

Consumption Strikes Back?: Measuring Long Run Risk *

Lars Peter Hansen[†]

John C. Heaton[‡]

Nan Li[§]

March 18, 2005

*We thank Fernando Alvarez, Ravi Bansal, John Cochrane, Jonathan Parker, Tano Santos, Chris Sims and Pietro Veronesi for valuable comments. We also thank workshop participants at Chicago, Emory, INSEAD, LBS, MIT, Montreal, NBER and Northwestern. Hansen gratefully acknowledges support from the National Science Foundation, Heaton from the Center for Research in Securities Prices, and Li from the Olin Foundation.

[†]Department of Economics, University of Chicago and NBER

[‡]Graduate School of Business, University of Chicago and NBER

[§]National University of Singapore

1 Introduction

Economically grounded models of asset pricing feature a role for information and risk as a device for explaining heterogeneity in asset values, including dividend price ratios and returns. Much of the applied time series research in the past has questioned the empirical relevance of risk identified from economic aggregates such as consumption. Recently, however, this has changed. Several investigations have documented at least a qualitative role for uncertainty about consumption and other macroeconomic aggregates in explaining heterogeneity in asset values. Much of the supporting evidence relies on statistical models of low frequency or long run growth rate uncertainty in consumption and financial cash flows.

Our paper aims to explore the resulting empirical challenges. How sensitive are risk-measures to details in the specification of the time series evolution? How accurately can we measure these components? When should we expect these components to play a fundamental role in valuation? To address these questions we are compelled to address formally the role of specification, measurement and pricing theory in quantifying risk.

We find it most valuable to address these issues in the context of a well specified, albeit highly stylized, model. Following Epstein and Zin (1989b), Weil (1990), Tallarini (1998), Bansal and Yaron (2004), and many others we use a recursive utility framework of Kreps and Porteus (1978). For these preferences, the intertemporal composition of risk matters to the decision maker. Risk cannot simply be reduced or averaged out. Instead the timing of when information is revealed about intertemporal consumption lotteries matters in the implied preference ordering. As emphasized by Epstein and Zin (1989b), these preferences also offer a convenient and appealing way to break the preference link between risk aversion and intertemporal substitution. As featured by Bansal and Yaron (2004) long run risk components can amplify the risk premia in security market prices. Using these preferences to model investors in equilibrium pricing models, requires measuring how consumption risk unfolds over time.

While we focus on a recursive utility specification, the intertemporal timing of risk matters in other models as well, including models that feature habit persistence (*e.g.* Constantinides (1990), Heaton (1995), and Sundaresan (1989)) and models of staggered decision-making (*e.g.* see Lynch (1996) and Gabaix and Laibson (2002).)

In this paper we study the intertemporal composition risk using (log) linear vector autoregressive (VAR) models of consumption and cash flows. These models are designed to accommodate transient dynamics in a flexible way. They are convenient time series models that allow us to explore the statistical accuracy of the risk measurements along with the sensitivity of these measurements to changes in the model specification. By using an explicit model of investor preferences, we are able to evaluate the economic importance of the statistical and model specification uncertainty confronted by an econometrician. The focus on long-run risk stretches these methods beyond their ability to capture transient dynamics. Risk measurements can depend on the extent to which the growth properties of alternative time series such as consumption and financial cash flows are explicitly linked. Growth restrictions can alter the implied riskiness over long time horizons in ways that are economi-

cally important. As a consequence, the flexibility of vector autoregressive methods does not necessarily extend to the measuring long run risk. This paper provides an examination of the sensitivity to estimation and model uncertainty confronted in econometrics.

For our sensitivity analysis of VAR's to have economic meaning, we are led to suggest and apply two other methods in our analysis. The first method is an expansion of the equilibrium prices in the intertemporal substitution parameter. We use this expansion to study the pricing implications of Kreps and Porteus (1978) preferences over intertemporal consumption lotteries. This expansion applies and extends an approach suggested by Kogan and Uppal (2001). It is related to but distinct from the log-linear approximations commonly used in the asset pricing literature.

The second method characterizes formally long run cash flow risk by exploiting a mathematical formulation of pricing developed in Hansen and Scheinkman (2003). This method is model based and computes a long run dominant pricing component for cash flows that grow stochastically. This pricing component dominates as the initial cash flow payments are stripped from the security, leaving only the tail payoffs to be valued. It isolates value movements due to cash-flow riskiness far into the future and it allows us to ascertain what model ingredients have important influences on the valuation of riskiness.

In section 2 we study a familiar model of asset prices to show why the intertemporal composition of risk might matter to an investor. This model illustrates why long run consumption risk might matter. We also develop a general approximation to the model's solution. In section 3 we identify several important aggregate shocks that affect long run consumption. The implications of these shocks for wealth and risk prices is also examined. Section 4 develops a notion of risk based on the low frequency properties of cash flows and consumption. Starting with the assumption of cointegration between cash flows and consumption we construct a decomposition of security prices that displays the contribution of long-run risk to returns and prices. In section 5 we provide statistical evidence for the long-run relationship between consumption and the dividends of portfolios of stocks. Section 6 includes portfolio prices in the analysis and section 7 concludes.

2 Asset Pricing

Models of asset pricing link investor preferences and opportunities to deduce equilibrium relations for returns and prices. These models explain return heterogeneity by the existence of risk premia. Investors require larger expected returns as compensation for holding riskier portfolios. Alternative asset pricing models imply alternative risk-return tradeoffs. Equivalently [*e.g.*, see Hansen and Richard (1987)] they imply an explicit model of stochastic discount factors, the market determined variables $S_{t+1,t}$ used by investors to value one-period and hence multiple period assets.

There remains considerable controversy within the asset pricing literature about the feasibility of constructing an economically meaningful model of stochastic discount factors and hence risk premia. Nevertheless in this section we find it useful to consider one such model that, by design, leads to tractable restrictions on economic time series. This model is

rich enough to help us examine return heterogeneity as it relates to risk and to understand better the intertemporal values of equity.

2.1 Preferences

We follow Epstein and Zin (1989b) Weil (1990) by depicting preferences recursively. As we show below, this model of preferences provides a simple justification for examining a long-run relationship between consumption and returns. In addition it provides a convenient separation between risk aversion and the elasticity of intertemporal substitution [see Epstein and Zin (1989b)]. This separation allows us to examine the effects of changing risk exposure with modest consequences for the risk-free rate. Many of the measurement challenges that emerge in this economic model carry over to others as well, including any model that features the intertemporal composition of risk, including models in which investor preferences display intertemporal complementarity or “habit persistence.”

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

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

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

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

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

The parameter $\frac{1}{\rho}$ is a measure of the intertemporal substitution implied by the preferences. A measure of risk aversion depends on the details of the gamble being considered. As featured by Kreps and Porteus (1978), preferences like these relax the restriction that intertemporal compound lotteries can be *reduced* by simply integrating out the uncertainty conditioned on current information. Instead the intertemporal composition of risk matters. As we will see, this will be reflected explicitly in the equilibrium asset prices that we characterize. On the other hand, the aversion to simple wealth gambles is given by θ .

Under a Cobb-Douglas specification ($\rho = 1$), recursion (1) becomes:

$$V_t = (C_t)^{(1-\beta)} \mathcal{R}_t(V_{t+1})^\beta.$$

In what follows, the $\rho = 1$ will receive special attention because of its analytical tractability.

To include stochastic growth in consumption we study an alternative recursion that scales continuation values by consumption:

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

Since consumption and continuation values are positive, we find it convenient to work with logarithms instead. Let v_t denote the logarithm of the continuation value relative to the logarithm of consumption, and let c_t denote the logarithm of consumption. We rewrite recursion (1) as

$$v_t = \frac{1}{1-\rho} \log \left((1-\beta) + \beta \exp \left[(1-\rho) \mathcal{Q}_t (v_{t+1} + c_{t+1} - c_t) \right] \right), \quad (2)$$

where \mathcal{Q}_t is the so-called *risk-sensitive recursion*:

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

(See Hansen and Sargent (1995) and Tallarini (1998) the relations they show to the risk sensitive control literature.) The risk sensitive recursion is convenient for our subsequent characterizations.

2.2 Shadow Valuation

Consider the shadow valuation of a given consumption process. The utility recursion gives rise to a corresponding valuation recursion and implies stochastic discount factors used to represent this valuation. In light of the intertemporal budget constraint, the valuation of consumption in equilibrium coincides with wealth.

The first utility recursion (1) is homogeneous of degree one in consumption and the future continuation utility. Use Euler's Theorem to write:

$$V_t = (MC_t)C_t + E \left[(MV_{t+1})V_{t+1} \middle| \mathcal{F}_t \right] \quad (3)$$

where

$$\begin{aligned} MC_t &= (1-\beta)(V_t)^\rho (C_t)^{-\rho} \\ MV_{t+1} &= \beta(V_t)^\rho [\mathcal{R}_t(V_{t+1})]^{\theta-\rho} (V_{t+1})^{-\theta} \end{aligned}$$

The right-hand side of (3) measures the shadow value of consumption today and the continuation value of utility tomorrow.

Let consumption be numeraire, and suppose for the moment that we value claims to the future continuation value V_{t+1} as a substitute for future consumption processes. Divide both sides of (3) by MC_t and use marginal rates of substitution to compute shadow values. The shadow value of a claim to a continuation value is priced using $\frac{MV_{t+1}}{MC_t}$ as a stochastic discount factor. A claim to next period's consumption is valued using

$$S_{t+1,t} = \frac{MV_{t+1}MC_{t+1}}{MC_t} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\rho} \left(\frac{V_{t+1}}{\mathcal{R}_t(V_{t+1})} \right)^{\rho-\theta} \quad (4)$$

as a stochastic discount factor. There are two contribution (typically highly correlated) contributions to the stochastic discount factor in formula (4). One is the direct consumption

growth contribution familiar from the Lucas (1978) and Breeden (1979) model of asset pricing. The other is the continuation value relative to its risk adjustment. The contribution is forward-looking and is present provided that ρ and θ differ.

Given the homogeneity in the recursion used to depict preferences, equilibrium wealth is given by $W_t = \frac{V_t}{MC_t}$. Substituting for the marginal utility of consumption, the wealth-consumption ratio is:

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

Taking logarithms, we find that

$$\log W_t - \log C_t = -\log(1 - \beta) + (1 - \rho)v_t$$

When $\rho = 1$ we obtain the well known result that the wealth consumption ratio is constant.

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 (1). One approach is to use the relationship between wealth and the continuation value, $W_t = V_t/MC_t$, to construct 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. An aggregate stock market return is sometimes used to proxy for this return as in Epstein and Zin (1989a), for example; or other components can be included such as human capital with assigned market or shadow values. In addition to requiring the use of a market measure of wealth, this approach precludes the special case in which $\rho = 1$. Since the consumption wealth ratio is constant when $\rho = 1$, we cannot infer the continuation value from wealth and consumption. Moreover, when ρ is close to one any volatility in the stochastic discount factor attributed to wealth should also be reflected in consumption volatility. This implication is typically ignored even when consumption and wealth are used simultaneously.

In this investigation, like that of Restoy and Weil (1998), we maintain the direct link between the continuation value V_t stochastic process governing future consumption. In the case of logarithmic intertemporal preferences the link between future consumption and the continuation value easily can be calculated as we demonstrate in the next section. It is well understood that $\rho = 1$ leads to substantial simplification in the equilibrium prices and returns (*e.g.* see Schroder and Skiadas (1999).)

Approximate characterization of equilibrium pricing for recursive utility have been produced by Campbell (1994) and Restoy and Weil (1998). In what follows we use a distinct but related approach. While Campbell (1994) and Restoy and Weil (1998) use log-linear approximation of budget constraints, we follow Kogan and Uppal (2001) by approximating around an explicit equilibrium computed when $\rho = 1$. Our approximation is in the parameter ρ .¹ Campbell and Viceira (2002) (chapter 5) show the close connection between

¹Strictly speaking, $\rho = 1$ is ruled out in the parameterization considered by Restoy and Weil (1998) including the return-based Euler equation exploited in their calculations. The economy we study is different from that Kogan and Uppal (2001), but they suggest that extensions in the directions that interest us would be fruitful.

approximation around the utility parameter $\rho = 1$ and approximation around a constant consumption-wealth ratio for portfolio problems. Chacko and Viceira (2003) apply these approximation methods to portfolio problems with incomplete markets.

Our application in what follows is to the study of a simple model of equilibrium price determination. We find some useful and intriguing contrasts between approximation methods. There are interesting differences in the implied risk prices and the implied consumption-wealth ratios. Before turning to these, we consider the special case in which $\rho = 1$.

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

As in many papers in asset pricing, we use a $\rho = 1$ specification as a convenient benchmark. Campbell (1996) argues for less intertemporal substitution and Bansal and Yaron (2004) assume more. We will explore such deviations subsequently. The $\rho = 1$ is convenient for our purposes because when consumption has a log linear time series evolution, we can solve for the continuation value. This feature gives us the flexibility to include important low frequency time series components in the model solution.

The $\rho = 1$ limit in recursion (2) for continuation values:

$$v_t = \beta \mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t). \quad (5)$$

The stochastic discount factor in this special case is:

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

Notice that the term associated with the risk-adjustment satisfies

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

and can thus be thought of as distorting the probability distribution. As is familiar for logarithmic preferences, consumption and wealth are proportional.²

Recursion (5) was used by Tallarini (1998) in his study of *risk sensitive* business cycles. An important limiting case occurs when $\theta = 1$. In this case preferences are logarithmic and separable over time and states of the world with discount factor β .

To make our formula for the marginal rate of substitution operational, we need a formula for V_{t+1} computed using the equilibrium consumption process. Suppose that the first-difference of the logarithm of equilibrium consumption has a moving-average representation:

$$c_t - c_{t-1} = \gamma(L)w_t + \mu_c$$

²The constancy of the consumption-wealth when $\rho = 1$ prevents our use of this model to interpret directly the findings of Lettau and Ludvigson (2001b) and Lettau and Ludvigson (2001a).

where $\{w_t\}$ is a vector, *iid* standard normal process and

$$\gamma(z) = \sum_{j=0}^{\infty} \gamma_j z^j$$

where γ_j is a row vector and

$$\sum_{j=0}^{\infty} |\gamma_j|^2 < \infty.$$

This linear times series representation is adopted to help us interpret some of the time series evidence that we will discuss subsequently. Log-linear approximations are often used in macroeconomic modelling, although in what follows we will take the log-linear specification as being correct.

Guess a solution:

$$v_t = v(L)w_t + \mu_v.$$

Rewrite recursion (5) as:

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

Thus v must solve:

$$zv(z) = \beta[v(z) - v(0) + \gamma(z) - \gamma(0)],$$

which in particular implies that

$$v(0) + \gamma(0) = \gamma(\beta).$$

Solving for v and μ_v :

$$\begin{aligned} v(z) &= \beta \frac{\gamma(z) - \gamma(\beta)}{z - \beta} \\ \mu_v &= \frac{\beta}{1-\beta} \left[\mu_c + \frac{(1-\theta)}{2} \gamma(\beta) \cdot \gamma(\beta) \right]. \end{aligned}$$

The formula $v(z)$ is the solution to the forecasting problem:

$$v(L)w_t = \sum_{j=1}^{\infty} \beta^j E(c_{t+j} - c_{t+j-1} - \mu_c | \mathcal{F}_t)$$

familiar from the rational expectations literature on the permanent income model of consumption. The risk parameter θ enters only the constant term of continuation value process. The term $\gamma(\beta)$, which enters the formulas for v and μ_v is the *discounted impulse response* of consumption growth rate to a shock.

