/2$.

Step 4. Solve Bellman equation. If we substitute the certainty equivalent into (4) and line up coefficients, we have

$$\begin{aligned} u &= \kappa_0 + \kappa_1 \left[u + g + p_v(1-\varphi_v)v + (\alpha/2)p_v^2 bb^\top \right] \\ p_x^\top &= e^\top(\kappa_1 A)(I - \kappa_1 A)^{-1} \\ p_v &= \kappa_1(1-\kappa_1\varphi_v)^{-1} \left[(e+p_x)^\top a + (\alpha/2)(e+p_x)^\top BB^\top(e+p_x) \right]. \end{aligned}$$

The constant u plays no role in the pricing that follows. The coefficient p_x has the form

$$p_x^\top = e^\top \left[(\kappa_1 A) + (\kappa_1 A)^2 + (\kappa_1 A)^3 + \cdots \right].$$

We think of it as capturing the Bansal-Yaron effect. x_t affects expected future consumption growth. Utility (scaled) is the valuation of that future consumption with κ_1 as the "discount factor" (recall $\kappa_1 = \beta$ when $\rho = 0$.) For example, with white noise consumption growth where $A = 0$ in equation (6), then x_t has zero effect on utility and $p_x = 0$. The p_v term determines how innovations to the conditional volatility change scaled utility. The $\kappa_1/(1-\kappa_1\varphi_v)$ again comes from discounting. The shock to volatility decays at rate φ_v and the "discount rate" is κ_1. The shock has two effects on utility. The first (the one involving the interaction parameter a) comes from the impact of volatility v_t on future values of x_t and expected future consumption growth. The second (the one involving $\alpha/2$) summarizes the impact of v_t on the through volatility. The sign of this term is determined by the relative magnitude of utility from growth and disutility from volatility. To calculate p_v, note

$$(e+p_x)^\top = e^\top(I-\kappa_1 A)^{-1}.$$

The solution for p_v follows immediately.

Step 5. Derive pricing kernel. With these inputs, we can calculate the pricing kernel. The (scaled) utility enters the pricing kernel through the term

$$\begin{aligned} \log(g_{t+1}u_{t+1}) - \log \mu_t(g_{t+1}u_{t+1}) &= -(\alpha/2)\left[p_v^2 bb^\top + (e+p_x)^\top BB^\top(e+p_x)v_t \right] \\ &\quad + v_t^{1/2}(e+p_x)^\top Bw_{t+1} + p_v b w_{t+1}. \end{aligned}$$

The right-hand side has two kinds of terms: innovations (the terms with w_{t+1}) and penalties for risk (those with $\alpha/2$). The pricing kernel (5) follows as

$$\log m_{t+1} = \log\beta + (\rho-1)\log g_{t+1} + (\alpha-\rho)\left[\log(g_{t+1}u_{t+1}) - \log\mu_t(g_{t+1}u_{t+1})\right]$$

$$
\begin{aligned}
&= \log\beta + (\rho-1)(g - e^\top a v) - (\alpha-\rho)(\alpha/2)p_v^2 bb^\top \\
&\quad + (\rho-1)e^\top A x_t + [(\rho-1)e^\top a - (\alpha-\rho)(\alpha/2)(e+p_x)^\top BB^\top(e+p_x)]v_t \\
&\quad + v_t^{1/2}[(\rho-1)e + (\alpha-\rho)(e+p_x)]^\top B w_{t+1} + (\alpha-\rho)p_v b w_{t+1} \\
&= \delta_0 + \delta_x^\top x_t + \delta_v v_t + v_t^{1/2}\lambda_x^\top w_{t+1}, + \lambda_v^\top w_{t+1}, \qquad (9)
\end{aligned}
$$

with the implicit definitions of $(\delta_0, \delta_x, \delta_v, \lambda_x, \lambda_v)$. You may recognize this as a close relative of so-called affine models of bond pricing. The main difference is that the parameters are not free: they're tied to preferences and the consumption growth process.

The pricing kernel illustrates the interaction of recursive preferences, predictability of consumption growth, and stochastic volatility. The state variables x_t play a role only to the extent they help to forecast future consumption growth. If the state variables do not forecast growth (in other words, when $A = 0$), they do not appear in the pricing kernel ($\delta_x = 0$). If the state variables x_t do enter the pricing kernel, their impact is governed by the intertemporal substitution parameter ρ. Volatility v_t appears here for two reasons. Volatility helps predict future consumption growth (via the a volatility-growth interaction term). The impact of this is controlled by intertemporal substitution (ρ). The volatility term also enters directly because it controls the conditional variance (the second term in δ_v). The size of this effect depends on risk aversion (through $\alpha/2$) and the departure from additive preferences (the difference $\alpha - \rho$). When either is zero, the term is also zero, and volatility affects the pricing kernel only through its impact on expected future consumption growth.

In what sense is our solution an approximation? The only relation that isn't exact is (8), which is exact when $\rho = 0$. Moreover, relative to standard methods (linearize around the deterministic steady state), uncertainty plays a central role.

Asset returns

Given a pricing kernel m, (gross) asset returns r satisfy the pricing relation

$$
E_t(m_{t+1} r_{t+1}) = 1.
$$

An asset, for our purposes, is a claim at date t to a dividend stream $\{d_{t+j}\}$ for $j \geq 1$. We use the pricing kernel (9) to derive prices and returns for a number of common assets, whose properties can then be compared to those we documented in Sections 2. For bonds, we will look at the short rate and multiperiod default-free bonds. For equity, we will look at a claim to next periods endowment, a consumption strip, and a claim to the endowment at all future dates, a consumption stream. We skip the pricing of more complicated securities, like a claim to a levered dividend stream that is imperfectly correlated with consumption growth. Solving for the returns on this sort of security would involve an additional approximation. For now, we will stick to securities whose returns we can solve exactly. (Despite the log-linear structure, not all these expressions that follow are transparent. The example that follows in Section 4.1 solves an example with a single state variable.)

Short rate. The dividend on a short-term default-free security is one unit of the consumption good next period: $d_{t+1} = 1$. The price of this 1-period (real) bond is $q_t^1 = E_t m_{t+1}$ and the return is $r_{t+1}^1 = 1/q_t^1 = 1/E_t m_{t+1}$. Thus

$$\log r_{t+1}^1 \quad = \quad -\log q_t^1 \quad = \quad -(\delta_0 + \lambda_v^\top \lambda_v/2) - \delta_x^\top x_t - (\delta_v + \lambda_x^\top \lambda_x/2) v_t,$$

The short rate, as you would expect, is a loglinear function of the state (x_t, v_t). This is not obvious from the expression, but in the examples we look at below, we'll see that the dynamics of the short rate are dominated by the volatility term (a similar flavor as Atkeson-Kehoe (2008)). In addition, we use the familiar preference calibration where α and $\alpha - \rho$ are negative, the short rate is decreasing in volatility. The example in Section 4.1 illustrates this better.

Consumption Strip. A consumption "strip" pays the consumption at (the single) date $t+s$; $d_{t+s} = c_{t+s}$. This isn't a real asset, but it illustrates how the various ingredients interact. Let's focus on the one-period ahead strip for now. Define q_t^s as the price-dividend ratio for strip with maturity s at date t. For $s = 1$,

$$q_t^s \quad = \quad E_t \left(m_{t+1} g_{t+1} \right)$$

and the return is $r_{t+1}^s = g_{t+1}/q_t^s$. The log growth rate is

$$\log g_{t+1} \quad = \quad g + e^\top x_{t+1} \quad = \quad g + e^\top [A x_t + a(v_t - v) + v_t^{1/2} B w_{t+1}].$$

The price is

$$\log q_t^s \quad = \quad (\delta_0 + g - e^\top a v + \lambda_v^\top \lambda_v/2) + (\delta_x^\top + e^\top A) x_t + \ [\delta_v + e^\top a + (\lambda_x^\top + e^\top B)(\lambda_x + B^\top e)/2] v_t$$

and the return is

$$\log r_{t+1}^s \quad = \quad -(\delta_0 + \lambda_v^\top \lambda_v/2) - \delta_x^\top x_t - [\delta_v + (\lambda_x^\top + e^\top B)(\lambda_x + B^\top e)/2] v_t + v_t^{1/2} e^\top B w_{t+1}.$$

The excess return is therefore

$$\log r_{t+1}^s - \log r_{t+1}^1 \quad = \quad (1/2)[\lambda_x^\top \lambda_x - (\lambda_x^\top + e^\top B)(\lambda_x + B^\top e)] v_t + v_t^{1/2} e^\top B w_{t+1}.$$

This expression gives us a look at how the cross correlation function for excess returns with consumption growth might work. Note the excess return here does not depend on the predictable component of consumption growth. Neither x_t state nor the volatility-growth interaction term, a, appear. The predictable component has an identical effect on all returns; hence it does not show up in excess returns. This is a general result that comes from the afine pricing kernel. This suggests the cross correlation function of excess returns with consumption growth will largely reflect cross correlation function of volatility with consumption growth (we re-visit this with the example in Section 4.1). Second, note that innovations to volatility (the bw_{t+1} term) does not show up in excess returns of the consumption strip since we are looking at a single period security (v_{t+1} does not affect g_{t+1}).

Consumption Stream. The consumption stream is an asset whose dividend is $d_{t+j} = c_{t+j}$ for all $j \geq 1$; that is s a claim to consumption from next period on. For now, we will think of this as "equity" (leaving issues like leverage and the imperfect correlation between dividends and consumption for later). The reason we focus on this asset is that the return has a simple form,

$$r^c_{t+1} \quad = \quad \beta^{-1} \left[g_{t+1} u_{t+1} / \mu_t (g_{t+1} u_{t+1}) \right]^\rho g_{t+1}^{1-\rho}. \tag{10}$$

The derivation of this depends only on the constant elasticity form of the time aggregator; it does not reflect any of the structure we've given to consumption growth or the certainty equivalent (other than linear homogeneity). See Appendix B for the details. We use the same loglinear approximation for scaled utility $\log u_{t+1}$ as equation (8). Everything else in the return, we can solve for exactly.

$$\begin{aligned} \log r^c_{t+1} \quad = \quad & -\log \beta + (1-\rho)(g - e^\top a v) - (\rho\alpha/2) p_v^2 b b^\top \\ & + (1-\rho) e^\top A x_t + [(1-\rho) e^\top a - (\rho\alpha/2)(e + p_x)^\top B B^\top (e + p_x)] v_t \\ & + v_t^{1/2} (e + \rho p_x)^\top B w_{t+1} + \rho p_v b w_{t+1}. \end{aligned}$$

The excess return is

$$\begin{aligned} \log r^c_{t+1} - \log r^1_{t+1} \quad = \quad & (1/2)[(\alpha-\rho)^2 - \alpha^2] p_v^2 b b^\top \\ & + [(\rho - 1)e + (\alpha - \rho)(e + p_x)]^\top B B^\top [(\rho - 1)e + (\alpha - \rho)(e + p_x)] v_t \\ & - (\alpha^2/2)(e + p_x)^\top B B^\top (e + p_x) v_t \\ & + v_t^{1/2} (e^\top + \rho p_x^\top) B w_{t+1} + \rho p_v b w_{t+1}. \end{aligned}$$

Notice, again, that it does not depend on x_t nor the interaction term a: all of the variation in the conditional mean of the excess return comes from v_t. The volatility term is a little complicated, but if you consider (as we do) situations in which α is large relative to ρ, then the volatility term is dominated by

$$(\alpha^2/2)(e+p_x)^\top BB^\top (e+p_x)v_t.$$

The quadratic form is the conditional variance of next-period utility, whose impact is governed largely by risk aversion (α) squared. Thus we see that risk aversion affects not only the average excess return, but also its variation. The Bansal-Yaron term p_x also plays a role: if it's small or even negative, the impact of volatility is also small. If $p_x = 0$, the volatility term is

$$[(\alpha-1)^2 - \alpha^2/2]e^\top BB^\top e\, v_t,$$

so, again, risk aversion is central.

