

Consumption Strikes Back?: Measuring Long-Run Risk¹

Lars Peter Hansen John C. Heaton
Department of Economics Graduate School of Business
University of Chicago University of Chicago
Email:l-hansen@uchicago.edu Email:john.heaton@uchicago.edu

Nan Li
School of Business
National University of Singapore
Email:biznl@nus.edu.sg

January 3, 2008

¹We thank Fernando Alvarez, Ravi Bansal, Jaroslav Borovička, John Cochrane, Ken Judd, Anil Kashyap, Mark Klebanov, Junghoon Lee, Johnathan Parker, Tano Santos, Chris Sims, Grace Tsiang, Pietro Veronesi, Kenji Wada, two anonymous referees and especially Monika Piazzesi and Tom Sargent for valuable comments. 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 risk prices of cash flows that are realized far into the future but are reflected in asset values. We apply this analysis to claims on aggregate cash flows, as well as to 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 unmodeled 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 to 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 comovement 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) advocate using 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 and these differences are most dramatic in the long run.

Given our choice of models and evidence, we devote part of our analysis to evaluating 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 provide a clear understanding of where historical data are 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 the long-run implications of models 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 component and a long-term risk premium. This long-term 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

The state of the economy is given by a 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 important determinants of the stochastic discount factor. We assume that differences in the logarithm of aggregate consumption are a linear function of the state vector.⁴

$$c_{t+1} - c_t = \mu_c + U_c x_t + \lambda_0 w_{t+1}. \quad (2)$$

Under this assumption, in section 3 we show that the logarithm of the stochastic discount factor $s_{t+1,t} \equiv \log S_{t+1,t}$ for a version of the recursive utility model is 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 $\xi_0 = -\gamma\lambda_0$ where γ 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. With recursive utility the impact of the vector of shocks w_{t+1} on the discount factor is modified. For example, when the intertemporal elasticity of substitution is equal to one, the weighting on the current shock becomes:

$$\xi_0 = -\lambda_0 + (1 - \gamma)\lambda(\beta)$$

where $\log(\beta)$ is the subjective rate of time discount in preferences,

$$\begin{aligned} \lambda(\beta) &= \lambda_0 + U_c \sum_{j=1}^{\infty} \beta^j G^{j-1} H \\ &= \lambda_0 + \beta U_c (I - \beta G)^{-1} H, \end{aligned}$$

⁴See Hansen, Heaton, Lee, and Roussanov (2006) for a generalization with stochastic volatility.

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

2.2 The risk-return tradeoff

Our decomposition of long-run returns requires a specification of the long-run components of cash flows. In our application these cash flows are dividends flowing to those holding the stocks in the portfolios. 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 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 exposure to long-run risk, and our aim is to assign prices to this exposure. 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 affect the risk prices in the limit.

In our analysis we consider the cash flows from several portfolios. Figure 1 displays three of our cash flow series: a portfolio of growth stocks, a portfolio of value stocks and the value-weighted market portfolio. The stocks included in the growth portfolio are those with a low ratio of book equity to market equity and the stock included in the value portfolio are those with a high ratio of book equity to market equity. The portfolios are re-balanced as in Fama and French (1992).⁵ In figure 1, the cash flows are depicted relative to aggregate consumption with the initial cash flows normalized to equal aggregate consumption. Notice that the cash flows of the growth portfolio grow much more slowly than those of value portfolio. The difference 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.

⁵Details of the construction of the portfolios and cash flows can be found in Hansen, Heaton, and Li

Cash Flows Relative to Consumption for Two Portfolios

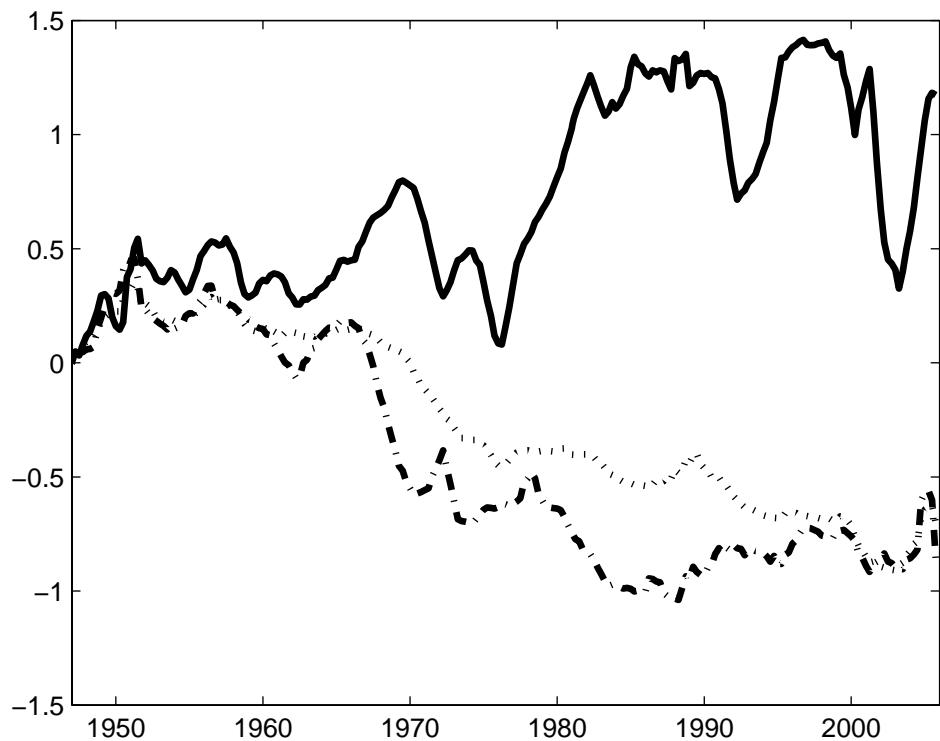


Figure 1: The quarterly logarithms of the ratios of portfolio cash flows to consumption are depicted. The solid line depicts the high book-to-market portfolio the dotted line the value-weighted market portfolio and the dash-dotted line the high book-to-market portfolio of stocks.

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 incorporate this process into the conditional expectations operator and create a one-period *valuation operator*:

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

This operator is much like the conditional expectations operator, but it differs in important ways. It is not representable using a transition density function that integrates to one because of the contribution of the stochastic discount factor and the stochastic growth. Our valuation operator allows us to fix the long-run cash flow dynamics but consider alternative transient components given by different choices of the function f . The date t price of the cash flow D_{t+1} is $D_t^* \mathcal{P}f(x_t)$.

Before proceeding, notice that pricing is recursive so that prices of cash flows multiple periods in the future are inferred from the 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. The prices of these cash flows eventually decline as the horizon j increases. The rate of decline or decay in these values depends on the expected growth in cash flows relative to discount rate. Our first result characterizes this limiting rate of decay in value.

When the function $f(x)$ is assumed to be a log-linear function of the state x , the functions $\{\mathcal{P}^j f(x), j = 1, 2, \dots\}$ are also log-linear 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:

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

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: $-\nu \equiv \mu_s + \zeta + \frac{|\bar{\omega}H + \xi_0 + \pi|^2}{2}$. 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 of repeated iteration of the above relations is summarized in the following:

(2005) and at <http://www.bschool.nus.edu.sg/staff/biznl/papers/bmdata.html>.

Result 2.1. *Equation:*

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

has a strictly positive solution e given by $e(x) = \exp[U_s(I - G)^{-1}x]$. The corresponding value of $-\nu$ is $-\nu = \mu_s + \zeta + \frac{|\bar{\omega}H + \xi_0 + \pi|^2}{2}$.

The 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. In what follows we will refer to e as the principal eigenfunction, and it will be used to represent some of the limits that follow.

While these iterations 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 to growth rate risk given by the vector π .⁶

We use this characterization of the limit to investigate long-run risk. As j gets larger, although $\mathcal{P}^j f(x)$ approaches zero it does so eventually at a rate that is approximately constant. The value of ν gives this asymptotic rate of decay of the values. It 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 $\{D_t^*\}$ is a geometric random walk with drift, the long-run growth rate is

$$\eta = \zeta + \frac{1}{2}\pi \cdot \pi.$$

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.⁷ 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

$$\pi^* \equiv -\xi_0 - U_s(I - G)^{-1}H$$

⁶We could 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 us to include “share models” with nonlinear share evolution equations as in Santos and Veronesi (2001).

⁷Formally, 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 $\exp(\eta)$.

$$\varsigma^* \equiv -\mu_s - \frac{\pi^* \cdot \pi^*}{2}.$$

The vector π^* prices exposure to long-run risk. It depends on the assumed consumption dynamics and the preferences of the representative consumer. The vector π measures the extent to which cash flows are exposed to long-run risk. 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: ς^* . 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.