The logarithm of the stochastic discount factor can now be depicted as:

$$s_{t+1,t} \equiv \log S_{t+1,t} = -\delta - \gamma(L)w_{t+1} - \mu_c + (1 - \theta)\gamma(\beta)w_{t+1} - \frac{(1 - \theta)^2\gamma(\beta) \cdot \gamma(\beta)}{2}$$

where $\beta = \exp(-\delta)$. The term $\gamma(\beta)w_{t+1}$ is the solution to

$$(1 - \beta) \sum_{j=0}^{\infty} \beta^j [E(c_{t+j}|\mathcal{F}_{t+1}) - E(c_{t+j}|\mathcal{F}_t)].$$

It is a geometric average of current and future consumption responses to a shock at a fixed date (say date $t + 1$). The discount factor dictates the importance of future responses in this weighted average. As the subjective discount factor β tends to unity, $\gamma(\beta)$ converges to $\gamma(1)$ which is cumulative growth rate response or equivalently the limiting consumption response in the infinite future.

The stochastic discount factor includes both the familiar contribution from contemporaneous consumption plus a forward-looking term that discounts the *impulse responses* for consumption growth. The innovation to the logarithm $s_{t+1,t}$ of the stochastic discount factor is:

$$[-\gamma_0 + (1 - \theta)\gamma(\beta)]w_{t+1},$$

which shows how a shock at date $t + 1$ alters the stochastic discount factor. This term determines the magnitude of the risk premium. For instance, the price of payoff $\phi(w_{t+1})$ is given by:

$$\begin{aligned} E[\exp(s_{t+1})\phi(w_{t+1})|\mathcal{F}_t] &= E[\exp(s_{t+1})|\mathcal{F}_t] \frac{E[\exp(s_{t+1})\phi(w_{t+1})|\mathcal{F}_t]}{E[\exp(s_{t+1})|\mathcal{F}_t]} \\ &= E[\exp(s_{t+1})|\mathcal{F}_t] \tilde{E}[\phi(w_{t+1})|\mathcal{F}_t]. \end{aligned}$$

The first term is pure discount term and the second the is the expectation of $\phi(w_{t+1})$ under the so-called risk neutral probability distribution. The first term is:

$$\log E[\exp(s_{t+1})|\mathcal{F}_t] = -\delta - \gamma^1(L)w_t - \frac{(1 - \theta)^2\gamma(\beta) \cdot \gamma(\beta)}{2}$$

where

$$\gamma^1(L) = \sum_{j=0}^{\infty} \gamma_{j+1}L^j.$$

The innovation to the logarithm $s_{t+1,t}$ of the stochastic discount factor:

$$[-\gamma_0 + (1 - \theta)\gamma(\beta)]w_{t+1}$$

is the mean of the normally distributed shock w_{t+1} under the risk-neutral distribution. The covariance matrix remains an identity matrix. The adjustment $-\gamma_0$ is familiar from the paper Hansen and Singleton (1983) and the term $(1 - \theta)\gamma(\beta)$ is the adjustment for the

intertemporal composition of consumption risk implied by the Kreps and Porteus (1978) specification of recursive utility. Large values of the risk parameter θ enhance the importance of this component. This latter effect is featured in the analysis of Bansal and Yaron (2004).³

Example 2.1. *Suppose that consumption evolves according to:*

$$c_{t+1} - c_t = \mu_c + U_c \cdot x_t + \gamma_0 w_{t+1}$$

where z_t evolves according to first-order vector autoregression:

$$x_{t+1} = Gx_t + Hw_{t+1}.$$

The matrix G has strictly stable eigenvalues (eigenvalues with absolute values that are strictly less than one), and $\{w_{t+1} : t = 0, 1, \dots\}$ is iid normal with mean zero and covariance matrix I . Then for $j > 0$,

$$\gamma_j = U'_c(G^{j-1})H,$$

and

$$v_t = U_v \cdot x_t + \mu_v$$

where

$$U_v \doteq \beta(I - G'\beta)^{-1}U_c,$$

$$\mu_v \doteq \frac{\beta}{1 - \beta} \left[\mu_c + \frac{(1 - \theta)}{2} \gamma(\beta) \cdot \gamma(\beta) \right],$$

and

$$\gamma(\beta) = \gamma_0 + \beta U'_c(I - G\beta)^{-1}H.$$

The logarithm of the stochastic discount factor is:

$$s_{t+1,t} = -\delta - \mu_c - U_c \cdot x_t - \gamma_0 w_{t+1} + (1 - \theta)\gamma(\beta)w_{t+1} - \frac{(1 - \theta)^2 \gamma(\beta) \cdot \gamma(\beta)}{2}$$

This model presents a measurement challenge for an econometrician. More than just the one-period response of consumption to underlying economic shocks matters. In addition the discounted response of consumption in underlying economic shocks is what is required to quantify the risk that matters to investors. For discount factors close to unity, this challenge is known to be more acute.

While this model has a simple and usable characterization of how temporal dependence in consumption growth alters risk premia, it has the counterfactual implication of risk premia that are time invariant. Other authors, including Campbell and Cochrane (1999) argue that

³Anderson, Hansen, and Sargent (2003) suggest a different interpretation for the parameter θ . Instead of risk, this parameter may reflect model misspecification that investors confront by not knowing the precise riskiness that they must confront in the marketplace. As argued by Anderson, Hansen, and Sargent (2003), under this alternative interpretation, $|(1 - \theta)\gamma(\beta)|$ is measure of model misspecification that investors have trouble disentangling because this misspecification is disguised by the underlying shocks that impinge on investment opportunities.

risk premia vary over the business cycle. Time varying risk premia could be added to the model by allowing for stochastic variation in volatility as in Bansal and Yaron (2004). We see no reason why this complexity, however, simplifies the measurement or approximation problem.

Because of the logarithmic nature of preferences, wealth in this economy is proportional to consumption

$$W_t = \frac{C_t}{1 - \beta}.$$

As noted by Gibbons and Ferson (1985), we may use the return on the wealth portfolio as a proxy for the consumption growth rate. In particular, the return on a claim to wealth is:

$$R_{t+1}^w = \frac{W_{t+1}}{\frac{\beta C_t}{1 - \beta}} = \frac{C_{t+1}}{\beta C_t}.$$

Thus

$$r_{t+1}^w = c_{t+1} - c_t - \log \beta$$

This leads Campbell and Vuolteenaho (2003) and Campbell, Polk, and Vuolteenaho (2003) to use a market wealth return as a proxy for consumption growth. With this proxy, these papers take γ_0 to be the familiar (conditional) CAPM risk adjustment and $(1 - \theta)\gamma(\beta)$ as an additional adjustment where γ is now measured using a market return.⁴ In this paper we instead follow Hansen and Singleton (1983), Restoy and Weil (1998), Bansal and Yaron (2004) and others by focus on consumption dynamics. This justifies our interest in computing continuation values for equilibrium consumption plans.

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

While $\rho = 1$ is a convenient benchmark, we are also interested in departures from this specification. To assess these departures, we consider an expansion for the continuation value around the point $\rho = 1$. Our aim is to compute a derivative Dv_t^1 to use in a first-order approximation:

$$v_t \approx v_t^1 + (\rho - 1)Dv_t^1$$

where v_t^1 is the continuation value for the case in which $\rho = 1$. In appendix 7, we derive the following recursion for the derivative:

$$Dv_t^1 = -\frac{(1 - \beta)(v_t^1)^2}{2\beta} + \beta \tilde{E}(Dv_{t+1}^1 | \mathcal{F}_t)$$

where \tilde{E} is the distorted expectation operator associated with the density

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

⁴Campbell and Vuolteenaho (2003) refer to this second term as the *bad* β term.

For the log-normal model of consumption, this distorted expectation appends a mean to the shock vector w_{t+1} . The distorted distribution of w_{t+1} remains normal, but instead of mean zero, it has a risk adjusted mean of $(1 - \theta)\gamma(\beta)$. The derivative Dv_t^1 is negative because it is the (distorted) expectation of the sum of negative random variables.

When ρ is different from, the wealth-consumption ratio is not constant. A first-order expansion of the continuation value implies a second-order expansion of the consumption-wealth ratio. This can be seen directly from (2.2):

$$\begin{aligned}\log W_t - \log C_t &= -\log(1 - \beta) + (1 - \rho) [v_t^1 + (\rho - 1)Dv_t^1] \\ &= -\log(1 - \beta) - (\rho - 1)v_t^1 - (\rho - 1)^2 Dv_t^1.\end{aligned}$$

The term v_t^1 is very similar (but not identical to) the term typically used when taking log-linear approximations.⁵ Recall that this term is the expected discounted value of consumption growth with an additive term constant term that adjusts for variability. In the first-order approximation of the wealth-consumption ratio, v_t^1 shows how the wealth-consumption ratio is altered with the intertemporal substitution parameter ρ . When consumption growth rates are predictable, growth rate forecasts alter the consumption-wealth ratio. Forecasts that a geometric average of future consumption will be higher than average imply a higher wealth-consumption ratio when ρ exceeds one and a lower one ρ is less than unity. In contrast, the risk parameter θ alters the constant term in v_t^1 . This implication of intertemporal substitution is familiar from previous literature (*e.g.* see Campbell (1996) and Restoy and Weil (1998)). By construction, the second-order term adjusts the wealth consumption ratio in a manner that is symmetric about $\rho = 1$. When ρ deviates from one, this second-order correction is positive.

The corresponding expansion for the logarithm of the stochastic discount factor is:

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

where⁶

$$Ds_{t+1,t}^1 = v_{t+1}^1 - \frac{1}{\beta}v_t^1 + (1 - \theta) \left[Dv_{t+1}^1 - \tilde{E} (Dv_{t+1}^1 | \mathcal{F}_t) \right].$$

Recall that in Example 2.1, $c_{t+1} - c_t$ has conditional mean: $\mu_c + U_c x_t$ and a shock contribution: $\gamma_0 w_{t+1}$. Using the parameterization, the logarithm of the continuation value/consumption ratio is:

$$v_{t+1}^1 = U_v \cdot x_{t+1} + \mu_v$$

⁵In log-linear approximation the discount rate in this approximation is linked to the mean of the wealth consumption ratio. In the ρ expansion, the subjective rate of discount is used instead.

⁶There are two ways we could use this formula to approximate one period pricing. We use $\exp [s_{t+1,t}^1 + (\rho - 1)Ds_{t+1,t}^1]$, as an approximate discount factor, but instead we could have used

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

The second approximation is not necessarily positive, but it produces the first-order expansion of the one-period pricing operator.

$$= U_v' H w_{t+1} + U_v' G x_t + \mu_v.$$

In the appendix we show that

$$Dv_{t+1}^1 = -\frac{1}{2} x_{t+1}' \Upsilon_d x_{t+1} + U_d \cdot x_{t+1} + u_d$$

where formulas for the matrices Υ_d and U_d are given in the appendix.

We can use this expansion to produce approximations to equilibrium prices. We compute implied risk neutral prices as a way to characterize risk premia. The risk neutral distribution alters the standard normal distribution for w_{t+1} . In the example economy the first-order approximation of the stochastic discount factor implies that the risk neutral distribution remains normal but with an enhanced covariance matrix and an alternative mean. In a continuous-time approximation, only the mean adjustment is present. The first-order expansion of the altered mean is:

$$-\gamma_0 + (1 - \theta)\gamma(\beta) + (\rho - 1)[H'U_v + (1 - \theta)(U_d - H'\Upsilon_d G x_t)]. \quad (6)$$

The term $(\rho - 1)(1 - \theta)(U_d - H'\Upsilon_d G x_t)$ is new relative to the more typical log-linear approximation. It is time varying when ρ and θ are distinct from unity. This is so even though the consumption process in the example economy is homoskedastic.

3 Shocks and Vector Autoregressions

We use vector autoregressive (VAR) models to both identify interesting aggregate shocks and to estimate $\gamma(L)$. As we discuss below we also use these methods to identify important long-run risks. We consider a specification that is rich enough to allow experimentation with different long-run assumptions and different variables that may be important in identifying the long run consequences of macroeconomic shocks.

3.1 Identifying Aggregate Shocks

In our initial model we let consumption be the first element of y_t and corporate earnings be the second element:

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

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

$$A_0 y_t + A_1 y_{t-1} + A_2 y_{t-2} + \dots + A_\ell y_{t-\ell} + B_0 = w_t, \quad (7)$$

The vector B_0 two-dimensional, and similarly 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 . We normalize A_0 to be lower triangular with positive entries on the diagonals. Form:

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

We are interested in a specification in which $A(z)$ is nonsingular for $|z| < 1$. Given this model, the discounted response of consumption to shocks is given by:

$$\gamma(\beta) = (1 - \beta)u_1' A(\beta)^{-1} \quad (8)$$

where u_1 is a column vector with a one in the first position and a zero in the second entry.

For our measure of aggregate consumption we use aggregate consumption of nondurables and services taken from the National Income and Product Accounts. This measure is quarterly from 1947 Q1 to 2002 Q4, is in real terms and is seasonally adjusted. Our inclusion of corporate earnings in the VAR is motivated by the work of Lettau and Ludvigson (2001b) and Santos and Veronesi (2001). This variable is meant to capture aggregate exposure to stock market cash flows. We measure corporate earnings from NIPA and convert this series to real terms using the implicit price deflator for nondurables and services.

We consider two specifications of the evolution of y_t . In one case the model is estimated without additional restrictions, and in the other we restrict the matrix $A(1)$ to have rank one:

$$A(1) = \alpha \begin{bmatrix} 1 & -1 \end{bmatrix}.$$

where the column vector α is freely estimated. This parameterization imposes two restrictions on the $A(1)$ matrix. We refer to the first specification as the *without cointegration* model and second as the *with cointegration* model.

The second system imposes a unit root in consumption and earnings, but restricts these series to grow together. In this system both series respond in the same way to shocks in the long run. Specifically, the limiting response of consumption and earnings to a shock at date 0 is the same. Since the cointegration relation we consider is prespecified, the *with cointegration* model can be estimated as a vector autoregression in the first-difference of the log consumption and the difference between the log earnings and log consumption.

Our use of a second time series is to identify additional sources of long run risk beyond just a single “consumption innovation.” Whereas Bansal and Yaron (2004) consider multivariate specifications 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.

In our analysis, we will not be concerned with the usual shock identification familiar from the literature on structural VAR’s (vector autoregressions). This literature is concerned with the assignment of structural labels to the underlying shocks and imposes *a priori* restrictions to make this assignment. While we have restricted C to be lower triangular, we view this as a normalization. This leads to the identification of two shocks, but other shock configurations are constructed by taking linear combinations of the two shocks we identify. Sometimes we will form the linear combination that captures fully the long-run contribution to consumption and earnings variability. What interests us is the intertemporal composition of consumption risk and not the precise labels attached to individual shocks.

We report impulse responses for estimates of the VAR with and without the cointegration restriction in figure 1. When cointegration is imposed, corporate earnings relative to consumption identifies an important long-run response to both shocks. The long run impact of the first “consumption shock” is twice that of the impact on impulse. While on impact, the second “earnings shock” is normalized to no impact on consumption, its long run impact is sizeable. As demonstrated in the recursive utility model, the geometrically weighted average response of consumption to the underlying shocks is a key ingredient in the stochastic discount factor. As the subjective discount rate converges to zero, this average coincides with the limiting consumption response.

Notice from the impulse responses in figure 1, that when the cointegration restriction is not imposed, the estimated long run consumption responses are substantially smaller. The imposition of the cointegration restriction is critical to locating an important low frequency component in consumption. Moreover, in the absence of this restriction, the overall feedback from earnings shocks to consumption is substantially weakened. The earnings shocks have little impact on consumption for the *no cointegration* specification.