Multiperiod bonds. An n-period bond is a claim to one unit of the consumption good n periods in the future: $d_{t+j} = 1$ for $j = n$, zero otherwise. Since the return on an $n+1$-period bond is $r_{t+1}^{n+1} = q_{t+1}^n / q_t^{n+1}$, the pricing relation implies that prices satisfy

$$q_t^{n+1} \;=\; E_t\left(m_{t+1} q_{t+1}^n\right).$$

We guess that prices are loglinear functions of the state:

$$\log q_t^n \;=\; \delta_0^n + \delta_x^{n\top} x_t + \delta_v^n v_t.$$

The coefficients satisfy the recursions

$$\begin{aligned}
\delta_0^{n+1} &= \delta_0 + \delta_0^n + [\delta_v^n(1-\varphi_v) - \delta_x^{n\top} a]v + (\delta_v^n b + \lambda_v^\top)(\delta_v^n b^\top + \lambda_v)/2 \\
\delta_x^{n+1\top} &= \delta_x^\top + \delta_x^{n\top} A \;=\; \delta_x^\top (I + A + \cdots + A^n) \\
\delta_v^{n+1} &= \delta_x^{n\top} a + \delta_v + \delta_v^n \varphi_v + (\delta_x^{n\top} B + \lambda_x^\top)(B^\top \delta_x^n + \lambda_x)/2.
\end{aligned}$$

Excess returns are therefore

$$\begin{aligned}
r_{t+1}^{n+1} - r_{t+1}^1 &= \log q_{t+1}^n - \log q_t^{n+1} + \log q_t^1 \\
&= (\delta_0^n - \delta_0^{n+1} - \delta_0^n) + \delta_x^{n\top} x_{t+1} + \delta_v^n v_{t+1} + (\delta_x^1 - \delta_x^{n+1})^\top x_t + (\delta_v^1 - \delta_v^{n+1}) v_t.
\end{aligned}$$

This is a little ugly, but the expressions can can be used to compute excess returns in the model, and thus their properties.

4 Examples – Cyclical behavior of theoretical excess returns

We find a couple of examples helpful to sort out how the preferences and the process for endowment growth lead to the cross correlation patterns we see in the data. The first example begins with a simpler process. The second example uses an endowment process that builds from Bansal and Yaron (2004) and generates asset return features we see in the data.

4.1 Stochastic volatility and no long-run risk

Here is a simpler scalar process for endowment growth where the algebra is a bit more transparent. Define consumption growth as $g_{t+1} = c_{t+1}/c_t$ with dynamics

$$\log g_{t+1} = g + a(v_t - v) + v_t^{1/2}\epsilon_{t+1} \quad (11)$$

$$v_{t+1} = (1-\varphi_v)v + \varphi_v v_t + \sigma\eta_{t+1} \quad (12)$$

where v is the unconditional mean of v_t, g is the unconditional mean of g_t and $0 < \varphi_v < 1$ is the autocorrelation of v_t. The $\epsilon_t, \eta_t \sim NID(0,1)$ and uncorrelated. This specification has predictable variation in the conditional variance (stochastic volatility). The conditional mean of consumption growth is correlated with the volatility through the (scalar) parameter a (our numerical illustration has $a \geq 0$). The only state variable, s_t, here is the conditional volatility v_t. The conditional mean of endowment growth has simple AR(1) dynamics and is too simple to generate realistic asset prices. It is missing the "long-run risk" component (since we have set $A = 0$ in equation (6)). However, it does illustrate how preferences and endowment link to the cross-correlation functions.

We can solve this example with exactly the same steps as in the previous section or just plug $A = 0$ into the previous solutions (or, of course, do both just to be inefficient). To get started, the pricing kernel is

$$\log m_{t+1} = \delta_0 + \delta_v v_t + \lambda_x v_t^{1/2}\epsilon_{t+1} + \lambda_v\sigma\eta_{t+1}$$

with

$$\delta_o = \log\beta + (\rho-1)(g-av) - (\alpha-\rho)(\alpha/2)p_v^2\sigma^2$$

$$\delta_v = (\rho-1)a - (\alpha-\rho)(\alpha/2)$$

$$\lambda_x = \alpha - 1$$

$$\lambda_v = (\alpha-\rho)p_v\sigma$$

$$p_v = \kappa_1(1-\kappa_1\varphi_v)^{-1}(a+\alpha/2)$$

Since v_t is the only state variable in this example, $\delta_x = 0$. As we have seen before, the volatility appears for two reasons: because it predicts future consumption growth (a in δ_v) and its effect on risk (the$-(\alpha - \rho)(\alpha/2)$ term). The sign of δ_v will depend on the relative magnitude of these two effects. We will come back to the importance of this term and what it implies about the parameters when we look at the cross correlation functions below. (The p_v term comes from the log-linear approximation to scaled utility. It determines how scaled-utility depends on the sole state variable, v_t.)

The short rate in this example is given by

$$\log r^1_{t+1} \quad = \quad r + [(1-\rho)a + 0.5(\alpha - 1 + \alpha(1-\rho))]v_t$$

The short rate, as you might expect, varies with the single state variable in the model, v_t volatility. The term $(1-\rho)a$ reflects the predictable component of consumption growth. The remaining term reflects the risk contribution of volatility to the short rate. The typical calibration has α and $\alpha - \rho$ negative, so this later term is negative. We will see in a moment, that to match the cyclical pattern of the risk free rate we need the coefficient on volatility here to be negative (high volatility implies low risk-free rates). For this to be the case, ρ needs to be near one so the $(1-\rho)a$ term is small.

Recall that we defined a consumption "strip" as a security that pays the consumption at (the single) date $t+s$; $d_{t+s} = c_{t+s}$ and we focus on the short-horizon strip with $s = 1$ whose return is r^s_{t+1}. Similarly, for consumption stream, r^c_{t+1} is the return on an asset whose dividend is all future consumption ($d_{t+j} = c_{t+j}$ for all $j \geq 1$). The excess return for these two equity securities is

$$\begin{aligned}
\log r^s_{t+1} - \log r^1_{t+1} \quad &= \quad 0.5[1-2\alpha]v_t + v_t^{1/2}\epsilon_{t+1} \\
\log r^c_{t+1} - \log r^1_{t+1} \quad &= \quad 0.5[1-2\alpha]v_t + v_t^{1/2}\epsilon_{t+1} + 0.5((\alpha-\rho)^2 - \alpha^2)p_v^2\sigma^2 + \rho p_v \sigma \eta_{t+1}.
\end{aligned}$$

Again, the excess returns do not depend on the predictable component of consumption growth. The predictable component, through the parameter a in this case, has an identical effect on all returns; hence it does not show up in excess returns. The volatility state variable only shows up with its risk effect (note $1 - 2\alpha > 0$ for the usual values of $\alpha < 0$). The cross correlation functions below will largely reflect the cross correlation of volatility with growth.

The expressions for the consumption strip and stream are both very similar with identical coefficients on v_t and ϵ_{t+1}. These terms reflect the risk in the immediate future outcome of g_{t+1}. The next two terms for the consumption stream, both multiplicative in σ, are the contribution of the stochastic volatility on longer horizon consumption stream dividends. The term last term determines how excess returns react to innovations to volatility. The p_v is the coefficient determines how innovations to volatility affect scaled utility. Substituting it in, we get

$$\rho p_v \sigma = \rho \frac{\kappa_1}{1-\kappa_1 \varphi_v}(a+\alpha/2)\sigma$$

The sign depends on the sign of ρ. A $\rho > 0$, as we use in our numerical examples below, corresponds to an EIS greater than one. Assuming $\rho > 0$, a positive shock to volatility ($\eta_{t+1} > 0$) is associated with higher excess returns if $a > (-\alpha/2)$. That is; the increase in future consumption via the parameter a has a larger utility gain than the increase to future risk.

Cross Correlations

The advantage of the simple endowment structure we look at here is that we can explicitly calculate the cross correlation functions to see what sort of parameters are needed to reproduce the data. The cross-correlation function depends on the covariance of the short rate and consumption growth. For lag k, we can compute this

$$\begin{aligned} \operatorname{cov}(\log r^1_{t+1}, \log g_{t+1-k}) &= \operatorname{cov}(\Big[[(1-\rho)a + 0.5(\alpha-\rho)\alpha - 0.5(\alpha-1)^2]v_t\Big], \Big[av_{t-k} + v^{1/2}_{t-k}\epsilon_{t+1-k}\Big]) \\ &\approx \Big[(1-\rho)a^2 - 0.5(\alpha(\rho-2)+1)a\Big] \operatorname{cov}(v_t, v_{t-k}), \end{aligned}$$

Using that ϵ_{t+1-k} is uncorrelated with v_t and if we ignore the $v_t^{1/2}$ term that modifies the iid shocks. For the consumption strip

$$\begin{aligned} \operatorname{cov}([\log r^s_{t+1} - \log r^1_{t+1}], \log g_{t+1-k}) &= \operatorname{cov}\left(\Big[0.5[1-2\alpha]v_t + v_t^{1/2}\epsilon_{t+1}\Big], \Big[av_{t-k} + v^{1/2}_{t-k}\epsilon_{t+1-k}\Big]\right) \\ &\approx 0.5a[1-2\alpha] \operatorname{cov}(v_t, v_{t-k}). \end{aligned}$$

And, similarly, for the consumption stream

$$\begin{aligned} & \operatorname{cov}([\log r^c_{t+1} - \log r^1_{t+1}], \log g_{t+1-k}) \\ = & \operatorname{cov}\left(\Big[0.5[1-2\alpha]v_t + v_t^{1/2}\epsilon_{t+1}\rho p_v \sigma \eta_{t+1}\Big], \Big[av_{t-k} + v^{1/2}_{t-k}\epsilon_{t+1-k}\Big]\right) \\ \approx & \ 0.5a[1-2\alpha] \operatorname{cov}(v_t, v_{t-k}) + \rho\frac{\kappa_1}{1-\kappa_1\varphi_v}[a^2 + 0.5a\alpha]\sigma \operatorname{cov}(\eta_{t+1}, v_{t-k}) \end{aligned}$$

Section 2 tells us our model needs to account for two facts: (1) The short rate leads growth with negative correlation (low short rate precedes boom); and (2) Excess returns on equity leads growth (high excess return precedes boom). For the short rate, since volatility has positive auto-covariance, we need $[(1-\rho)a^2 + a(\alpha - 1 - \rho\alpha)] < 0$. The first term reflects the effect of predictable growth on the short rate. The second term reflects the risk impact of persistent volatility shocks. This suggests that a $\rho > 0$ (an EIS larger than one) is needed. Given α needs to be negative for a realistic equity premium, the volatility and growth need to have positive correlation and we need to set $a > 0$. (We need to cautious here, the data comes from the nominal rate and our model gives us only the real rate. More later.)

For the consumption strip, with $a > 0$ and $\alpha < 0$, the term $0.5a[1 - 2\alpha]$ is positive and the consumption strip returns leads growth through the higher equity premium. The consumption stream is a bit more complicated since the innovation to volatility, η_{t+1}, has a direct effect on the consumption stream return. $\text{cov}(\eta_{t+1}, v_{t-k})$ is positive, by construction. So the sign depends on the sign of $\rho(a + \alpha/2)$. Maintaining the assumption that $\rho > 0$, we need $a > (-\alpha/2)$. A shock to volatility has two offsetting effects: higher future growth and higher future risk. In order for high equity returns to be associated with a shock to volatility, we need the growth component to be larger than the risk component.

4.2 Numerical Examples

We present three numerical example to show how this works. This is illustrative: we haven't worked out all the implications for means, variances, and autocorrelations of returns. But since cross-correlations do not depend on magnitudes, it's likely we can match the unconditional features about as well as our starting point, Bansal and Yaron (2004).

Parameter values

We start with the Bansal-Yaron (2004) parameter values and vary them as needed. Our version of (6) is a scalar ARMA(1,1) approximation to the two-component Bansal-Yaron process:

$$\log g_{t+1} \quad = \quad (1-\varphi_g)g + \varphi_g \log g_t + a(v_{t-1} - v) - \theta v_{t-1}^{1/2}\epsilon_t + v_t^{1/2}\epsilon_{t+1}$$

Volatility follows the same AR(1) process we have assumed throughout.

$$v_{t+1} \quad = \quad (1-\varphi_v)v + \varphi_v v_t + b\eta_{t+1},$$

The $\epsilon_t, \eta_t \sim NID(0,1)$ and uncorrelated. Again, the (scalar) parameter a controls the effect of volatility on consumption growth. Note that here, even when $a = 0$, consumption growth can have a persistent component. The parameters φ_g (autoregressive term) and θ (moving-average term) control the persistence of the shocks. The long-run risk characteristic of the Bansal Yaron framework is that shocks have a small very persistent component. This is achieved in this setting with $\varphi_g \approx \theta$. These parameters characterize the matrix A in the general specification (6). (More details of how the ARMA(1,1) is calibrated to the Bansal-Yaron specification are in the Appendix. Also includes is the algebra for computing the cross-correlations in our setting.)

