

Consumption Strikes Back?: Measuring Long-Run Risk¹

Lars Peter Hansen
Department of Economics
University of Chicago
Email:l-hansen@uchicago.edu

John C. Heaton
Graduate School of Business
University of Chicago
Email:john.heaton@uchicago.edu

Nan Li
School of Business
National University of Singapore

First draft: May 2004
Current draft: October 2007

¹We thank Fernando Alvarez, Ravi Bansal, John Cochrane, Ken Judd, Anil Kashyap, Mark Klebanov, Junghoon Lee, Jonathan Parker, Tano Santos, Tom Sargent, Chris Sims, Grace Tsiang, Pietro Veronesi, Kenji Wada, and two anonymous referees for valuable comments. We also thank workshop participants at Chicago, Columbia, Duke, Emory, INSEAD, LBS, MIT, Montreal, NBER, NYU, and Northwestern. Hansen gratefully acknowledges support from the National Science Foundation under award number SES0519372, Heaton from the Center for Research in Securities Prices, and Li from the Olin Foundation.

Abstract

We characterize and measure a long-run risk return tradeoff for the valuation of cash flows exposed to fluctuations in macroeconomic growth. This tradeoff features the risk prices of cash flows that are realized far into the future but are reflected in asset values. We apply this analysis to a claims on aggregate cash flows, as well as to the cash flows from value and growth portfolios. Based on vector autoregressions, we characterize the dynamic response of cash flows to macroeconomic shocks and document that there are important differences in the long-run responses. We isolate those features of a recursive utility model and the consumption dynamics needed for the long run valuation differences among these portfolios to be sizable. Finally, we show how the resulting measurements vary when we alter the statistical specifications of cash flows and consumption growth.

Key words: Risk-return tradeoff, long run, asset pricing, macroeconomic risk

1 Introduction

In this paper we ask: how is risk exposure priced in the long run? Current period values of cash flows depend on their exposure to macroeconomic risks, risks that cannot be diversified. The risk exposures of cash flows are conveniently parameterized by the gap between two points in time: the date of valuation and the date of the payoff. We study how such cash flows are priced, including an investigation of the limiting behavior as the gap in time becomes large. While statistical decompositions of cash flows are necessary to the analysis, we supplement such decompositions with an economic model of valuation to fully consider the pricing of risk exposure in the long run.

Long-run contributions to valuation are of interest in their own right, but there is a second reason for featuring the long run in our analysis. Highly stylized economic models, like the ones we explore, are misspecified when examined with full statistical scrutiny. Behavioral biases or transactions costs, either economically grounded or metaphorical in nature, challenge the high frequency implications of pricing models. Similarly, while un-modeled features of investor preferences such as local durability or habit persistence alter short run value implications, these features may have transient consequences for valuation.¹ One option is that we repair the valuation models by appending *ad hoc* transient features, but instead we accept the misspecification and seek to decompose the implications.

Characterizing components of pricing that dominate over long horizons helps us understand better the implications of macroeconomic growth rate uncertainty for valuation. Applied time series analysts have studied extensively a macroeconomic counterpart to our analysis by characterizing how macroeconomic aggregates respond in the long run to underlying economic shocks.² The unit root contributions measured by macroeconomists are a source of long-run risk that should be reflected in the valuation of cash flows. We measure this impact on financial securities.

Our study considers the prices of exposures to long run macroeconomic uncertainty, and the implications of these prices for the values of cash flows generated by portfolios studied previously in finance. These portfolios are constructed from stocks with different ratios of book value to market value of equity. It has been well documented that the one period average returns to portfolios of high book-to-market stocks (value portfolios) are substantially larger than those of portfolios of low book-to-market stocks (growth portfolios).³ We find that the cash flows of value portfolios exhibit positive co-movement in the long run with macroeconomic shocks while the growth portfolios show little covariation with these shocks. Equilibrium pricing reflects this heterogeneity in risk exposure: risk averse investors must be compensated more to hold value portfolios. We quantify how this compensation depends on investor preferences and on the cash flow horizon.

¹Analogous reasoning led Daniel and Marshall (1997) to use an alternative frequency decomposition of the consumption Euler equation.

²For instance, Cochrane (1988) uses time series methods to measure the importance of permanent shocks to output, and Blanchard and Quah (1989) advocated uses restrictions on long run responses to identify economic shocks and measure their importance.

³See, for example, Fama and French (1992)

The pricing question we study is distinct from the more common question in empirical finance: what is the short-run tradeoff between risk and return measured directly from returns? Even when equities are explored it is common to use the one period return on equity as an empirical target. Instead we decompose prices and returns by horizon. For instance, the one-period return to a portfolio is itself viewed as the return to a portfolio of claims to cash flows at different horizons. Moreover the price of a portfolio reflects the valuation of cash flows at different horizons. We use these representations to ask: when will the cash flows in the distant future be important determinants of the one-period equity returns and how will the long-run cash flows be reflected in portfolio values? From this perspective we find that there are important differences in the risks of value and growth portfolios that are most dramatic in the long run.

Given our choice of models and evidence, we devote part of our analysis to measuring the estimation accuracy and to assessing the sensitivity of our risk measurements to the dynamic statistical specification. Both tasks are particularly germane because of our consideration of long-run implications. Our purpose in making such assessments is to give a clear understanding of where sample information is informative and where long-run prior restrictions are most relevant.

In section 2 we present our methodology for log-linear models and derive a long-run risk return tradeoff for cash flow risk. In section 3 we use the recursive utility model to show why the intertemporal composition of risk that is germane to an investor is reflected in both short run and long run risk return tradeoffs. In section 4 we identify important aggregate shocks that affect consumption in the long run. Section 5 constructs the implied measures of the risk-return relation for portfolio cash flows. Section 6 concludes.

2 Long run risk

Characterization of long run implications through the analysis of steady states or their stochastic counterparts is a familiar tool in the study of dynamic economic models. We apply an analogous idea for the long-run valuation of stochastic cash flows. The resulting valuation allows us to decompose long-run expected returns into the sum of a risk-free return and a long-run risk premium. This long-run risk premium is further decomposed into the product of a measure of long-run exposure to risk and the price of long-run risk. Unlike approaches that examine the relationship between one-period expected returns and preferences that feature a concern about long-run risk (e.g. (Bansal and Yaron 2004) and (Campbell and Vuolteenaho 2004)) our development focuses on the intertemporal composition of risk prices, and in particular on the implied risk prices for cash flows far into the future. The result we establish for long-run expected returns has the same structure as the standard decomposition of one-period expected returns into a risk-free component plus the product of the price of risk and the risk exposure.

2.1 Stochastic discount factors

Uncertainty in the economy is given by the dynamics of a state vector x_t which evolves according to a first-order vector autoregression:

$$x_{t+1} = Gx_t + Hw_{t+1}. \quad (1)$$

The matrix G has eigenvalues with absolute values that are strictly less than one. $\{w_{t+1} : t = 0, 1, \dots\}$ is a vector of normal random variables that are independently and identically distributed over time with mean zero and covariance matrix I . Although we consider a first-order system, higher-order systems are accommodated by augmenting the state vector.

The time t price of an asset payoff at time $t + 1$ is determined by a stochastic discount factor $S_{t+1,t}$. For example, let $f(x_{t+1})$ be a claim to consumption at time $t + 1$. The time t price of this claim is $E[f(x_{t+1})S_{t+1,t}|x_t]$. Multi-period claims are valued using multiples of the stochastic discount factor over the payoff horizon.

As we develop in section 3 the stochastic discount factor is determined by a representative agent's intertemporal marginal rate of substitution. We feature two important specifications for the preferences of this agent: CRRA utility with a power utility function and the recursive utility model of Kreps and Porteus (1978), Epstein and Zin (1989b) and Weil (1990).

Since the representative agent's utility is defined over aggregate consumption, the dynamics of consumption are an important determinant of the stochastic discount factor. We assume that differences in the logarithm of aggregate consumption are a linear function of the state vector:⁴

$$c_{t+1} - c_t = \mu_c + U_c x_t + \gamma_0 w_{t+1}. \quad (2)$$

Under this assumption, in section 3 we show that the logarithm of the stochastic discount factor $s_{t+1,t} \equiv S_{t+1,t}$ are also linked to the state vector by:

$$s_{t+1,t} = \mu_s + U_s x_t + \xi_0 w_{t+1}. \quad (3)$$

In the case of CRRA utility $\xi_0 = -\theta\gamma_0$ where θ is the coefficient of relative risk aversion. As in the work of Hansen and Singleton (1983), shocks to aggregate consumption have a negative price so that assets with payoffs that are exposed to these shocks have higher average returns. When the intertemporal elasticity of substitution is equal to one and the preferences of the representative agent are of the recursive form the weighting on the current shock becomes:

$$\xi_0 = -\gamma_0 + (1 - \theta)\gamma(\beta)$$

where $\log(\beta)$ is the pure rate of time preferences,

$$\gamma(\beta) = \gamma_0 + \beta U_c (I - G\beta)^{-1} H,$$

and θ is a measure of risk aversion. The vector $\gamma(\beta)$ is the discounted impulse response of consumption to each of the respective components of the standardized shock vector w_{t+1} .

⁴See ?) for a generalization with stochastic volatility

As emphasized by (Bansal and Yaron 2004), the contribution of the discounted response to the stochastic discount factor makes consumption predictability a potentially potent way to enlarge risk prices, even over short horizons. Further the term $\theta\gamma(\beta)$ captures the “bad beta” of Campbell and Vuolteenaho (2004) except that they measure shocks using the market return instead of aggregate consumption.

The linear specification of the discount factor (3) assumes that the intertemporal elasticity of substitution is equal to one. We explore perturbations of this assumption and alternative assumptions about the risk aversion parameter θ .

2.2 The risk-return tradeoff

Our decomposition of long-run returns requires a specification of the long-run components of cash flows. We first consider a growth process modeled as the exponential of a random walk with drift:

$$D_t^* = \exp \left[\zeta t + \sum_{j=1}^t \pi w_j \right] . \quad (4)$$

Observed cash flows to securities have additional transient or stationary components. We let $\{D^t\}$ be an observed cash flow which is linked to the growth process via:

$$D_t = D_t^* f(x_t) . \quad (5)$$

To price D_t we value both the transient component $f(x_t)$ and the growth component D_t^* . The vector π measures the long-run risk exposure, and our aim is to assign risk prices to the long-run risk exposure vector π . An important result is that the effect of the growth component on the long-run risk-return tradeoff is invariant to the specification of f . The parameter ζ and the transient component $f(x_t)$ contribute to the implied asset values, but they do not effect the risk prices in the limit.

In our analysis we consider the cash flows from several portfolios. Figure 1 displays two of our cash flow series. Portfolio 1 is a growth portfolio and portfolio 5 is a value portfolio. The portfolios are re-balanced as in Fama and French (1992).⁵ In figure 1, both cash flows are depicted relative to aggregate consumption with the initial cash flows normalized to equal aggregate consumption. Notice that the cash flows of portfolio 1 grow much slower than those of portfolio 5. These differences in growth rates imply that the two portfolios are characterized by different values of ζ and/or π . Our goal is to understand how different assumptions about the long-run are reflected in expected returns.

To do this we consider fixing the growth process (4) and examine pricing for arbitrary choices of the function f . Since pricing is given by the conditional expectation of the stochastic discount factor times the asset payoff, fixing the growth process means that we sweep this process into the conditional expectations operator and create a one-period *value operator*:

$$\mathcal{P}f(x) = E [\exp (s_{t+1,t} + \zeta + \pi w_{t+1}) f(x_{t+1}) | x_t = x] .$$

⁵Details of the construction of the portfolios and cash flows can be found in Hansen, Heaton, and Li (2005) and at <http://www.bschool.nus.edu.sg/staff/biznl/bmdata.html>

Cash Flows Relative to Consumption for Two Portfolios

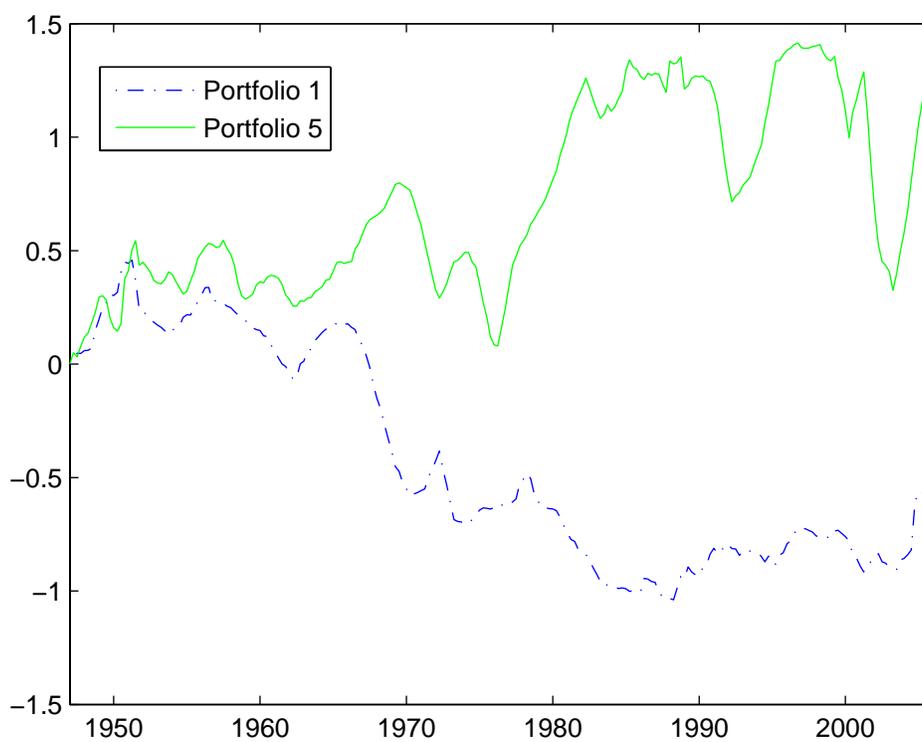


Figure 1: Natural logarithms of the ratios of portfolio cash flows to consumption. The plot given by — is for the cash flows from a portfolio of high book-to-market stocks (portfolio 5). The plot given by - · - · is for flows from a portfolio of low book-to-market stocks (portfolio 1). Quarterly.