Using the cointegration specification, we explore the statistical accuracy of the estimated responses. Following suggestions of Sims and Zha (1999) and Zha (1999) we impose Box-Tiao priors on the coefficients of each equation and simulate histograms for the parameter estimates. This provides approximation for a Bayesian posterior with a relatively diffuse (and improper) prior distribution. These “priors” are chosen for convenience, but they give us a convenient way to depict the sampling uncertainty associated with the estimates.

In the model of Hansen and Singleton (1983), it is the immediate innovation in consumption in consumption that matters for pricing one-period securities. Figure 2 gives a histogram for the standard deviation of this estimate. In other words it gives the histogram for the estimate of the (1, 1) entry of A_0 .

For comparison we also report the histogram for a long-run response. Both shocks have long run consequences. Similar to Blanchard and Quah (1989), we construct a *temporary shock* as a linear combination of the two original shocks that has no long run consequences. Similarly, we define a *permanent shock* to be the orthogonal linear combination of shocks. Mechanically we build an orthonormal matrix C such that

$$\lim_{\beta \rightarrow 1} (1 - \beta)(u_1)' A(\beta) C u_2 = 0$$

where u_2 is vector with a zero in the first position and a one in the second position. The permanent shock is the first entry of Cw_t and the transitory shock is the second entry of this transformed shock vector. Figure 2 also gives a histogram for the long run consumption response to a long run shock. The permanent shock is normalized to have unit standard deviation, so that we can compare magnitudes across the long-run and short run responses.

As might be expected, the short run response is much more accurate than the long run response. Notice that the horizontal scales of histogram differ by a factor of ten. In particular, while the long run response is centered at higher value and it also has a substantial right tail. Consistent with impulse response functions, the median long-run response is about double that of the short-term response, but in addition nontrivial probabilities are given to

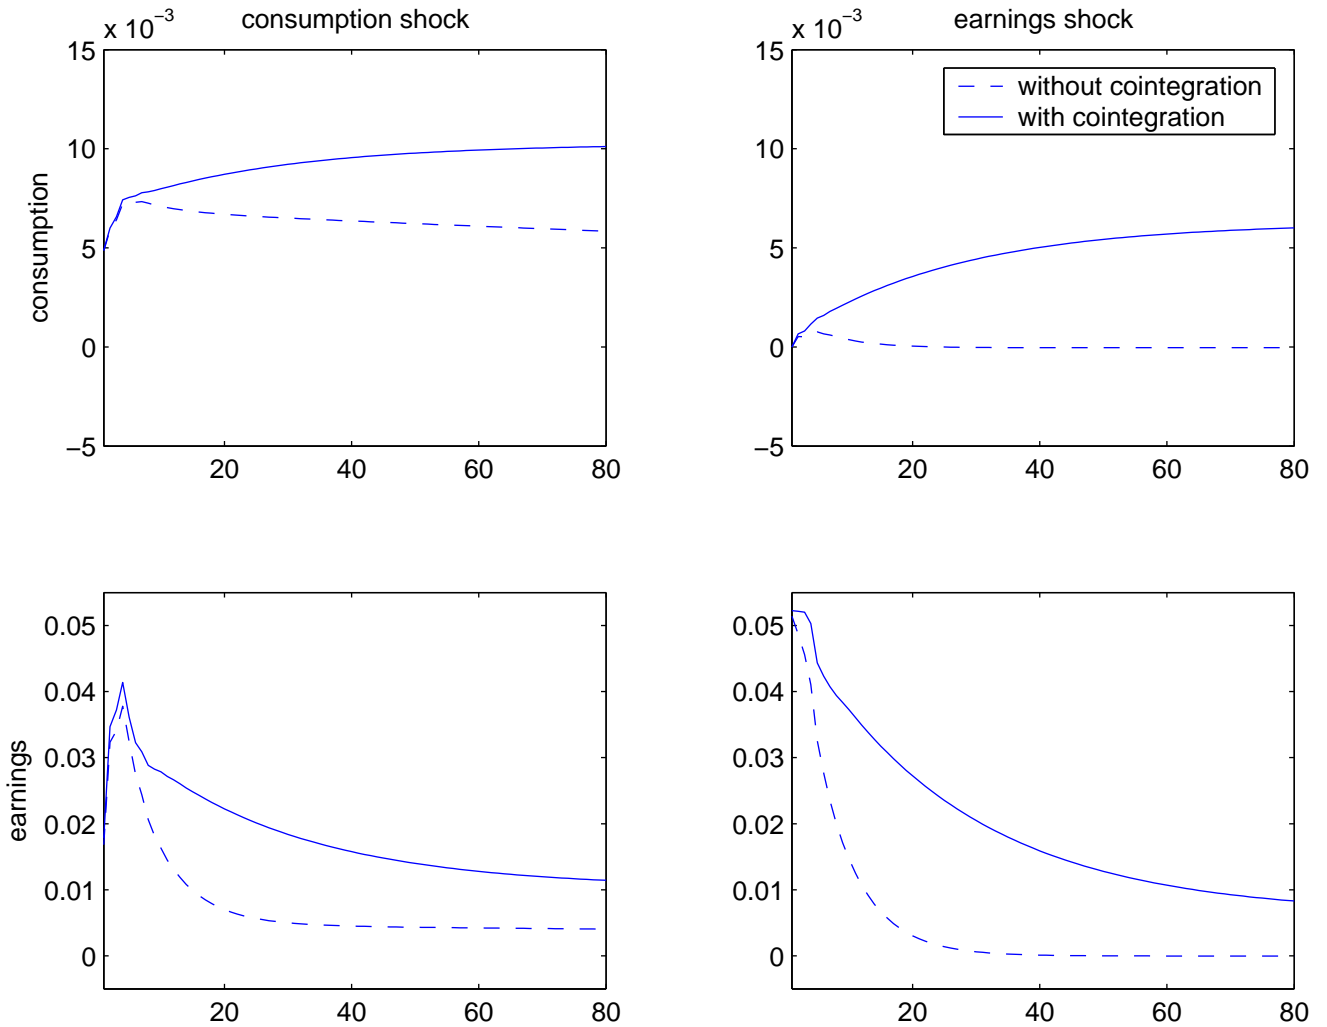


Figure 1: *Impulse responses of consumption and earnings.* The impulse responses without imposing cointegration were constructed from a bivariate VAR with entries c_t , e_t . These responses are given by the dashed lines $- - -$. Solid lines $—$ are used to depict the impulse responses estimated from a cointegrated system. The impulse response functions are computed from a VAR with $c_t - c_{t-1}$ and $c_t - e_t$ as time series components.

substantially larger responses.⁷ Thus from the standpoint of sampling accuracy, the long run response could be even more than double that of the immediate consumption response. Because discounted future consumption enters the pricing model this low-frequency component of consumption is potentially important. For example, Bansal and Yaron (2004) argue that this component of consumption is important for understanding the equity premium.

The cointegrated specification with a known cointegrating coefficient imposes a restriction on the VAR. To explore the statistical plausibility of this restriction, we free up the cointegration relation by allowing consumption and earnings to have different long run responses. We introduce a freely estimated cointegrating coefficient λ , and explore the cointegrating restriction $A(1) = \alpha F'$ where $F' = [-\lambda \ 1]$. We concentrate the likelihood function by doing a repeated maximization for each choice of parameter λ . The concentrated likelihood function is reported in figure 3. We complement this calculation with a simulation of the posterior distribution for the cointegrating coefficient imposing a Box-Tiao prior for each VAR conditioned on the cointegrating coefficient. The resulting histogram is depicted in figure 4. For sake of computation, we used a uniform prior over the interval $[-2, 2]$ for the cointegrating coefficient. Both figures suggest that the “balanced growth” coefficient of unity is plausible.

To explore more generally the cointegrating restriction, figure 5 compares the implied spectral densities for growth rates when the cointegration restriction is imposed (with $\lambda = 1$) and when it is not. Under the cointegration restriction, there is an important low frequency component to the consumption growth rate. The spectral densities in figure 5, however, are very close except at a narrow frequency range, suggesting that these two models are hard to distinguish on statistical grounds.⁸

3.2 Implications for Wealth and Risk Prices

The inclusion of corporate earnings in the cointegrated VAR for consumption appears to identify an important long-run shock to consumption. There are issues with the statistical reliability of the estimation of this effect, however. A potential avenue for improving the inference can be found in the model of section 2. Here we consider the effects of the alternative specifications for the dynamics of consumption as well as the importance of the expansion of section 2.4.

First consider figure 6 which gives the fitted values for the equilibrium prices of the two stocks as given by (6). This is based on the parameter estimates of the model of consumption and corporate earnings with cointegration. For this example the utility function parameters

⁷The 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. 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.

⁸The model with cointegration imposes two restrictions on the matrix $A(1)$. Twice the likelihood ratio for the two models is 5.9. As a consequence, the Bayesian information or Schwarz criterion selects the restricted model, although it is hard to defend the formal use of this particular decision criterion.

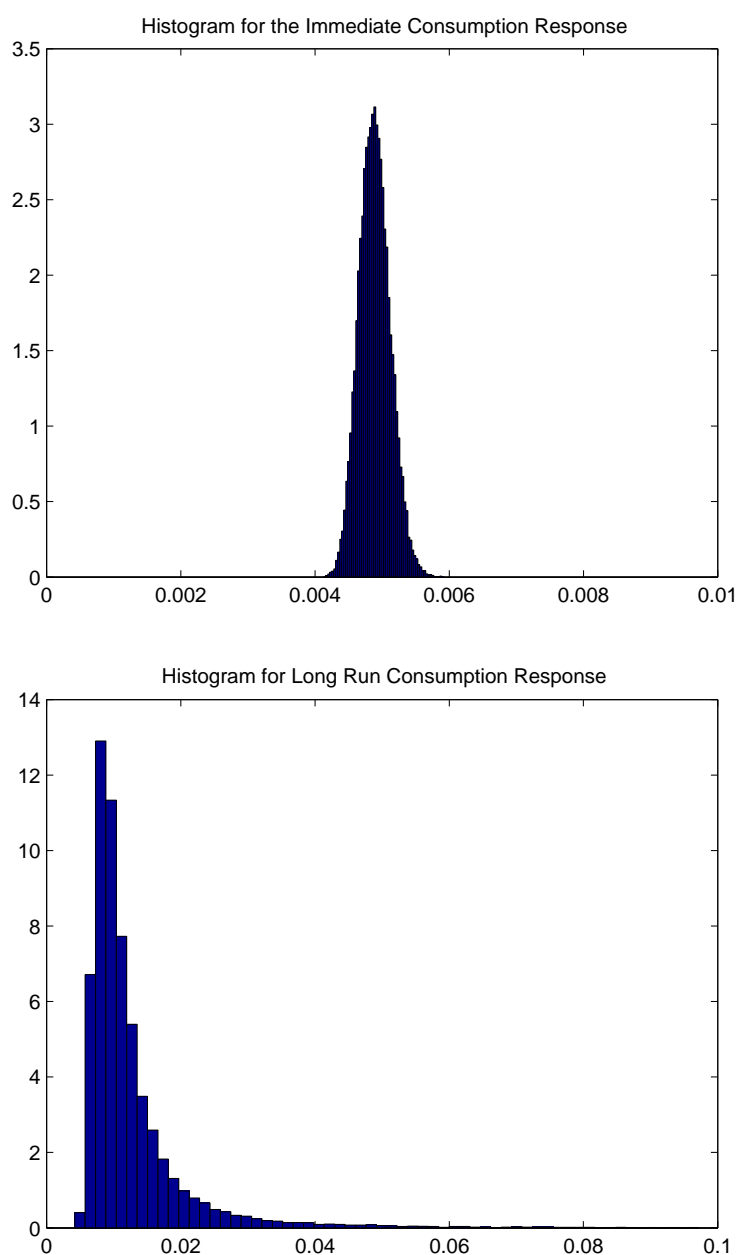


Figure 2: *Approximate posterior distributions for responses.* The top panel gives the approximate posterior for the immediate response to consumption and the bottom panel the approximate posterior for the long-run response of consumption to the permanent shock. The histograms have sixty bins with an average bin height of unity. They were constructed using using Box-Tiao priors for each equation. Vertical axes are constructed so that on average the histogram height is unity.

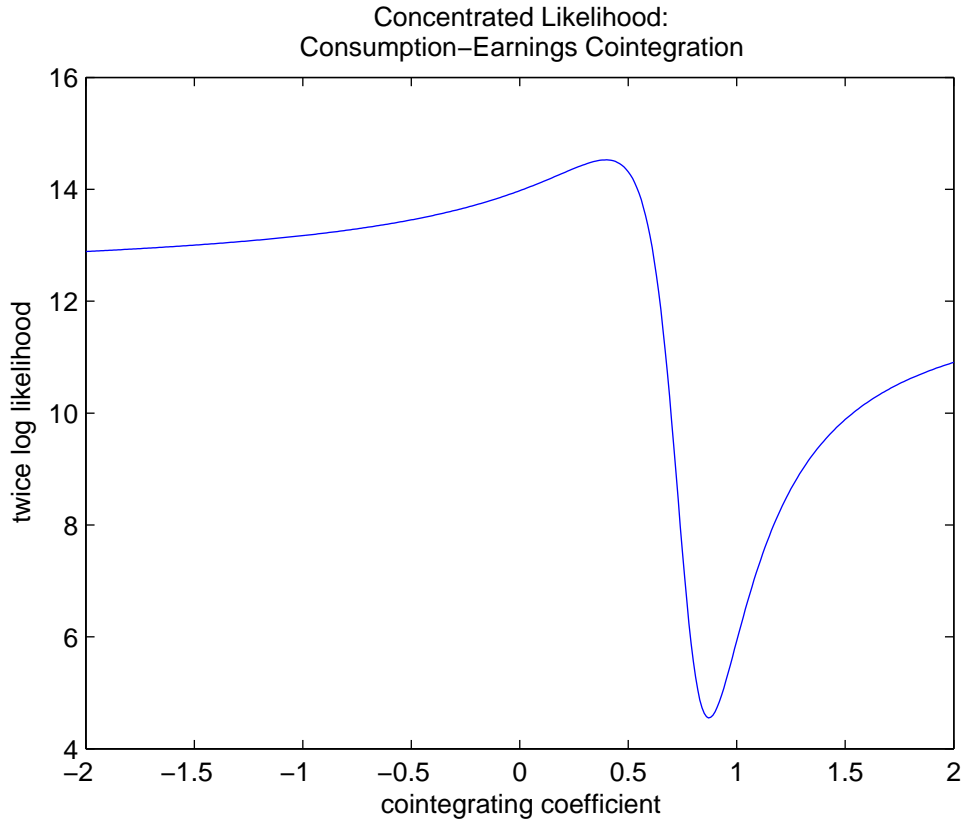


Figure 3: This figure reports twice the concentrated conditional log-likelihood function constructed from a bivariate VAR with entries c_t , e_t . The horizontal axis is the value of the cointegrating coefficient as it scales c_t . The log-likelihood is depicted relative to an unrestricted log-level VAR. The maximum likelihood estimate is obtained by minimizing the objective.




Figure 4: *Approximate posterior probabilities for the cointegrating coefficient.* Box-Tiao priors are imposed on the regression coefficients and innovation variances conditioned on the cointegrating coefficient. Posterior probabilities are computed by simulating from a Markov chain constructed from the conditional likelihood function.