In this setting we look at three versions of the model: (1) The Bansal-Yaron setting with long run risk (φ_g and θ non-zero) but no interaction between the stochastic volatility and growth ($a = 0$); (2) The no-long run risk example we looked at in the previous section (Section 4.1) that has no long run risk ($\varphi_g = \theta = 0$) but a volatility-growth connection ($a > 0$); (3) An example calibrated to the cross correlation functions we see in the data that features both long-run risk (φ_g and θ non-zero) and a volatility-growth link $a > 0$. Each of these calibrations is a different endowment process. To make things roughly comparable, the autocorrelation for log consumption growth ($\text{corr}(\log g_t, \log g_{t-1})$) in all three examples is set to 0.12. All the parameter assumptions are in Table 1. You can see the key assumptions in Figures 13, 14, and 15. Part a in each of these figures show the auto-correlation function of endowment growth (i.e., the cross correlation of log growth with itself) and part b shows the cross-correlation function of endowment growth and volatility.

Preference parameters are relatively standard with risk aversion at -9 (high, as usual in asset pricing models). As we saw above, ρ, needs to be positive and this implies and IES greater than one. This is common, but not without controversy, in asset pricing settings. It turns out that in our setting, a higher-than-typical value for this parameter, $\rho = 0.7$ is needed to get the short-rate cross correlation function to match the data. This is a direct implication of the discussion in the preceding section.

Cross correlation of asset returns

This is all a lot of work. Now, lets look at the results. The Bansal-Yaron calibration is in Figure 13. Note the small but persistence effects of shocks in Figure 13a. In this setting, we have set the parameter $a = 0$, so there is no correlation between volatility and consumption growth in Figure 13b. The result here is that the cross correlations in the short rate and equity excess returns look nothing like the data of Section 2. In particular, the short-rate lags the cycle and the equity excess return moves (almost) completely contemporaneously with the endowment growth.

Next, Figure 14 is a numerical version of the example we looked at in Section 4.1. Here, you can see that the endowment growth has much less persistence, (Figure 14a) but does have a correlation between volatility and growth (Figure 14b). The result is that the short rate and equity excess returns lead the cycle in a direction that mimics the signs we see in the data. This is encouraging, but the magnitude and timing is off here since the lead is too short.

Finally, Figure 15 puts all these ingredients together. Figure 15a shows the long persistence of shocks to the growth rate and Figure 15b shows that volatility is correlated with endowment growth. Given the parameters we have attached to our recursive preferences, this volatility pattern translates into equity premiums and the cyclical properties of equity returns in Figure 15d mimic those we see in the data (Figure 7, for example). The exception is the contemporaneous correlation, which has a sharp positive spike at lag zero. This is a direct result of the dividend in our model being consumption itself. In real life this isn't true: the contemporary contemporaneous correlation between consumption growth and dividend growth is small.

In addition, with the high value we have chosen for ρ, the short rate also in Figure 15 captures what we see in the data. The cyclical patterns for excess returns in Section 2 are quite robust to inflation. The inflation component nets out of realized excess returns. For the short rate we need to be more cautious. The data for short rate cross correlations comes from the nominal rate and our model gives us only the real rate. In our model, however, it easy to produce a (real) short rate that has positive (or even zero) correlation with growth with a small change to the parameter ρ without noticeably changing the cross correlations for equity excess returns. However, we leave a more careful investigation of the role of nominal and real rates for later.

Figure 16 shows that this calibration also captures the bond portfolio excess return that we saw in Figure 10. Figure 16 shows the 60 month bonds; but all maturities show a similar pattern. In our model, bond excess returns lead the cycle. However, they do not display the same contemporaneous correlation we see in the date. More generally, capturing a richer set of assets in our setting is an interesting direction to pursue. In particular, understanding how we could generate portfolios with different average equity premia, like seen in the Fama French portfolios, that maintain similar cross-correlation functions?

5 Discussion

We have some work to do to nail down the details, but the numerical example shows that this kind of mechanism can replicate the shape of cross-correlation functions between excess returns and economic growth. There remain some open issues. Most of the open issues are familiar in asset pricing and unique to our setting.

Risk and risk aversion. We've mimicked the cyclical behavior of excess returns in a model in which expected excess returns stem from variations in risk with constant risk aversion. We could have addressed the same issue by allowing risk aversion to vary across states, as it does in Campbell and Cochrane (1999) and Routledge and Zin (2010), or by letting the price of risk vary with the distribution of wealth across individuals, as in Lustig and Van Nieuwerburgh (2005) or Guvenen (2009). We have no particular reason to prefer our approach to these alternatives. Our point is simply that the data implies cyclical variation in excess returns. Exactly how different source of cyclical variation, beyond stochastic volatility, can be correlated with growth might be interesting to consider.

A related issue is the magnitude of risk aversion used in our example. A common argument against risk aversion parameters this large is that when we extrapolate them to large risks (aggregate risks are small), the extent of risk aversion seems unreasonable. Extrapolation of this sort depends a lot on the form of risk preference, including the power form used here and expected utility in general. A natural resolution is a form of risk preference that exhibits different aversion to small and large risks. One such is disappointment aversion. Campanale, Castro, and Clementi (2006) show that the first-order risk aversion exhibited by such preferences exhibits substantial aversion to the small risks we see in the aggregate yet has modest aversion to the large risks faced by individuals. We could model that explicitly in this case, but at some cost of computational complexity. It's simpler to think of our risk preferences as a local approximation for the small risks present in this model.

An alternative is to interpret the risk aversion parameter as aversion to uncertainty about the model's structure. Barillas, Hansen, and Sargent (2008) show that a modest amount of uncertainty about the model's structure (the stochastic process for consumption, for example) can look like extreme risk aversion.

Consumption and returns. Empirical work by Canzoneri, Dumby, and Diba (2007) and Parker and Julliard (2005) shows that the contemporaneous correlation between returns and consumption growth is small, but increases as we expand the time interval. The evidence is Sections 2 is similar: since the correlation is larger with a lag of several months, it's not hard to imagine that the correlation will increase with the time interval. Our theoretical model suggestions an explanation: that the pricing kernel contains an additional factor that is correlated with consumption growth only with a lag.

Alvarez-Jermann bound. Our example has a relatively small bound. Is there enough variability in the pricing kernel with our parameter values?

Endogenous consumption. In a more complete model, variation in the conditional variance of (say) productivity shocks will generate an endogenous response from consumption. Naik (1993) and Primiceri, Schaumburg, Tambalotti (2006) are examples. It remains to be seen whether this produces the interaction we have specified for volatility and consumption growth. Some details on how the lead-lag relationship might work in a production setting with recursive preferences is in Backus, Routledge, and Zin (2007).

Cross-section of returns. We've looked at the possibility that aggregate volatility might account for the cyclical behavior of excess returns. Koijen, Lustig, and Van Nieuwerburgh (2008) use a similar approach to account for the cross section of asset returns.

.

Solution method. There is a growing literature on perturbation methods, in which uncertainty doesn't appear until the second-order approximation. Collard and Juillard (2001) is an elegant example, and Van Binsbergen, Fernandez-Villaverde, Koijen, and Rubio-Ramirez (2008) show how such methods can be extended to models with recursive preferences. Our approach differs in two respects: recursive preferences lead to more complex equilibrium conditions, and we use methods similar to those in finance in which variances appear even with first-order (linear) approximations. This isn't a substitute for high-order approximations, but it allows us to generate reasonably accurate solutions without giving up the convenience of linearity.

6 Conclusions

In short: The returns data suggest that excess returns lead the business cycle by 6-9 months. We replicate this feature in an exchange economy that has a few key features: Recursive preferences; an endowment process that has predictable variation in both consumption growth and its conditional variance; and where predictability in consumption growth depends on past volatility. All of these features are necessary to account for the cyclical behavior of the excess returns, and in particular, the equity premium.

A Data sources

[Later.]

B Theoretical results

The Kreps-Porteus pricing kernel

The pricing kernel in a representative agent model is the marginal rate of substitution between (say) consumption at date t [c_t] and consumption in state s at $t+1$ [$c_{t+1}(s)$]. Here's how that works with recursive preferences. With this notation, the certainty equivalent (3) might be expressed less compactly as

$$\mu_t(U_{t+1}) = \left[\sum_s \pi(s) U_{t+1}(s)^\alpha\right]^{1/\alpha},$$

where $\pi(s)$ is the conditional probability of state s and $U_{t+1}(s)$ is continuation utility. Some derivatives of (2) and (3):

$$\begin{aligned}
\partial U_t/\partial c_t &= U_t^{1-\rho}(1-\beta)c_t^{\rho-1} \\
\partial U_t/\partial \mu_t(U_{t+1}) &= U_t^{1-\rho}\beta\mu_t(U_{t+1})^{\rho-1} \\
\partial \mu_t(U_{t+1})/\partial U_{t+1}(s) &= \mu_t(U_{t+1})^{1-\alpha}\pi(s)U_{t+1}(s)^{\alpha-1}.
\end{aligned}$$

The marginal rate of substitution between consumption at date t and consumption in state s at $t+1$ is

$$\begin{aligned}
\frac{\partial U_t/\partial c_{t+1}(s)}{\partial U_t/\partial c_t} &= \frac{[\partial U_t/\partial \mu_t(U_{t+1})][\partial \mu_t(U_{t+1})/\partial U_{t+1}(s)][\partial U_{t+1}(s)/\partial c_{t+1}(s)]}{\partial U_t/\partial c_t} \\
&= \pi(s)\,\beta\left(\frac{c_{t+1}(s)}{c_t}\right)^{\rho-1}\left(\frac{U_{t+1}(s)}{\mu_t(U_{t+1})}\right)^{\alpha-\rho}.
\end{aligned}$$

The pricing kernel (5) is the same with the probability $\pi(s)$ left out and the state left implicit.

Equity prices and returns

We define equity at t as a claim to consumption from $t+1$ on. The return is the ratio of its value at $t+1$, measured in units of $t+1$ consumption, to the value at t, measured in units of t consumption. The value at $t+1$ is U_{t+1} expressed in c_{t+1} units:

$$\begin{aligned}
U_{t+1}/\left(\partial U_{t+1}/\partial c_{t+1}\right) &= U_{t+1}/\left[(1-\beta)U_{t+1}^{1-\rho}c_{t+1}^{\rho-1}\right] \\
&= (1-\beta)^{-1}u_{t+1}^{\rho}c_{t+1}.
\end{aligned}$$

The value at t is the certainty equivalent expressed in c_t units:

$$\begin{aligned} q_t^c c_t &= \frac{\partial U_t/\partial \mu_t(U_{t+1})}{\partial U_t/\partial c_t}\,\mu_t(U_{t+1}) \;=\; \frac{\beta \mu_t(U_{t+1})^\rho}{(1-\beta)c_t^\rho}\,c_t \\ &= \beta(1-\beta)^{-1}\mu_t(g_{t+1}u_{t+1})^\rho c_t. \end{aligned}$$

The return is the ratio:

$$\begin{aligned} r_{t+1}^c &= \beta^{-1}\left[u_{t+1}/\mu_t(g_{t+1}u_{t+1})\right]^\rho g_{t+1} \\ &= \beta^{-1}\left[g_{t+1}u_{t+1}/\mu_t(g_{t+1}u_{t+1})\right]^\rho g_{t+1}^{1-\rho}. \end{aligned}$$

Check to see if this satisfies the Euler equation:

$$\begin{aligned} E_t\left(m_{t+1}r_{t+1}^c\right) &= E_t\left[g_{t+1}u_{t+1}/\mu_t(g_{t+1}u_{t+1})\right]^\alpha \\ &= \mu_t(g_{t+1}u_{t+1})^\alpha/\mu_t(g_{t+1}u_{t+1})^\alpha \;=\; 1. \end{aligned}$$

ARMA(1,1) Specification of Bansal Yaron

[later]

Computing cross correlations

Recall that the state is $s_t = (x_t, v_t)$ and the "expanded state" is $s_t^* = (s_t, s_{t-1})$. The latter has the law of motion

$$s_{t+1}^* = A_* s_t^* + B_* w_{t+1},$$

with

$$A_s = \begin{bmatrix} A & a \\ 0 & \varphi_v \end{bmatrix}, \quad B_s = \begin{bmatrix} B \\ b \end{bmatrix}$$

and

$$A_* = \begin{bmatrix} A_s & 0 \\ I & 0 \end{bmatrix}, \quad B_* = \begin{bmatrix} B_s \\ 0 \end{bmatrix}.$$

The unconditional variance is

$$G(0) = E\left(s_t^* s_t^{*\top}\right) = A_* G(0) A_*^\top + B_* B_*^\top.$$

We compute $G(0)$ iteratively using Hansen and Sargent's (2005) Matlab program `doublej.m`. Autocovariances follow from

$$G(k) = E\left(s_t^* s_{t-k}^{*\top}\right) = \begin{cases} A_*^k G(0) & k > 0 \\ G(0)(A_*^k)^\top & k < 0. \end{cases}$$

Since $G(-k) = G(k)^\top$, positive k is sufficient.