This operator is much like the conditional expectations operator. Here our valuation operator allows us to fix the long-run components but consider alternative transient components given by different values of the function f . Notice that the date t price of the cash flow D_{t+1} is $D_t^* \mathcal{P}f(x_t)$.

Before moving on, notice that pricing is recursive so that prices of cash flows multiple periods in the future are inferred from this one-period pricing operator through iteration. For example, the time t value of date $t + j$ cash flow (5) is given by:

$$D_t^* [\mathcal{P}^j f(x_t)] = D_t^* E \left(\exp \left[\sum_{\tau=1}^j (s_{t+\tau, t+\tau-1} + \pi w_{t+\tau}) + j\zeta \right] f(x_{t+j}) | x_t = x \right)$$

where the notation \mathcal{P}^j denotes the application of the one-period valuation operator j times.

When the function $f(x)$ is assumed to be an exponential function of the Markov state, the functions $\{\mathcal{P}^j f(x), j = 1, 2 \dots\}$ are also exponential functions of the state. To see this let $f(x) = \exp(\omega x + \kappa)$ for some row vector ω and some number κ . Using the properties of the lognormal distribution:

$$\exp(\omega^* x + \kappa^*) = \mathcal{P}f(x) = \mathcal{P}[\exp(\omega x + \kappa)]$$

where

$$\omega^* = \omega G + U_s \tag{6}$$

and

$$\kappa^* = \kappa + \mu_s + \zeta + \frac{|\omega H + \xi_0 + \pi|^2}{2}. \tag{7}$$

Iteration of (6) and (7) j times yields the coefficients for the function $\mathcal{P}^j f(x)$.

Repeated iteration of (6) converges to a limit that is a fixed point of this equation:

$$\bar{\omega} = U_s(I - G)^{-1}.$$

The differences in the κ 's from (7) converge to:

$$-\nu \equiv \mu_s + \zeta + \frac{|\bar{\omega} H + \xi_0 + \pi|^2}{2}. \tag{8}$$

We include the minus sign in front of ν because the right-hand side will be negative in our applications. In our present-value calculations the contribution to value from cash flows in the distant future becomes arbitrarily small. The limit is the solution to:

Result 2.1. *A solution to the equation.*⁶

$$\mathcal{P}e = \exp(-\nu)e$$

for a strictly positive function e is given by $e(x) = \exp[U_s(I - G)^{-1}x]$ and $-\nu$ by (8).

⁶The equation in Result 2.1 is in form of an eigenvalue problem, and e is the unique (up to scale) solution that is strictly positive and satisfies a stability condition developed in the appendix.

While these iteration can be characterized simply for exponential functions of the Markov state, the same limits are obtained for a much richer class of functions. (See appendix A for a characterization of these functions.) Moreover, the limits do not depend on the starting values for ω and κ , but ν in particular depends on the exposure vector π to growth rate risk.⁷

We use this characterization of the limit to investigate long-run risk. As j gets larger, $\mathcal{P}^j(f)(x)$ approaches zero. The value of ν gives the asymptotic rate of decay of the values. The rate of decay reflects two competing forces, the asymptotic rate of growth of the cash flow and the asymptotic, risk adjusted rate of discount.

To isolate the rate of discount or long-run rate of return, we compute the limiting growth rate. Given that $\{G_t^*\}$ is a geometric random walk with drift, the long-run growth rate is

$$\eta = \zeta + \frac{1}{2}\pi \cdot \pi. \quad (9)$$

The variance adjustment, $\pi \cdot \pi$, reflects the well known Jensen's inequality adjustment. The transient components of cash flows do not alter the long-run growth rate.⁸ The asymptotic rate of return is obtained by subtracting the growth rate η from ν . The following theorem summarizes these results and gives a well defined notion of the price of long-run cash flow risk.

Theorem 1. *Suppose that the state of the economy evolves according to (1) and the stochastic discount factor is given by (3), then the asymptotic rate of return is:*

$$\eta + \nu = \zeta^* + \pi^* \cdot \pi$$

where

$$\begin{aligned} \pi^* &\doteq -\xi_0 - U_s(I - G)^{-1}H \\ \zeta^* &\doteq -\mu_s - \frac{\pi^* \cdot \pi^*}{2}. \end{aligned}$$

The term π^* is the price of exposure to long-run risk as measured by π . By setting $\pi = 0$ we consider cash flows that do not grow over time and are stationary. An example is a discount bond, whose asymptotic pricing is studied by Alvarez and Jermann (2005). The asymptotic rate of return for such a cash flow with no long run risk exposure is: ζ^* . Thus $\pi \cdot \pi^*$ is the contribution to the rate of return coming from the exposure of cash flows to long run risk. Since π measures this exposure, π^* is the corresponding price vector.

⁷In fact we could to represent these transient components with a larger state vector provided that this state vector does not Granger cause $\{x_t\}$ in the sense of nonlinear prediction. This allows to include “share models” with nonlinear share evolution equations as in Santos and Veronesi (2001).

⁸Formally, a unit function is the eigenfunction of the *growth operator*:

$$\mathcal{G}f(x) = E[\exp(\zeta + \pi w_{t+1}) f(x_{t+1}) | x_t = x]$$

with an eigenvalue given by $\exp(\eta)$.

Theorem 1 gives the long-horizon counterpart to a risk-return tradeoff. The price of growth rate risk exposure parameterized by π is π^* . In the case of the power utility model:

$$\pi^* = \theta\gamma(1)$$

where $\gamma(1)$ is the long-run (undiscounted) response vector for consumption to the underlying shocks. In the recursive utility model with a unitary elasticity of substitution, this price is:

$$\pi^* = \gamma_0 + (\theta - 1)\gamma(\beta)$$

which is approximately the same for β close to unity. The period counterparts will differ provided that consumption is predictable (see Kocherlakota (1990) and Bansal and Yaron (2004)).

The top panel of figure 2 displays estimates of risk-adjusted returns for each portfolio based on statistical models of cash flows and a pricing model described below. Expected returns are given as a function of the horizon of future cash flows. The expected rates of return start at a similar level for both portfolios and then significantly separate as the horizon increases. In particular, the expected return to the value portfolio increases with horizon in contrast to the growth portfolio. This effect is due to important exposure to long-run macroeconomic risk in the value portfolio. Short-run risk exposures are not significantly different across the portfolios, however. This is reflected in similar expected short-horizon returns for each portfolio.

2.3 Risk premia over alternative horizons

While we have characterized the limiting expected rate of return, it is of interest more generally to see how returns depends on the horizon of the payoffs. Consider the expected return to holding a claim a single cash flow D_{t+j} . This return is given by the ratio of expected cash flow to current price. We scale this by the horizon and take logarithms to yield:

$$\frac{1}{j} [\log \mathcal{G}^j f(x_t) - \log \mathcal{P}^j f(x_t)] .$$

This expected return depends on the transitory cash flows $f(x_{t+j})$. When a corresponding risk-free return is subtracted from this return, this formula provides a measure of the risk-premia by horizon. The risk premia reflect both risk exposure and risk prices associated with the different horizons.

The top panel of 2 displays estimates of risk-adjusted returns for the cash flows produced by the growth (portfolio 1) and value portfolio (portfolio 5). These calculation assume recursive preferences.⁹ As a basis of comparison when considering this figure, note that the observed average returns to these portfolios are substantially different. For example, as

⁹These calculations assume a risk aversion parameter of 20. The large number is used to amplify the effects, and we will have more to say about this parameter subsequently.

reported in table 1 below, the expected one-period returns to portfolios 1 and 5 are 6.79% and 11.92% respectively.

The pattern of risk premia across horizons are intriguing. The expected return to the value portfolio increases with horizon in contrast to the growth portfolio. This effect is due to important exposure to long-run macroeconomic risk in the value portfolio. Exposure to short-run risk are not significantly different across the two portfolios, however. This is manifested in the very similar expected short-horizon returns for each portfolio. As the horizon increases the expected returns settle approximate their long-run limits as given in 1 and the limiting differences are reflected in the long-run prices of exposure to growth-rate risk. In section 5.5 we investigate how sensitive these measurements are to errors in the statistical specification and to estimation accuracy.

2.4 One-period returns

The one period return to a portfolio of cash flows is a weighted average of one-period returns to holding each cash flow. The holding period return to a security that pays off $\psi(x_{t+j})$ in period j is given by:

$$R_{t+1,t}^j = \exp(\zeta + \pi w_{t+1}) \frac{\mathcal{P}^{j-1} f(x_{t+1})}{\mathcal{P}^j f(x_t)}.$$

The logarithm of the expected gross returns for alternative j are reported in the bottom panel figure 2 for the growth and value portfolios. As j gets large these returns are approximately equal to:

$$R_{t+1,t}^d = \exp(\nu) \exp(\zeta + \pi w_{t+1}) \frac{e(x_{t+1})}{e(x_t)} \quad (10)$$

which is the holding period return to a security that pays off the dominant eigenfunction or pricing factor e over *any* horizon j . Thus for a given π the holding period returns become approximately the same as the horizon increases. The weighting of these returns is dictated by the relative magnitudes of $\mathcal{P}^j f$, which will eventually decay asymptotically at a rate ν . Thus, ν gives us a measure of duration, the importance of holding period returns far into the future relative to holding period returns today. When ν is closer to zero, the holding-period returns to cash flows far into the future are more important contributors to the portfolio decomposition of one period returns.¹⁰

The logarithm of the return $R_{t+1,t}^d$ has two components: a cash flow component: $\zeta + \pi w_{t+1}$ determined by the reference growth process and a valuation component $\nu + \log e(x_{t+1}) - \log e(x_t)$ determined by the dominant eigenvalue and eigenfunction. While ν and the cash flow component change as we alter the cash flow risk exposure vector π , $\log e(x_{t+1}) - \log e(x_t)$ remains the same.

¹⁰Lettau and Wachter (2005) also consider the decomposition of returns into the holding period returns of the component cash flows. Their focus is different because they feature a single aggregate return with portfolio dynamics meant to capture differences in average returns across portfolios instead of differences in observed cash flow dynamics.

Expected Returns Rates by Horizon

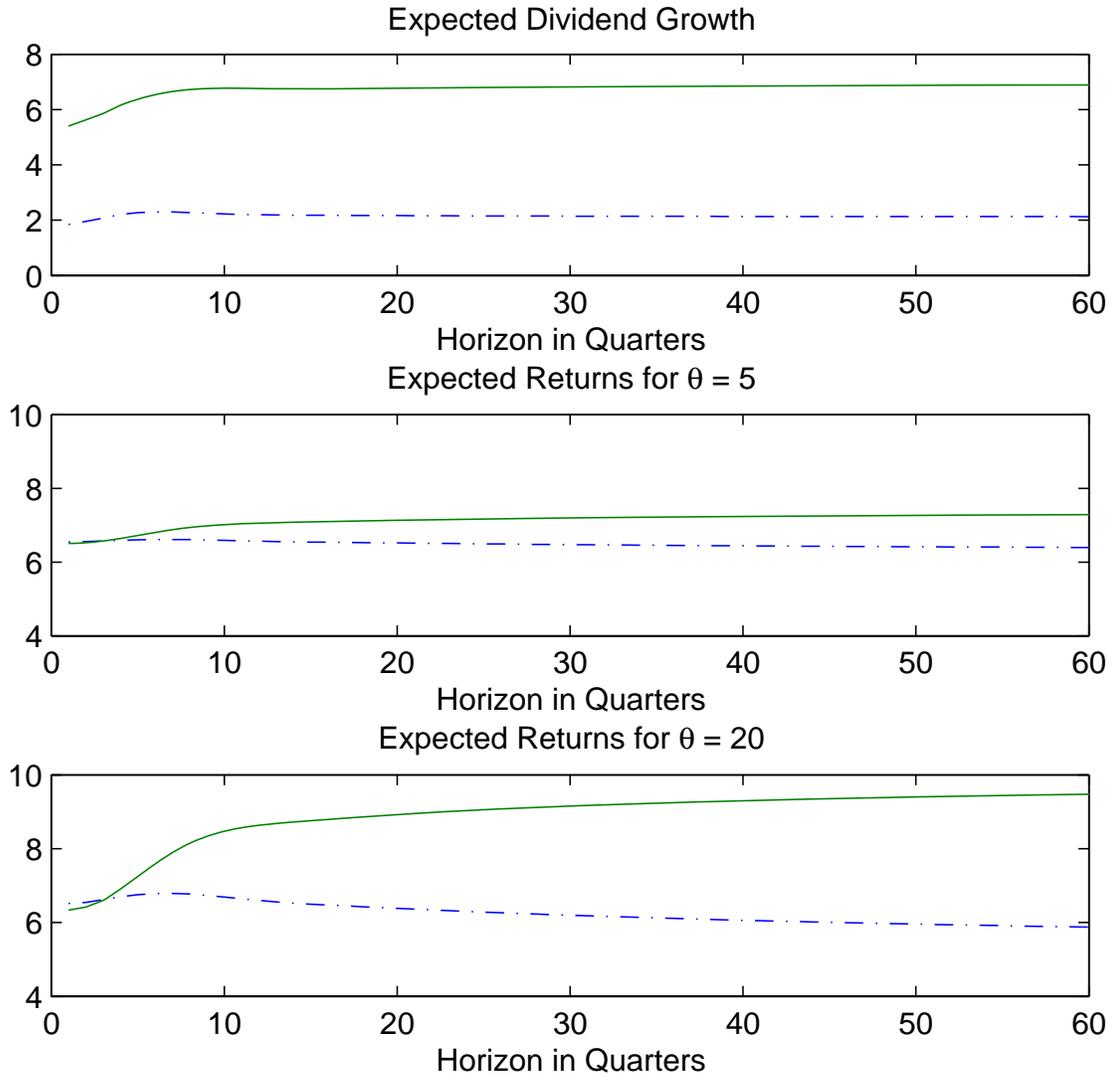


Figure 2: The plots given by — are for portfolio 5, a value portfolio and the plots given by -. are for portfolio 1, a growth portfolio. Rates of return are given in annual percentage rates.