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:

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

where $\lambda(1)$ 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:

$$\pi^* = \lambda_0 + (\gamma - 1)\lambda(\beta)$$

which is approximately the same for β close to unity. The single period counterparts will differ provided that consumption is predictable (see Kocherlakota (1990) and Bansal and Yaron (2004)). Bansal, Dittmar, and Kiku (2007) study a limiting version of a risk-return return relation under log-linearity and power utility. They focus on cumulative returns and study consumption betas in a model in which dividends and consumption are cointegrated with a coefficient that can differ across portfolios. The limiting portfolio beta is determined by the cointegrating coefficient, which they use as a measure of long-run cash flow risk exposure. Cointegration is not featured in our analysis. Our limiting risk prices can be used in general log-linear settings including their environment.

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 to 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 figure 2 displays estimates of risk-adjusted returns for the cash flows produced by the growth and value portfolios (portfolios 1 and 5 respectively). These calculations assume recursive preferences.⁸ For comparison estimated returns for the market portfolio are also reported. When considering the expected returns for the growth and value portfolios, note that the observed average returns to these portfolios are substantially different. As reported in table 1 below, the expected one-period returns to the growth and values portfolios are 6.8% and 11.9% respectively.

The pattern of risk premia across horizons is intriguing. The expected returns to the value portfolio increase with horizon in contrast to the market portfolio and especially the growth portfolio. This effect is due to important exposure to long-run macroeconomic risk in the cash flows from the value portfolio. In contrast the risk premia for all portfolios are similarly small for short horizons. As the horizon increases the expected returns approximate their long-run limits given in theorem 1 and the limiting differences are reflected in the prices of growth-rate risk. In subsection 5.5 we investigate how sensitive these measurements are to errors in the specification of growth and to estimation accuracy.

2.4 One-period returns

The expected one-period returns for the growth and value portfolios are substantially different. Each of these one-period returns is a weighted average of one-period returns to holding the corresponding cash flows at alternative horizons. The gross holding period return to a security that pays off $f(x_{t+j})$ 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:

$$R_{t+1,t}^d = \exp(\nu) \exp(\zeta + \pi w_{t+1}) \frac{e(x_{t+1})}{e(x_t)} \quad (8)$$

which is the holding period return to a security that pays off the pricing factor e over *any* horizon j . This pricing factor is the principal eigenfunction of result 2.1. 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.⁹

⁸The risk aversion parameter is assumed to be 20. This large value is used to amplify the effects of risk. We will say more about this parameter subsequently.

⁹In a paper presented at the same NBER Summer Institute (2004) as our paper, Lettau and Wachter (2005) also considered the decomposition of one-period returns into the holding-period returns of the compo-

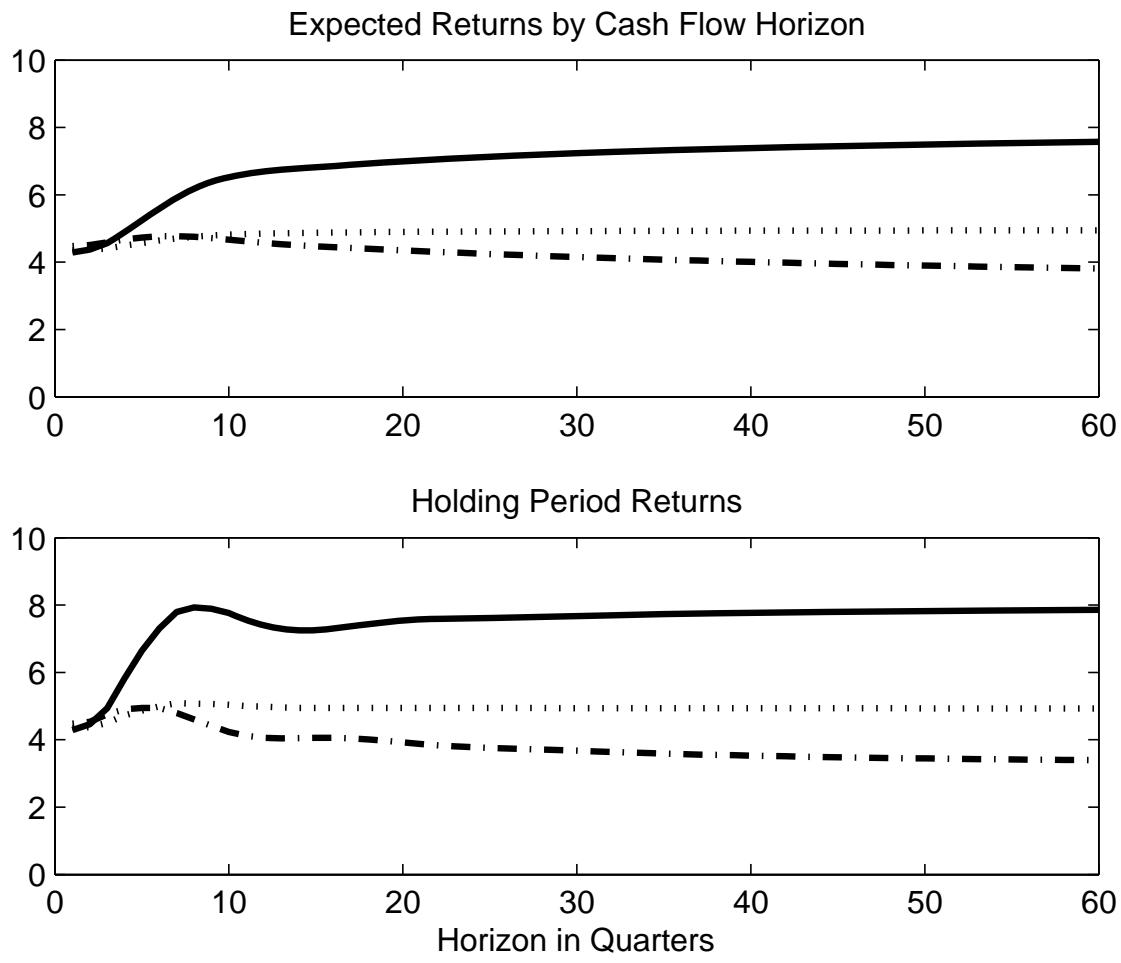


Figure 2: The plots given by a solid line are for portfolio 5, a value portfolio, the plots given by a dash-dotted line are for portfolio 1, a growth portfolio, and the plots given by the dotted line are for the market portfolio. Rates of return are given in annual percentage rates. The expected rates of return in the top panel are computed over the investment horizon, and the returns in the bottom panel are one-quarter holding period returns on payoffs at the respective horizons.

The logarithm of the return R_{t+1}^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 principal eigenfunction and its associated eigenvalue. 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.

The bottom panel of figure 2 depicts this decomposition of expected returns for the growth, value and market 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 values in these figures are also good approximations to the entire figure after about two years.

2.5 Other models of the stochastic discount factor

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.¹⁰ While Bansal and Lehmann (1997) focus on logarithmic bounds on stochastic discount factors, the long-term risk-return tradeoff of theorem 1 is invariant across a similar variety of models. Such models differ only in their transient implications for valuation.¹¹ Thus the limiting one-period return (8) will be altered by the inclusion of a common state dependent contribution, but the long-term tradeoff remains the same. See Hansen (2008) for a more extensive discussion of these issues.

3 Pricing under recursive utility

In what follows we develop more fully a recursive utility model of investor preferences. As we will illustrate, this model provides an important role for the intertemporal composition of consumption risk for valuation at short as well as long horizons. The resulting specification of the stochastic discount factor gives us a tractable characterization of long-run implications that is rich enough to imply differences in expected returns as they relate to long-run risk.

next cash flows. Their focus is different because they feature the decomposition of a single aggregate return in a model in which the expected holding-period returns are larger for shorter horizons than longer ones. While they build a simple model of portfolio cash flows, they do not match their model to actual cash flows from portfolios. In contrast our focus is on the observed behavior of the cash flows, and we find interesting differences in the return decompositions for the alternative book-to-market portfolios even without the aggregate decomposition they advocate.

¹⁰Their analysis extends to some recent models of social externalities or preference shocks such as Campbell and Cochrane (1999) and Menzly, Santos, and Veronesi (2004).

¹¹Formally, the eigenvalue of result 2.1 remains the same but the eigenfunction is altered.

3.1 Preferences and the stochastic discount factor

The time t utility of the representative consumer is given by the CES recursion:

$$V_t = [(1 - \beta) (C_t)^{1-\rho} + \beta [\mathcal{R}_t(V_{t+1})]^{1-\rho}]^{\frac{1}{1-\rho}}. \quad (9)$$

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:

$$\mathcal{R}_t(V_{t+1}) \equiv (E [(V_{t+1})^{1-\gamma} | \mathcal{F}_t])^{\frac{1}{1-\gamma}}$$

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 .