Figure 5: *Spectral Densities of consumption and earnings growth*

were set at: $\beta = 0.98^{1/4}$, $\theta = 20$ and $\rho = 1/1.5$. The value of ρ is consistent with the assumptions of Bansal and Yaron (2004) and others and implies an intertemporal elasticity of substitution of 1.5.

Notice the dynamics of consumption along with our approximation imply substantial variation in risk prices even though there is no variation in the conditional volatility of consumption. Our expansion of the model around the case of $\rho = 1$ explicitly allows for this possibility.

The assumed dynamics of consumption are quite important, however. To see this, consider figure 7 which gives the corresponding plot for the without cointegration case. In this case the prices are essentially constant. There is some variability but when plotted on the same scale as figure 6 these prices appear to be constant. These contrasting results under the two specifications of the dynamics of consumption emphasize the potential importance of long-run risk.

To further investigate the importance of the specification of consumption dynamics the fitted values of the log wealth-consumption ratio for each model of consumption are plotted in figure 8. Notice that without cointegration there is substantially less variation in the wealth-consumption ratio. The second order terms in the expansion developed in section 2.4 are very important in producing the variation that is seen in figure 8. Without these terms both plots in figure 8 show much less variation.

Although the estimated dynamics of consumption with cointegrated earnings produce some variation in the consumption wealth ratio this variation is much too small if we were to directly link stock market returns to the claim on consumption. The variation in W_t/C_t is much smaller than the observed variation in the price dividend ratio from stock markets. For this reason it is important to allow for a distinction between consumption and the cash flows to holders of equity. In the next section we investigate the potential implications of this distinction.

4 Long Run Cash Flow Risk

We have seen some evidence for an important long run component in consumption when combined with the preference specification of section 2.1. We now investigate how long run risk is encoded in asset prices where the long run dynamics of cash flows are not linked one for one with consumption. Specifically, we consider when riskiness about long run cash flow growth can have an important contribution to current value.

To think about this issue first consider a stationary Markov specification for $\{x_t\}$, a process used to depict the underlying valuation. The logarithm of consumption evolves according to:

$$c_t - c_{t-1} = \mu_c(x_{t-1}) + \sigma_c(x_{t-1}) \cdot w_t.$$

This model nests the specifications we have considered so far as special cases.

Cash flows or dividends to risky securities are allowed to be “levered” claims on consumption in the long run. Following Bansal, Dittmar, and Lundblad (2002), Lettau, Ludvigson,

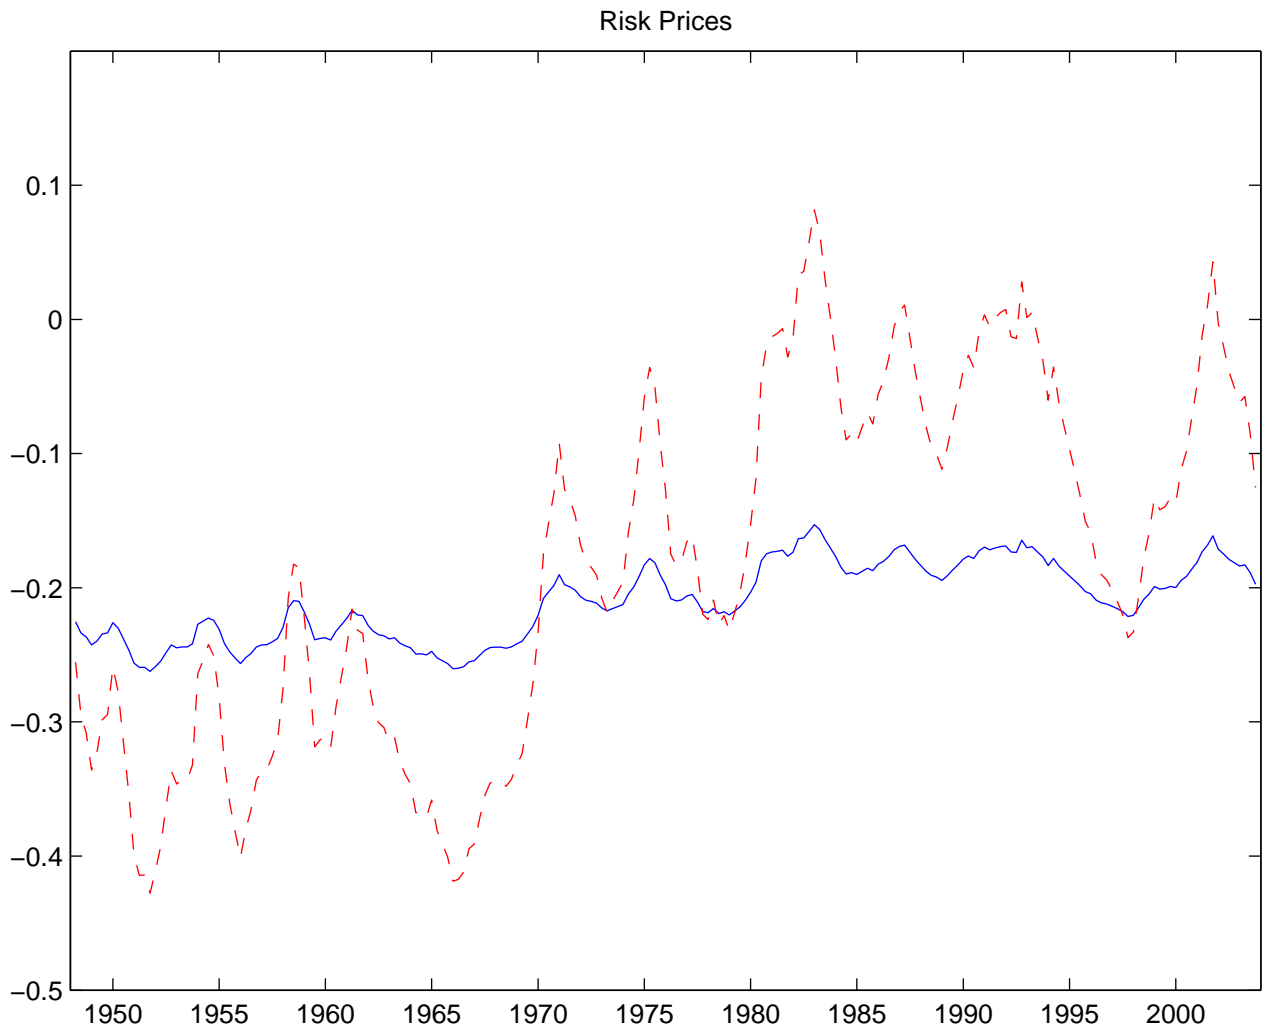


Figure 6: *Fitted prices of the two shocks using the model with cointegration. The solid line — gives the price of the consumption shock (the first shock) and the dashed line — — — gives the prices of the earnings shock (the second shock)*

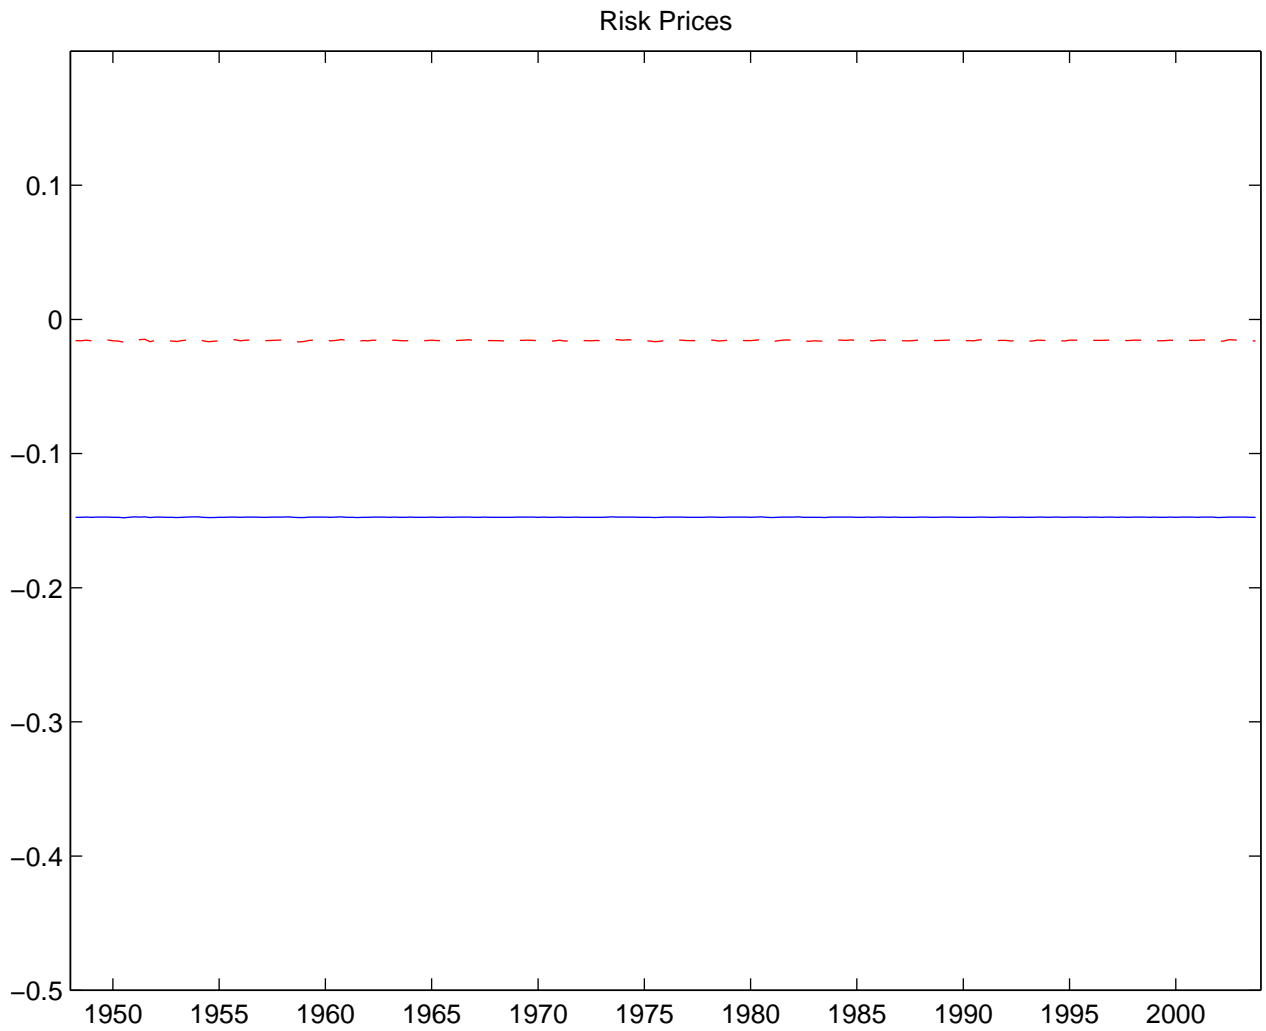


Figure 7: *Fitted prices of the two shocks using the model without cointegration. The solid line — gives the price of the consumption shock (the first shock) and the dashed line - - - gives the prices of the earnings shock (the second shock)*

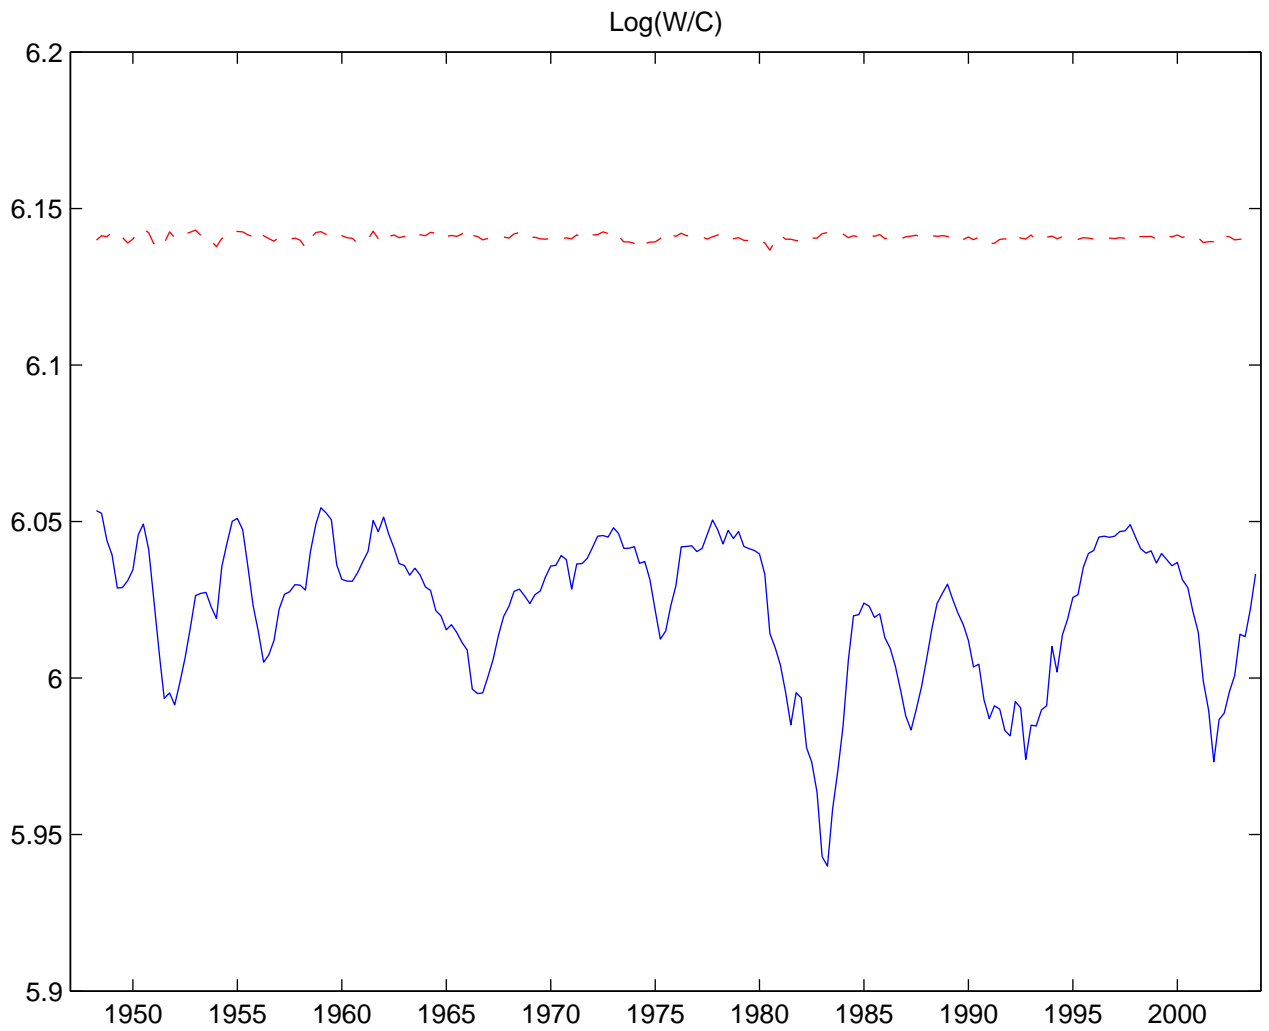


Figure 8: *Fitted log wealth-consumption ratio. The solid line — gives the values for the model with cointegration and the dashed line - - - gives the values for the model without cointegration.*

and Wachter (2004), and others, as an example we study the valuation of cash flows that are cointegrated with consumption:

$$d_t = \lambda c_t + \zeta t + \phi(x_t) \quad (9)$$

where d_t is the logarithm of the cash flow. Since the Markov process is stationary, growth is governed by the parameter pair (λ, ζ) . For the time being we allow for nonlinearities in the time series model, although in our computations we will revert to the log-linear specification used previously. Also we will examine empirically other examples of the long-run characteristics of cash flows.