The bottom panel of figure 2 depicts this decomposition of expected returns for the growth and value portfolios described as a function of horizon. The expected rate of return is much larger for the value portfolio once we look at the returns to holding portfolio cash flows beyond two years into the future. The limiting value of these figures are also good approximations to the entire figure after about five to size years.

2.5 Transient components to stochastic discount factors

Bansal and Lehmann (1997) have shown that a variety of asset pricing models imply common bounds on the expected growth rate in logarithms of the stochastic discount factors. These include asset pricing models with forms of habit persistence and social externalities. Their analysis extends to some recent models of social externalities or preference shocks such as Menzly, Santos, and Veronesi (2004).

While Bansal and Lehmann (1997) focus on specific stochastic discount factor bounds, the long-term risk return tradeoff of theorem 1 is invariant to many of these same changes. These various models differ only in their transient implications for valuation. Formally, the eigenvalues of result 2.1 remain the same but the eigenfunctions are altered. Thus the limiting one-period return (10) will be altered by the inclusion of a common state dependent contribution, but the long-term tradeoff remains the same. See Hansen (2006) for a more extensive discussion of this.

3 Pricing under recursive utility

In what follows we feature a recursive utility model of investor preferences. This model provides an important role for long-run consumption risk. The resulting specification of the stochastic discount factor gives us a tractable characterizing of long-run implications that is rich enough to imply differences in expected returns as they relate to long-run risk.

3.1 Preferences and the stochastic discount Factor

We follow Kreps and Porteus (1978), Epstein and Zin (1989b) and Weil (1990) in choosing to examine recursive preferences. As we will see below, this specification of preferences provides a simple justification for examining the temporal composition of risk in consumption.

In our specification of these preferences, we use a CES recursion:

$$V_t = [(1 - \beta) (C_t)^{1-\rho} + \beta \mathcal{R}_t(V_{t+1})^{1-\rho}]^{\frac{1}{1-\rho}}. \quad (11)$$

The random variable V_{t+1} is the continuation value of a consumption plan from time $t + 1$ forward. The recursion incorporates the current period consumption C_t and makes a risk adjustment $\mathcal{R}_t(V_{t+1})$ to the date $t + 1$ continuation value. We use a CES specification for this risk adjustment as well:

$$\mathcal{R}_t(V_{t+1}) \doteq \left[E(V_{t+1})^{1-\theta} | \mathcal{F}_t \right]^{\frac{1}{1-\theta}}$$

where \mathcal{F}_t is the current period information set. The outcome of the recursion is to assign a continuation value V_t at date t .

The preferences provide a convenient separation between risk aversion and the elasticity of intertemporal substitution [see Epstein and Zin (1989b)]. For our purposes, this separation allows us to examine the effects of changing risk exposure with modest consequences for the risk-free rate. When there is perfect certainty, the value of $1/\rho$ determines the *elasticity of intertemporal substitution* (EIS). A measure of risk aversion depends on the details of the gamble being considered. As emphasized by Kreps and Porteus (1978), with preferences like these intertemporal compound consumption lotteries cannot necessarily be *reduced* by simply integrating out future information about the consumption process. Instead the timing of information has a direct impact on preferences and hence the intertemporal composition of risk matters. As we will see, this is reflected explicitly in the equilibrium asset prices we characterize. On the other hand, the aversion to simple wealth gambles is given by θ . Since we will explore “large values” of this parameter we also consider other interpretations of it related to investor concerns about model misspecification.

In a frictionless market model, one-period stochastic discount factors are given by the intertemporal marginal rates of substitution between consumption at date t and consumption at date $t + 1$. For simplicity, we assume an endowment economy but more generally this consumption process is the outcome of an equilibrium with production.¹¹ Aggregate consumption evolves according to:

$$c_{t+1} - c_t = \mu_c + U_c x_t + \gamma_0 w_{t+1} \quad (12)$$

where $c_{t+1} - c_t$ is the first difference in the logarithm of aggregate consumption. Preferences are common across consumers and in equilibrium they equate their intertemporal marginal rates of substitution. Since we are using a recursive specification with two CES components, it is straightforward to show that the implied stochastic discount factor is:

$$S_{t+1,t} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\rho} \left(\frac{V_{t+1}}{\mathcal{R}_t(V_{t+1})} \right)^{\rho-\theta} . \quad (13)$$

(For instance, see ?.) There are two contributions to the stochastic discount factor. One is the direct consumption growth contribution familiar from the Rubinstein (1976), Lucas (1978) and Breeden (1979) model of asset pricing. The other is the continuation value relative to its risk adjustment. This second contribution is forward-looking and is present only when ρ and θ differ.

A challenge in using this model empirically is to measure the continuation value, V_{t+1} , which is linked to future consumption via the recursion (11). One possible approach to the measurement problem is to use the link between the continuation value and wealth defined as the value of the aggregate consumption stream in equilibrium. A direct application of

¹¹The inclusion of production implies additional restrictions and can suggest alternative specifications for the consumption dynamics.

Euler's theorem for constant returns to scale functions implies that

$$\frac{W_t}{C_t} = \frac{1}{1-\beta} \left(\frac{V_t}{C_t} \right)^{1-\rho}.$$

where W_t is wealth at time t . When $\rho \neq 1$ this link between wealth, consumption and the continuation value implies a representation of the stochastic discount factor based on consumption growth and the return to a claim on future wealth. In general this return is unobservable. The return to a stock market index is sometimes used to proxy for this return as in Epstein and Zin (1989a); or other components can be included such as human capital with assigned market or shadow values (see Campbell (1994)).

In this investigation, like that of Restoy and Weil (1998) and Bansal and Yaron (2004), we base the analysis on a well specified stochastic process governing consumption and avoid the need to construct a proxy to the return on wealth. This is especially important in our context because we are interested in risk determined by the long-run effects of shocks the aggregate quantities. These shocks may not be reflected in the variation of a proxy for the return to aggregate wealth such as a stock index. In contrast to Restoy and Weil (1998) and Bansal and Yaron (2004), we begin with the case of $\rho = 1$ since it is understood that logarithmic intertemporal preferences lead to substantial simplification of the calculation of equilibrium prices and returns [*e.g.* see Schroder and Skiadas (1999)]. When $\rho = 1$ the wealth to consumption ratio is a constant and the construction of the stochastic discount factor using the return to the wealth portfolio breaks down.¹² We then explore sensitivity of pricing implications as we change the elasticity of intertemporal substitution.¹³ For example, Campbell (1996) argues for less less elasticity than the log case and Bansal and Yaron (2004) argue for more.

To develop our approach to calculating the continuation value scale V_t in (11) by consumption.¹⁴

$$\frac{V_t}{C_t} = \left[(1-\beta) + \beta \mathcal{R}_t \left(\frac{V_{t+1} C_{t+1}}{C_{t+1} C_t} \right)^{1-\rho} \right]^{\frac{1}{1-\rho}}.$$

Let v_t denote the logarithm of the continuation value relative to the logarithm of consumption, and let c_t denote the logarithm of consumption and rewrite recursion (11) as:

$$v_t = \frac{1}{1-\rho} \log \left((1-\beta) + \beta \exp \left[(1-\rho) \mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t) \right] \right), \quad (14)$$

¹²Log-linear methods typically approximate around a constant consumption-wealth ratio. Setting $\rho = 1$, justifies this. The approximation method we explore is very similar to log-linear approximation, but we employ this one because of its explicit link to a local theory of approximation and because it allows us to impose stochastic dynamics with an extensive amount of persistence in the limit economy.

¹³Log-linear methods typically approximate around a constant consumption-wealth ratio. Setting $\rho = 1$, justifies this. The approximation method we explore is very similar to log-linear approximation. We employ it in part because of its explicit link to a local theory of approximation and because it allows us to impose stochastic dynamics with an extensive amount of persistence in the limit economy.

¹⁴Recursion (15) is used by Tallarini (1998) in his study of *risk sensitive* business cycles and asset prices.

where \mathcal{Q}_t is:

$$\mathcal{Q}_t(v_{t+1}) = \frac{1}{1-\theta} \log E(\exp[(1-\theta)v_{t+1}] | \mathcal{F}_t).$$

3.2 The special case in which $\rho = 1$

The $\rho = 1$ limit in recursion (14) is:

$$\begin{aligned} v_t &= \beta \mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t) \\ &= \frac{\beta}{1-\theta} \log E(\exp[(1-\theta)(v_{t+1} + c_{t+1} - c_t)] | \mathcal{F}_t). \end{aligned} \quad (15)$$

Recursion (15) is used by Tallarini (1998) in his study of *risk sensitive* business cycles and asset prices. For the log-linear stochastic specification, the solution for the continuation value is

$$v_t = \mu_v + U_v x_t$$

where:

$$\begin{aligned} U_v &\doteq \beta U_c (I - \beta G)^{-1}, \\ \mu_v &\doteq \frac{\beta}{1-\beta} \left[\mu_c + \frac{(1-\theta)}{2} |\gamma_0 + U_v H|^2 \right]. \end{aligned}$$

In this formula $U_v x_t$ is the discounted sum of expected future growth rates of consumption constructed using the subjective discount factor β . The term $\gamma_0 + U_v H$ is the shock exposure vector of the continuation value for consumption.

The stochastic discount factor when $\rho = 1$ is:

$$S_{t+1,t} \equiv \beta \left(\frac{C_t}{C_{t+1}} \right) \left[\frac{(V_{t+1})^{1-\theta}}{\mathcal{R}_t (V_{t+1})^{1-\theta}} \right].$$

Notice that the term of $S_{t+1,t}$ associated with the risk-aversion parameter θ satisfies:

$$E \left[\frac{(V_{t+1})^{1-\theta}}{\mathcal{R}_t (V_{t+1})^{1-\theta}} | \mathcal{F}_t \right] = E \left(\frac{(V_{t+1})^{1-\theta}}{E[(V_{t+1})^{1-\theta} | \mathcal{F}_t]} \right) = 1. \quad (16)$$

The resulting formula for the stochastic discount factor is:

$$s_{t+1,t} = \mu_s + U_s x_t + \xi_0 w_{t+1}$$

where:

$$\begin{aligned} \mu_s &= \log \beta - \mu_c - \frac{(1-\theta)^2 |\gamma_0 + U_v H|^2}{2} \\ U_s &= -U_c \\ \xi_0 &= -\gamma_0 + (1-\theta)(\gamma_0 + U_v H). \end{aligned}$$

The coefficient vector ξ_0 on the shock w_{t+1} has the following interpretation. From the consumption dynamics (12), the initial response of consumption at date $t + 1$ to a shock w_{t+1} is $\gamma_0 w_{t+1}$ and the response of c_{t+j} for $j > 1$ is $\gamma_j \doteq U_c G^{j-1} H$. The discounted (by the subjective rate of discount) value of these response is:

$$\gamma(\beta) \doteq \gamma_0 + \beta U_c (I - \beta G)^{-1}.$$

Thus $\xi_0 = -\gamma_0 + (1 - \theta)\gamma(\beta)$ as claimed in section 2. The term $\gamma(\beta)$ is a target of measurement even for one-period pricing. This the impact of predictability in consumption growth that is featured in Bansal and Yaron (2004). It reflects the intertemporal composition of consumption risk, and it gives an important measurement challenge for implementation. Long-term risk can have important implications for even one-period pricing. The impact persists as we infer risk prices over longer-horizons as is conveyed by the limiting pricing formulas in section 2.

Since the term (16) in the one-period period stochastic discount factor is positive, and it has conditional expectation equal to unity and can thus be thought of as distorting the probability distribution. The presence of this distortion is indicative of a rather different interpretation of the parameter θ . Anderson, Hansen, and Sargent (2003) argue that this parameter may reflect investors' concerns about not knowing the precise riskiness that they confront in the marketplace instead of incremental risk aversion applied to continuation utilities. Under this view, the original probability model is viewed as a statistical approximation, but investors are concerned that this model may be misspecified. Although we continue to refer to θ as a risk-aversion parameter, this alternative interpretation is germane to our analysis because we will explore sensitivity of our measurements to the choice of θ . Changing the interpretation of θ alters what we might view as reasonable values of this parameter.