Returns and excess returns are linear functions of the expanded state: $r_t = Hs_t^*$ say for a vector of returns and excess returns. Autocovariances are

$$E\left(r_t r_{t-k}^\top\right) \;=\; hG(k)h^\top.$$

Cross-covariances are off-diagonal elements and cross-correlations are scaled by standard deviations.

References

Alvarez, Fernando, and Urban Jermann, 2005, "Using asset prices to measure the persistence of the marginal utility of wealth," *Econometrica* 73, 1977-2016.

Ang, Andrew, Monika Piazzesi, and Min Wei, 2006, "What does the yield curve tell us about GDP growth?" *Journal of Econometrics* 131, 359-403.

Atkeson, Andrew, and Patrick Kehoe, 2008, "On the need for a new approach to analyzing monetary policy," *NBER Macroeconomics Annual*, forthcoming.

Backus, David K., Bryan R. Routledge, and Stanley E. Zin, 2007, "Asset prices in business cycle analysis," manuscript, November.

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

Barillas, Francisco, Lars Peter Hansen, and Thomas J. Sargent, 2008, "Doubts or variability," manuscript, July.

Campanale, Claudio, Rui Castro, and Gian Luca Clementi, 2006, "Asset pricing in a general equilibrium production economy with Chew-Dekel risk preferences," manuscript, December.

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

Canzoneri, Matthew B., Robert E. Cumby, and Behzad T. Diba, 2007, "Euler equations and money market interest rates: a challenge for monetary policy models," *Journal of Monetary Economics* 54, 1863-1881.

Casassus, Jaime, Peng Liu and Ke Tang, 2010, Long-term economic relationships and correlation structure in commodity markets, Pontificia Universidad Catolica de Chile Working Paper, February, 2010.

Constantinides, George M. and Anisha Ghosh, 2010, "The Predictability of Returns with Regime Shifts in Consumption and Dividend Growth," Carnegie Mellon Working Paper, January, 2010

Collard, Fabrice, and Michel Juillard, 2001, "Accuracy of stochastic perturbation methods: the case of asset pricing models," *Journal of Economic Dynamics and Control* 25, 979-999.

Epstein, Larry G., and Stanley E. Zin, 1989, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework," *Econometrica* 57, 937-969.

Estrella, Arturo, and Gikas A. Hardouvelis, 1991, "The term structure as a predictor of real economic activity," *Journal of Finance* 46, 555-576.

Fama, Eugene F., and Kenneth R. French, 1989, "Business conditions and expected returns on stocks and bonds," *Journal of Monetary Economics* 25, 23-49.

Fama, Eugene F., and Kenneth R. French, 1992, "The cross-section of expected stock returns," *Journal of Finance* 47, 427-467.

Gallmeyer, Michael F., Burton Hollifield, and Stanley E. Zin, 2005, "Taylor rules, McCallum rules and the term structure of interest rates," *Journal of Monetary Economics* 52, 921-950.

Guvenen, Fatih 2009, A Parsimonious Macroeconomic Model for Asset Pricing* *Econometrica*, November 2009, Vol. 77, No 6, pp. 1711-1750.

Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, "Consumption strikes back? Measuring long-run risk," *Journal of Political Economy* 116, 260-302.

Hansen, Lars Peter, and Thomas J. Sargent, 2005, *Recursive Models of Dynamic Linear Economies*, manuscript, September.

Parker, Jonathan A., and Christian Julliard, 2005, "Consumption risk and the cross section of expected returns," *Journal of Political Economy* 113, 185-222.

Kandel, Shmuel, Robert F. Stambaugh, 1990, "Expectations and volatility of consumption and asset returns," *Review of Financial Studies* 3, 207-32.

King, Robert G., and Mark W. Watson, 1996, "Money, prices, interest rates and the business cycle," *Review of Economics and Statistics* 78, 35-53.

Koijen, Ralph, Hanno Lustig, and Stijn Van Nieuwerburgh, 2008, "The bond risk premium and the cross-section of expected stock returns," manuscript.

Kreps, David M., and Evan L. Porteus, 1978, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica* 46, 185-200.

Lustig, Hanno, and Stijn Van Nieuwerburgh, 2005, "Housing collateral, consumption insurance and risk premia: an empirical perspective," *Journal of Finance* 60, 1167-1219.

Mueller, Philippe, 2008, "Credit spreads and real activity," manuscript, London School of Economics, January.

Naik, Vasanttilak, 1994, "Asset prices in time-varying production economies with time-varying risk," *Review of Financial Studies* 7, 781-801.

Primiceri, Giorgio E., Ernst Schaumburg, and Andrea Tambalotti, 2006, "Intertemporal disturbances," NBER Working Paper 12243, May.

Routledge, Bryan R., and Stanley E. Zin, 2010, "Generalized disappointment aversion and asset prices," *Journal of Finance* , Volume 65, 4

Rouwenhorst, K. Geert, 1995, "Asset pricing implications of equilibrium business cycle models," in T.F. Cooley, editor, *Frontiers of Business Cycle Research*, Princeton, NJ: Princeton University Press.

Sargent, Thomas J., and Christopher A. Sims, 1977, "Business cycle modeling without pretending to have too much *a priori* economic theory," in C. Sims et al., eds., *New Methods in Business Cycle Research*, Minneapolis: Federal Reserve Bank of Minneapolis.

Stock, James H., and Mark W. Watson, 1989, "New indexes of coincident and leading economic indicators," in *NBER Macroeconomics Annual*, O.J. Blanchard and S. Fischer, eds., 352-394.

Stock, James H., and Mark W. Watson, 2003, "Forecasting output and inflation: the role of asset prices," *Journal of Economic Literature* 41, 788-829.

Stock, James H., and Mark W. Watson, 2005, "Implications of dynamic factor models for VAR analysis," NBER Working Paper No. 11467, June.

van Binsbergen, Jules, Jesus Fernandez-Villaverde, Ralph Koijen, and Juan Rubio-Ramirez, 2008, "Likelihood estimation of DSGE models with Epstein-Zin preferences," manuscript, March.

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

Table 1: Numerical Examples Parameters for Section 4.2

$$\log g_{t+1} = (1-\varphi_g)g + \varphi_g \log g_t + a(v_{t-1} - v) - \theta v_{t-1}^{1/2}\epsilon_t + v_t^{1/2}\epsilon_{t+1}$$
$$v_{t+1} = (1-\varphi_v)v + \varphi_v v_t + \sigma\eta_{t+1}$$

		(1) Bansal-Yaron	(2) Section 4.1	(3) Bansal-Yaron $a > 0$
AR(1)	φ_g	0.9400	0	0.9400
MA(1)	θ	-0.8604	0	-0.9400
Volatility-growth connection	a	0	24.5	30
Endowment autocorrelation		0.12	0.12	0.12
See Figure		13	14	15,16
		Common Endowment Parameters		
Mean endowment growth	g		0.0015	
Mean volatility	v		0.080^2	
Volatility autocorrelation	φ_v		0.987	
Volatility of volatility	σ		0.23×10^{-5}	
		Common Preference Parameters		
Risk parameter	α		-9	
IES parameter	ρ		0.7	
Discount	$\kappa_1 \approx \beta$		0.997	

Parameters for the examples Section 4.2 and Figures 13 , 14 , 15, and 16.

Figure 1: Cross Correlation: Equity Returns and Industrial Production

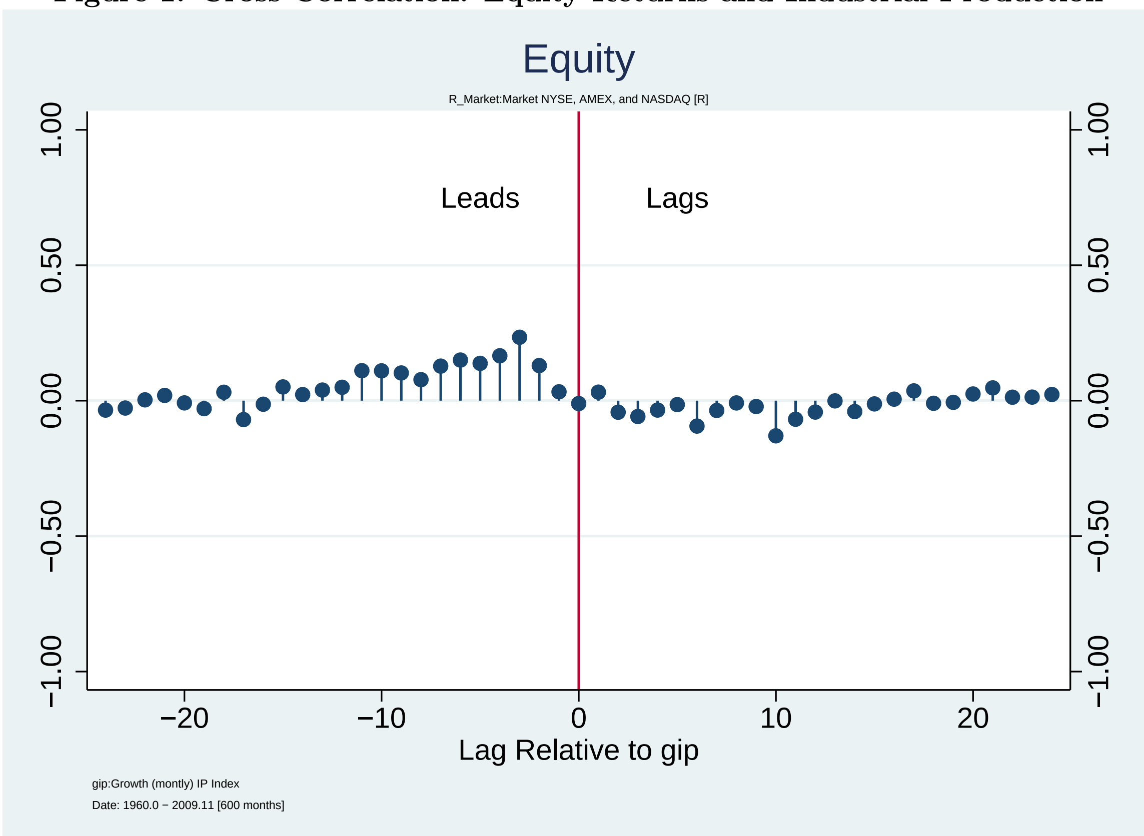

The figure depicts the cross-correlation function for the return on an aggregate equity portfolio and the monthly growth rate of industrial production. On the left side of the figure, the return leads growth, on the right side it lags. The sample period is 1960-present.