To study the effects of our long-run specification of risk, we use eigenvalue methods. Prior to developing these methods, consider a square matrix M raised to a power j . For simplicity suppose the matrix has distinct eigenvalues, and write the eigenvalue decomposition as:

$$M = T\Delta T^{-1}$$

where Δ is a diagonal matrix of eigenvalues. Then

$$M^j = T\Delta^j T^{-1}.$$

Suppose that the largest eigenvalue in absolute value is positive, and call this eigenvalue δ . Then the sequence, $\{\delta^{-j}\Delta^j : j = 1, 2, \dots\}$ converges to a matrix of zeros except in one position where there is a one. Thus

$$\delta^{-j}M^j = T(\delta^{-j}\Delta^j)T^{-1}$$

converges to a constant matrix as j gets arbitrarily large. The eigenvalue δ determines the asymptotic growth factor of the matrix sequence M^j and $\log \delta$ is the corresponding asymptotic growth rate.

We use this same approach, but applied to operators. Markov valuation and conditional expectation operators over multiple time intervals can be depicted as iterates of their single-period counterparts. Logarithms of dominant eigenvalues give us a characterization of long-run growth in values and expectations. In what follows, we construct the operators of interest and apply them to characterize long-run behavior.

4.1 Operator Valuation

One counterpart to the matrix M , is a one-period valuation operator given by:

$$\mathcal{P}_\lambda \psi(x) = E(\exp[s_{t+1,t} + \lambda(c_{t+1} - c_t)]\psi(x_{t+1})|x_t = x).$$

Formally, we view this operator as a mapping from L^2 into L^2 where L^2 is the space of (Borel measurable) functions ψ for which

$$E\psi(x_t)^2 < \infty.$$

This operator takes a payoff at date $t + 1$ of the form:

$$\exp(\lambda c_{t+1}) \psi(x_{t+1})$$

and maps it into a price today scaled by $\exp(\lambda c_t)$. Since payoffs and prices are scaled, the valuation operator depends on the choice of λ used in the scaling. Consistent with formula (9), consumption provides the only source of growth in this specification.

Multi-period prices can be inferred from this one-period pricing operator through iteration. The value of a date $t + j$ cash flow:

$$\exp[\zeta(t + j) + \lambda c_{t+j}] \psi(x_{t+j}) \tag{10}$$

is:

$$\exp[\zeta(t + j) + \lambda c_t] (\mathcal{P}_\lambda)^j \psi(x_t).$$

The notation $(\mathcal{P}_\lambda)^j$ denotes the application of the one-period valuation operator j times.

If we take this cash flow to be a dividend process, the date t price-dividend ratio is:

$$\frac{P_t}{D_t} = \frac{\sum_{j=1}^{\infty} \exp(\zeta j) [(\mathcal{P}_\lambda)^j \psi(x_t)]}{\psi(x_t)}. \tag{11}$$

The term

$$\frac{\exp(\zeta j) [(\mathcal{P}_\lambda)^j \psi(x_t)]}{\psi(x_t)} \tag{12}$$

is the contribution of the date $t + j$ derivative to the price-dividend ratio, and the price-dividend ratio adds over these objects. Computing these individual terms gives a value decomposition of the price-dividend ratio by time horizon.

Since we allow for the growth rates in the cash flows to vary over time, we shall also have need to define operators that we use to measure these rates and the limiting growth behavior. Let

$$\mathcal{G}_\lambda \psi(x) = E(\exp[\lambda(c_{t+1} - c_t)] \psi(x_{t+1}) | x_t = x).$$

By iterating on this growth operator, we can study expected cash flow growth over multi-period horizons. In particular, the expected value of the cash flow (10):

$$\exp[\zeta(t + j) + \lambda c_t] (\mathcal{G}_\lambda)^j \psi(x_t).$$

4.2 Limiting Behavior

For positive cash flows we can characterize the limiting or tail contribution by studying the limit:

$$\lim_{j \rightarrow \infty} \frac{1}{j} \log \left[\frac{\exp(\zeta j) [(\mathcal{P}_\lambda)^j \psi(x_t)]}{\psi(x_t)} \right] = \zeta + \lim_{j \rightarrow \infty} \frac{\log [(\mathcal{P}_\lambda)^j \psi(x_t)]}{j}$$

which will depend on λ but often not depend on ψ . This calculation gives us an asymptotic decay rate that depends on both cash flow growth through its dependence on λ and ζ and on

the economic value associated with that growth. It measures how long-run prospects about dividends contribute to the value.

Analogously, the asymptotic cash-flow growth is measured as:

$$\lim_{j \rightarrow \infty} \frac{1}{j} \log \left[\frac{\exp(\zeta j) [(\mathcal{G}_\lambda)^j \psi(x_t)]}{\psi(x_t)} \right] = \zeta + \lim_{j \rightarrow \infty} \frac{\log [(\mathcal{G}_\lambda)^j \psi(x_t)]}{j}$$

As in the case of matrices, the asymptotic decay rate is determined by the dominate eigenvalue of the valuation operator \mathcal{P}_λ . An eigenvalue solves the equation:

$$\mathcal{P}_\lambda \psi_\lambda = \exp(-\nu_\lambda) \psi_\lambda$$

where ψ is an eigenfunction and $\exp(-\nu)$ is an eigenvalue. Then the invariance property of an eigenfunction implies that

$$(\mathcal{P}_\lambda)^j \psi = \exp(-j\nu_\lambda) \psi_\lambda$$

for any j . The dominant eigenvalue (when it exists) is associated with a strictly positive eigenfunction. Moreover,

$$\lim_{j \rightarrow \infty} \frac{\log [(\mathcal{P}_\lambda)^j \psi(x_t)]}{j} = -\nu_\lambda$$

provided that ψ is strictly positive. This limiting rate is invariant to the specific choice of positive ψ used in defining the cash flow. It does of course depend on λ .

Similarly, we compute the limiting growth rate of the cash flow by solving the eigenfunction problem:

$$\mathcal{G}_\lambda \varphi_\lambda = \exp(\epsilon_\lambda) \varphi_\lambda.$$

Then

$$\lim_{j \rightarrow \infty} \frac{\log [(\mathcal{G}_\lambda)^j \psi(x_t)]}{j} = \epsilon_\lambda$$

for strictly positive ψ . This limit is of interest because cash flow growth is a contributor to manner in which future payoffs contribute to current value.

We aim to approximate returns to a security with dividends payments only far off into the future. We construct these objects by supposing we have a security with a dividend:

$$D_{t+j} = \exp[(t+j)\zeta + \lambda c_{t+j}] \psi_\lambda(x_{t+j}).$$

This security has a constant price/dividend ratio. Using the eigenvalue property and formula (11)

$$\frac{P_t}{D_t} = \frac{\exp(\zeta - \nu_\lambda)}{1 - \exp(\zeta - \nu_\lambda)}.$$

The corresponding return is:

$$R_{t+1} = \left(\frac{D_{t+1}}{D_t} \right) \exp(-\zeta + \nu_\lambda).$$

Thus a feature of payoffs constructed from the dominant eigenfunction is that the expected returns can be inferred directly from the expected dividend growth. While gross returns are proportional to dividend growth factors, the proportionality factor depends in part on riskiness of the security.

Let R_{t+j}^j denote the return compounded over j time periods. Expected return growth can be inferred directly from expected dividend growth. In particular,

$$\lim_{j \rightarrow \infty} \frac{\log E(R_{t+j}^j | x_t)}{j} = \zeta + \epsilon_\lambda - \zeta + \nu_\lambda = \epsilon_\lambda + \nu_\lambda.$$

So far we have shown how to compute the expected value of a long-horizon tail return implied by an asset pricing model. What matters is the value of λ . Not surprisingly, the deterministic component of the cash flow growth vanishes in our return calculation. To produce a risk-return tradeoff, we must compare this return with a risk-free counterpart. Given our interest in long-horizon risk, the natural benchmark is obtained by setting $\lambda = 0$. The dominant eigenvalue of \mathcal{G}_0 is one and is associated with a constant eigenfunction. Thus $\xi_0 = 0$. When $\lambda = 0$ the only mechanism for cash-flow growth is a positive value of ζ , which has no impact on returns. Thus the long-horizon counterpart to an expected (logarithm of a) risk-free return is ν_0 , and thus risk-premium associated with λ :

$$\epsilon_\lambda + \nu_\lambda - \nu_0.$$

The $\lambda = 0$ return turns out to be the maximal growth return of Bansal and Lehmann (1997). This follows from the work of Alvarez and Jermann (2001) and Hansen and Scheinkman (2003). Alvarez and Jermann (2001) study the holding period returns to long-horizon discount bonds and show that in the limit these holding period returns approximate the maximal growth return of Bansal and Lehmann (1997). Hansen and Scheinkman (2003) show that this limiting return is the return on the dominant eigenfunction for the ($\lambda = 0$) pricing operator. As a consequence it is approximately the long-horizon return on any security with a terminal payoff of the form $\psi(x_{t+j})$ not just a discount bond.

We compute long-run average excess rates of return $\epsilon_\lambda + \nu_\lambda - \nu_0$ implied by the asset pricing model in example 2.1. In appendix B we discuss how these calculations are done. We used the two alternative specifications of the evolution of consumption of section 3.1. For purposes of illustration, we set $\beta = 0.98^{1/4}$, but we consider values for θ of 5, 10 and 20. We also examine how the parameter λ affects the returns.

The results are reported in figure 9. The upper panel of the figure provides results for the model of consumption and corporate earnings with cointegration and the lower panel for the case without cointegration. Notice that in each case the long-run return rises with λ and that the slope of the curve increases with the risk aversion parameter θ . If substantial variation in λ is allowed the model, can predict large differences in long-run returns especially for the case of the model with cointegration between consumption and corporate earnings. In this model there are significant long run shocks to consumption and if the cash flows being priced leverage this effect then expected returns are substantially affected.

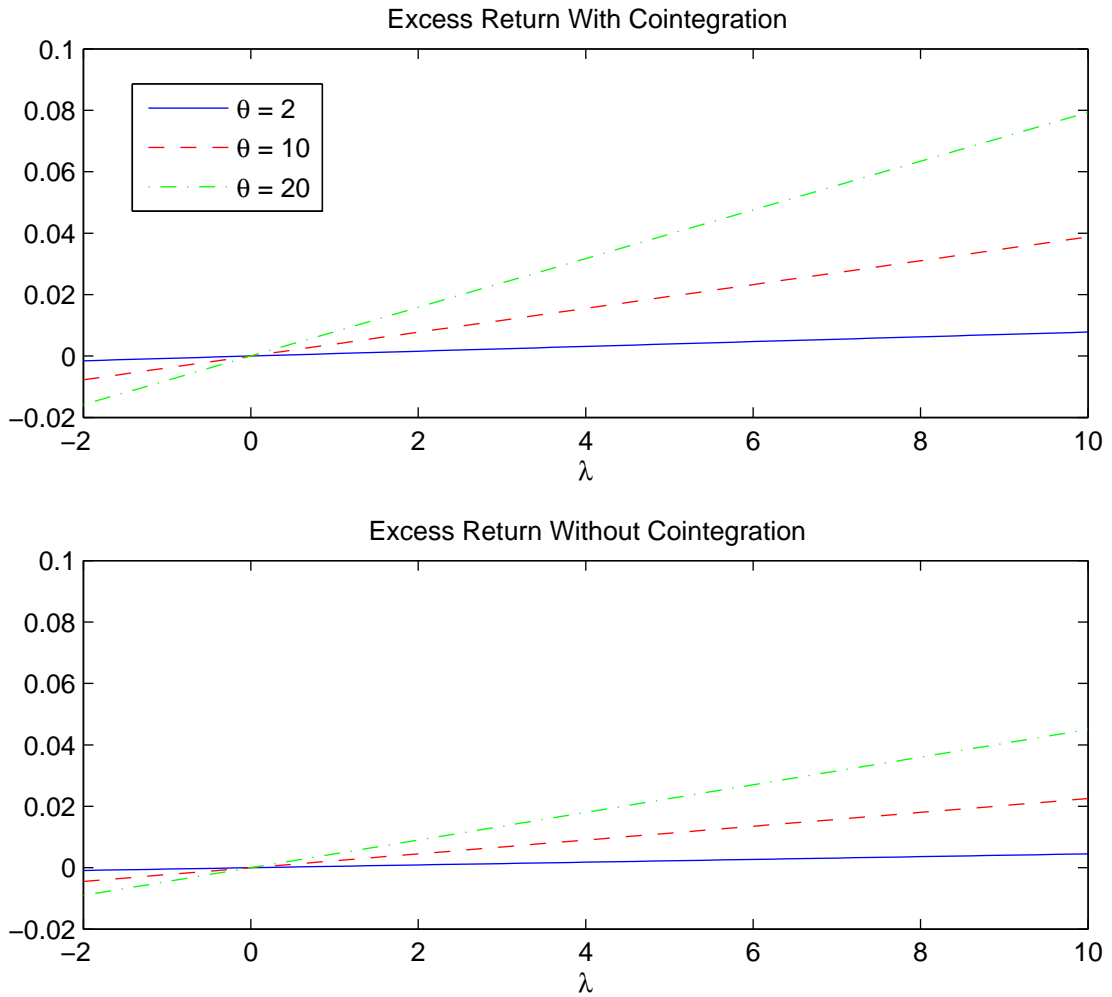


Figure 9: Model implications for long-run return. The plots give curves of the form $\epsilon_\lambda + \nu_\lambda - \nu_0$ as a function of λ . The curves are computed using $\beta = 0.98^{1/4}$ and $\rho = 1$.

The model predicts substantial differences in long run returns across difference measured cash flows if there are corresponding differences in the exposure of cash flows to long run shocks to consumption. We now turn to measurement of this long run risk.

5 Portfolio Dividends

In this section we report estimates of long run dividend growth and the risk associated with the growth. We use the 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.

5.1 Book to Market Portfolios

We follow Fama and French (1993) and construct portfolios of returns by sorting stocks according to their book-to-market values. We use a coarser sort into 5 portfolios to make our analysis tractable. Summary statistics for the basic portfolios are reported in table 1. Notice that 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. Average returns also follow this sort. Portfolio 1 has the lowest average return and portfolio 5 has the highest return. It is well known that the differences in average returns are not well explained by exposure to contemporaneous covariance with consumption. This is reflected in the last row of 1 which reports the correlation between consumption growth and each return. Notice that there is little variation in this measure of risk.

In this section we are particularly interested in the behavior of dividends from the constructed portfolios. The constructed dividend processes accommodate the changes in the classification of the primitive assets and depend on the relative prices of the new and old asset in the book-to-market portfolios. The construction of portfolio dividends is detailed in Hansen, Heaton, and Li (2004) and follows the work of Bansal, Dittmar, and Lundblad (2002), Menzly, Santos, and Veronesi (2004).

Even with this relative price adjustment, the sorting used in constructing the portfolios can induce permanent differences in dividend growth. Indeed the dividends from financial portfolios do not appear to grow one-to-one with consumption. This has been documented in a variety of different places and is evident in figure 10, where we report the logarithms of portfolio dividends relative to aggregate consumption.⁹ The first three portfolios appear to grow slower than consumption, and even market dividends display this same pattern. Portfolios four and five show more pronounced growth than consumption. These features of the time series lead us to explore some alternative specifications of dividend growth.