This specification of investor 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 γ . 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. 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:

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

(For instance, see Hansen, Heaton, Lee, and Roussanov (2006).) 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 γ 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 (9). 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 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 to 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 logarithmic intertemporal preferences substantially simplify 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. We then explore sensitivity of pricing implications as we change the elasticity of intertemporal substitution.¹² For example, Campbell (1996) argues for less elasticity than the log case and Bansal and Yaron (2004) argue for more.

To calculate the continuation value, first scale V_t in (9) by consumption:

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

Next let v_t denote the logarithm of the ratio of the continuation value to consumption, and let c_t denote the logarithm of consumption and rewrite recursion (9) as:

$$v_t = \frac{1}{1-\rho} \log ((1-\beta) + \beta \exp [(1-\rho) \mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t)]), \quad (10)$$

where \mathcal{Q}_t is:

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

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 a 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.

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

The $\rho = 1$ limit in recursion (10) is:

$$\begin{aligned} v_t &= \beta Q_t(v_{t+1} + c_{t+1} - c_t) \\ &= \frac{\beta}{1-\gamma} \log E(\exp[(1-\gamma)(v_{t+1} + c_{t+1} - c_t)] | \mathcal{F}_t). \end{aligned} \quad (11)$$

Recursion (11) 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:

$$\begin{aligned} U_v &\equiv \beta U_c(I - \beta G)^{-1}, \\ \mu_v &\equiv \frac{\beta}{1-\beta} \left[\mu_c + \frac{(1-\gamma)}{2} |\lambda_0 + U_v H|^2 \right]. \end{aligned}$$

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 $\lambda_0 + U_v H$ is the shock exposure vector of the continuation value for consumption.

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

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

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

$$E \left[\frac{(V_{t+1})^{1-\gamma}}{[\mathcal{R}_t(V_{t+1})]^{1-\gamma}} | \mathcal{F}_t \right] = E \left(\frac{(V_{t+1})^{1-\gamma}}{E[(V_{t+1})^{1-\gamma}] | \mathcal{F}_t} \right) = 1. \quad (12)$$

As we asserted in section 2 the stochastic discount factor is a linear function of the lagged state and the shock vector w_{t+1} as in (3). The coefficients of this function are:

$$\begin{aligned} \mu_s &= \log \beta - \mu_c - \frac{(1-\gamma)^2 |\lambda_0 + U_v H|^2}{2} \\ U_s &= -U_c \\ \xi_0 &= -\lambda_0 + (1-\gamma) (\lambda_0 + U_v H). \end{aligned}$$

The matrix of coefficients ξ_0 weights the shock vector. From the consumption dynamics (2), the initial response of consumption to a date $t+1$ shock w_{t+1} is $\lambda_0 w_{t+1}$ and the response of c_{t+j} for $j > 1$ is $\lambda_j \equiv U_c G^{j-1} H$. The discounted (by the subjective rate of discount) value of these responses is:

$$\lambda(\beta) \equiv \lambda_0 + \beta U_c (I - \beta G)^{-1} H = \lambda_0 + U_v H.$$

Thus $\xi_0 = -\lambda_0 + (1-\gamma)\lambda(\beta)$ as claimed in section 2. The term $\lambda(\beta)$ is a target of measurement even for one-period pricing. This is the impact of predictability in consumption growth that is featured in Bansal and Yaron (2004). It reflects the intertemporal composition of consumption risk, and creates an important measurement challenge in implementation. Long-run risk can have important implications even for one-period pricing. The impact persists over longer-horizons as is conveyed by the limiting pricing formulas of section 2.

Since the term (12) in the one-period period stochastic discount factor is positive, and it has conditional expectation equal to unity, it can be thought of as distorting the probability distribution. The presence of this distortion is indicative of a rather different interpretation of the parameter γ . Instead of incremental risk aversion applied to continuation utilities, Anderson, Hansen, and Sargent (2003) argue that γ may reflect investors' concerns about not knowing the precise riskiness that they confront in the marketplace. In this case the original probability model is viewed as a statistical approximation and investors are concerned that the model may be misspecified. Although we continue to refer to γ 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 γ . Changing the interpretation of γ alters what might be viewed as reasonable values of this parameter.

To be concrete, under the alternative interpretation suggested by Anderson, Hansen, and Sargent (2003), $(\gamma - 1)\lambda(\beta)$ is the contribution to the induced prices 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 out easily with historical data. This lead us to ask how large is $-\gamma(1-\beta)\lambda(\beta)$ in a statistical sense. To gauge this, when $\gamma = 10$ and $|(1-\beta)\lambda(\beta)| = .01$ 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 γ changes this probability to about 6%.¹³ In this sense $\gamma = 10$ is in an interesting range of statistical ambiguity while $\gamma = 20$ leads to an alternative model that is 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 γ in our empirical work, perhaps part of the large choice of γ can be ascribed to statistical ambiguity on the part of investors.

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. Hansen, Heaton, Lee, and Roussanov (2006) 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

¹³These 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.

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 $Ds_{t+1,t}^1$ is the derivative of the discount factor with respect to ρ evaluated at $\rho = 1$. This derivative is a stochastic process that is a quadratic function of the shock vector w_{t+1} . The approximation of the discount factor allows us to calculate the derivative of the asymptotic rate of return for any cash flow process. The details of the calculation and implementation of this approximation are given in appendix A.

4 Long-run consumption risk

We now describe the estimation of the consumption dynamics needed to characterize how risk 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 dynamics. For our initial model we let log consumption be the first element of y_t and log corporate earnings be the second element. Our use of corporate earnings in the VAR is important as a predictor of consumption and as an additional source of aggregate risk.¹⁴ Changes in corporate earnings signal changes in aggregate productivity which will have long-run consequences for consumption.

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:

$$A_0 y_t + A_1 y_{t-1} + A_2 y_{t-2} + \dots + A_\ell y_{t-\ell} + B_0 = w_t, \quad (13)$$

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:

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

By inverting the matrix polynomial $A(z)$ for the autoregressive representation, we obtain the power series expansion for the moving-average coefficients. We are interested in a specification in which $A(z)$ is nonsingular for $|z| < 1$. 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:

$$\lambda(\beta) = (1 - \beta) u_c A(\beta)^{-1} \quad (14)$$

as required by our model. When $A(1)$ is singular, there are unit root components to the times series. While $A(1)$ cannot necessarily be inverted, (14) will still have a well defined

¹⁴Whereas 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.

limit as β tends to one provided that the limiting response of the logarithm of consumption to a shock is finite.

Following Hansen, Heaton, and Li (2005), we impose two restrictions on the matrix $A(1)$. We impose a unit root in consumption and earnings, but we restrict these series to grow together.¹⁵ 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, we estimate the model as a vector autoregression in the first-difference of the log consumption and the difference between log earnings and log consumption. In section 5.5 we explore other assumptions about growth.

As is known from the literature on structural VARs, ideally a macroeconomic model assigns economic interesting “labels” to shocks by imposing *a priori* restrictions to make this assignment. Macroeconomic labeling is not featured in our analysis. Instead we use two identification schemes depending on the purpose.

Our first assignment simplifies the representation and interpretation of our results. Given our focus on the analysis of long-run risk, we normalize the shocks so that only one shock has long-run consequences. We achieve this by first constructing a temporary shock to consumption that has no long-run impact on consumption and corporate earnings. We construct the second shock so that it is uncorrelated with the first one and has equal permanent effects on consumption and earnings.¹⁶ By design, exposure to this shock dominates long-run valuation.¹⁷

For reporting the accuracy of measurements, we use a recursive scheme to identify shocks. This makes the Bayesian method of inference for impulse response functions proposed by Sims and Zha (1999) and Zha (1999) directly applicable. Under this scheme the second shock is restricted not to influence the growth rate of consumption in the initial period. The likelihood function for the two equation system factors into two separate pieces, one coming from the consumption growth equation and the other from an equation with the log of the ratio of corporate earnings to consumption on the left-hand side and the contemporaneous growth rate of consumption on the right-hand side along with the appropriate number of lagged values of each of the variables. We impose separable noninformative priors for the regression coefficients conditioned on the regression error variances and on the marginals for the regression error variances as in Box and Tiao (1973). These “priors” are chosen for convenience, but they give us a simple way to depict the uncertainty associated with the estimates. We use these priors in computing posterior distributions for the short-run and long-run responses of the permanent shock to consumption. We consider only the region of the posterior distribution for which the transformed VAR system (expressed in terms of consumption growth rates and the difference in logarithms of consumption and corporate earnings) has stable dynamics.¹⁸

¹⁵Formally, we restrict $A(1) = A(1) = \alpha [1 \ -1]$ where the column vector α is freely estimated.

¹⁶This construction is much in the same spirit as Blanchard and Quah (1989).

¹⁷This is formally true in the power utility model and approximately true in the recursively utility model.

¹⁸We used only simulation draws for which the absolute values of the eigenvalues of the transformed VAR were all less than .999.