Figure 2: Cross Correlation: Equity Returns and Industrial Production

a. Year-on-year growth and monthly returns

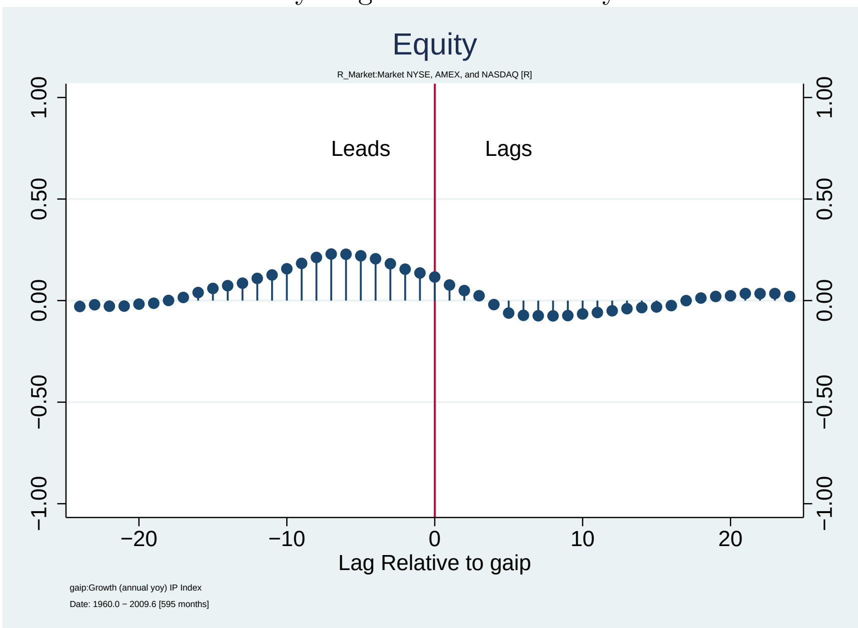

b. Year-on-year growth and year-on-year returns

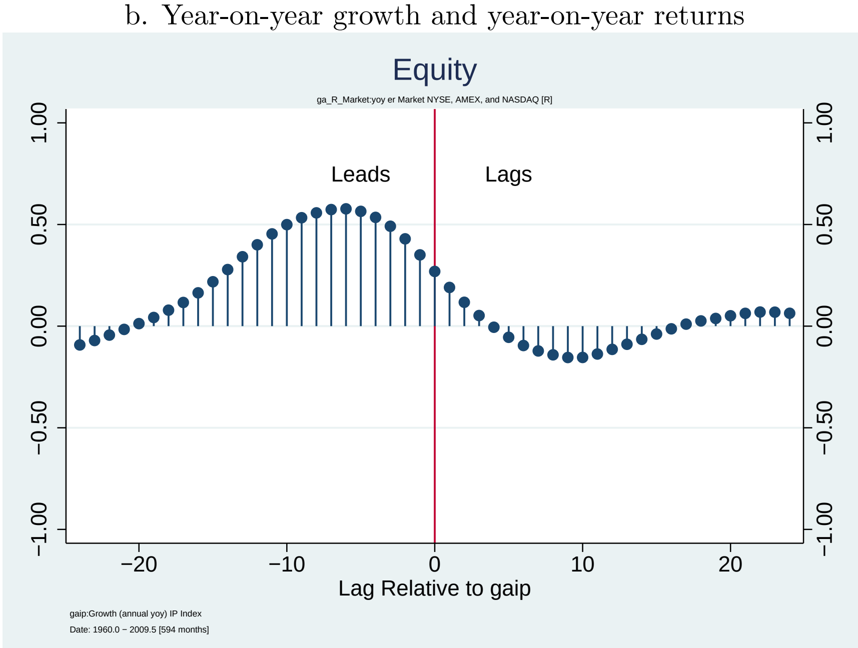

The figure depicts the cross-correlation function for the return on an aggregate equity portfolio and the monthly growth rate of industrial production. (a) Shows year-on-year growth and monthly equity returns. (b) Shows the year-on-year growth and year-on-year equity returns. The sample period is 1960-present.

Figure 3: Cross Correlation: Equity Returns

a. Employment growth

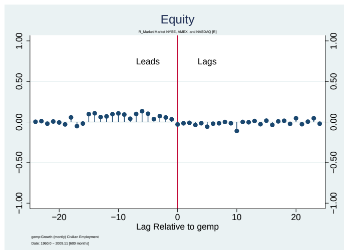

b. Year-on-Year Employment growth

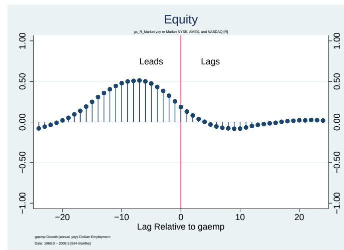

c. Consumption growth

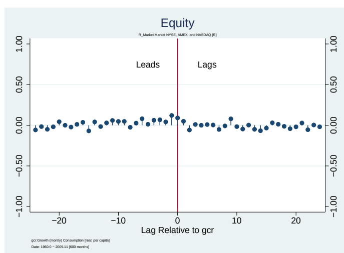

d. Year-on-Year Consumption growth

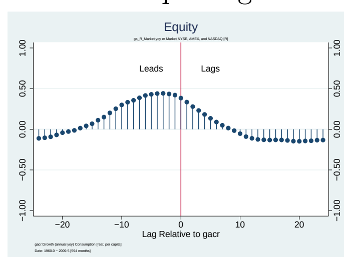

e. S&P 500 earnings growth

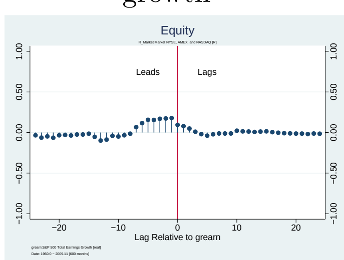

f. Year-on-year S&P 500 earnings growth

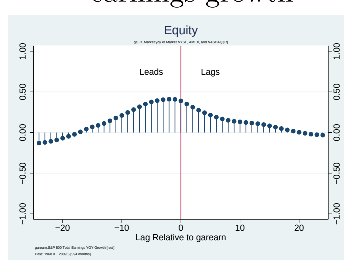

g. S&P 500 dividend growth

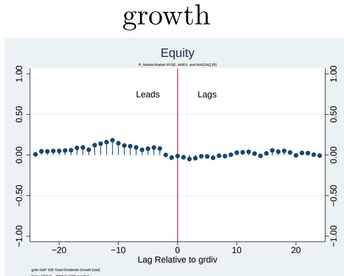

h. Year-on-year S&P 500 dividend growth

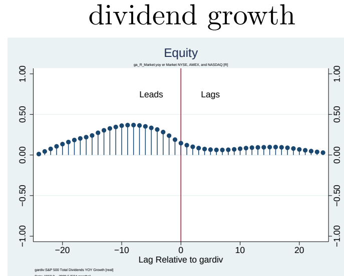

The figure depicts the cross-correlation function for the return on an aggregate equity portfolio and the growth rate of various measures of economic activity. Figures on the left depcit monthly growth and monthly equity returns. Figures on the right are year-on-year growth and year-on-year equity returns. The sample period is 1960-present.

Figure 4: Time Series Equity Returns

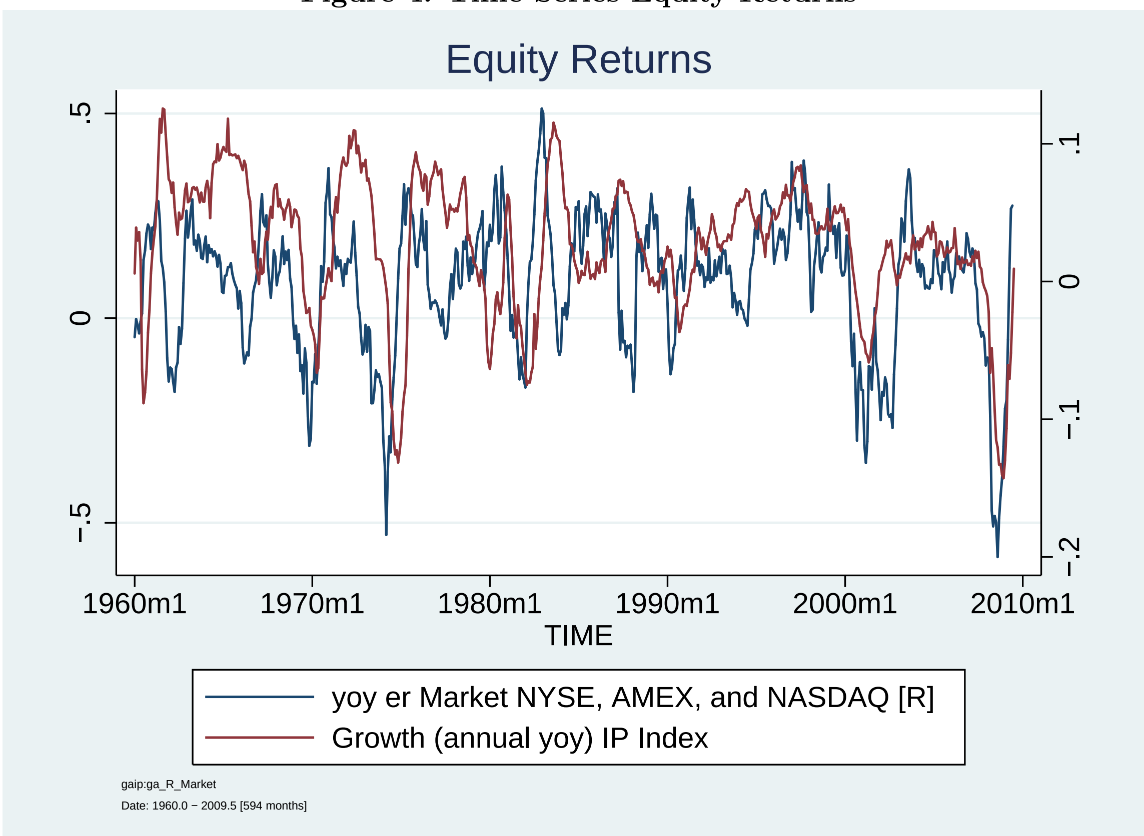

Time-series plot of industrial produciton year-on-year growth rate (right scale) and year-on-year equity returns (left scale). Data is 1960 to present.

Figure 5: Cross Correlation: Interest Rates and Industrial Production

[a]Yield Spread; Monthly IP growth

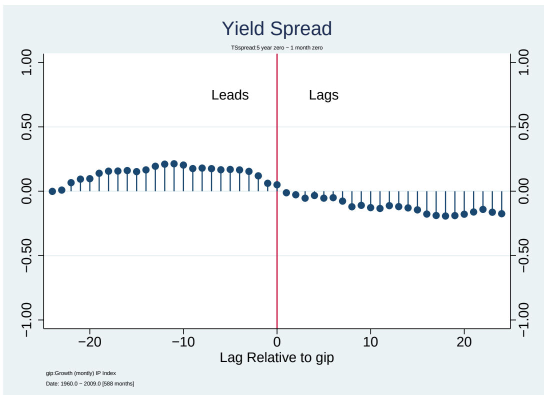

[b] Yield Spread; Year-onyear IP growth

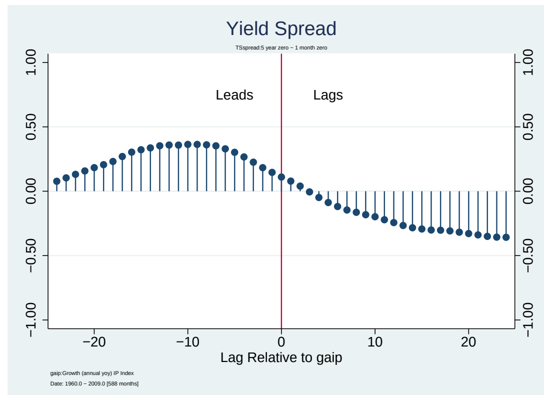

[a]Yield Short Bond; Monthly IP growth

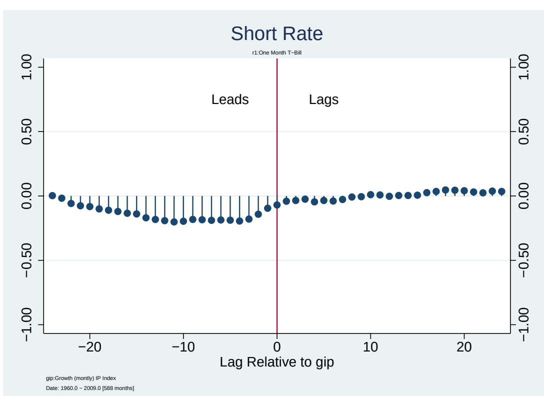

[b] Yield Short Bond; Year-onyear IP growth

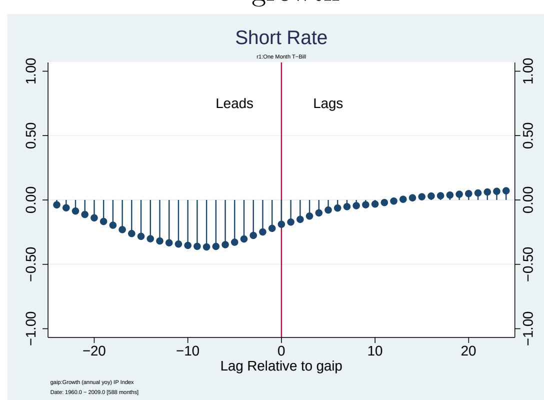

The figure depicts the cross-correlation function with industrial production. The top figures are for the yield spread (the difference in the five-year and the one-month default-free bond). The bottom figures are for the one-month default-free bond. Figures on the left depcit monthly growth and yields. Figures on the right are year-on-year growth and yields. The sample period is 1960-present. The sample period is 1960-present.