⁹In an attempt to construct consumption-dividend ratios that are stationary, Menzly, Santos, and Veronesi (2004) divide consumption by population but not dividends. While population is not a simple time trend, its time series trajectory is much smoother than either consumption or dividends.

Table 1: Properties of Portfolios Sorted by Book-to-Market

| | Portfolio | | | | | Market |
|------------------------------|-----------|------|------|------|-------|--------|
| | 1 | 2 | 3 | 4 | 5 | |
| Avg. Return (%) | 6.48 | 6.88 | 8.90 | 9.32 | 11.02 | 7.23 |
| Std. Return (%) | 18.8 | 16.4 | 14.8 | 15.8 | 17.8 | 16.5 |
| Avg. B/M | 0.32 | 0.62 | 0.84 | 1.12 | 2.00 | 0.79 |
| Sharpe Ratio | 0.18 | 0.20 | 0.28 | 0.28 | 0.30 | 0.21 |
| Correlation with Consumption | 0.20 | 0.18 | 0.20 | 0.20 | 0.21 | 0.20 |

Portfolios formed by sorting stocks into 5 portfolios using NYSE breakpoints from Fama and French (1993). Portfolios are ordered from lowest to highest average book-to-market value. Data are quarterly from 1947 Q1 to 2002 Q4 for returns and annual from 1947 to 2001 for B/M ratios. Returns are converted to real units using the implicit price deflator for nondurable and services consumption. Average returns are converted to annual units using the natural logarithm of quarterly gross returns multiplied by 4. The standard deviation of returns is also put in annual units by multiplying the standard deviation of quarterly log gross returns by 2. This assumes that returns are independently distributed over time. “Avg. B/M” for each portfolio is the average portfolio book-to-market over the period computed from COMPUSTAT. The Sharpe Ratio is based on quarterly observations. Correlation with consumption is measured as the contemporaneous correlation between log returns and log consumption growth.




Figure 10: *Portfolio Dividends Relative to Consumption*

5.2 Adding Dividends to the VAR's

Consider a VAR with three variables: consumption, corporate earnings and dividends (all in logarithms). Consumption and corporate earnings are modelled as before in a cointegrated system. We use the cointegrated system because it “identifies” a long run consumption risk component that is distinct from the one-step-ahead forecast error of consumption. In addition to the consumption and earnings we include in sequence the dividend series from each of the five book-to-market portfolios and from the market. Thus the same two shocks as were identified previously remain shocks in this system because consumption and corporate earnings remain an autonomous system. An additional shock is required to account for the remaining variation in dividends beyond what is explained by consumption and corporate earnings.

We consider three different specifications of growth to assess sensitivity to model specification. These three specifications differ in how the dividend evolution equation is specified. We append a dividend equation

$$A_0^* y_t^* + A_1^* y_{t-1} + A_2^* y_{t-2} + \dots + A_\ell^* y_{t-\ell} + B_0^* + B_1^* t = w_t^*, \quad (13)$$

to equation system (7). In this equation the vector of inputs is

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

and the shock w_t^* is scalar with mean zero and unit variance. This shock is uncorrelated with the shock w_t that enters (7). The third entry of A_0^* is normalized to be positive. We refer to (13) as the dividend equation, and the shock w_t^* as the dividend shock. As in our previous estimation, we set $\ell = 5$.

The first specification adds the restriction that the trend coefficient B_1^* equal zero. We use specification (7) in which consumption and earnings are restricted to grow together (cointegration is imposed). Given our interest in measuring long run risk, we measure the permanent response of dividends to the permanent shock. In the long run, both consumption and corporate earnings respond to this shock in the same manner, but the dividend response is left unconstrained. We let λ^* denote the ratio of the long run dividend response to the long run consumption response. We measure this for each of the five portfolios. In this case we allow the matrix:

$$\begin{bmatrix} A(1) & 0 \\ A^*(1) \end{bmatrix} \quad (14)$$

to have rank two where

$$A^*(1) \doteq \sum_{j=0}^{\ell} A_j^* z^j.$$

The cointegrating vector $(1, 1, \lambda^*)$ is in the null space of this rank two matrix.

The second specification includes a time trend by freely estimating B_1^* . We refer to this as the time trend specification. The third specification again restricts B_1^* to be zero, but it introduces an additional growth component in the dividend evolution by restricting:

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

for some scalar α^* . The matrix (14) is now restricted to be of rank one. This allows for an additional stochastic growth component in dividends. We identify a counterpart to λ^* as the long run response of dividends to permanent consumption/earnings shock. We refer to this as the dividend growth rate specification.

We estimate λ^* for each of the five portfolios and each of the dividend growth specifications. The results are given table 2. This table includes quantiles obtained via simulation using the Box-Tiao priors for each of the three equations. For each of the three growth rate specifications, the median estimates of λ^* are larger for the larger book to market portfolios. While this qualitative pattern is roughly preserved across growth specifications, the estimates of λ^* are very sensitive to which model of growth is presumed. The larger estimates are obtained for the dividend growth rate specification and smallest for the time trend specification. Even conditioning on a growth configuration, the sampling variability for λ^* is substantial, particularly for the time trend specification and the dividend growth rate specification.

In the case of the time trend specification, the potentially potent cointegration restriction is offset of the challenge of estimating simultaneously a trend rate of growth in conjunction with λ^* . In the dividend growth rate specification, even though consumption and earnings are restricted to respond in the long run to a permanent shock in identical ways, it remains a challenge to measurement to infer the dividend response to this same shock. Both specification uncertainty and sampling uncertainty make the measurement of λ^* particularly challenging. The role of specification uncertainty is illustrated in the impulse response figure 11. This figure features the responses of portfolio one and five to a permanent shock. For each portfolio, the measured responses obtained for each of the three growth configurations are quite close up to about three to four years and then they diverge. Both portfolios initially respond positively to this shock with peak responses occurring in about seven time periods. The response of portfolio one is much larger in this initial phase. The tail responses differ substantially depending on the growth configuration that is imposed in estimation. The estimated response of portfolio one is eventually negative when time trends are included or an additional stochastic growth factor is included.

To better understand the importance of time trends, figure 12 plots both the level of dividends and the fitted values implied by the “aggregate” innovations to consumption and corporate earnings alone. The presence of a deterministic trends in a log levels specification allows the VAR model to fit the low frequency movements of dividends for portfolios 1 and 2 much better. This low frequency effect is reflected in figure 13 which presents the spectral densities of dividend growth implied by the VAR model with time trends. The dashed lines give the spectral densities implied by all of the shocks and the solid lines give the spectral densities implied by the aggregate innovations alone. For comparison figure 14 presents the

Table 2: Long Run Responses

| No Time Trends | | | | | |
|-----------------------|----------|------|------|------|------|
| Portfolio | Quantile | | | | |
| | .05 | .25 | .5 | .75 | .95 |
| 1 | -1.35 | 0.08 | 0.35 | 0.53 | 0.83 |
| 2 | 0.18 | 0.47 | 0.66 | 0.98 | 2.46 |
| 3 | 0.54 | 0.69 | 0.78 | 0.88 | 1.09 |
| 4 | 0.55 | 0.65 | 0.70 | 0.77 | 0.89 |
| 5 | 0.65 | 1.21 | 1.44 | 1.71 | 3.00 |
| market | 0.36 | 0.44 | 0.49 | 0.53 | 0.61 |

| Time Trend Included | | | | | |
|----------------------------|----------|-------|-------|-------|-------|
| Portfolio | Quantile | | | | |
| | .05 | .25 | .5 | .75 | .95 |
| 1 | -6.24 | -4.26 | -3.42 | -2.72 | -1.77 |
| 2 | -16.47 | -7.28 | -4.70 | -3.15 | -1.52 |
| 3 | -0.64 | 1.21 | 2.09 | 2.84 | 3.87 |
| 4 | 1.10 | 1.70 | 2.07 | 2.44 | 3.05 |
| 5 | -1.66 | 1.66 | 3.87 | 7.32 | 21.18 |
| market | -1.60 | -0.43 | 0.10 | 0.53 | 1.12 |

| Dividend Growth Rate | | | | | |
|-----------------------------|----------|-------|-------|-------|------|
| Portfolio | Quantile | | | | |
| | .05 | .25 | .5 | .75 | .95 |
| 1 | -4.84 | -2.61 | -1.35 | -0.20 | 1.60 |
| 2 | -2.13 | -0.68 | 0.18 | 1.03 | 2.43 |
| 3 | 0.34 | 1.82 | 2.97 | 4.32 | 7.05 |
| 4 | 1.82 | 3.05 | 3.97 | 5.01 | 6.94 |
| 5 | 1.42 | 3.64 | 5.18 | 6.87 | 9.93 |
| market | -0.20 | 0.84 | 1.49 | 2.20 | 3.49 |

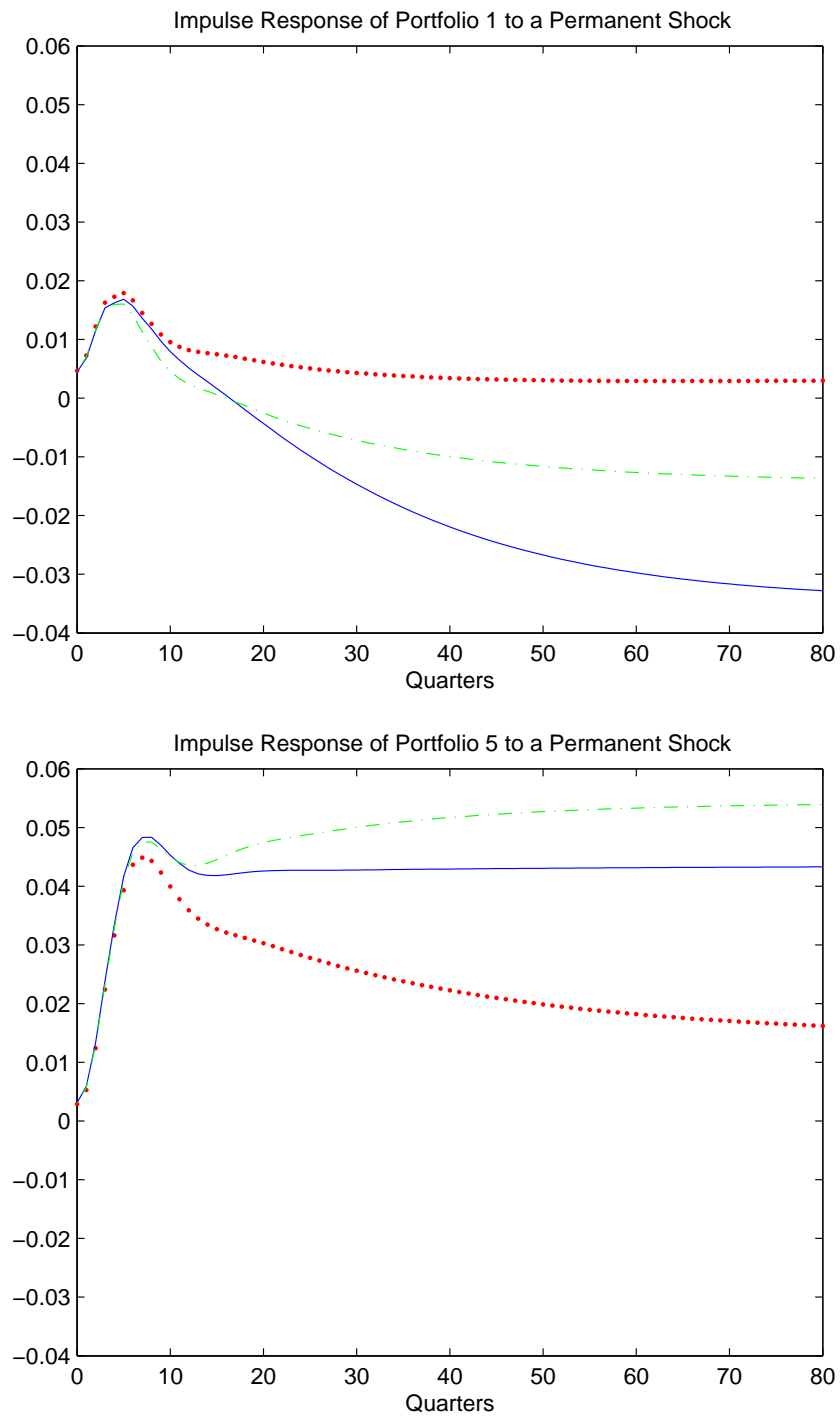


Figure 11: *Impulse response functions for two portfolios.* The \cdots curve is generated from the level specification for dividends; the $—$ is generated from the level specification with time trends included; and the $- - -$ curve is generated from the first difference specification.

analogous figure for the case where the models are restricted to have no time trend. When time trends are included, the aggregate shocks do a much better job of matching the low frequency dynamics of dividend growth for portfolios 1, 2 and 4. In the presence of a time trend, the cointegrating relationship between consumption and dividends better captures the low frequency movements in dividends beyond 16 years.

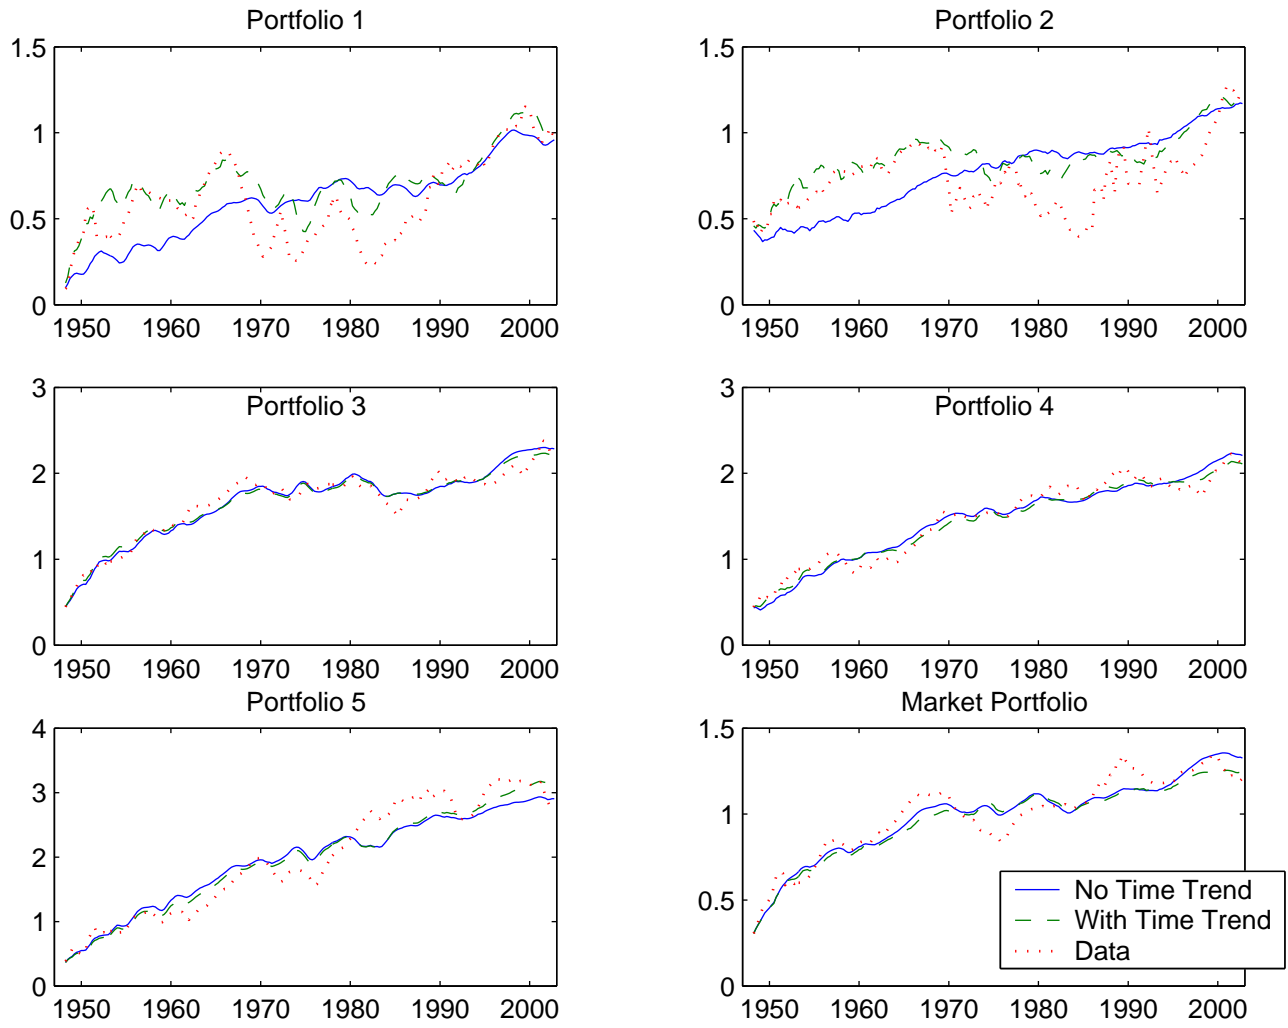


Figure 12: Portfolio dividends relative to fitted values based on aggregate innovations alone.