To be concrete, under the alternative interpretation suggested by Anderson, Hansen, and Sargent (2003), $(\theta - 1)\gamma(\beta)$ is the contribution to prices induced because investors cannot identify potential model misspecification that is disguised by shocks that impinge on investment opportunities. An investor with this concern explores alternative shock distributions including ones with a distorted mean. He uses a penalized version of a max-min utility function. In considering how big the concern is about model misspecification, we ask if it could be ruled easily with historical data. This lead us to ask how large is $-\theta(1 - \beta)\gamma(\beta)$ in a statistical sense. To gauge this, when $\theta = 10$ and $|(1 - \beta)\gamma(\beta)| = .01$ a hypothetical decision maker asked to tell the two models apart would have about 24% chance of getting the correct answer given 250 observations. Doubling θ changes this probability to about 6%.¹⁵ In this sense $\theta = 10$ is in an interesting range of statistical ambiguity while $\theta = 20$ leads to an alternative model that considerably easier to discriminate based on historical data. See Hansen (2007) for a more extensive discussion of such calculations. As we explore large values of θ in our empirical work, perhaps part of large choice of θ can be ascribed to statistical ambiguity on the part of investors.

¹⁵These numbers are essentially the same if the prior probability across models is the same or if the min-max solution of equating the type I and type II errors is adopted.

3.3 Intertemporal substitution ($\rho \neq 1$)

Approximate characterization of equilibrium pricing for recursive utility have been produced by Campbell (1994) and Restoy and Weil (1998) based on a log-linear approximation of budget constraints. ?) use a distinct but related approach and follow Kogan and Uppal (2001) by approximating around an explicit equilibrium computed when $\rho = 1$ and then varying the parameter ρ . The stochastic discount factor is expressed as an expansion around the case of $\rho = 1$:

$$s_{t+1,t} \approx s_{t+1,t}^1 + (\rho - 1)Ds_{t+1,t}^1.$$

where:

$$Ds_{t+1,t}^1 = \frac{1}{2}w_{t+1}'\Theta_0w_{t+1} + w_{t+1}'\Theta_1x_t + \vartheta_0 + \vartheta_1x_t + \vartheta_2w_{t+1}.$$

Formulas for Θ_0 , Θ_1 , ϑ_0 , ϑ_1 and ϑ_2 are also given in Hansen, Heaton, Lee, and Rousanov (2006). When $\rho \neq 1$ the discount factor includes quadratic terms in the shock vector w_{t+1} . The approximation to the discount factor allows us to calculate the derivative of the asymptotic rate of return for any cash flow process. The details of the justification and implementation of these formulas are given in appendix A.

4 Long-Run Consumption Risk

Our first measurement task is to estimate the consumption dynamics needed to characterize how risk exposure is priced. As in much of the empirical literature in macroeconomics, we use vector autoregressive (VAR) models to identify interesting aggregate shocks and estimate the macroeconomic dynamics. In our initial model we let consumption be the first element of y_t and corporate earnings be the second element:

$$y_t = \begin{bmatrix} c_t \\ e_t \end{bmatrix}.$$

Our use of corporate earnings in the VAR is important for two reasons. First, it is used as a predictor of consumption and an additional source of aggregate risk.¹⁶ For example, changes in corporate earnings potentially signal changes in aggregate productivity which will have long-run consequences for consumption. Second, corporate earnings provide a broad-based measure of the ultimate source of the cash flows to capital. The riskiness of the equity claims on these cash flows provides a basis of comparison for the riskiness of the cash flows generated by the portfolios of stocks that we consider in section 5.

¹⁶Whereas Bansal and Yaron (2004) also consider multivariate specification of consumption risk, they seek to infer this risk from a single aggregate time series on consumption or aggregate dividends. With flexible dynamics, such a model is not well identified from time series evidence. On the other hand, while our shock identification allows for flexible dynamics, it requires that we specify *a priori* the important sources of macroeconomic risk.

The process $\{y_t\}$ is presumed to evolve as a VAR of order ℓ . In the results reported subsequently, $\ell = 5$. The least restrictive specification we consider is:

$$A_0 y_t + A_1 y_{t-1} + A_2 y_{t-2} + \dots + A_\ell y_{t-\ell} + B_0 = w_t, \quad (17)$$

The vector B_0 is two-dimensional, and the square matrices $A_j, j = 1, 2, \dots, \ell$ are two by two. The shock vector w_t has mean zero and covariance matrix I .

Form:

$$A(z) \doteq A_0 + A_1 z + A_2 z^2 + \dots + A_\ell z^\ell.$$

We are interested in a specification in which $A(z)$ is nonsingular for $|z| < 1$ implying that the process $\{y_t\}$ is stationary. By inverting the matrix polynomial $A(z)$ for the autoregressive representation, we obtain the power series expansion for the moving-average coefficients. The discounted consumption response is $u_c A(\beta)^{-1}$ where u_c selects the first row, the row consisting the consumption responses. Multiplying by $(1 - \beta)$ gives the geometric average response:

$$\gamma(\beta) = (1 - \beta) u_c A(\beta)^{-1}$$

as required by our model.

Following Hansen, Heaton, and Li (2005), we restrict the matrix $A(1)$ to have rank one:

$$A(1) = \alpha [1 \quad -1].$$

where the column vector α is freely estimated. This parameterization imposes two restrictions on the $A(1)$ matrix. It imposes a unit root in consumption and earnings, but restricts these series to grow together. In this system, both series respond in the same way to shocks in the long run, so they are cointegrated. Since the cointegration relation we consider is prespecified, the model can be estimated as a vector autoregression in the first-difference of the log consumption and the difference between the log earnings and log consumption. In section 5.5 we explore other growth restrictions specifications.

In our analysis, we will not be concerned with the usual shock identification familiar from the literature on structural VAR's. This literature assigns structural labels to the underlying shocks and imposes *a priori* restrictions to make this assignment. Our primary interest is the intertemporal composition of consumption risk and not the precise labels attached to individual shocks. We construct two uncorrelated shocks as follows. One is temporary, formed as a linear combination of shocks that has no long run impact on consumption and corporate earnings. The second is permanent which effects consumption and earnings equally in the long run.¹⁷ Thus we impose $\gamma(1) = [0 \quad 1]$ as a convenient identifying restriction for the shocks.

For our measure of aggregate consumption we use aggregate consumption of nondurables and services taken from the National Income and Product Accounts. This measure is quarterly from 1947 Q1 to 2005 Q4, is in real terms and is seasonally adjusted. We measure

¹⁷This construction is much in the same spirit as Blanchard and Quah (1989).

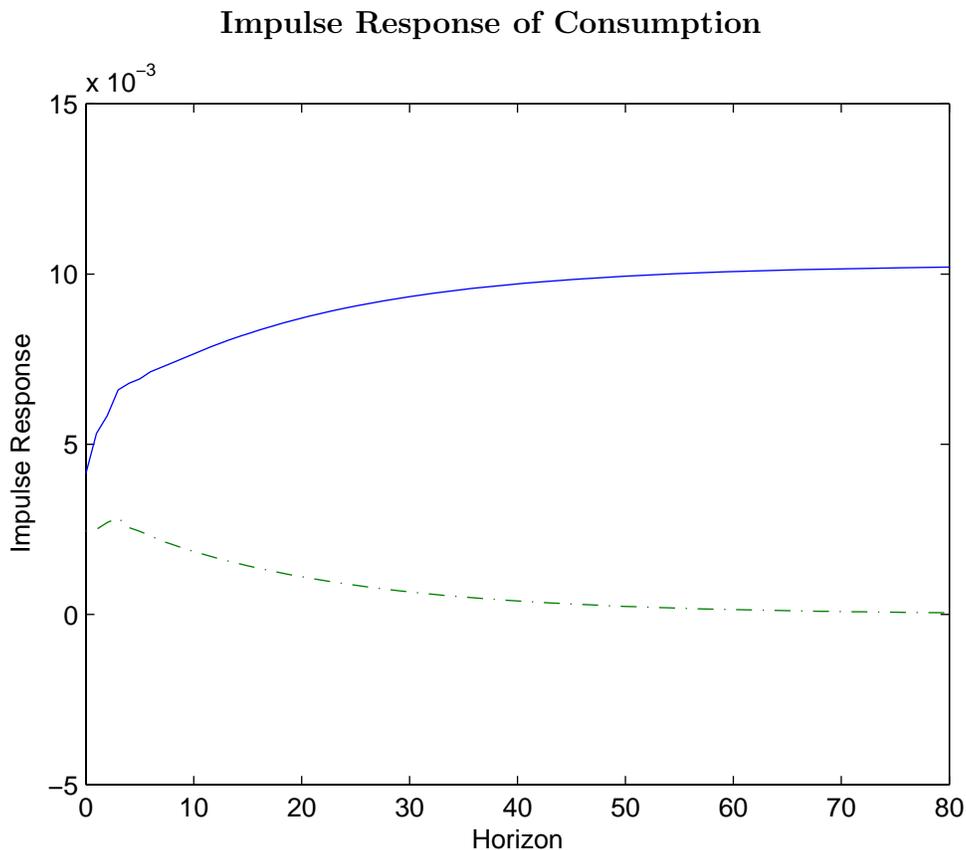


Figure 3: The two curves are impulse responses of consumption to shocks implied by bivariate VAR's where consumption and earnings are assumed to be cointegrated. — depicts the impulse response to a permanent shock. - . depicts the impulse response to a temporary shock. Each shock is given a unit impulse. Responses are given at quarterly horizons

corporate earnings from NIPA and convert this series to real terms using the implicit price deflator for nondurables and services. Using these series, we estimate the system with cointegration.

In figure 3 we report the response of consumption to permanent and temporary shocks. The immediate response of consumption to a permanent shock is approximately twice that of the response to a temporary shock. Permanent shocks are an important feature of aggregate consumption. The full impact of the permanent shock is slowly reflected in consumption and ultimately accumulates to a level that is more than twice the on-impact response.

4.1 Estimation accuracy

With recursive utility, the geometrically weighted average response of consumption to the underlying shocks affects both short-run and long-run risk prices. For this reason, the pre-

dictable responses of consumption to shocks identified by the VAR with cointegration, affect risk prices at all horizons. Using the cointegration specification, we explore the statistical accuracy of the estimated responses.

Following suggestions of Sims and Zha (1999) and Zha (1999), we impose conjugate priors from regression analysis on the coefficients of each equation in the VAR and simulate histograms for the parameter estimates. This provides an approximation for Bayesian posteriors with a relatively diffuse (and improper) prior distribution. These “priors” are chosen for convenience, but they give us a simple way to depict the sampling uncertainty associated with the estimates. We use these priors in computing posterior distributions for the immediate response of consumption to a temporary shock and for the long-run response of the permanent shock to consumption. This long-run response is $|\gamma(1)|$ and the short-run response is γ_0 . Figure 4 gives the posterior histogram for both responses.

As might be expected, the short-run response estimate is much more accurate than the long-run response. Notice that the horizontal scales of the histograms differ by a factor of ten. In particular, while the long-run response is centered at a higher value, it also has a substantial right tail. Consistent with the estimated impulse response functions, the median long-run response is about double that of the short-term response. In addition nontrivial probabilities are given to substantially larger responses.¹⁸ Thus, from the standpoint of sampling accuracy, the long-run response could be even more than double that of the immediate consumption response. Our interest is in how these measurement challenges carry over to risk pricing.

4.2 Specification sensitivity

Cointegration plays an important role in identifying both the long-run impact of the permanent shock depicted in figure 3 and the temporal pattern of the responses to both shocks. In figure 5 we depict $|\gamma(\beta)|$ as a function of β for the baseline model and for two alternative specifications, a VAR estimated in log-levels and a VAR estimated in first differences. The log-level VAR is estimated to be stable, and as a consequence the implied $|\gamma(1)| = 0$. This convergence is reflected in the figure, but only for values of β very close to unity. For more moderate levels of β , the log level specification reduces the measure of $|\gamma(\beta)|$ by a third. The first difference specification gives results that are intermediate relative to the baseline specification and the log-level specification. In summary, our restriction that consumption and earnings respond to permanent shocks in the same way ensures that a larger value of $|\gamma(\beta)|$.

¹⁸The accuracy comparison could be anticipated in part from the literature on estimating linear time series models using a finite autoregressive approximation to an infinite order model (see Berk (1974)). The *on impact* response is estimated at the parametric rate, but the *long-run* response is estimated at a considerably slower rate that depends on how the approximating lag length increases with sample size. Our histograms do not confront the specification uncertainty associated with approximating an infinite order autoregression, however.