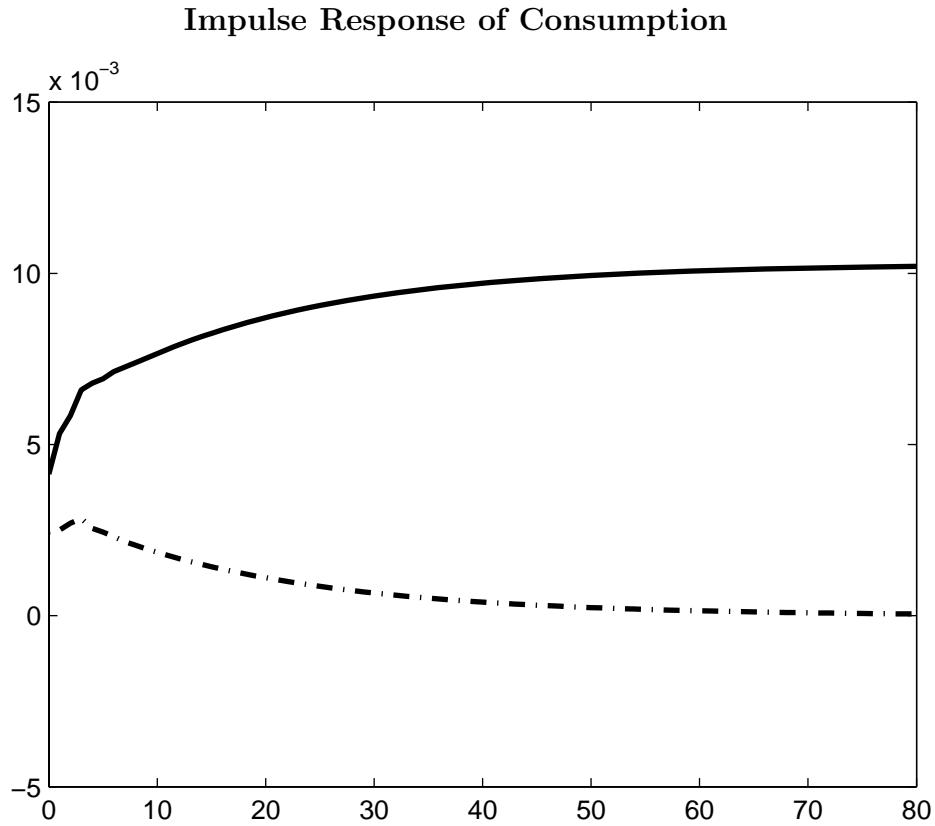


Figure 3: Impulse responses of consumption to shocks implied by bivariate VARs where consumption and earnings are assumed to be cointegrated. The solid line depicts the impulse response to a permanent shock. The dash-dotted line depicts the impulse response to a temporary shock. Each shock is given a unit impulse. Responses are given at quarterly intervals.

We use aggregate consumption of nondurables and services taken from the National Income and Product Accounts as our measure of consumption. This measure is quarterly from 1947 Q1 to 2005 Q4, is in real terms and is seasonally adjusted. We measure 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 responses of future consumption to the underlying shocks affects both short-run and long-run risk prices. For this reason, the predictable responses of consumption to shocks identified by the VAR with cointegration, affect risk prices at all horizons. The estimated responses are subject to statistical error especially over the long run. To compute posterior distributions, we imposed the priors described previously on each equation in the VAR system and simulated the posterior histograms for the parameter estimates. While these priors are chosen for convenience, they give us a simple way to depict the statistical uncertainty associated with the estimates. We display the implied posterior distributions for the short-run and long-run responses in figure 4. The magnitude of the long-run response is $|\lambda(1)|$ and the magnitude of the short-run response is $|\lambda_0|$. The vertical lines in each plot are located at the posterior medians.

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 posterior probabilities are given to substantially larger responses.¹⁹ Thus, from the standpoint of statistical accuracy, the long-run response could be more than double that of the immediate consumption response. When $\beta \approx 1$, long-run risk prices are approximately equal to $\gamma\lambda(1)$. These prices are expressed as required additions to expected rates of returns for an exposure to a shock with a unit standard deviation. Depending on the choice of γ , long-run risk prices could be quite substantial when accounting for statistical uncertainty.

4.2 Specification sensitivity

In the long-run risk model of Bansal and Yaron (2004), low-frequency shocks to consumption are driven by an unobserved latent variable. In contrast we identify the long-run impact of shocks using corporate earnings and the assumption that these earnings are cointegrated with consumption. Cointegration plays an important role both in identifying the long-run impact of the permanent shock depicted in figure 3 and in determining the temporal pattern of the responses to both shocks. The impact of these identified patterns on prices is given by $\gamma\lambda(\beta)$ when $\rho = 1$.

To assess the importance of the assumption of cointegration, in figure 5 we depict $|\lambda(\beta)|$ as a function of β for the baseline model and for two alternative specifications of the relationship between consumption and corporate earnings: a VAR estimated in log-levels

¹⁹The 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.

Posterior Distributions for Immediate and Long-run Responses

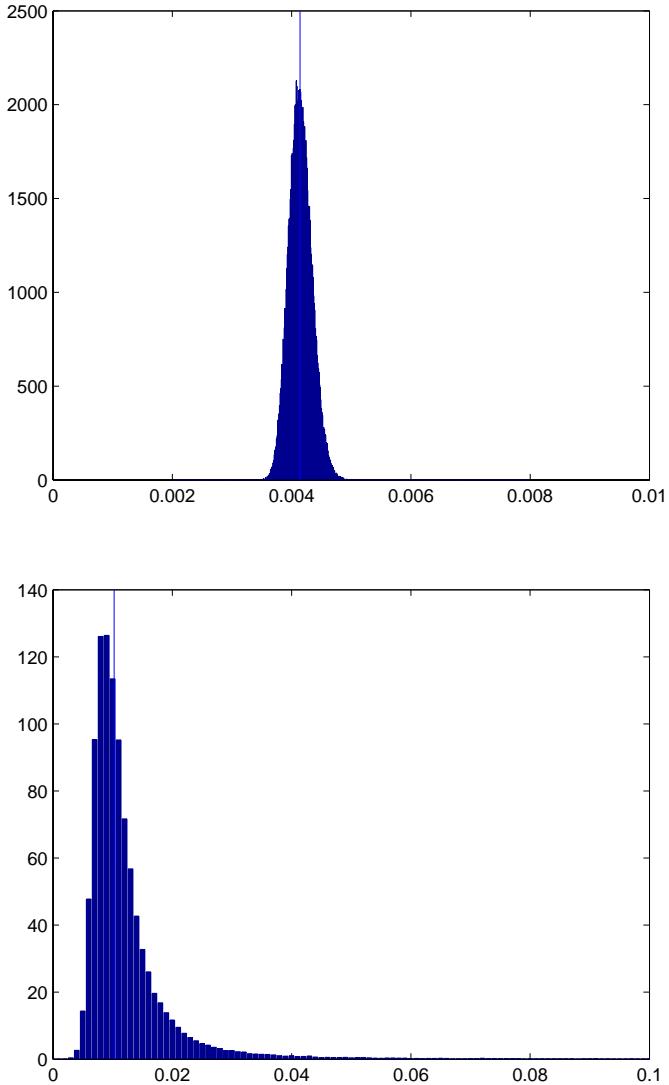


Figure 4: Top figure gives the posterior histogram for the magnitude $|\lambda_0|$ of immediate response of consumption to shocks. The bottom figure gives posterior histogram for the magnitude $|\lambda(1)|$ of the long-run response of consumption to the permanent shock. The vertical axis is constructed so that the histograms integrate to unity. Vertical lines are located at the posterior medians.