Up until now, we have taken the linear cointegration model with time trends literally. Is it realistic to think of these as deterministic time trends in studying the economic components of long-run risk? We suspect not. While there may be important components to the cash flows for portfolios 1 and 2 that are very persistent, it seems unlikely that these are literally deterministic time trends known to investors. Within the statistical model, the time trends for these portfolios in part offset the negative growth induced by the cointegration. We suspect that the substantially negative estimates of λ^* probably are not likely to be the true limiting measures of how dividends respond to consumption and earnings shocks. While the long-run risk associated with portfolios 1 and 2 looks very different from that of portfolio five, a literal interpretation of the resulting co-integrating relation is hard to defend.

There is a potential pitfall in using maximum likelihood methods conditioned on initial data points as we have here. Sims (1991) and Sims (1996) warn against the use of such methods because the resulting estimates might imply that

... the first part of the sample behavior of the data is dominated by a large “transient”. That is, the estimates imply that the initial data points are very far from the deterministic trend line or steady state, in the sense that the estimated model implies that future deviations as great as the initial deviation will be extremely rare.

This seems to be particularly true for the models we fit to portfolios 1 and 2. In figure 15 we report the time series trajectory implied by the initial conditions alone with time trends included. These trajectories are depicted by the dot-dashed lines and are generated using the conditional maximum likelihood estimates. The initial conditions appear to be far from the trend lines for portfolios 1 and 2, and as a consequence have a nontrivial trajectory. By conditioning, the maximum likelihood estimates allow for this feature of the model fit. Do we really believe that investors have confidence at the beginning of the sample in such a trajectory? We suspect not.

We investigate what happens when we include initial information into the likelihood function when we approximate posteriors. By assumption, linear combinations of the time series are asymptotically stationary. We now restrict these components to be drawn from the stationary distribution. This gives us additional information to build into the the likelihood function. Table 3 updates the quantiles for portfolios one, two and five. Some of the extreme quantiles are altered and brought closer to zero, but the approximate median of the posteriors are similar with and without this additional information.¹⁰ In particular the first quantile for portfolio 2 is increased from -16.47 to -5.29 and the 5 quantile for portfolio 5 is reduced from 21.18 to 9.57.

In summary, while there is some intriguing heterogeneity in the long run cash flow responses, the measured long run responses are typically not well estimated and sensitive to the growth configuration of the time series.

¹⁰Sims (1991) suggests incorporating additional prior information to prevent over-fitting of initial conditions when trends are included.

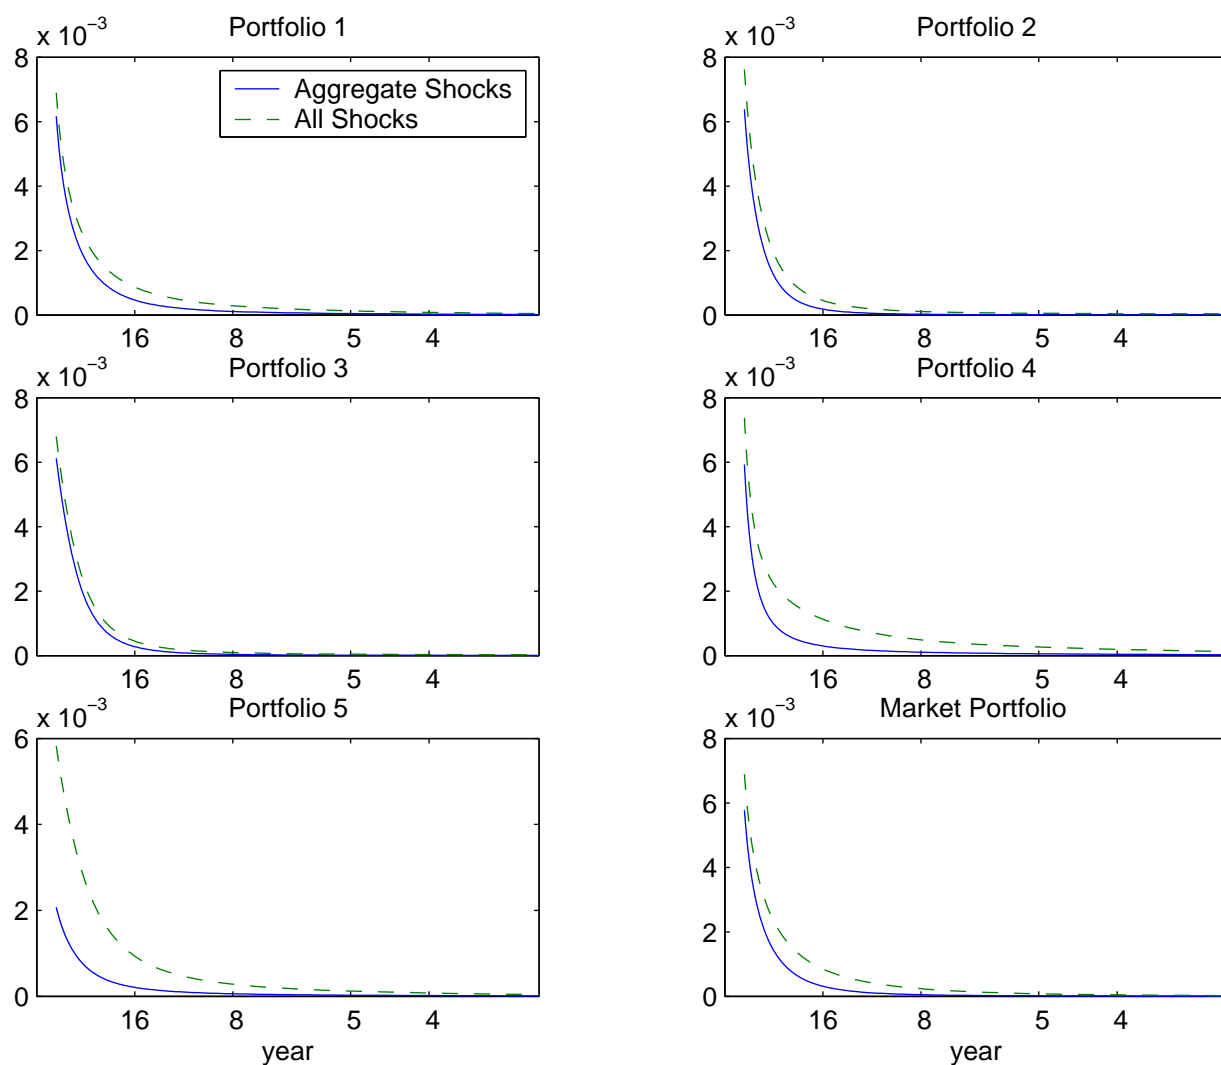


Figure 13: Implied spectral density for dividend growth from VAR's with consumption and aggregate earnings and a time trend in dividends. Densities are normalized by the standard deviation of dividend growth implied by the model. The solid lines give the spectral densities implied by the aggregate shocks to consumption and corporate earnings alone. The dashed lines give the implied spectral densities from the complete model. The x-axes give the number of years that correspond to each frequency.

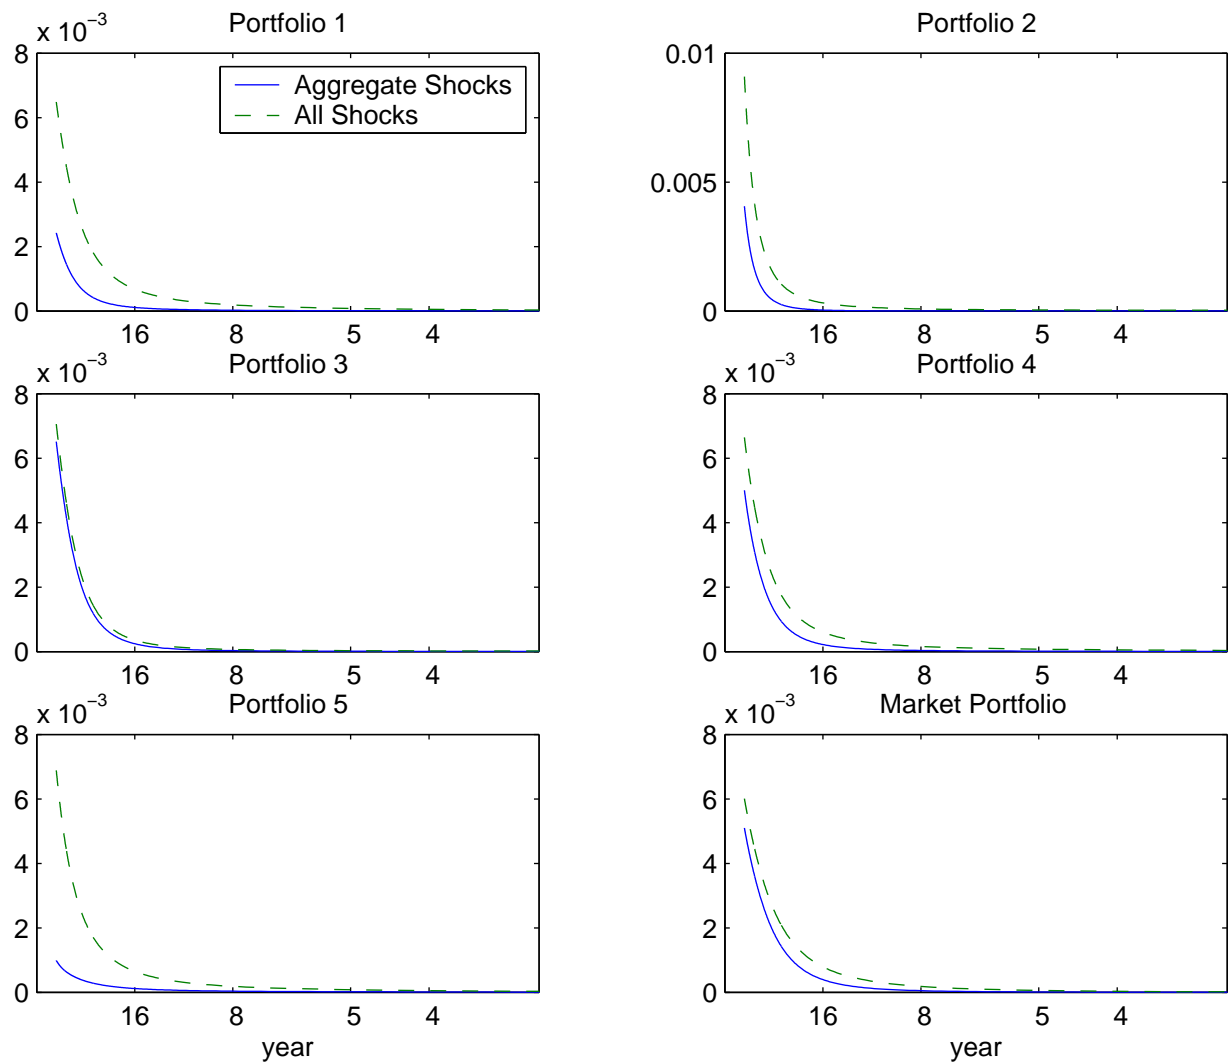


Figure 14: Implied spectral density for dividend growth from VAR's with consumption and aggregate earnings and no time trend in dividends. Densities are normalized by the standard deviation of dividend growth implied by the model. The solid lines give the spectral densities implied by the aggregate shocks to consumption and corporate earnings alone. The dashed lines give the implied spectral densities from the complete model. The x-axes give the number of years that correspond to each frequency.

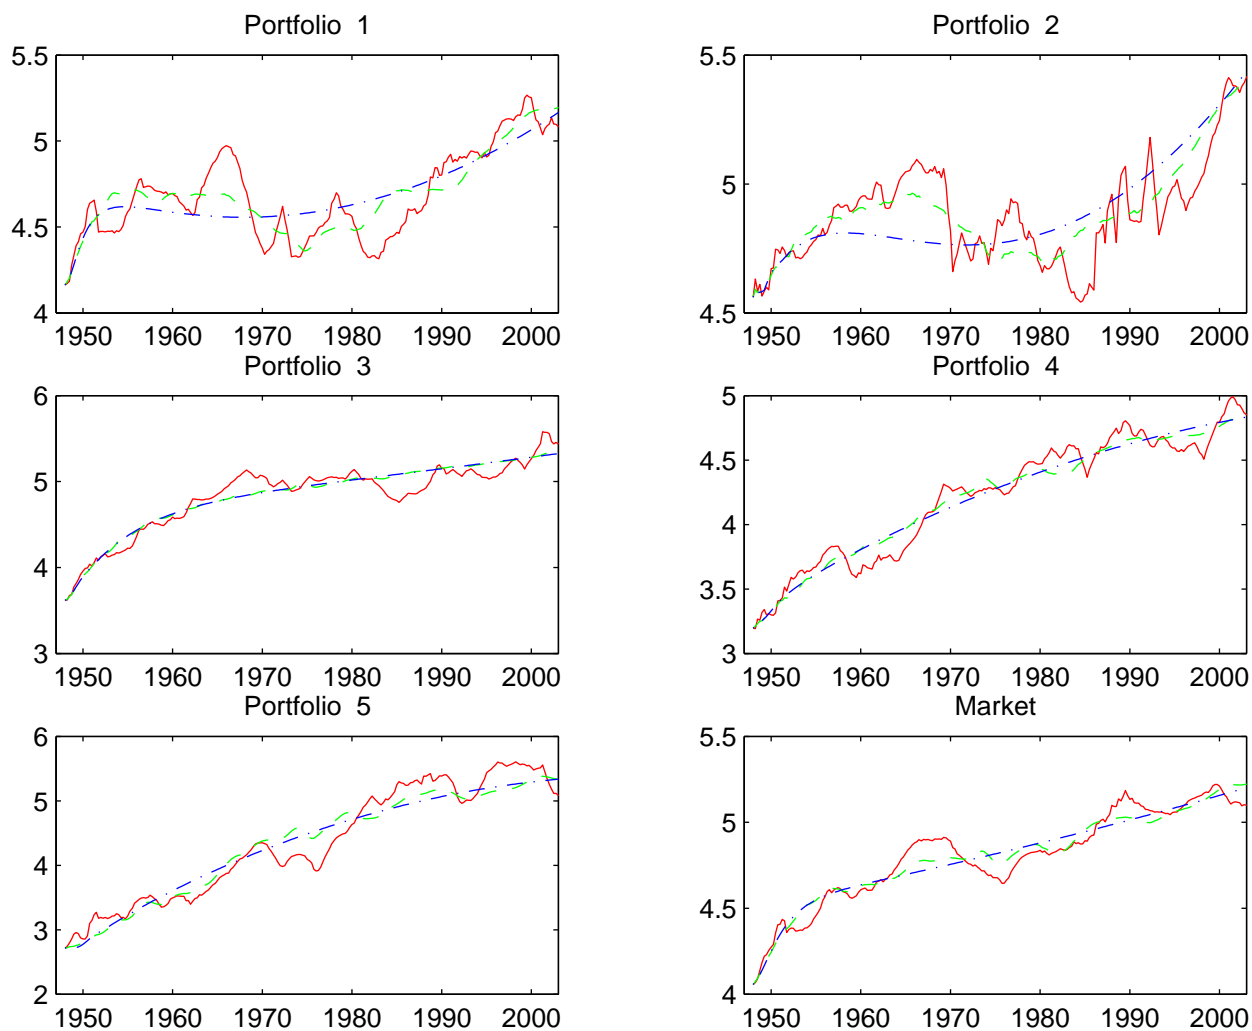


Figure 15: *Portfolio Dividends and Fitted Values*. Solid Lines — display the data. Dashed lines — — — are the fitted values based on consumption shocks alone. Dot-dashed lines — · — · are fitted values with all shocks set to zero .