Approximate posterior distributions for immediate and long-run responses

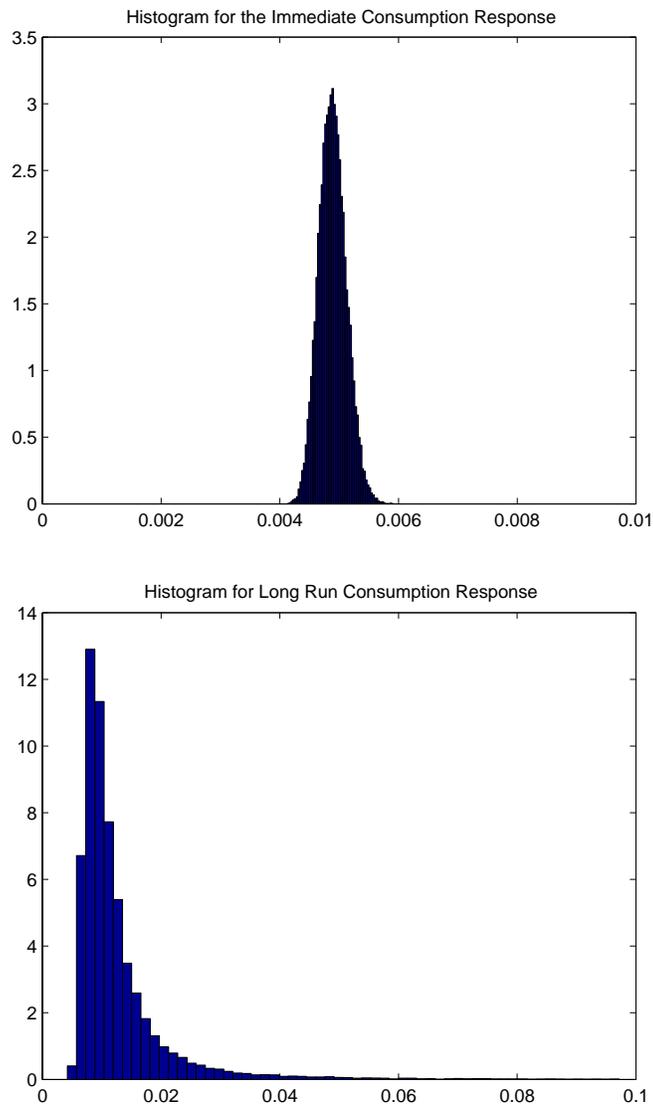


Figure 4: Top figure gives the posterior histogram for the magnitude $|\gamma_0|$ of immediate response of consumption to shocks. Bottom figure gives the long-run response of consumption to the permanent shock. The histograms have sixty bins with an average bin height of unity. They were constructed using conjugate regression priors for each equation. The vertical axis is constructed so that on average the histogram height is unity.

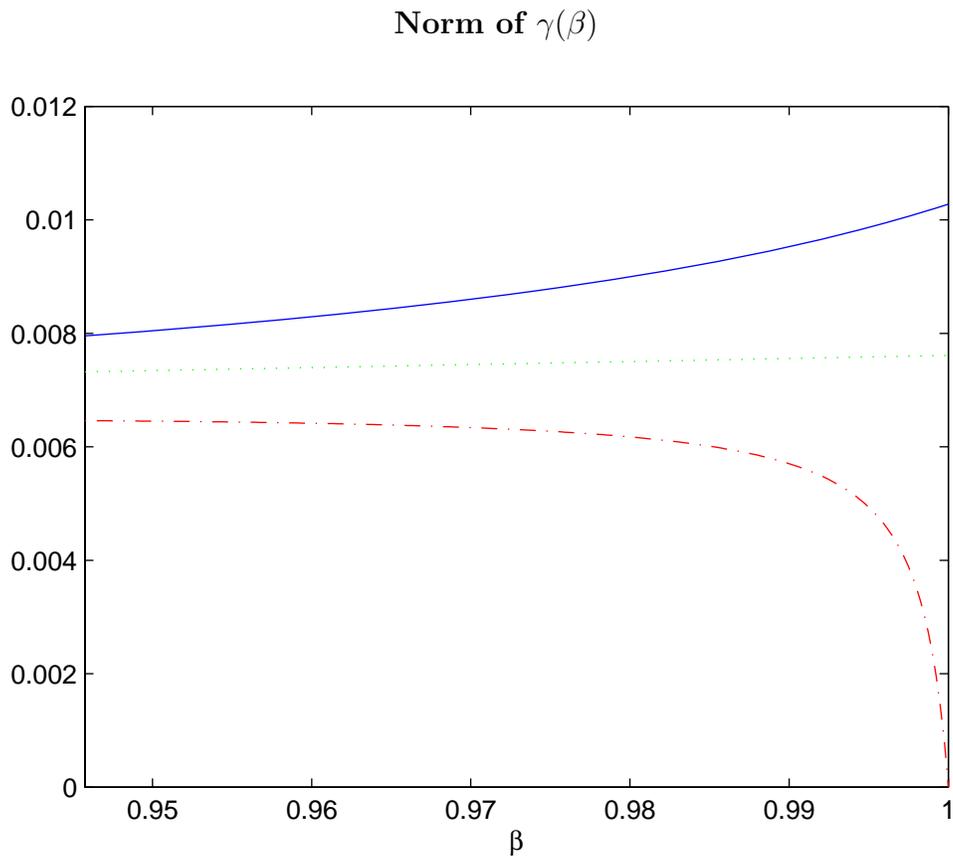


Figure 5: Norm of $\gamma(\beta)$ for different values of β and three different VAR systems. The solid line — is for the cointegrated system, the dashed-dot line - · -· is for the system without cointegration, the dotted line ··· is for the specification with first differences used for all variables.

4.3 Pricing implications for aggregate consumption

To examine the long-run risk components of aggregate consumption, consider pricing a claim to aggregate consumption. Such pricing calculations are commonly performed, but are not directly relevant in equity pricing. In this case π is equal to the long-run exposure of consumption to the two shocks: $\gamma(1)$. With recursive preferences and $\rho = 1$, the asymptotic rate of return of Theorem 1 reduces to:

$$-\log \beta + \mu_c + \frac{\gamma(1) \cdot \gamma(1)}{2}$$

which does not depend on the risk-aversion parameter θ . The excess of the asymptotic return to the consumption claim over the riskless return is:

$$\gamma(1) \cdot [\gamma(1) + (\theta - 1)\gamma(\beta)] .$$

The expected excess return is essentially proportional to θ due to the dependence of the risk-free benchmark on θ .¹⁹

Even in the long-run, the consumption claim is not very risky. The VAR system implies that $\gamma(1) \cdot \gamma(1) = 0.0001$. Hence when β is near unity, increases in θ has a very small impact on the excess return to the consumption claim. For example, even when $\theta = 20$ the expected excess return, in annual units, is .8% ($= 20 \times 0.0001 \times 4$). A similar conclusion holds when we use aggregate stock market dividends instead of consumption as a cash flow measure, but we obtain a different conclusion when we consider cash flows from portfolios.

5 Long-Run Cash Flow Risk

Theorem 1 characterizes the limiting factor risk prices π^* for cash flow exposure to long-run risk. This relation in conjunction with our economic model 3 allow us to compute limiting valuation and risk adjustments. Our task in this section is to measure the long-run risk exposure of the cash flows from some portfolios familiar from financial economics and to consider the implied differences in values and the risk premia of returns.

Previously, Bansal, Dittmar, and Lundblad (2005) and Campbell and Vuolteenaho (2004) have related measures of long-run cash flow risk to one period returns using a log-linearization of the present value relation. Our aim is different, but complementary to their study. As we described in section 2, we study how long run cash flow risk exposure is priced in the context of economic models of valuation.

We consider cash flows that may not grow proportionately with consumption. This flexibility is consistent with the models of Campbell and Cochrane (1999), Bansal, Dittmar, and Lundblad (2005), Lettau, Ludvigson, and Wachter (2004), and others. It is germane to our empirical application because the sorting method we use in constructing portfolios can

¹⁹Notice that when $\beta = 1$, the expected excess return reduces to $\theta\gamma(1) \cdot \gamma(1)$ and the proportionality is exact.

induce permanent differences in dividend growth.²⁰ While physical claims to resources may satisfy balanced growth restrictions, financial claims of the type we investigate need not as reflected in the long run divergence displayed in figure 1.

Consistent with our use of VAR methods, we consider a log-linear model of cash flow growth:

$$d_{t+1} - d_t = \mu_d + U_d x_t + \iota_0 w_{t+1}.$$

where d_t is the logarithm of the cash flow. This growth rate process has a moving-average form:

$$d_{t+1} - d_t = \mu_d + \iota(L)w_{t+1}.$$

where:

$$\iota(z) = \sum_{j=0}^{\infty} \iota_j z^j$$

and:

$$\iota_j = \begin{cases} \iota_0 & \text{if } j = 0 \\ U_d G^{j-1} H & \text{if } j > 0 \end{cases}$$

5.1 Martingale extraction

In section 2, we considered benchmark growth processes that were geometric random walks with drifts. Empirically our cash flows are observed to have stationary components as well. This leads us to construct the random walk components to the cash flow process. Specifically, we represent the log dividend process as the sum of a constant, a martingale with stationary increments and the first difference of a stationary process.²¹ Write:

$$d_{t+1} - d_t = \mu_d + \iota(1)w_{t+1} - U_d^* x_{t+1} + U_d^* x_t$$

where:

$$\begin{aligned} \iota(1) &= \iota_0 + U_d(I - G)^{-1}H \\ U_d^* &= U_d(I - G)^{-1} \end{aligned}$$

Thus d_t has a growth rate μ_d and a martingale component with increment: $\iota(1)w_t$. To relate this to the development in section 2, $\iota(1) = \pi$, $\mu_d = \zeta$ and $f(x_t) = \exp(U_d^* x_t)$ in the cash flow representation (5). We will fit processes to cash flows to obtain estimates of $\iota(1)$ and μ_d .

²⁰This restriction is at odds with many models in the literature. For example, a discrete-time version of the share model of Santos and Veronesi (2001) can be depicted as in (??) with $D_t^* = C_t$. The long-run cash flow risk price is the same as that of consumption.

²¹Martingale approximations are commonly used in establishing central limit approximations (*e.g.* see Gordin (1969) or Hall and Heyde (1980)), and are not limited to log-linear processes. For scalar linear time series, it coincides with the decomposition of ?).

5.2 Empirical Specification of Dividend Dynamics

We identify dividend dynamics and, in particular, the martingale component $\iota(1)$ using VAR methods. Consider a VAR with three variables: consumption, corporate earnings and dividends (all in logarithms). Consumption and corporate earnings are modeled as before in a cointegrated system. In addition to consumption and earnings, we include in sequence the dividend series from each of the five book-to-market portfolios and from the market. Thus the same two shocks as were identified previously remain shocks in this system because consumption and corporate earnings remain an autonomous system. An additional shock is required to account for the remaining variation in dividends beyond what is explained by consumption and corporate earnings. As is evident from figures 1, these series have important stochastic low frequency movements relative to consumption. Cash flow models that feature substantial mean reversion or stochastically stable shares relative to aggregate consumption are poor descriptions of these data.

Formally, we append a dividend equation:

$$A_0^*y_t^* + A_1^*y_{t-1} + A_2^*y_{t-2} + \dots + A_\ell^*y_{t-\ell} + B_0^* = w_t^*, \quad (18)$$

to equation system (17). In this equation the vector of inputs is

$$y_t^* \doteq \begin{bmatrix} y_t \\ d_t \end{bmatrix} = \begin{bmatrix} c_t \\ e_t \\ d_t \end{bmatrix}$$

and the shock w_t^* is scalar with mean zero and unit variance. This shock is uncorrelated with the shock w_t that enters (17). The third entry of A_0^* is normalized to be positive. We refer to (18) as the dividend equation, and the shock w_t^* as the dividend shock. As in our previous estimation, we set $\ell = 5$. We presume that this additional shock has a permanent impact on dividends by imposing the linear restriction:

$$A^*(1) = [\alpha^* \quad -\alpha^* \quad 0].$$

In the next section we will explore sensitivity of our risk measures to alternative specifications of long-run stochastic growth in the cash flows.

A stationary counterpart to this log level specification can be written in terms of the the variables $(c_t - c_{t-1}), (e_t - c_t), (d_t - d_{t-1})$. We estimate the VAR using these transformed variables with four lags of the growth rate variables and five lags of the logarithmic differences between consumption and earnings.

5.3 Book to Market Portfolios

We use five portfolios constructed based on a measure of book equity to market equity, and characterize the time series properties of the dividend series as it covaries with consumption and earnings. We follow Fama and French (1993) and construct portfolios of returns by sorting stocks according to their book-to-market values. We use a coarser sort into 5 portfolios

Properties of Portfolios Sorted by Book-to-Market

	Portfolio					
	1	2	3	4	5	Market
One-period Exp. Return (%)	6.79	7.08	9.54	9.94	11.92	7.55
Long-Run Return (%)	8.56	8.16	10.72	10.84	13.01	8.77
Avg. B/M	0.32	0.61	0.83	1.10	1.80	0.65
Avg. P/D	51.38	34.13	29.02	26.44	27.68	32.39

Table 1: Data are quarterly from 1947 Q1 to 2005 Q4 for returns and annual from 1947 to 2005 for B/M ratios. Returns are converted to real units using the implicit price deflator for nondurable and services consumption. Average returns are converted to annual units using the natural logarithm of quarterly gross returns multiplied by 4. “One-period Exp. Return,” we report the predicted quarterly gross returns to holding each portfolio in annual units. The expected returns are constructed using a separate VAR for each portfolio with inputs: $(c_t - c_{t-1}, e_t - c_t, r_t)$ where r_t is the logarithm of the gross return of the portfolio. We imposed the restriction that consumption and earnings are not Granger caused by the returns. One-period expected gross returns are calculated conditional on being at the mean of the state variable implied by the VAR. “Long-Run Return” reports the limiting value of the logarithm of the expected long-horizon return from the VAR divided by the horizon. “Avg. B/M” for each portfolio is the average portfolio book-to-market over the period computed from COMPUSTAT. “Avg. P/D” gives the average price-dividend for each portfolio where dividends are in annual units.

to make our analysis tractable. In addition we use the value-weighted CRSP return for our “market” return.