Figure 6: Cross Correlation: Excess Equity Returns and Industrial Production

a. Monthly growth and monthly returns

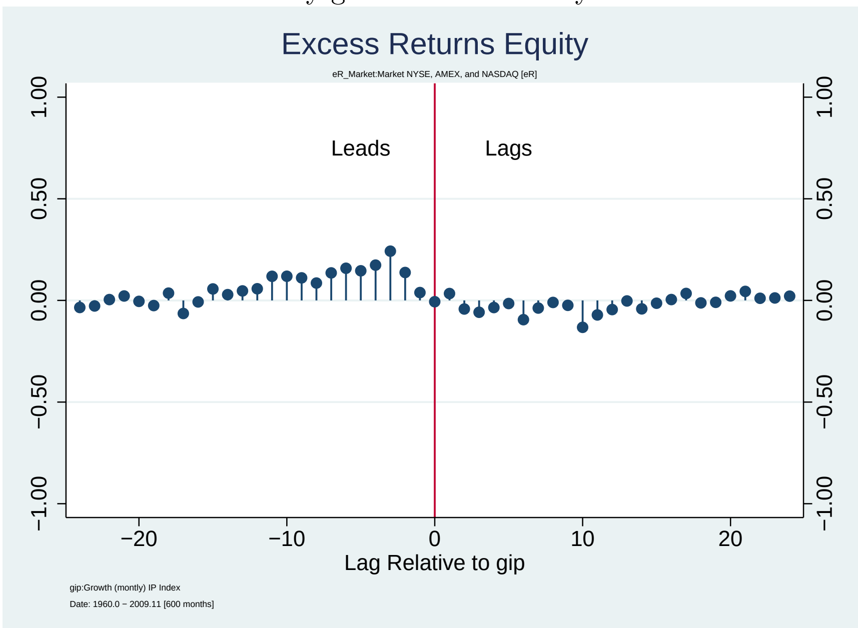

b. Year-on-year growth and year-on-year returns

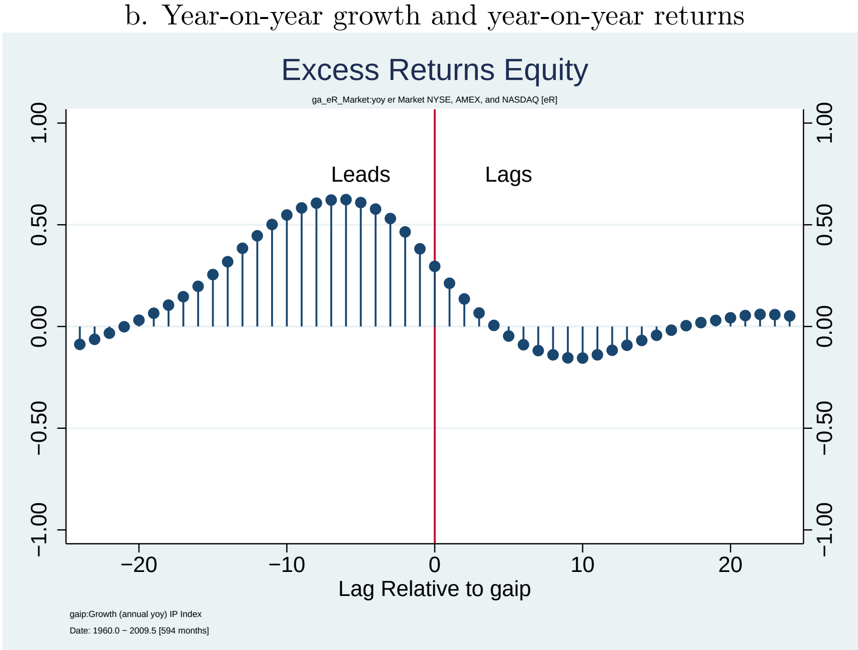

The figure depicts the cross-correlation function for the return on an aggregate equity portfolio in excess of the one-mooth risk free rate (i.e. excess returns) and the growth rate of industrial production. The top figure is monthly growth and monthly returns. The bottom is year-on-year growth and year-on-year excess returns. The sample period is 1960-present.

Figure 7: Cross Correlation: Excess Equity Returns by Industry and Industrial Production

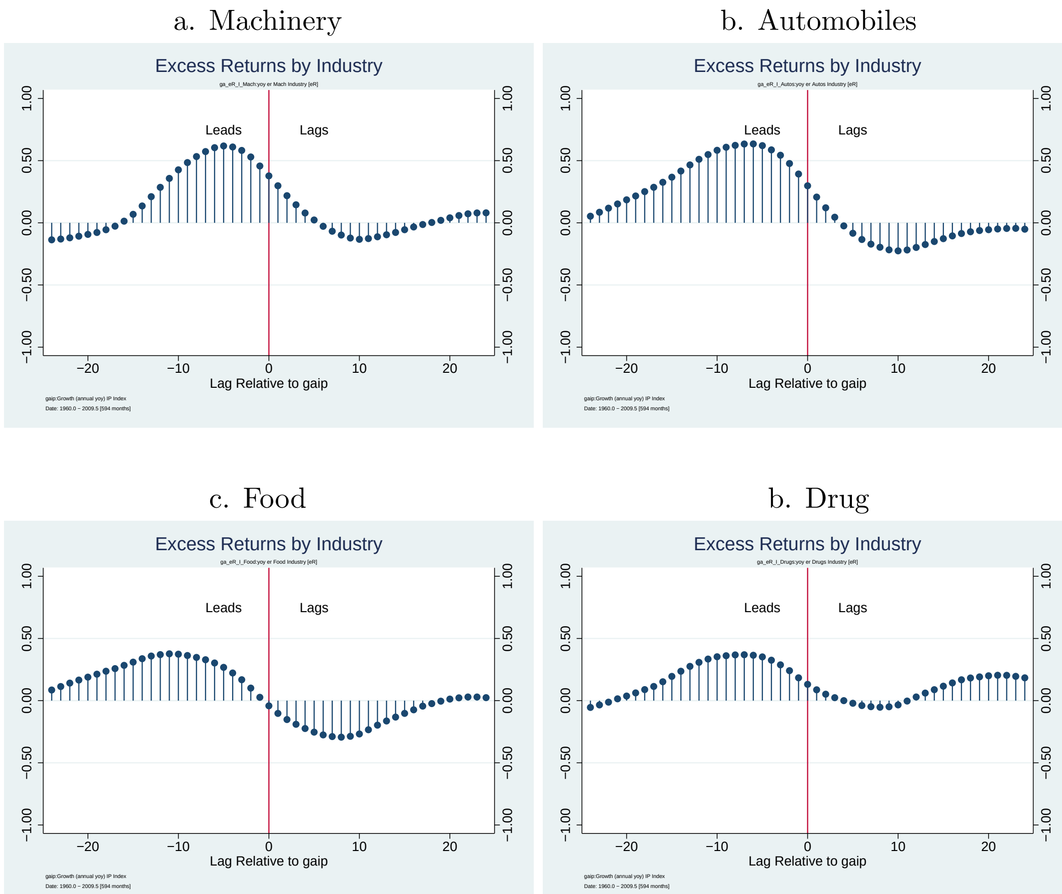

The figure depicts the cross-correlation function for the return on equity portfolios by industry in excess of the one-mooth risk free rate (i.e. excess returns) and the growth rate of industrial production. See Ken French web site for industry SIC definitions. Figures are all year-on-year growth and year-on-year excess returns. The sample period is 1960-present.

Figure 8: Cross Correlation: Excess Equity Returns by Sorted Portfolios and Industrial Production

a. Size-sorted – smallest decile

b. Size-sorted – largest decile

c. BEME-sorted – smallest decile

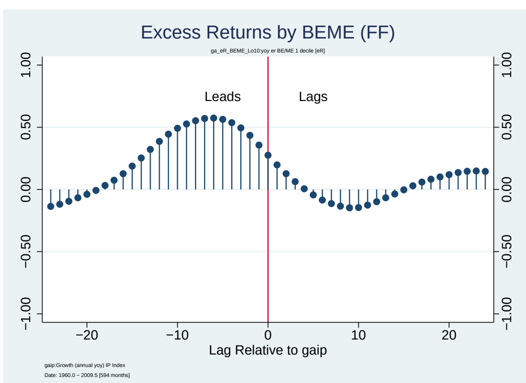

d. BEME-sorted – largest decile

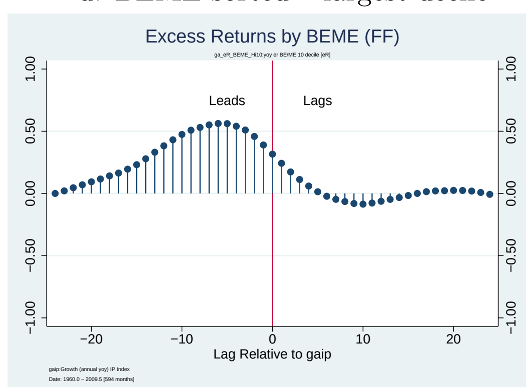

The figure depicts the cross-correlation function for the return on equity portfolios by size and book-to-market characteristics in excess of the one-mooth risk free rate (i.e. excess returns) and the growth rate of industrial production. See Ken French web site for procedures for constructing size and book-to-market sorted portfolios. Figures are all year-on-year growth and year-on-year excess returns. The sample period is 1960-present.

Figure 9: Cross Correlation: Excess Bond Returns and Industrial Production

a. Bond portfolio: 1-6 months

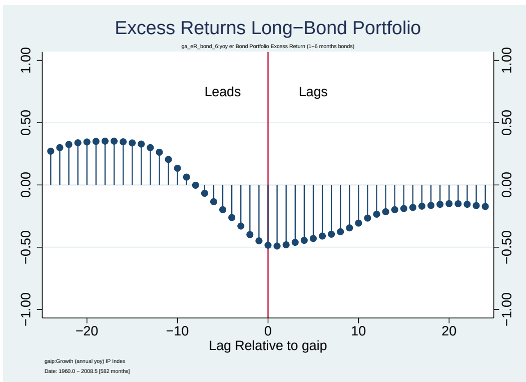

b. Bond portfolio: 19-24 months

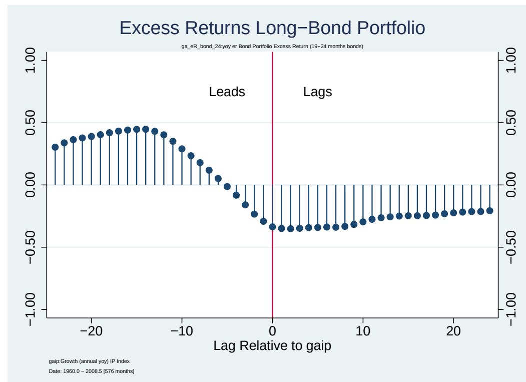

c. Bond portfolio: 55-60 months

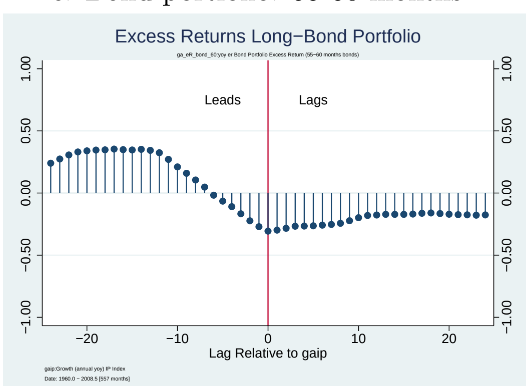

d. Bond portfolio: 115-120 months

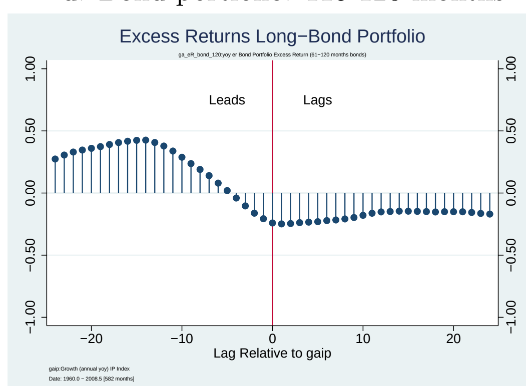

The figure depicts the cross-correlation function for the return on portfolios of bonds in excess of the one-mooth risk free rate (i.e. excess returns) and the growth rate of industrial production. The bond portfolios are from CRSP/Fama and hold bonds of fixed maturity (rebalancing montly). Figures are all year-on-year growth and year-on-year excess returns. The sample period is 1960-present.