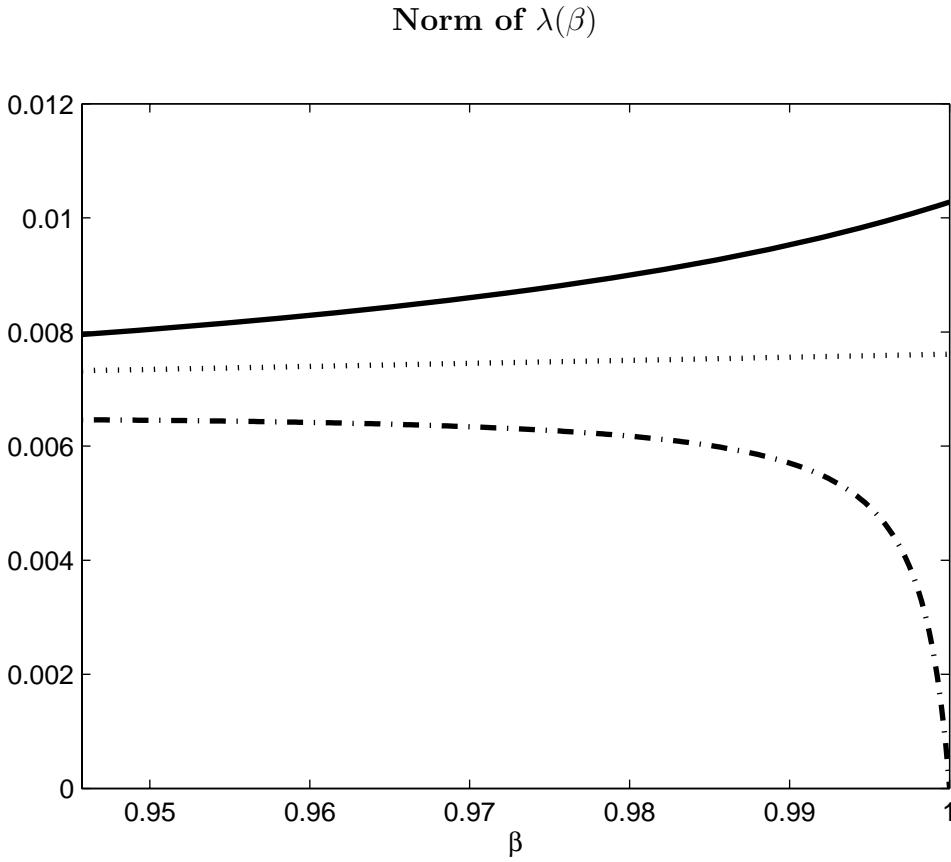


Figure 5: Norm of $\lambda(\beta)$ for different values of β and three different VAR systems. The solid line depicts $\lambda(\beta)$ for the cointegrated system, the dash-dotted line for the the system without cointegration, and the dotted line for the specification with first differences used for all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

and a VAR estimated in first differences. The log-level VAR is estimated to be stable, and as a consequence the implied $|\lambda(1)| = 0$. 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 $|\lambda(\beta)|$ 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 a larger value of $|\lambda(\beta)|$ and hence larger risk prices for any given value of γ .

4.3 Implications for pricing aggregate consumption

Although consumption is not equal to dividends it is still instructive to examine the price of aggregate risk as represented by a claim on aggregate consumption. In this case π is equal to

the long-run exposure of consumption to the two shocks: $\lambda(1)$. With recursive preferences and $\rho = 1$, the excess of the asymptotic return to the consumption claim over the riskless return is:

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

The expected excess return is essentially proportional to γ due to the dependence of the risk-free benchmark on γ when β is close to one.

Even in the long-run, the consumption claim is not very risky. The point estimates of the VAR system imply that $\lambda(1) \cdot \lambda(1) = 0.0001$. Hence when β is near unity, increases in γ has a very small impact on the expected excess return to the consumption claim. For example, even when $\gamma = 10$ the expected excess return, in annual units, is .4% ($= 10 \times 0.0001 \times 4$).

Notice, however that because of significant sampling uncertainty this excess return could be much larger. For example, assuming that $\gamma = 10$, the posterior distribution illustrated in figure 4 implies that there is a 10% chance that the long-run excess return to holding the consumption claim could be larger than 1.65% annually. By way of contrast, the same posterior distribution for the short-run excess return (given by $\gamma\lambda_0 \cdot \lambda(1)$) has 10% of its mass above a much smaller value of 0.4%. A notable price for long-run risk cannot be ruled out by this data once a large value of γ is assumed.

5 Long-run cash flow risk

We now ask whether exposure to long-run risk can help to explain differences in returns for particular portfolios of stocks familiar from financial economics. 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-term cash flow risk exposure is priced.

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 and is suggested by figure 1. It is germane to our empirical application because the sorting method we use in constructing portfolios can induce permanent differences in dividend growth. 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.²⁰ Write:

$$\begin{aligned} d_{t+1} - d_t &= \mu_d + U_d x_t + \iota_0 w_{t+1} \\ &= \mu_d + \iota(1) w_{t+1} - U_d^* x_{t+1} + U_d^* x_t \end{aligned}$$

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 .

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 separately 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 we restrict consumption and corporate earnings to be jointly autonomous. 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 figure 1, these series have important low frequency movements relative to consumption. Cash flow models that feature substantial

²⁰A martingale decomposition is commonly used in establishing central limit approximations (*e.g.* see Gordin (1969) or Hall and Heyde (1980)). For a scalar linear time series, it coincides with the decomposition of Beveridge and Nelson (1981).

mean reversion or stochastically stable shares relative to aggregate consumption are poor descriptions of these data.

Formally, we append a dividend equation:

$$A_0^*y_t^* + A_1^*y_{t-1}^* + A_2^*y_{t-2}^* + \dots + A_\ell^*y_{t-\ell}^* + B_0^* = w_t^*, \quad (15)$$

to equation system (13). In this equation the vector of inputs is

$$y_t^* \equiv \begin{bmatrix} y_t \\ d_t \end{bmatrix}$$

and the shock w_t^* is scalar with mean zero and unit variance. This shock is uncorrelated with the shock w_t that enters (13). The third entry of A_0^* is normalized to be positive. We refer to (15) as the dividend equation, and the shock w_t^* as the dividend shock. As in our previous estimation, we set $\ell = 5$. Initially, we presume that this additional shock has a permanent impact on dividends,²¹ but subsequently we will explore sensitivity of our risk measures to alternative specifications of long-run stochastic growth in the cash flows. We estimated the VAR using the transformed variables: $(c_t - c_{t-1})$, $(e_t - c_t)$, $(d_t - d_{t-1})$ to induce stationarity 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 consider 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 build portfolios by sorting stocks according to their book-to-market values. A coarse sort into 5 portfolios makes our analysis tractable. Our market portfolio is the CRSP value-weighted portfolio.

Summary statistics for returns on 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. Book-to-market is an index of “growth” versus “value”. Stocks with low book-to-market values are growth stocks because of their high market value relative to book value. Conversely stocks with high values of book-to-market are value stocks. Portfolios 1 through 5 are therefore ordered from growth to value. We call portfolio 1 the growth portfolio and portfolio 5 the value portfolio.

Notice that the expected one-period rates of return increase from growth to value stocks. The difference in expected returns to holding the value portfolio versus the growth portfolio is substantial. It is well documented that the differences in expected returns across the portfolios cannot be explained by differences in the contemporaneous covariances of the returns with consumption growth.

²¹This imposes the linear restriction:

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

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
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: “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 given by the first differences in log consumption, the difference between log consumption and log corporate earnings and 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. “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.

In this section we are particularly interested in the behavior of the cash flows from the portfolios and how they are priced. The constructed cash flow processes accommodate changes in the classification of the underlying stocks and depend on the relative prices of the new and old stocks that move in and out of 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 Heaton (1995) and Bansal, Dittmar, and Lundblad (2005). The geometric averaging induces a transient distortion to our cash flows, but does not distort the long run stochastic behavior.

5.4 Investor preferences and intertemporal pricing

Our measurements of the risk prices depend on parameters that govern investor preferences: the parameter ρ that governs intertemporal substitution, the parameter γ that contributes to risk aversion, and the subjective discount factor β . We now explore how the intertemporal pricing implications differ as we change these parameters. In our analysis aggregate consumption is held fixed at the process we estimated from historical data. In a model with explicit production, the dynamics for consumption would be altered as we change investor preferences, in ways that may be empirically implausible. Consideration of production is interesting because it may imply additional model implications. It is still revealing, however, to explore how prices and risk premia are altered for a given consumption process as in the theoretical analysis of Lucas (1978).

5.4.1 Risk and return for alternative horizons

We first consider expected returns to holding claims to portfolio cash flows at different horizons. As in figure 2 and section 2.3, we take logarithms of the expected returns and scale them by horizon. To study sensitivity to changes in the parameters that govern investor preferences, we find it convenient to split the results into two parts: risk-free returns by horizon (figure 6) and expected excess returns by horizon (figure 7). The results in both figures are computed assuming that the Markov state is set to its unconditional mean.