Summary statistics for these portfolios are reported in table 1. The portfolios are ordered by average book to market values where portfolio 1 has the lowest book-to-market value and portfolio 5 has the highest. Both one-period and long-run average returns generally follow this sort. For example, portfolio 1 has much lower average returns than portfolio 5. It is well documented that the differences in these average returns are not explained by exposure to contemporaneous covariance with consumption.

In this section we are particularly interested in the behavior of the cash flows from the constructed portfolios. The constructed cash flow processes accommodate changes in the classification of the primitive assets and depend on the relative prices of the new and old asset in the book-to-market portfolios. The monthly cash flow growth factors for each portfolio are constructed from the gross returns to holding each portfolio with and without dividends. The difference between the gross return with dividends and the one without dividends times

the current price-dividend ratio gives the cash flow growth factor. Accumulating these factors over time gives the ratio of the current period cash flow to the date zero cash flow. We normalize the date zero cash flow to be unity. The measure of quarterly cash flows in quarter t that we use in our empirical work is the geometric average of the cash flows in quarter $t - 3$, $t - 2$, $t - 1$ and t . This last procedure removes the pronounced seasonality in dividend payments. Details of this construction are given in Hansen, Heaton, and Li (2005), which follows the work of Bansal, Dittmar, and Lundblad (2005). The geometric averaging induces a transient distortion to our cash flows, but will not distort the long run stochastic behavior.

We estimate $\iota(1)$ from the VAR inclusive of portfolio dividends which gives us a measure of π . We the limiting rates of returns using the methods described in section 2. Table 2 gives long-run average rates of return for the five book-to-market portfolios. We explore formally sensitivity to the risk aversion parameter θ and report derivatives with respect to the intertemporal elasticity parameter ρ .

We first consider the implied logarithm of the expected return decomposed and scaled by horizon. These are reported in figure 6 and discussed previously. The figures are computed assuming that the Markov state is set to its unconditional mean. This figure reproduces the decompositions depicted in the upper panel of 2, but here we include sensitivity to changes in the parameter ρ . Using the horizon counterpart to the ρ derivatives discussed previously, we compute approximations for $\rho = 1.5$ and $\rho = .5$. Changing ρ leads to a roughly parallel shift in the curves, with a larger value of ρ increasing the overall returns.

In table 2 we report the limiting cash flow discount rates or long-run expected returns. These adjust the asymptotic decay in valuation for dividend growth. Recall from table 1 that one-period and long-horizon reinvestment expected returns are similar for each portfolio. The asymptotic cash flow discount rates only achieve comparable dispersion for large values of θ , say $\theta = 20$. While the discount rates in table 2 are lower, common changes in these rates can be achieved by simply altering the subjective discount factor β .

Portfolio 1 has low long-run cash flow covariation with consumption relative to portfolio five. This results in larger risk adjustments for the high book-to-market portfolios. Complementary to many other asset pricing studies, differences in the average rates of return on long-run cash flows are small except for large values of the risk aversion parameter θ , say $\theta = 20$. In contrast to aggregate securities, the implied heterogeneity in the the limiting expected returns are now substantial when θ is large. For the reasons we gave earlier, changing θ alters the expected excess returns almost proportionately.

The derivatives with respect to ρ are similar across securities so that modest movements in ρ have very little impact on the excess long-run returns, but rather substantial impact on the returns. Notice that for large values of θ , increasing ρ above one reduces some of the expected excess long-run rates of return. Recall that under recursive utility, the temporal pattern of risk matters. Shifts in intertemporal preferences change the way these patterns are evaluated and therefore affect attitudes toward risk and risk premia.

Up until now, we have focused on the return implications of the cash flows. In the remainder of this section we consider implications for portfolio values. We consider two

Expected Returns by Horizon

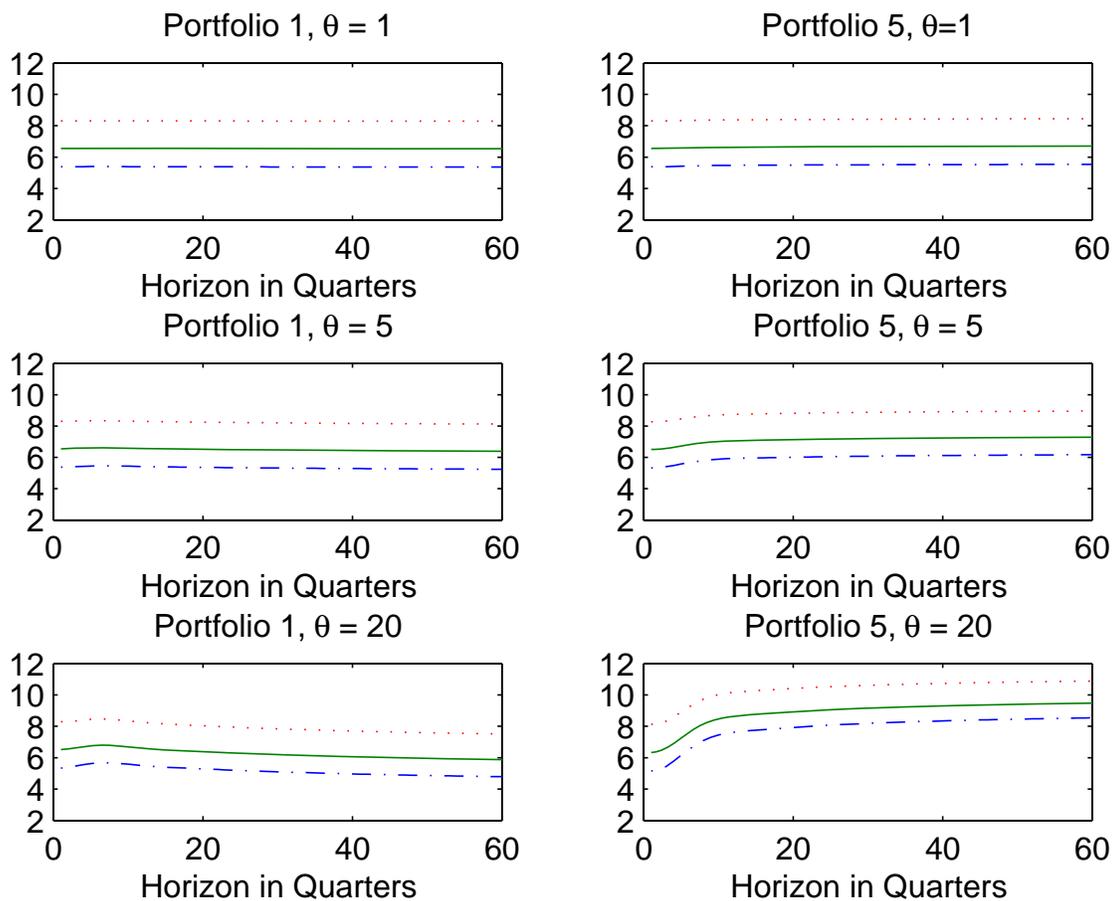


Figure 6: Expected returns to holding cash flows from portfolios 1 and 5 at different horizons. $-.$ assumes $\rho = 1/1.5$, $—$ assumes $\rho = 1$, \cdots assumes $\rho = 1.5$. Expected returns are in annualized percentages.

Limiting Cash Flow Discount Rates

Portfolio	Rate of Return	Excess Return	Return Derivative	Excess Return Derivative
$\theta = 1$				
1	6.51	-0.02	3.50	-0.00
2	6.54	0.01	3.50	0.00
3	6.67	0.14	3.52	0.01
4	6.70	0.17	3.52	0.01
5	6.75	0.22	3.52	0.02
market	6.60	0.06	3.51	0.01
$\theta = 5$				
1	6.27	-0.10	3.45	0.01
2	6.42	0.05	3.44	-0.00
3	7.03	0.66	3.35	-0.09
4	7.16	0.79	3.36	-0.08
5	7.42	1.05	3.33	-0.11
market	6.68	0.30	3.41	-0.03
$\theta = 20$				
1	5.39	-0.39	3.24	0.05
2	5.98	0.21	3.18	-0.01
3	8.37	2.59	2.75	-0.44
4	8.89	3.12	2.81	-0.37
5	9.92	4.15	2.67	-0.52
market	6.98	1.20	3.02	-0.17

Table 2: Limiting expected returns to holding the cash flows of portfolio 1 through 5. Excess returns are measured relative to the return on a long horizon discount bond. The derivative entries in columns four and five are computed with respect to ρ and evaluated at $\rho = 1$. Returns are reported in annualized percentages.

Limiting Decay Rates and Price-Dividend Ratios

Preference Parameters	Portfolio				
	1	2	3	4	5
	Limiting Decay Rates (annual %)				
$\theta = 5, \rho = 1$	4.16	4.48	2.71	3.14	0.40
$\theta = 20, \rho = 1$	3.27	4.04	4.04	4.87	2.90
$\theta = 30, \rho = 1$	2.68	3.74	4.93	6.02	4.57
$\theta = 30, \rho = 1/1.5$	1.66	2.74	4.13	5.19	3.83
	Price-Dividend Ratios				
$\theta = 5, \rho = 1$	23.6	22.1	36.9	32.0	251.4
$\theta = 20, \rho = 1$	28.5	23.8	26.6	21.9	35.9
$\theta = 30, \rho = 1$	33.6	25.1	22.8	18.4	23.4
$\theta = 30, \rho = 1/1.5$	54.4	34.4	27.7	21.7	28.4

Table 3: The limited decay rates are multiplied by 400 to produce annual percentages. The reported price-dividend ratios use the mean value of the state vector $x_t = 0$. These calculations assumed a subjective factor of $\beta = 0.97^{1/4}$.

calculations for alternative parameter configurations. First we report the limiting expected decay rate in value, ν , for each of the five portfolios. Under a continuous-time approximation, $1/\nu$ is the corresponding limiting price dividend ratio.²² Moreover, in the one-period return decomposition reported in bottom panel of figure 2, $\frac{\exp(-j\nu)}{1-\exp(-\nu)}$ gives an approximation to the weighting of the holding period returns in the total return. For smaller values of ν , the one-period holding period return on a cash flow that pays off far into the future is a more important contributor to the overall one-period return. We also report the implied price-dividend ratios when $x_t = 0$.²³

In table 3 we report the value of decay rates and price-dividend ratios for several different parameter configurations. We consider three alternative values of θ . Interestingly, the value decay rate portfolio 1 is much higher than that of portfolio 5 when $\theta = 5$. Bigger values of θ are required to ensure that price-dividend ratios are larger for low book-to-market portfolios than for high ones. Recall that low book-to-market portfolios generate less cash flow growth than high book-to-market portfolios. The differential risk adjustments have to

²²The discrete time counterpart is $\frac{\exp(-\nu)}{1-\exp(-\nu)}$.

²³These are not the same as the average price-dividend ratios even though they are evaluated at the average value of x_t . This is because the price-dividend ratios are nonlinear functions of x_t .

more than compensate for the growth rate disparity. For example, the second and third rows of each panel consider larger values of θ . Even $\theta = 20$ is not large enough for limiting price-dividend ratio of the low book-to-market portfolios to exceed those of the high book-to-market portfolios. By further increasing $\theta = 30$ (fourth row of each panel), the value decay for portfolio one is noticeably less than that of portfolio five. Recall that lower decay rates in value imply that the contributions of the holding period returns for cash flows far into the future are more pronounced.

Next, we consider the impact of changing the intertemporal substitution parameter ρ . As we saw in table 2 a decrease in ρ lowers all expected rates of returns. As a consequence decreasing ρ to $1/1.5$ (fourth row of each panel in table 3) reduces all of the decay rates. Correspondingly, price-dividend ratios are increased. The increase is particularly pronounced for the low book-to-market portfolios.

The decay rates and price-dividend ratios are depicted for a single value of the subjective discount factor β . Changing β gives an alternative mechanism for altering price-dividend ratios. Decreasing β results in a roughly parallel shift in the decay rates for all of the portfolios and corresponding increases the price-dividend ratios. Unfortunately, this change with low risk aversion will not induce relatively low decay rates for the low-book to market portfolios. Moreover, this decrease increases the one-period risk free rate, which is arguably high for our baseline value of $\beta = .97^{1/4}$.²⁴

All of this discussion abstracts from errors in estimating the cash flow growth rates and risk exposure. We address estimation accuracy and sensitivity to specification in the next two sections.

5.4 Sampling uncertainty in long run risk prices

We consider sampling uncertainty in some of the inputs used to measure long run risk. Recall that these inputs are based in part on low frequency extrapolation of VAR systems fit to match transition dynamics. As in the related macroeconomics literature, we expect a substantial degree of sampling uncertainty. We now quantify how substantial this is for our application.

When $\rho = 1$, the expected excess returns are approximately equal to:

$$\theta\gamma(1) \cdot \pi.$$

We now investigate the statistical accuracy of $\gamma(1) \cdot \pi$ for the five portfolios, and for the difference between portfolios 1 and 5. The vector π is measured using $\iota(1)$. In table 4 we report the approximate posterior distribution for $\gamma(1) \cdot \pi$ computed using the approach of Sims and Zha (1999) and Zha (1999). As before we use Box-Tiao priors. We scale the values of $\gamma(1) \cdot \pi$ by 400 just as we did when reporting predicted annualized average returns in percentages. While there is a considerable amount of statistical uncertainty in these risk measures, there are important differences in the risk relative risk exposures of portfolios 1

²⁴For instance, the quarterly risk-free rate is 6.6% in annual units with $\theta = 1$.