Figure 10: Cross Correlation: Excess Oil Returns and Industrial Production

a. Oil 3 month

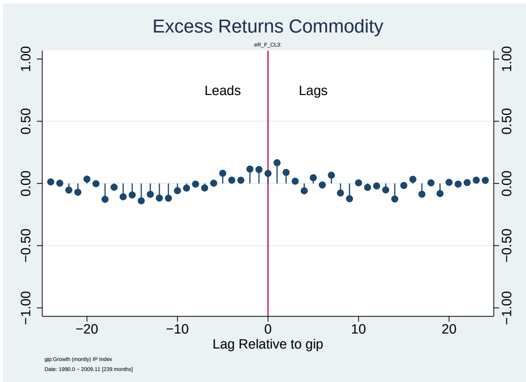

b. Oil 3 month (year-on-year)

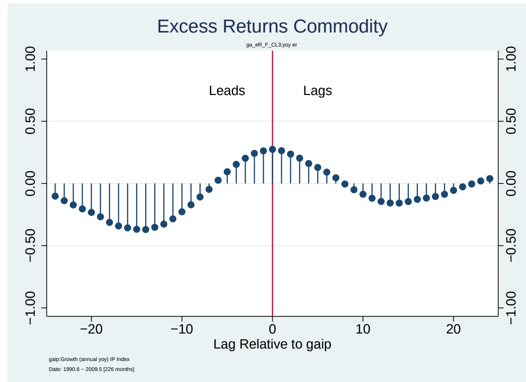

c. Oil 12 month

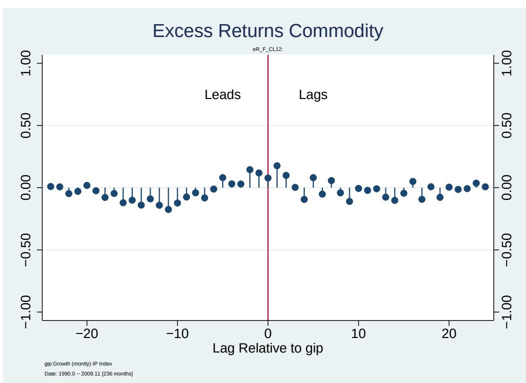

d. Oil 12 month (year-on-year)

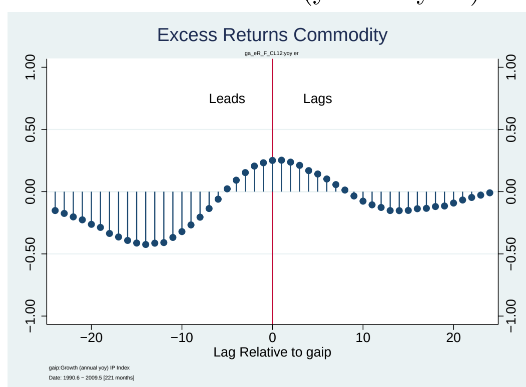

e. Oil 24 month

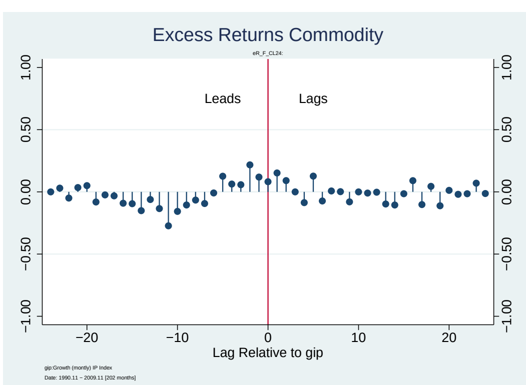

f. Oil 24 month (year-on-year)

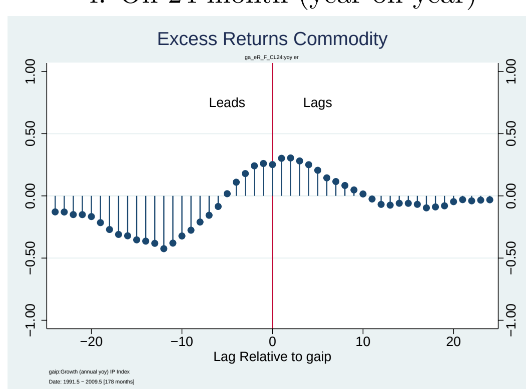

The figures depicts the cross-correlation function for the return on a portfolios of commodities in excess of the one-mooth risk free rate (i.e. excess returns) and the growth rate of industrial production. Returns are one-month (log) excess returns on fully-collaterialized futures contract. Figures on the left depcit monthly growth and monthly equity returns. Figures on the right are year-on-year growth and year-on-year equity returns. Oil is Crude Oil (West Texas Intermediate) futures price from NYMEX. The sample period is 1990-2010.

Figure 11: Cross Correlation: Excess Commodities Returns and Industrial Production

a. OJ 3 month

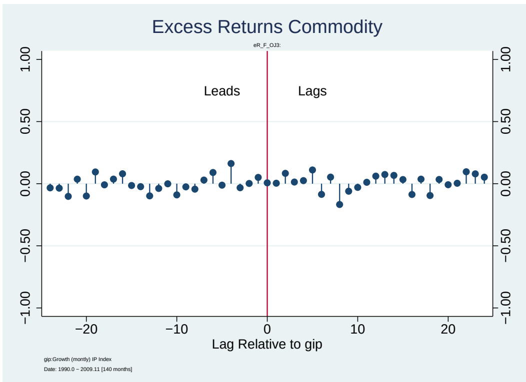

b. OJ 3 month (year-on-year)

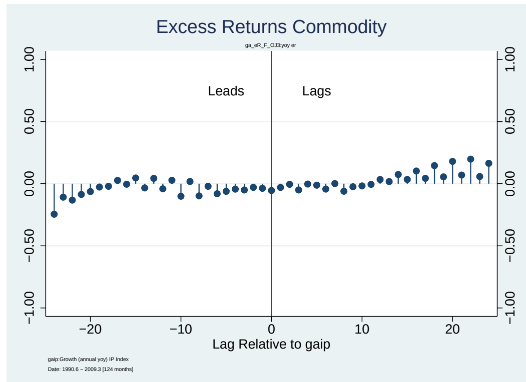

c. Wheat 3 month

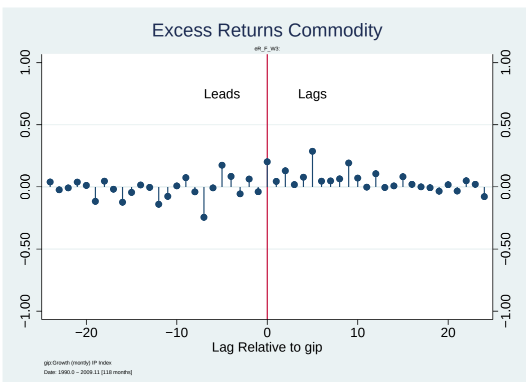

d. Wheat 3 month (year-on-year)

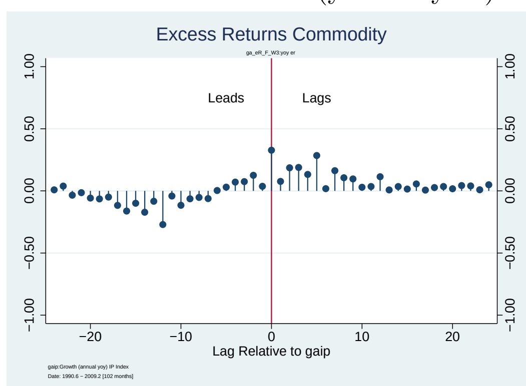

d. Copper 3 month

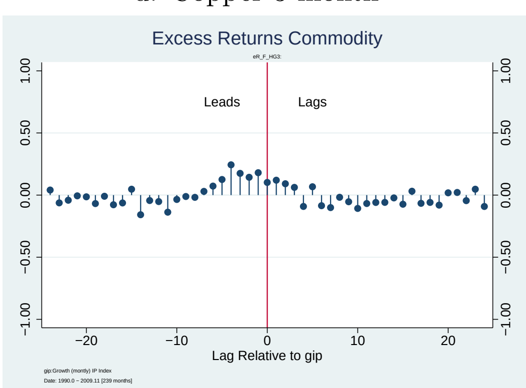

e. Copper 3 month (year-on-year)

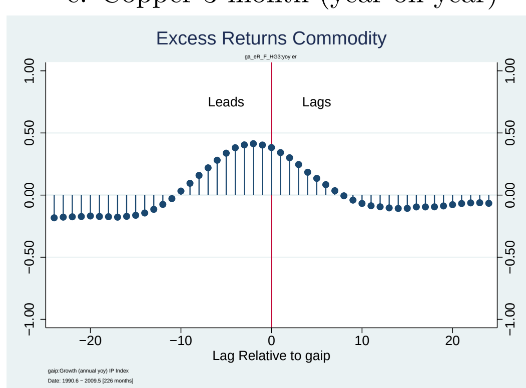

The figures depicts the cross-correlation function for the return on a portfolios of commodities in excess of the one-mooth risk free rate (i.e. excess returns) and the growth rate of industrial production. Returns are one-month (log) excess returns on fully-collaterialized futures contract. Figures on the left depcit monthly growth and monthly equity returns. Figures on the right are year-on-year growth and year-on-year equity returns. OJ is Florida Concentrated Orange Juice. Wheat is wheat. Copper is High-grade copper. Price from NYMEX. The sample period is 1990-2010.

Figure 12: Cross Correlation with endowment growth: Numerical Example "Bansal-Yaron"

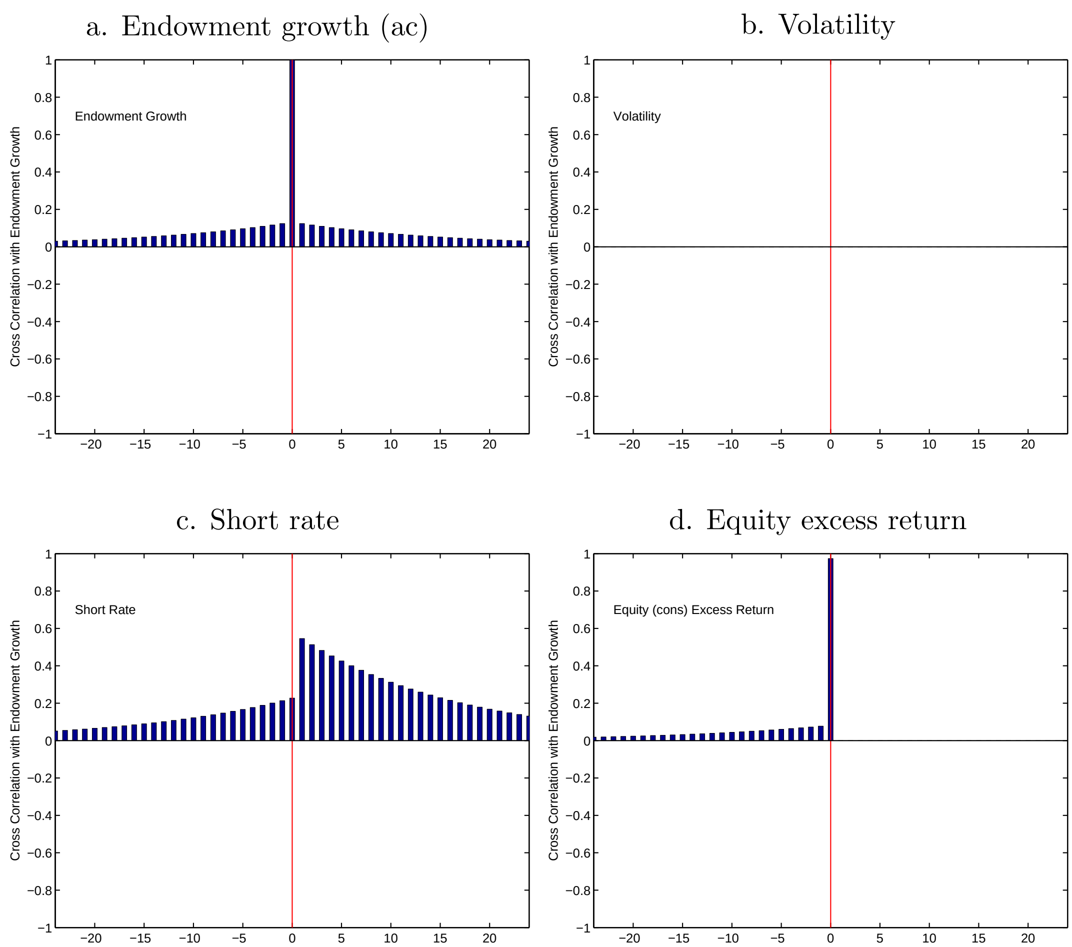

Cross correlation functions with endowment growth. The example shows the Bansal-Yaron setting with long run risk (φ_g, θ non-zero) but no interaction between the stochastic volatility and growth ($a = 0$). Note (a) is the endowment autocorrelation function. Parameters see Table 1.