These figures display the temporal counterpart to a pricing insight in Epstein and Zin (1989b). Even large changes in the risk aversion parameter γ have only a modest impact on implied risk-free returns across the entire term structure (figure 6). In contrast the parameter ρ (the reciprocal of the elasticity of substitution) has a big impact on risk-free returns. The impact of ρ could be offset or enhanced by changing the subjective discount factor β . In particular, changing β from 0.97 to 0.99 on an annual basis shifts the curves down by about

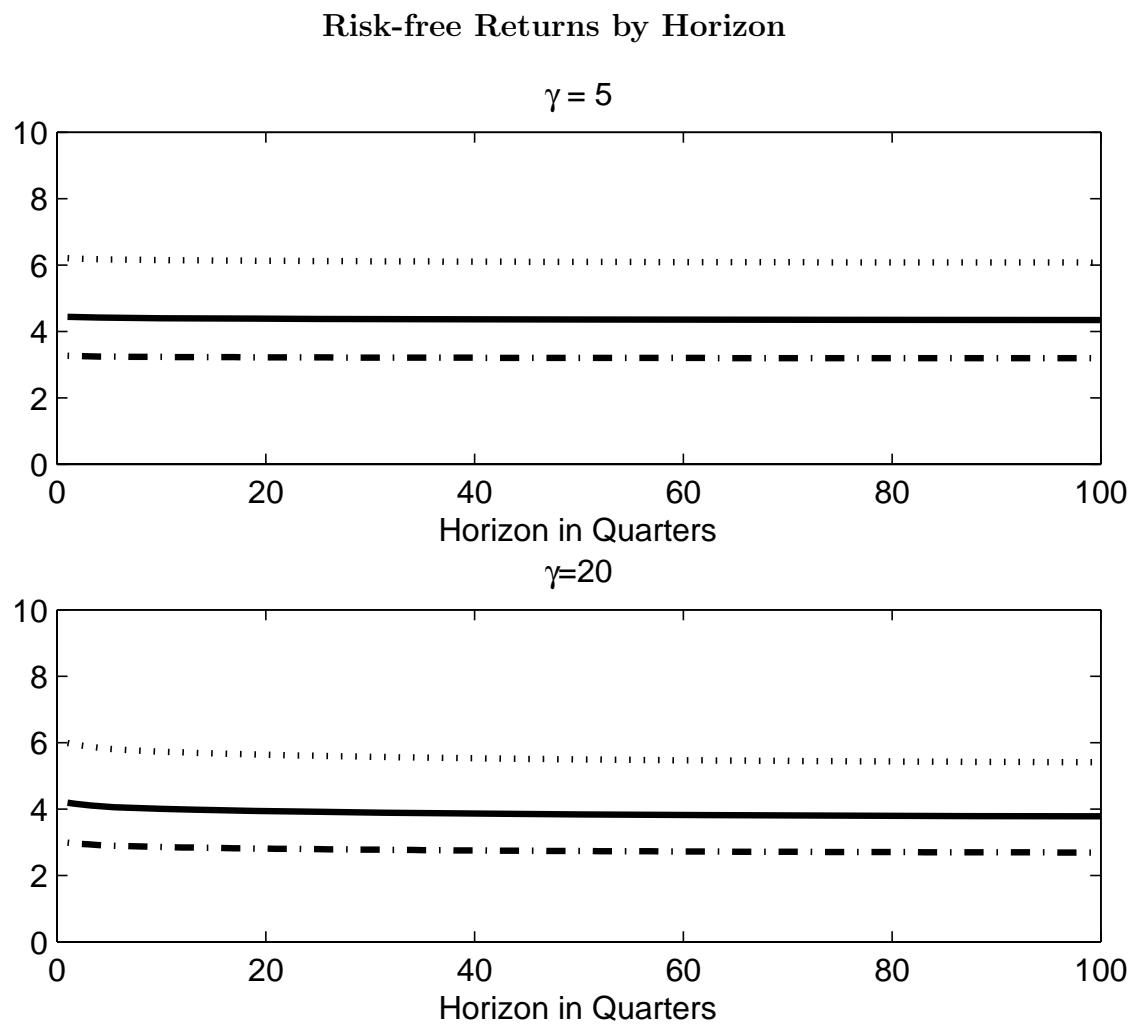


Figure 6: Logarithm of risk-free returns for alternative horizons. The dash-dotted line assumes $\rho = 2/3$, the solid line assumes $\rho = 1$ and the dotted line assumes $\rho = 3/2$. Returns are in annualized percentages. β is $0.97^{1/4}$.

two percent. Such a shift can be defended given the seemingly low short-term risk free rate but can have adverse consequences for dividend-price ratios as we will discuss in the next subsection. The real term structure in this model is rather “boring”, but papers by Piazzesi and Schneider (2007) and Kleshchelski and Vincent (2007) suggest ways to enrich this model to confront term structure evidence. Our primary interest in this paper is in the pricing of risk over longer horizons.

Figure 7 shows that the parameter γ has a substantial impact on the expected excess returns across different horizons as well as on the long-horizon limits. Changing ρ from unity to $3/2$ or $2/3$ has only a minor impact on the risk premia across the different horizons. The impact is only visible for very large values of γ . While our model solution is valid for arbitrarily large values of γ , it is local in ρ , which discourages us from exploring more extreme values of ρ .

The predicted risk premia obtained by holding the cash flows of portfolio 1 at alternative horizons are all close to zero. Moreover, the expected excess returns to holding these cash flows vary only slightly with horizon regardless of the value of γ . This occurs because portfolio 1 has low cash flow covariation with consumption at all horizons.

The parameter γ has a substantial impact on the predicted excess returns to holding the cash flows of portfolio 5 (a portfolio of high book-to-market stocks). These cash flows have much different exposure to consumption risk across horizons. The short-run exposure is similar to that of portfolio 1 but the long-run exposure is much higher. As in other studies, to magnify the importance of these differences we must assume that risk aversion is relatively high. For example when $\gamma = 20$ expected excess returns rise dramatically with horizon for portfolio 5. Recall that this portfolio has observed average returns that are quite high compared to those of portfolio 1. Figure 7 provides a possible explanation of the observed fact: portfolio 5 has cash flows with substantial exposure to consumption risk in the long run.

5.4.2 Value-based measures of duration

Up until now, we have focused on the return implications of the cash flows. Changes in preference assumptions also have implications for the contributions of future cash flows to current-period values. Recall from the Gordon growth model that it is the difference between the rate of return and the rate of cash-flow growth that determines the price-dividend ratio. The discrete-time counterpart to this formula states:

$$\frac{\text{price}}{\text{dividend}} = \frac{\exp(\text{growth rate})}{\exp(\text{return rate}) - \exp(\text{growth rate})}$$

As in the Gordon growth model, the difference between the long-term rates of return and growth gives a limiting measure of the *duration* of a cash flow. This difference is the parameter ν in result 2.1. When ν is small, cash flows far into the future remain important contributors to current-period values.²² As we argued in subsection 2.4, this duration mea-

²²Formal application of the Gordon growth model in this context gives the price-dividend ratio for a security with a transient component to the cash flow that is proportional the function e in result 2.1.