Accuracy of Risk Measures

Portfolio	Quantile				
	.05	.25	.5	.75	.95
1	-.63	-.13	-.05	-.01	.10
2	-.19	-.03	.01	.04	.22
3	.01	.06	.12	.28	1.42
4	.04	.10	.17	.32	1.46
5	.04	.12	.21	.42	1.88
market	-.01	.03	.06	.12	.58
5-1	.02	.07	.32	.67	2.81

Table 4: Accuracy of estimates of $\gamma(1) \cdot \pi$ scaled by 400. Quantiles were computed by simulating 100,000 times using Box-Tiao priors. The quantiles were computed using only simulation draws for which the absolute values of the eigenvalues were all less than .999. The fraction of accepted draws ranged from .986 to .987. The quantiles were computed using VARs that included consumption, corporate earnings and a single dividend series with one exception. To compute quantiles for the 5 – 1 row, dividends for both portfolios were included in the VAR.

and 5. We conclude that, from a long-run perspective, there is significant evidence that the portfolio of value stocks is riskier than the portfolio of growth stocks.

5.5 Specification sensitivity for cash flows

So far our measurements and inferences are conditioned on particular models of stochastic growth. In this section we explore the impact of changing the growth configurations in the cash flow dynamics. All of these specifications allow for the dividends from the portfolios to have growth components that are different from consumption to accommodate the growth heterogeneity that is evident in figure 1.

In our baseline model, we identified dividend dynamics and, in particular, the martingale component $\iota(1)$ using VAR methods. We used a VAR with three variables: consumption, corporate earnings and dividends (all in logarithms). Consumption and earnings were restricted to the same long-run response to permanent shocks. In addition, dividends had their own stochastic growth component.

We now consider two alternative specifications of dividend growth. Both are restrictions on the equation:

$$A_0^* y_t^* + A_1^* y_{t-1} + A_2^* y_{t-2} + \dots + A_\ell^* y_{t-\ell} + B_0^* + B_1^* t = w_t^*,$$

where the shock w_t^* is scalar with mean zero and unit variance and uncorrelated with the shock vector w_t that enters (17). The third entry of A_0^* is normalized to be positive. As in our previous estimation, we set $\ell = 5$.

The first alternative specification restricts that the trend coefficient B_1^* equal zero, and is the model used by Hansen, Heaton, and Li (2005). Given our interest in measuring long-run risk, we measure the permanent response of dividends to the permanent shock. While both consumption and corporate earnings continue to be restricted to respond to permanent shocks in the same manner, the dividend response is left unconstrained. There is no separate growth component for dividends in this specification. The second alternative specification includes a time trend by freely estimating B_1^* . A model like this, but without corporate earnings, was used by Bansal, Dittmar, and Lundblad (2005). We refer to this as the time trend specification. In this model the time trend introduces a second source of dividend growth.

The role of specification uncertainty is illustrated in the impulse responses depicted in figure 7. This figure features the responses of portfolios 1 and 5 to a permanent shock. For each portfolio, the measured responses obtained for each of the three growth configurations are quite close up to about 12 quarters (3 years) and then they diverge. Both portfolios initially respond positively to the shock with peak responses occurring in about seven quarters periods. The response of portfolio 5 is much larger in this initial phase than that of portfolio 1. The two alternative models for portfolio 5 give essentially the same impulse responses. The time trend is essentially zero for portfolio 5. The limiting response of the alternative models are much lower than that of the baseline specification.

For portfolio 1 there are important differences in the limiting responses of all three models. While the limiting response of the baseline model is negative, when a time trend is introduced in place of a stochastic growth component, the limit becomes substantially more negative. The time trend specification implies that portfolio 1 provides a large degree of consumption insurance in the long run in contrast to the small covariation measured when the additional growth factor is stochastic, as in our baseline dividend growth model. When consumption/earnings is the sole source of growth, the limiting response is positive but small. While the limiting responses are sensitive to the growth specification, the differences in the long-run responses between portfolios 1 and 5 are approximately the same for the time trend model and for our baseline dividend growth model.²⁵

While the use of time trends as alternative sources of cash flow growth alters our results, it requires that we take these trends literally in quantifying long run risk. Is it realistic to think of these secular movements, that are independent of consumption growth, as deterministic trends when studying the economic components of long-run risk? We suspect not. While there may be important persistent components to the cash flows for portfolio 1, it seems

²⁵Bansal, Dittmar, and Lundblad (2005) use their estimates with a time trend model as inputs into a cross sectional return regression. While estimation accuracy and specification sensitivity may challenge these regressions, the consistency of the ranking across methods is arguably good news, as emphasized to us by Ravi Bansal. As is clear from our previous analysis, we are using the economic model in a more formal way than the running of cross-sectional regressions.

Impulse Response Functions for Two Portfolios

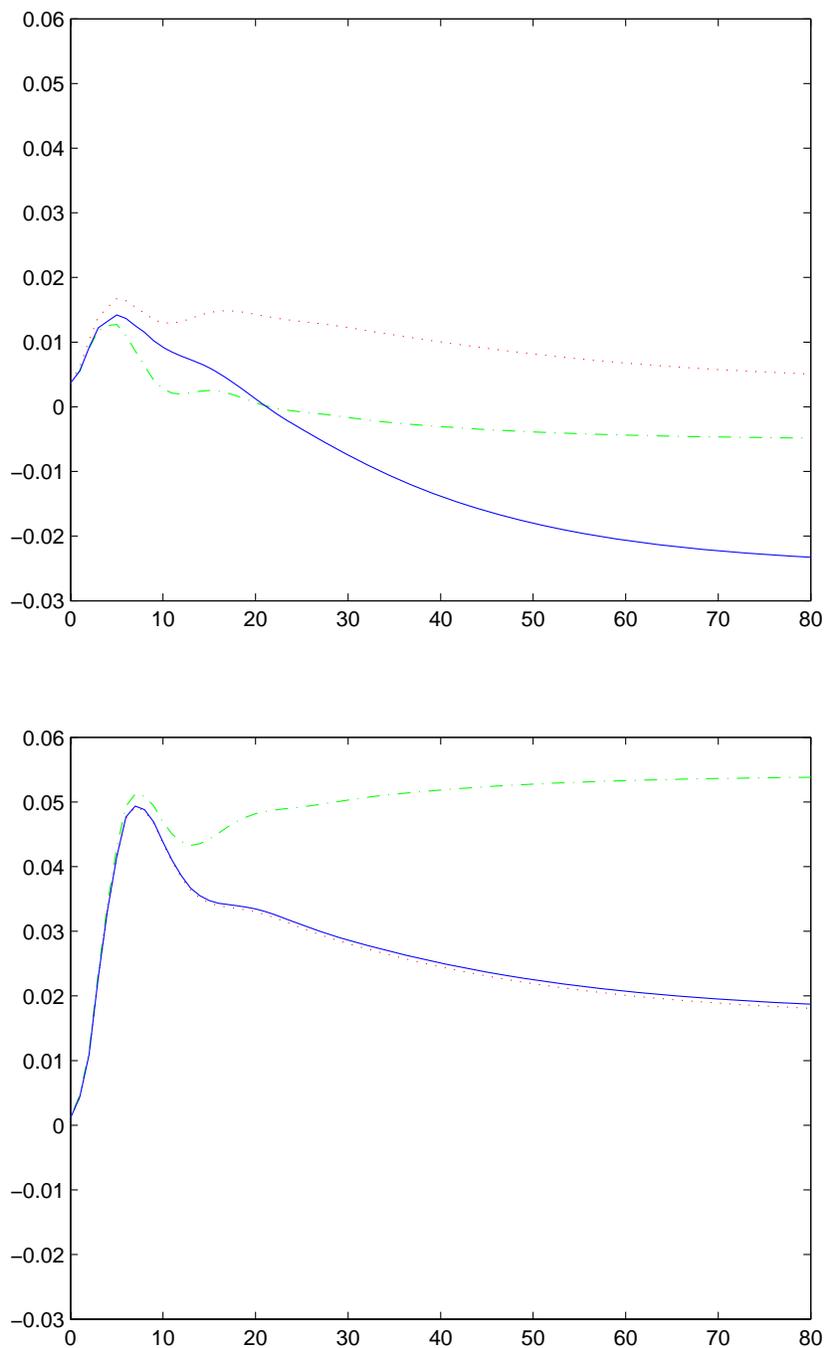


Figure 7: This figure depicts the impulse responses to a permanent shock to consumption of the cash flows to portfolios 1 and 5. The \cdots curve is generated from the level specification for dividends; the --- is generated from the level specification with time trends included; and the --- curve is generated from the first difference specification.

unlikely that these components are literally deterministic time trends known *a priori* to investors. The time trend for this portfolio is in part offset by the negative growth induced by cointegration. We suspect that the substantially negative limiting response for portfolio 1 is unlikely to be the true limiting measures of how dividends respond to consumption and earnings shocks.²⁶

In summary, while there is intriguing heterogeneity in the long run cash flow responses and implied returns, the implied risk measures are sensitive to the growth specification. Given the observed cash flow growth, it is important to allow for low frequency departures from a balanced growth restriction. The simple cointegration model introduces only one free growth parameter for each portfolio, but results in a modest amount of cash flow heterogeneity. The time trend growth models impose additional sources of growth. The added flexibility of the time trend specification may presume too much investor confidence in a deterministic growth component, however. The dividend growth specification that we used in our previous calculations, while *ad hoc*, presumes this additional growth component is stochastic and is a more appealing specification to us.²⁷

²⁶Sims (1991) and Sims (1996) warn against the use of time trends using conditional likelihood methods because the resulting estimates might over fit the the initial time series, ascribing it to a transient component far from the trend line.

²⁷In the specifications we have considered, we have ignored any information for forecasting future consumption that might be contained in asset prices. Our model of asset pricing implies a strict relationship between cash flow dynamics and prices so that price information should be redundant. Prices, however, may reveal additional components to the information set of investor and hence a long-run consumption risk that cannot be identified from cash flows. When we consider an alternative specification of the VAR where we include consumption, corporate earnings, dividends as well as prices, we obtain comparable heterogeneity.

6 Conclusion

Growth-rate variation in consumption and cash flows have important consequences for asset valuation. The methods on display in this paper formalize the long-run contribution to value of the stochastic components of discount factors and cash flows and quantify the importance of macroeconomic risk. We used these methods to isolate features of the economic environment that have important consequences for long-run valuation and heterogeneity across cash flows. We made operational a well defined notion of long-run cash flow risk and a well defined limiting contribution to the one-period returns coming from cash flows in the distant future.

In our empirical application we showed that the stochastic growth of growth portfolios has negligible covariation with consumption in the long run while the cash flow growth of value portfolios has positive covariation. For these differences to be important quantitatively for our long run risk-return calculations, investors must be either highly risk averse or highly uncertain about the probability models that they confront.²⁸

In this paper we used an *ad hoc* VAR model to identify shocks. In contrast to VAR methods, an explicit valuation model is a necessary ingredient for our analysis; and thus we analyzed the valuation implications through the lens of a commonly used consumption-based model. There are important reasons for extending the scope of our analysis in future work, and the methods we described here are amenable to such extensions. One next step is to add more structure to the macroeconomic model, structure that will sharpen our interpretation of the sources of long-run macroeconomic risk.

While the recursive utility model used in this paper has a simple and usable characterization of how temporal dependence in consumption growth alters risk premia in the long run, other asset models have interesting transient implications for the intertemporal composition of risk, including models that feature habit persistence (*e.g.* Constantinides (1990), Heaton (1995), and Sundaresan (1989)) and models of staggered decision-making (*e.g.* see Lynch (1996) and Gabaix and Laibson (2002).)

The model we explore here focuses exclusively on time variation in conditional means. Temporal dependence in volatility can be an additional source of long-run risk. Time variation in risk premia can be induced by conditional volatility in stochastic discount factors.²⁹ While the direct evidence from consumption data for time varying volatility in post war data is modest, the implied evidence from asset pricing for conditional volatility in stochastic discount factors is intriguing. For instance, Campbell and Cochrane (1999) and others argue that risk prices vary over the business cycle in ways that are quantitatively important. The models of Campbell and Cochrane (1999) and Lettau and Wachter (2005) are alternative ways to alter the long-run risk characterization through volatility channels.

While the methods we have proposed aid in our understanding of asset-pricing models, they also expose measurement challenges in quantifying the long-run risk-return tradeoff.

²⁸This latter conclusion can be made precise by using detection probabilities in the manner suggested by Anderson, Hansen, and Sargent (2003).

²⁹It can also be induced by time variation in risk exposure.

Important inputs into our calculations are the long-run riskiness of cash flows and consumption. As we have shown, these objects are hard to measure in practice. Statistical methods typically rely on extrapolating the time series model to infer how cash flows respond in the long-run to shocks. This extrapolation depends on details of the growth configuration of the model, and in many cases these details are hard to defend on purely statistical grounds. Also there is pervasive statistical evidence for growth rate changes or breaks in trend lines, but this statistical evidence is difficult to use directly in models of decision-making under uncertainty without some rather specific ancillary assumptions about investor beliefs. Many of the statistical challenges that plague econometricians presumably also plague market participants. Naive application of rational expectations equilibrium concepts may endow investors in these models with too much knowledge about future growth prospects. Learning and model uncertainty are likely to be particularly germane to understanding long run risk.