Figure 13: Cross Correlation with endowment growth: Numerical Example of Section 4.1

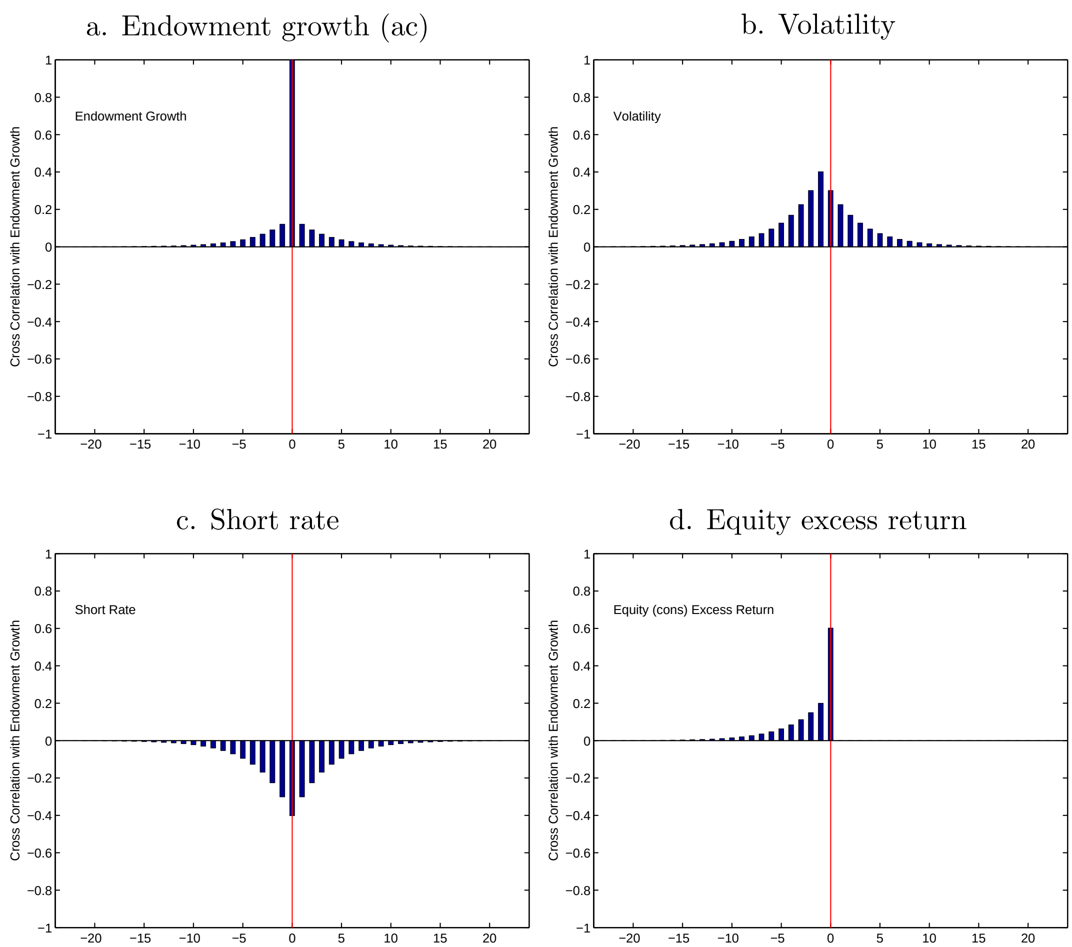

Cross correlation functions with endowment growth. This example shows the setting with no-long run risk ($\varphi_g = \theta = 0$) but a volatility-growth connection (a¿0) as in Section 4.1. Note (a) is the endowment autocorrelation function. Parameters see Table 1.

Figure 14: Cross Correlation with endowment growth: Numerical Example "bansal-Yaron" plus $a > 0$

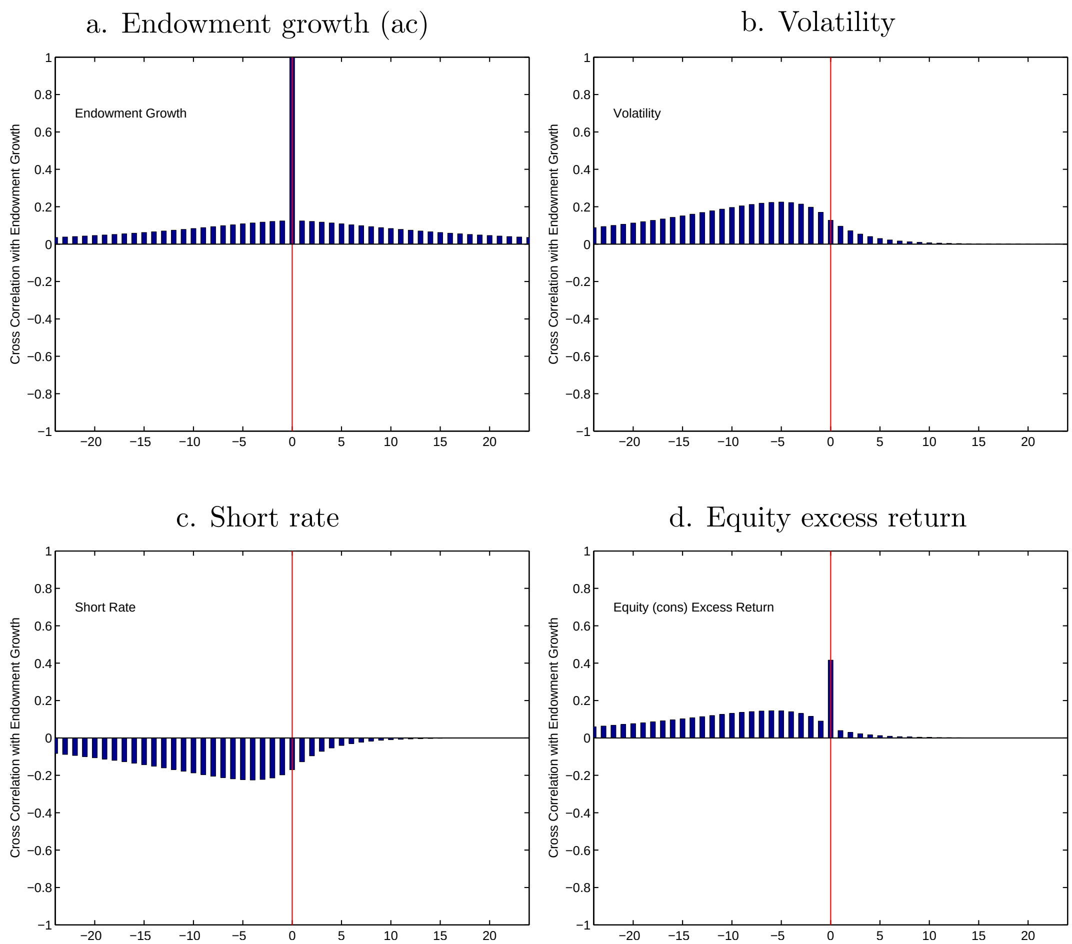

Cross correlation functions with endowment growth. The example is calibrated to the cross correlation functions we see in the data that features both long-run risk (φ_g, θ non-zero) and a volatility-growht link $a > 0$. Note (a) is the endowment autocorrelation function. Parameters see Table 1.

Figure 15: Cross Correlation with endowment growth: Numerical Example "bansal-Yaron" plus $a > 0$

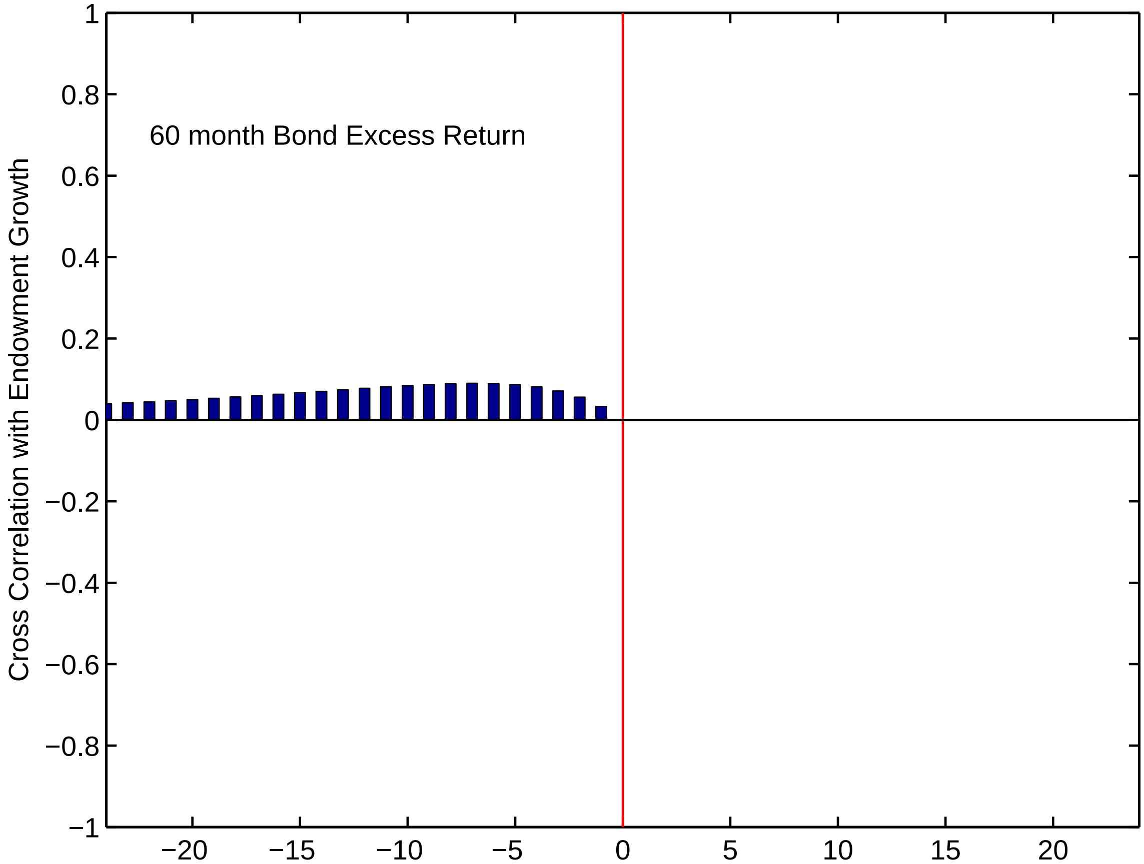

Cross correlation functions with endowment growth with the one-month holding period excees return on a portfolio of 60 month bonds. The example is calibrated to the cross correlation functions we see in the data that features both long-run risk (φ_g, θ non-zero) and a volatility-growht link $a > 0$. Parameters see Table 1.