Expected Excess Returns by Horizon

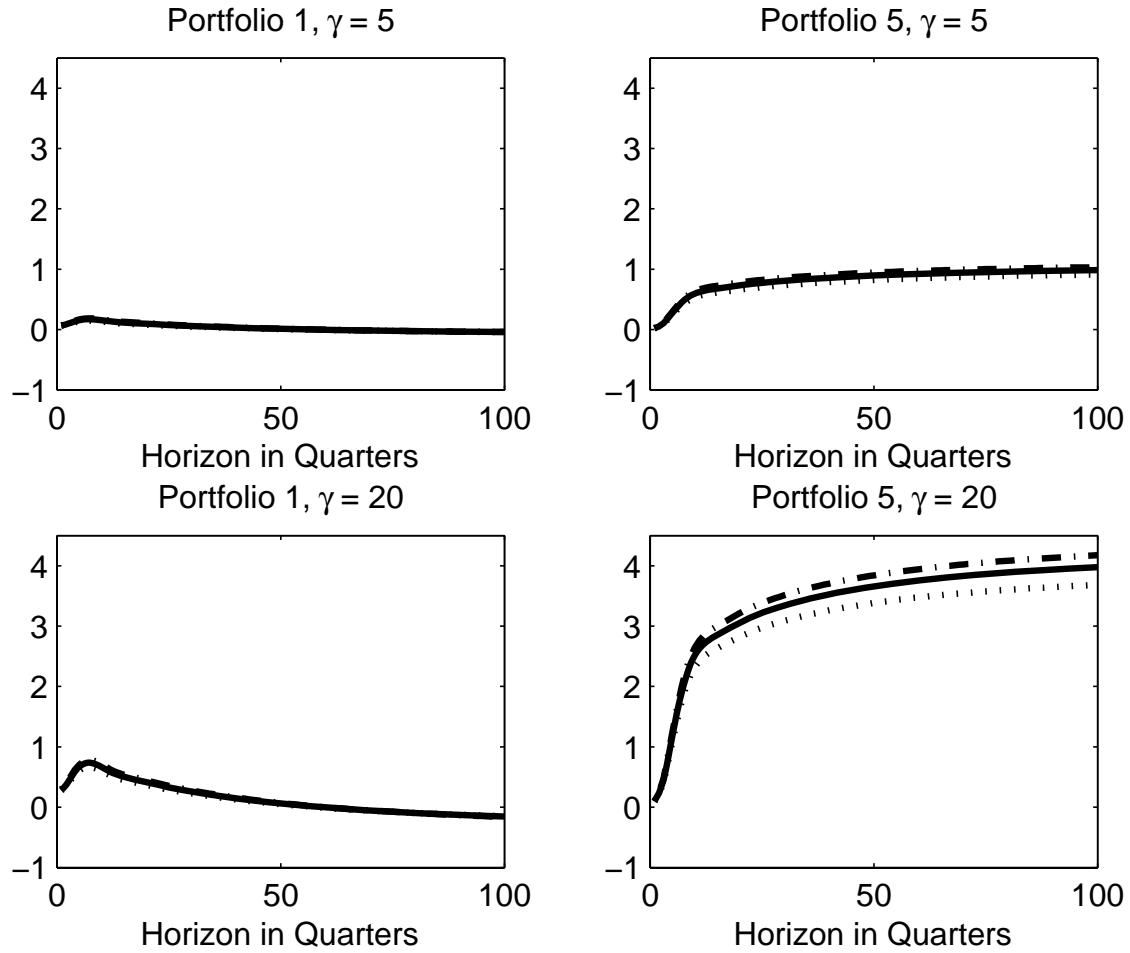


Figure 7: Logarithms of the ratio of expected returns to holding cash flows from portfolios 1 and 5 at different horizons to risk-free counterparts dividend by the horizon. The dash-dotted line assumes that $\rho = 2/3$, and the dotted line assumes that $\rho = 3/2$. Expected excess returns are in annualized percentages. β is set to $0.97^{1/4}$.

sure incorporates an adjustment for growth-rate risk exposure. To characterize the role of investor preferences, Table 2 splits ν into the two components that enter the Gordon growth model: a rate of return and a rate of dividend growth.

As we showed in figure 1 and in table 2, the long-run growth rates of the portfolios are substantially different. Low book-to-market portfolios have low limiting growth rates, but they also have high price-dividend ratios (see table 1). Naive application of the Gordon growth model with a common rate of return for all portfolios would suggest that the low book-to-market portfolios should have low price-dividend ratios. Of course these portfolios have different exposures to risk and hence measures of duration are potentially greatly affected by the price of this risk and not just by differences in cash-flow growth rates. This leads us to explore when the model-implied rates of return can offset the growth-rate differences.

For the low book-to-market portfolios to have comparable measures of duration relative to the high book-to-market portfolios, their limiting rates of return must be substantially lower. This can be attained by making γ sizeable. Changing ρ (as reflected by the derivatives) has an important impact on the value-based measures of duration, but this impact is almost the same across the various portfolios. Similarly, changing the subjective discount factor β alters the rates of return in the same way across the portfolios. Differential rates of return are achieved by making γ large. Notice that for $\rho = 1$ and $\gamma = 5$ our measure of duration for portfolio five is barely positive. Even small increases in β or reductions in ρ make this measure negative implying an infinite price-dividend ratio. This tension is less severe when we make γ larger.

5.5 Measurement accuracy of long-run risk prices

So far our discussion in this section abstracts from errors in estimating the cash flow growth rates and risk exposure. The results in table 2 are likely to be fragile from the standpoint of measurement accuracy, but we include them because they illustrate some important consequences of changes in investor preferences. In the next two subsections, we address formally estimation accuracy and sensitivity to specification in the next two subsections as they relate to risk prices.

5.5.1 Estimation uncertainty

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

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

We now investigate the statistical accuracy of $\lambda(1) \cdot \pi$ for the five portfolios, and for the difference between portfolios 1 and 5. The vector π is measured using $\iota(1)$. In table 3 we report the posterior distribution for $\lambda(1) \cdot \pi$ computed using the same Bayesian approach we described previously, except that we now include a third equation in the VAR. We scale the values of $\lambda(1) \cdot \pi$ by 400 just as we did when reporting predicted annualized average returns in percentages.

Limiting Cash Flow Discount and Growth Rates

Portfolio	Rate of return	Derivative	Rate of growth
$\gamma = 5$			
1	6.27	3.45	2.11
2	6.42	3.44	1.94
3	7.03	3.35	4.32
4	7.16	3.36	4.02
5	7.42	3.33	7.02
$\gamma = 20$			
1	5.39	3.24	2.11
2	5.98	3.18	1.94
3	8.37	2.75	4.32
4	8.89	2.81	4.02
5	9.92	2.67	7.02

Table 2: Limiting expected rates of return and growth rates for the cash flows of portfolio 1 through 5. The derivative entries in column three are computed with respect to ρ and evaluated at $\rho = 1$. Returns and growth rates are reported in annualized percentages.

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 3: Accuracy of estimates of $\lambda(1) \cdot \pi$ scaled by 400. Quantiles were computed by simulating 100,000 times using “noninformative” priors. 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.

Given that our measurements are based on the implied limiting behavior of the estimated VAR, we expect considerable amount of statistical uncertainty in these risk measures. Nevertheless, there are important differences in the relative risk exposures of portfolios 1 and 5. There is significant evidence that the cash flows of a portfolio of value stocks is riskier than the portfolio of growth stocks.

5.5.2 Specification uncertainty

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 for cash flow dynamics. The specifications we consider allow portfolio dividends to have growth patterns that are distinct from consumption and accommodate the heterogeneity evident in figure 1. These alternative models of dividend growth have antecedents in the prior empirical literature. As we will see, however, the various models have much different implications for the properties of long-run returns predicted by our model.

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:

$$A_0^*y_t^* + A_1^*y_{t-1} + A_2^*y_{t-2} + \dots + A_\ell^*y_{t-\ell} + B_0^* + B_1^*t = w_t^*, \quad (16)$$

where the shock w_t^* is scalar with mean zero and unit variance and uncorrelated with the shock vector w_t that enters (13). The third entry of A_0^* is normalized to be positive. As in our previous estimation, we set $\ell = 5$ and the third column of A_j^* for $j = 0, 1, \dots, \ell$ is restricted to have zeros in its first two entries. In other words, we continue to restrict the dividend process not to Granger-cause either consumption or corporate earnings or equivalently we may view consumption and earnings jointly as an autonomous stochastic system. We also continue to presume that we may configure the first two shocks so that one of them has a common permanent impact on consumption and corporate earnings while the other one has only a transient impact on both series.

The first alternative specification restricts the trend coefficient B_1^* to be equal zero, and is the model used by Hansen, Heaton, and Li (2005). The other coefficients in the last row of equation system (16) are unrestricted in our estimation. Given our interest in measuring long-run risk, we measure the permanent response of dividends to the permanent consumption 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. In contrast to our baseline specification, there is no separate growth component for dividends in this specification. Such a component could emerge in the estimation, but we do not restrict equation system (16) to have a second growth component in contrast to the baseline 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).²³ 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 8. This figure features the responses of the cash flows of portfolios 1 and 5 to a permanent shock to consumption. 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 in ways that are quantitatively important. 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. The limiting response of the baseline model is negative. Notice, however that 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

²³While they do not include macroeconomic predictors of consumption, Bansal, Dittmar, and Lundblad (2005) do allow for dividends to Granger-cause consumption.

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.²⁴

While the use of time trends in the second alternative specification as additional sources of cash flow growth alters our results, use of these results 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 unlikely that these components are literally deterministic time trends known *a priori* to investors. We suspect that the substantially negative limiting response for portfolio 1 is unlikely to be the true limiting measure of how dividends respond to a permanent shock to consumption.²⁵ 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.²⁶

²⁴Bansal, 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. We are using the economic model in a more formal way than the running of cross-sectional regressions, however.

²⁵Sims (1991) and Sims (1996) warn against the use of time trends using conditional likelihood methods because the resulting estimates might over fit the initial time series, ascribing it to a transient component far from the trend line.

²⁶In the specifications we have considered, we have ignored any information for forecasting future consumption that might be contained in asset prices. Since our model of asset pricing implies a strict relationship between cash flow dynamics and prices should be redundant. Prices, however, may reveal additional components to the information set of investor. When we consider an alternative specification of the VAR that includes prices (without imposing the pricing restrictions), we obtain comparable heterogeneity.