A Eigenfunction results

In what follows we use the notation:

$$M_{t+1,t} \doteq \exp(s_{t+1,t} + \pi w_{t+1})$$

A.1 Eigenfunctions and stability

We follow Hansen and Scheinkman (2006) by formalizing the approximation problem as a change in measure. Our analysis is in discrete time in contrast to their continuous-time analysis. Moreover, we develop some explicit formulas that exploit our functional forms.

Write the eigenfunction problem as:

$$E [M_{t+1,t} \phi(x_{t+1}) | x_t] = \exp(-\nu) \phi(x_t)$$

where $M_{t+1,t} = \exp(s_{t+1,t} + \pi w_{t+1})$. Then

$$\hat{M}_{t+1,t} = \exp(\nu t) M_{t+1,t} \begin{bmatrix} \phi(x_{t+1}) \\ \phi(x_t) \end{bmatrix}$$

satisfies:

$$E \left(\hat{M}_{t+1,t} | x_t \right) = 1.$$

As a consequence $\hat{M}_{t+1,t}$ induces a distorted conditional expectation operator. Recall our solution $\phi(x) = \exp(-\bar{\omega}x)$ to this problem. Then by the usual complete the square argument, $\hat{M}_{t+1,t}$ changes the distribution of w_{t+1} from being a multivariate standard normal to being a multivariate normal with mean:

$$\hat{\mu}_w = H' \bar{\omega}' + \pi' + \xi'_0 \quad (19)$$

and covariance matrix I . This adds a constant term to the growth rate of consumption. Let the implied distorted expectation operator \hat{E} .

We use this distorted shock distribution in our computations. For instance,

$$E [M_{t+1,t} f(x_{t+1}) | x_t] = \exp(\rho) \phi(x_t) \hat{E} \left[\frac{f(x_{t+1})}{\phi(x_{t+1})} | x_t \right].$$

Iterating, we obtain:

$$E [M_{t+j,t} f(x_{t+j}) | x_t] = \exp(\rho j) \phi(x_t) \hat{E} \left[\frac{f(x_{t+j})}{\phi(x_{t+j})} | x_t \right].$$

The limit that interests us is:

$$\lim_{j \rightarrow \infty} \hat{E} \left[\frac{f(x_{t+j})}{\phi(x_{t+j})} | x_t \right] = \hat{E} \left[\frac{f(x_t)}{\phi(x_t)} \right]$$

provided that $\{x_t\}$ has a well defined stationary distribution under the \hat{E} probability distribution and the conditional expectation operator converges the corresponding unconditional expectation operator.

Let q and \hat{q} denote the stationary densities of $\{x_t\}$ under E and the \hat{E} measures. Define $\varphi \doteq \hat{q}/(q\phi)$ implying that

$$\hat{E} \left[\frac{f(x_t)}{\phi(x_t)} \right] = E[\varphi(x_t)f(x_t)] .$$

The density q is normal with mean zero and covariance matrix:

$$\Sigma = \sum_{j=0}^{\infty} (G^j) H H' (G^j)',$$

which can be computed easily using a doubling algorithm. The density \hat{q} is normal with mean

$$\hat{\mu}_x \doteq (I - G)^{-1} H (-H' \omega' + \pi' + \xi_0'), \quad (20)$$

and the same covariance matrix as q .

Consider now a joint Markov process $\{(x_t, z_t) : t \geq 0\}$, and the equation:

$$E \left[M_{t+1,t} \left(\frac{z_{t+1}}{z_t} \right) \left(\frac{\phi(x_{t+1})}{z_{t+1}} \right) | x_t \right] = \exp(\rho) \left[\frac{\phi(x_t)}{z_t} \right] .$$

While this amounts to a rewriting of the initial eigenvalue equation, it has a different interpretation. The process $\{z_t\}$ is a transient contribution to the stochastic discount factor, and the eigenfunction equation is now expressed in terms of the composite state vector (x, z) with the same eigenvalue and an eigenfunction $\phi(x)/z$. The limit of interest is now:

$$\lim_{j \rightarrow \infty} \hat{E} \left[\frac{f(x_{t+j}) z_{t+j}}{\phi(x_{t+j})} | x_t \right] = \hat{E} \left[\frac{f(x_t) z_t}{\phi(x_t)} \right] .$$

To study this limit we require that the process $\{(x_t, z_t)\}$ be stationary under the distorted probability distribution and that $f(x_t)z_t$ have a finite expectation under this distribution.

In the special case in which $G = 0$, and $\phi = 1$ it suffices to study $f(x_t)z_t$.

A.2 Eigenvalue derivative

We compute this derivative using the approach developed in Hansen (2006). Suppose that $\hat{M}_{t+1,t}$ depends implicitly on a parameter ρ . Since each member of the parameterized family has conditional expectation equal to unity,

$$\hat{E} \left(\frac{\partial \log \hat{M}_{t+1,t}}{\partial \rho} | x_t \right) = E \left(\frac{\partial \hat{M}_{t+1,t}}{\partial \rho} | x_t \right) = 0.$$

Note that

$$\hat{E} \left(\frac{\partial \log \hat{M}_{t+1,t}}{\partial \rho} \Big| x_t \right) = \hat{E} \left(\frac{\partial \log M_{t+1,t}}{\partial \rho} \Big| x_t \right) - \frac{\partial \nu}{\partial \rho} + \hat{E} \left(\frac{\partial \log \phi(x_{t+1})}{\partial \rho} \Big| x_t \right) - \frac{\partial \log \phi(x_t)}{\partial \rho}$$

Since the left-hand side is zero, applying the Law of Iterated Expectation under the $\hat{\cdot}$ probability measure:

$$0 = \hat{E} \left(\frac{\partial \log M_{t+1,t}}{\partial \rho} \right) - \frac{\partial \nu}{\partial \rho} + \hat{E} \left(\frac{\partial \log \phi(x_{t+1})}{\partial \rho} \right) - \hat{E} \left(\frac{\partial \log \phi(x_t)}{\partial \rho} \right).$$

Since $\{x_t\}$ is stationary under the $\hat{\cdot}$ probability measure,

$$\frac{\partial \nu}{\partial \rho} = \hat{E} \left(\frac{\partial \log M_{t+1,t}}{\partial \rho} \right).$$

To apply this formula, write

$$\log M_{t+1,t} = s_{t+1,t} + \pi w_{t+1}$$

Differentiating with respect to ρ :

$$Ds_{t+1,t}^1 = \frac{1}{2} w_{t+1}' \Theta_0 w_{t+1} + w_{t+1}' \Theta_1 x_t + \vartheta_0 + \vartheta_1 x_t + \vartheta_2 w_{t+1}.$$

Recall that under the distorted distribution w_{t+1} has a constant mean $\hat{\mu}_w$ conditioned on x_t given by (19) and x_t has a mean $\hat{\mu}_x$ given by (20). Taking expectations under the distorted distribution:

$$\hat{E} (Ds_{t+1,t}^1) = \frac{1}{2} (\hat{\mu}_w)' \Theta_0 \hat{\mu}_w + \frac{1}{2} \text{trace}(\Theta_0) + (\hat{\mu}_w)' \Theta_1 \hat{\mu}_x + \vartheta_0 + \vartheta_1 \hat{\mu}_x + \vartheta_2 \hat{\mu}_w.$$

References

- Alvarez, F. and U. Jermann (2005). Using asset prices to measure the the persistence in the marginal utility of wealth. *Econometrica* 73, 1977–2016.
- Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). A quartet of semi-groups for model specification, detection, robustness and the price of risk. *Journal of European Economic Association* 1, 66–123.
- Bansal, R., R. Dittmar, and C. Lundblad (2005). Consumption, dividends, and the cross-section of equity returns. *Journal of Finance forthcoming*.
- Bansal, R. and B. N. Lehmann (1997). *Macroeconomic Dynamics* 1, 333–354.
- Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481–1509.
- Berk, K. N. (1974). Consistent autoregressive spectral estimates. *Annals of Statistics* 2, 289–502.
- Blanchard, O. J. and D. Quah (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review* 79, 655–673.
- Breeden, D. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics* 7, 265–296.
- Campbell, J. (1996). Understanding risk and return. *Journal of Political Economy* 104, 298–345.
- Campbell, J. and J. Cochrane (1999). By force of habit. *Journal of Political Economy* 107, 205–251.
- Campbell, J. and T. Vuolteenaho (2004). Bad beta, good beta. *American Economic Review* 94, 1249–1275.
- Campbell, J. Y. (1994). Intertemporal asset pricing without consumption data. *American Economic Review* 83, 487–512.
- Cochrane, J. H. (1988). How big is the random walk in gnp. *Quarterly Journal of Economics* 96, 893–920.
- Constantinides, G. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy* 104, 519–543.
- Daniel, K. and D. Marshall (1997). Equity-premium and risk-free-rate puzzles at long horizons. *Macroeconomic Dynamics* 1, 452–484.
- Epstein, L. and S. Zin (1989a). Substitution, risk aversion and the temporal behavior of stock returns: An empirical investigation. *Journal of Political Economy* 99, 263–286.
- Epstein, L. G. and S. E. Zin (1989b). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica* 57, 937–969.

- Fama, E. and K. French (1992). The cross-section of expected stock returns. *Journal of Finance* 47, 427–465.
- Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3–56.
- Gabaix, X. and D. Laibson (2002). The 6d bias and the equity premium puzzle. In B. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual*. MIT Press.
- Gordin, M. I. (1969). The central limit theorem for stationary processes. *Soviet Math. Dokl.* 10, 1174–1176.
- Hall, P. and C. C. Heyde (1980). *Martingale Limit Theory and its Application*. Academic Press.
- Hansen, L. P. (2006). Modeling the long run: Valuation in dynamic stochastic economies. Prepared for the Fisher-Schultz Lecture at the 2006 European Meetings of the Econometric Society.
- Hansen, L. P. (2007). Beliefs, doubts and learning: Valuing macroeconomic risk. *American Economic Review* 97, 1–30. Richard Ely Lecture to the American Economic Association.
- Hansen, L. P., J. Heaton, and N. Li (2005). Intangible risk. In C. Corrado, J. Haltiwanger, and D. Sichel (Eds.), *Measuring Capital in the New Economy*. University of Chicago Press.
- Hansen, L. P., J. C. Heaton, J. Lee, and N. Rousanov (2006). Intertemporal substitution and risk aversion. University of Chicago.
- Hansen, L. P. and J. Scheinkman (2006). Long term risk: An operator approach. University of Chicago.
- Hansen, L. P. and K. Singleton (1983). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy* 91, 249–265.
- Heaton, J. (1995). An empirical investigation of asset pricing with temporally dependent preference specifications. *Econometrica* 63, 681–717.
- Kocherlakota, N. R. (1990). Disentangling the coefficient of relative risk aversion from the elasticity of intertemporal substitution: An irrelevance result. 45, 175–190.
- Kogan, L. and R. Uppal (2001). Risk aversion and optimal portfolio policies in partial and general equilibrium economies. Sloan School of Management and London Business School.
- Kreps, D. M. and E. L. Porteus (1978). Temporal resolution of uncertainty and dynamic choice. *Econometrica* 46, 185–200.
- Lettau, M. and J. Wachter (2005). Why is long-horizon equity less risky? a duration-based explanation of the value premium. NBER working paper number 11144.

- Lettau, M. L., S. C. Ludvigson, and J. A. Wachter (2004). The declining equity premium: What role does macroeconomic policy play? New York University.
- Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica* 46, 1429–1445.
- Lynch, A. (1996). Decision frequency and synchronization across agents: Implications for aggregate consumption and equity returns. *Journal of Finance* 51, 1479–1497.
- Menzly, L., T. Santos, and P. Veronesi (2004). Understanding predictability. *Journal of Political Economy* 112(1), 1–47.
- Restoy, F. and P. Weil (1998, June). Approximate equilibrium asset prices. NBER Working Paper 6611.
- Rubinstein, M. (1976). The valuation of uncertain income streams and the price of options. *Bell Journal of Economics* 7, 407–425.
- Santos, T. and P. Veronesi (2001). Labor income and predictable stock returns. Graduate School of Business, University of Chicago.
- Schroder, M. and C. Skiadas (1999). Optimal consumption and portfolio selection with stochastic differential utility. *Journal of Economic Theory* 89, 68–126.
- Sims, C. A. (1991). Comment on ‘to criticize the critics’. *Journal of Applied Econometrics* 6, 423–434.
- Sims, C. A. (1996). Inference for multivariate time series models with trend. Originally presented at the August 1992 ASSA Meetings.
- Sims, C. A. and T. Zha (1999). Error bands for impulse responses. *Econometrica* 67, 1113–1156.
- Sundaresan, S. (1989). Intertemporally dependent preferences and volatility of consumption and wealth. *Review of Financial Studies* 3, 73–89.
- Tallarini, T. (1998). Risk sensitive business cycles. *Journal of Monetary Economics* 45, 507–532.
- Weil, P. (1990). Nonexpected utility in macroeconomics. *Quarterly Journal of Economics* 105, 29–42.
- Zha, T. (1999). Block recursion and structural vector autoregression. *Journal of Econometrics* 90, 291–316.