Impulse Response Functions for Portfolios One and Five: Alternative Long-Run Specifications

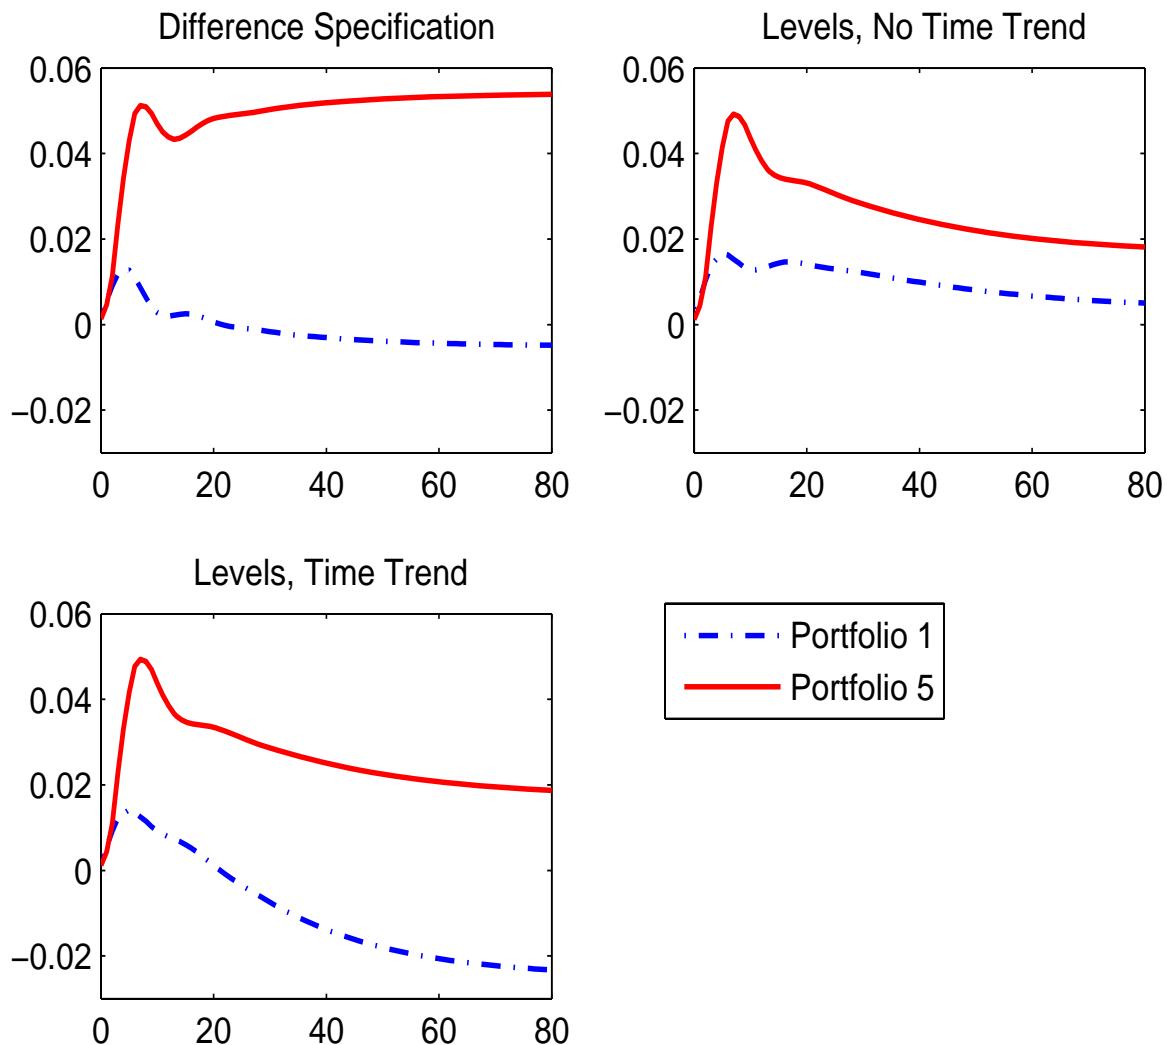


Figure 8: Impulse responses to a permanent shock to consumption of the cash flows to portfolios 1 and 5. Upper left figure is from the first-difference specification used as our baseline model. Upper right figure is from the level specification without time trends. Lower left figure is from the level specification with time trends.

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. Finally, we showed how valuation based measures of the duration of cash flows are linked explicitly to the long-run riskiness of the cash flows.

In our empirical application we showed that the cash flow growth of portfolios of growth stocks 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, investors in our model must be either highly risk averse or highly uncertain about the probability models they confront. Increasing the intertemporal substitution parameter ρ magnifies the differential of the long-run counterpart of price-dividend ratios.

There are three intriguing extensions of our work, a) providing a structural interpretation of shocks, b) exploring alternative models of investor preferences and constraints, and c) introducing time variation in local risk prices. In regards to a), for convenience we used an *ad hoc* VAR model to identify macro economic shocks to be priced. An important 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. In regards to b), other asset models have interesting transient implications for the intertemporal composition of risk prices and exposures. These include 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).) In regards to c), temporal dependence in volatility can be an additional source of long-run risk. Time variation in risk prices can be induced by conditional volatility in stochastic discount factors. It remains to explore implications of stochastic volatility for long-term valuation.

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. 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 the details of the growth configuration of the model. In many cases these details are hard to defend on purely statistical grounds. 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.²⁷

²⁷See Hansen (2007) for a discussion of learning and uncertainty and its impact for short-run pricing.

A Eigenfunction results

In what follows we use the notation:

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

A.1 Eigenfunctions and stability

We follow Hansen and Scheinkman (2007) 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.

We formalize the approximation problem as a change in measure.²⁸ Write the eigenfunction problem as:

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

Then

$$\hat{M}_{t+1,t} = \exp(\nu t) M_{t+1,t} \left[\frac{e(x_{t+1})}{e(x_t)} \right]$$

satisfies:

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

As a consequence $\hat{M}_{t+1,t}$ induces a distorted conditional expectation operator. Recall our solution $e(x) = \exp(\bar{\omega}x)$ 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:

$$\hat{\mu}_w = H' \bar{\omega}' + \pi' + \xi'_0 \quad (17)$$

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

To characterize the limiting behavior, we use this distorted shock distribution in our computations. For instance,

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

Iterating, we obtain:

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

The limit that interests us is:

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

²⁸See Hansen and Scheinkman (2007) for a justification of this change of measure in a continuous-time nonlinear environment. Our analysis is in discrete time and exploits our log-linear formulation.

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

Let q and \hat{q} denote the stationary densities of $\{x_t\}$ under E and the \hat{E} measures. 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 the same covariance matrix, but the nonzero mean for w_t induced the following nonzero mean for x_t :

$$\hat{\mu}_x \equiv (I - G)^{-1} H (H' \bar{\omega}' + \pi' + \xi_0'). \quad (18)$$

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

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

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 $e(x)/z$. The limit of interest is now:

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

To ensure this limit is well defined, we require that the joint process $\{(x_t, z_t)\}$ be stationary, ergodic and aperiodic under the distorted probability distribution and that $f(x_t)z_t/e(x_t)$ have a finite expectation under this distribution.

A.2 Eigenvalue derivative

We compute this derivative using the approach developed in Hansen (2008). 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

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

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

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

Differentiating with respect to ρ :

$$Ds_{t+1,t}^1 = \frac{1}{2} w_{t+1}' \Gamma_0 w_{t+1} + w_{t+1}' \Gamma_1 x_t + \vartheta_0 + \vartheta_1 x_t + \vartheta_2 w_{t+1}.$$

Recall that under the distorted distribution w_{t+1} has a constant mean $\hat{\mu}_w$ conditioned on x_t given by (17) and x_t has a mean $\hat{\mu}_x$ given by (18). Taking expectations under the distorted distribution:

$$\hat{E}(Ds_{t+1,t}^1) = \frac{1}{2} (\hat{\mu}_w)' \Gamma_0 \hat{\mu}_w + \frac{1}{2} \text{trace}(\Gamma_0) + (\hat{\mu}_w)' \Gamma_1 \hat{\mu}_x + \vartheta_0 + \vartheta_1 \hat{\mu}_x + \vartheta_2 \hat{\mu}_w.$$

References

- Alvarez, F. and U. Jermann (2005). Using asset prices to measure the 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 D. Kiku (2007). Cointegration and consumption risks in asset returns. *Review of Financial Economics* forthcoming.
- 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.
- Beveridge, S. and C. R. Nelson (1981). A new approach to the decomposition of economic time series into permanent and transitory components with particular attention to the measurement of the ‘business cycle’. *Journal of Monetary Economics* 7, 151–174.
- Blanchard, O. J. and D. Quah (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review* 79, 655–673.
- Box, G. E. P. and G. C. Tiao (1973). *Bayesian Inference in Statistical Analysis*. Reading, Massachusetts: Addison-Wesley.
- 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. (2007). Beliefs, doubts and learning: Valuing macroeconomic risk. *American Economic Review* 97, 1–30. Richard Ely Lecture.
- Hansen, L. P. (2008). 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., 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. Roussanov (2006). Intertemporal substitution and risk aversion. University of Chicago, Forthcoming in the Handbook of Econometrics.
- Hansen, L. P. and J. Scheinkman (2007). 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.
- Kleshchelski, I. and N. Vincent (2007). Robust equilibrium yield curves. Northwestern University and HEC Montreal.
- 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.
- Piazzesi, M. and M. Schneider (2007). Equilibrium yield curves. *NBER Macroeconomics Annual* 21, 389–442.
- 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.