Table 3: Long Run Responses with Time Trends

| Conditional Likelihood | | | | | |
|-------------------------------|----------|-------|-------|-------|-------|
| Portfolio | Quantile | | | | |
| | .05 | .25 | .5 | .75 | .95 |
| 1 | -6.24 | -4.26 | -3.42 | -2.72 | -1.77 |
| 2 | -16.47 | -7.28 | -4.70 | -3.15 | -1.52 |
| 5 | -1.66 | 1.66 | 3.87 | 7.32 | 21.18 |

| Unconditional Likelihood | | | | | |
|---------------------------------|----------|-------|-------|-------|-------|
| Portfolio | Quantile | | | | |
| | .05 | .25 | .5 | .75 | .95 |
| 1 | -5.49 | -3.97 | -3.23 | -2.59 | -1.70 |
| 2 | -8.64 | -5.59 | -4.05 | -2.86 | -1.46 |
| 5 | -0.90 | 1.67 | 3.36 | 5.42 | 9.57 |

The quantiles from unconditional likelihood were computed by using an accept/reject Markov chain monte carlo simulation based on the likelihood ratio from the initial data. The number of accepted simulations was 12% for portfolio 1, was 11% for portfolio 2, and 9% for portfolio 5.

6 Portfolio Returns

Many asset pricing papers focusing on the pricing of one-period returns. The aim of this literature is to test alternative models of one-period stochastic discount factors. As we have seen long-run consumption risk contributes to the stochastic discount factor under recursive utility. Under log utility this additional factor is the discounted consumption response, $\gamma(\beta)w_{t+1}$ for a discount factor β . In this section we explore the extent to which adding in discounting improves our ability to explain the risk premiums in the book-to-market portfolios.

As precursor to measuring return risk, we add price-dividend ratios to the VAR of consumption, earnings and dividends growth. Let

$$\tilde{y}_t = \begin{bmatrix} y_t^* \\ p_t \end{bmatrix}$$

where p_t is the market value of a portfolio.

$$\tilde{A}_0 \tilde{y}_t + \tilde{A}_1 \tilde{y}_{t-1} + \tilde{A}_2 \tilde{y}_{t-2} + \dots + \tilde{A}_\ell \tilde{y}_{t-\ell} + \tilde{B}_0 = \tilde{w}_t$$

where \tilde{w}_t is a multivariate standard normal. The matrix $\sum_{j=0}^{\ell} \tilde{A}_j$ is restricted to have reduced rank to accommodate two cointegrating vectors. Consistent with our earlier specification, $c_t - e_t$ is presumed to be stationary implying that $[1 \ -1 \ 0 \ 0]$ is a cointegrating vector. We also presume that the dividend-price ratio is stationary implying that $[0 \ 0 \ 1 \ -1]$ is a second cointegrating vector.¹¹ We include market values in the VAR for two reasons. The first is that we can approximate return risk and the second is that we can use dividend-price ratios as forecasters of future consumption. Second, we can deduce the implied return riskiness from the implied riskiness of dividends and capital gains. On the other hand, we are not using the full implications of the model to restrict the market values. The only restrictions imposed on the VAR are the reduced rank restrictions. We do not limit the feedback between series. As with dividends, we examine each portfolio separately to avoid dramatic parameter proliferation.

To derive an implication for returns, we follow Campbell and Shiller (1988) and use the approximation:

$$r_{t+1} = (d_{t+1} - d_t) + \chi + \varrho(p_{t+1} - d_{t+1}) - (p_t - d_t) \quad (15)$$

where ϱ is constructed from the average logarithm of the dividend price ratio:

$$\varrho \doteq \frac{1}{1 + \exp(\bar{d} - \bar{p})}.$$

Campbell and Shiller (1988) show this approximation is reasonably accurate in practice.¹² We use this formula to measure the one-period return response to the shock $\kappa \tilde{w}_{t+1}$ to the shock vector \tilde{w}_{t+1} .

¹¹We impose this restriction by estimating a VAR with $\ell - 1$ lags of $c_t - c_{t-1}$ and $d_t - d_{t-1}$ and ℓ of $e_t - c_t$ and $p_t - d_t$.

¹²Alternatively, we may look separately at the price appreciation and dividend growth components to returns.

Table 4: Return Risk

| $\beta = .98$ | | | | | |
|---------------|-------|-------|----------------|------|------|
| Portfolio | .05 | .25 | Quantile .5 | .75 | .95 |
| 1 | -3.89 | -1.40 | 0.41 | 2.24 | 4.73 |
| 2 | -1.71 | 0.46 | 1.93 | 3.34 | 5.24 |
| 3 | -1.55 | 0.61 | 1.96 | 3.18 | 4.82 |
| 4 | -0.51 | 1.30 | 2.55 | 3.8 | 5.61 |
| 5 | -0.58 | 1.14 | 2.40 | 3.71 | 5.59 |
| $\beta = .99$ | | | | | |
| Portfolio | .05 | .25 | Quantile .5 | .75 | .95 |
| 1 | -4.74 | -2.53 | -0.68 | 1.41 | 4.49 |
| 2 | -2.78 | -0.77 | 0.81 | 2.52 | 4.97 |
| 3 | -2.63 | -0.58 | 0.96 | 2.50 | 4.58 |
| 4 | -1.23 | 0.40 | 1.77 | 3.37 | 5.41 |
| 5 | -1.08 | 0.27 | 1.46 | 2.92 | 5.15 |
| $\beta = 1$ | | | | | |
| Portfolio | .05 | .25 | Quantile .5 | .75 | .95 |
| 1 | -5.16 | -3.09 | -1.64 | 0.18 | 4.23 |
| 2 | -3.25 | -1.53 | -0.44 | 1.17 | 4.37 |
| 3 | -3.12 | -1.32 | -0.27 | 1.27 | 4.09 |
| 4 | -1.51 | -0.31 | 0.64 | 2.55 | 5.05 |
| 5 | -1.24 | -0.22 | 0.39 | 1.84 | 4.60 |

Table 5: Return Risk Differences

| β | Quantile | | | | |
|---------|----------|------|------|------|------|
| | .05 | .25 | .5 | .75 | .95 |
| .98 | -1.03 | 0.30 | 1.20 | 2.13 | 3.56 |
| .99 | -0.70 | 0.54 | 1.33 | 2.18 | 3.60 |
| 1.00 | -0.33 | 0.49 | 1.19 | 2.03 | 3.49 |

Table 6: The quantiles are for the differences in the regression coefficients of portfolio return five and portfolio return one onto long run consumption. These quantile were computed using Box-Tiao priors for a six variate VAR.

Recall that the recursive utility model features the contribution of $\gamma(\beta)\tilde{w}_{t+1}$ to the stochastic discount factor for subjective discount rate β . This term is magnified increasing risk aversion through the choice of θ . The VAR implies a measure $\gamma(\beta)\tilde{w}_{t+1}$. We use this in conjunction with the return risk to measure the regression coefficient:

$$\frac{\kappa \cdot \gamma(\beta)}{\gamma(\beta) \cdot \gamma(\beta)}$$

for each of the five portfolios. We report the resulting measurements and measurement quantiles implied by Box-Tiao priors in table 4.

The medians (.5 quantiles) reported in this table confirm the qualitative findings of Parker and Julliard (2004) using a rather different approach to measurement. The high book-to-market portfolios do tend to me more highly correlated with the long run measure of consumption risk consistent with the measured excess returns. For subjective discount factors near one, the low book-to-market portfolios have one-period returns that are negatively correlated with long run consumption risk. The measurements are not very accurate as reflected by the quantiles, however; and they are sensitive to the choice of discount factor β . Obtaining statistically reliable and meaningful measurement of return risk based on long-run consumption may prove to be an elusive measurement challenge unless more structure is imposed on the economic model.

Fitting VAR's to each portfolio misses some cross correlation patterns. The estimated "regression" coefficients are likely to be correlated because of correlations in the underlying time series of dividends and dividend price ratios. To study accuracy of differences in the risk measures, we fit a six variable VAR by including simultaneously dividend growth and price-dividend ratios of for portfolios one and five along with consumption and earnings. The results are reported in table 5. The medians are positive for the three values of β as are the twenty-five percent quntiles. On the other hand, the .05 quantile is consistently negative.

7 Conclusion

Long run or growth rate variation in consumption or cash flows can have important consequences in asset valuation. Some recent time series evidence supports so called *consumption-based* models by appealing to long-run consumption risk.

Using statistical methods to measure directly long-run cash flow variation is a challenging endeavor, however. Statistical methods typically rely on extrapolating the time series model to infer how cash flows respond in the long-run to shocks. This extrapolation depends on details of the growth configuration of the model, and in many cases these details are defended primarily on statistical grounds. Moreover, the simple linear models we consider are likely to be misspecified. There is pervasive statistical evidence for growth rate changes or breaks in trend lines, but this statistical evidence is difficult to use directly in models of decision-making under uncertainty without some rather specific ancillary assumptions.

There are two complementary responses to this conundrum. One is to resort to the use of highly structured, but easily interpretable, models of long-run growth variation. The other is to exploit the fact that asset values encode information about long-run growth. To break this code requires a reliable economic model of the long-run risk-return relation. We suggest model-based methods for economic characterizations of this relation. These methods give us clues as to what types of models feature or amplify the role of long-run risk. Unfortunately, as yet there is not an empirically well grounded, and economically relevant model of asset pricing to use in deducing investors beliefs about the long-run from values of long-lived assets. Much progress has been made in our understanding of models, but less in understanding the precise nature of long-run growth rate risk in the underlying economy.

A Expansion

We compute the first-order expansion:

$$v_t \approx v_t^1 + (\rho - 1)Dv_t^1$$

where v_t^1 is the continuation value for the case in which $\rho = 1$. We base this calculation on the approximate recursion:

$$v_t \approx \beta \left[\mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t) + (1 - \rho)(1 - \beta^2) \frac{\mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t)^2}{2} \right].$$

Then

$$v_t^1 = \beta \mathcal{Q}_t(v_{t+1}^1 + c_{t+1} - c_t),$$

which is the $\rho = 1$ exact recursion and

$$\begin{aligned} Dv_t^1 &= -\beta(1 - \beta^2) \frac{\mathcal{Q}_t(v_{t+1}^1 + c_{t+1} - c_t)^2}{2} + \beta \tilde{E}(Dv_{t+1}^1 | \mathcal{F}_t) \\ &= -\frac{(1 - \beta^2)(v_t^1)^2}{2\beta} + \beta \tilde{E}(Dv_{t+1}^1 | \mathcal{F}_t) \end{aligned} \quad (16)$$

where \tilde{E} is the distorted expectation operator associated with the density

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

Consider example 2.1. Then

$$(v_t^1)^2 = (x_t)' U_v (U_v)' x_t + 2\mu_v U_v \cdot x_t + (\mu_v)^2.$$

Write:

$$Dv_t^1 = -\frac{1}{2} x_t' \Upsilon_d x_t + U_d \cdot x_t + \mu_d.$$

From (16),

$$\begin{aligned} \Upsilon_d &= \frac{(1 - \beta^2)}{\beta} U_v (U_v)' + \beta G' \Upsilon_d G \\ U_d &= -\frac{(1 - \beta^2)}{\beta} \mu_v U_v - \beta(1 - \theta) G' \Upsilon_d H \gamma(\beta)' + \beta G' U_d \\ \mu_d &= -\frac{(1 - \beta^2)}{2\beta} (\mu_v)^2 - \frac{\beta(1 - \theta)^2}{2} \gamma(\beta) H' \Upsilon_d H \gamma(\beta)' \\ &\quad + \beta(1 - \theta) U_d \cdot [H \gamma(\beta)] - \frac{\beta}{2} \text{trace}(H' \Upsilon_d H) + \beta \mu_d \end{aligned} \quad (17)$$

The first equation in (17) is a Sylvester equation and is easily solved. Given Υ_d , the solution for U_d is:

$$U_d = -(I - \beta G')^{-1} \left[\frac{1 - \beta^2}{\beta} \mu_v U_v + \frac{\beta(1 - \theta)}{2} G' \Upsilon_d H \gamma(\beta)' \right],$$

and given Υ_d and U_d the solution for μ_d is:

$$\mu_d = \frac{-\frac{(1-\beta^2)}{2\beta}(\mu_v)^2 - \frac{\beta(1-\theta)^2}{2}\gamma(\beta)H'\Upsilon_dH\gamma(\beta)' + \beta(1-\theta)U_d \cdot [H\gamma(\beta)] - \frac{\beta}{2}\text{trace}(H'\Upsilon_dH)}{1 - \beta}$$

Finally, consider the first-order expansion of the logarithm of the stochastic discount factor:

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

Recall that the log discount factor is given by:

$$\begin{aligned} s_{t+1,t} &= -\delta - \rho(c_{t+1} - c_t) + (\rho - \theta)[v_{t+1} + c_{t+1} - \mathcal{Q}_t(v_{t+1} + c_{t+1})] \\ &= -\delta - \rho(c_{t+1} - c_t) + (\rho - \theta)[v_{t+1} + c_{t+1} - c_t - \mathcal{Q}_t(v_{t+1} + c_{t+1} - c_t)] \end{aligned}$$

Differentiating with respect to ρ gives:

$$\begin{aligned} Ds_{t+1,t}^1 &= -(c_{t+1} - c_t) + [v_{t+1}^1 + c_{t+1} - c_t - \mathcal{Q}_t(v_{t+1}^1 + c_{t+1} - c_t)] \\ &\quad + (1 - \theta) \left[Dv_{t+1}^1 - \tilde{E}(Dv_{t+1}^1 | \mathcal{F}_t) \right] \\ &= v_{t+1}^1 - \frac{1}{\beta}v_t^1 + (1 - \theta) \left[Dv_{t+1}^1 - \tilde{E}(Dv_{t+1}^1 | \mathcal{F}_t) \right]. \end{aligned}$$

Note that

$$v_{t+1}^1 - \mathcal{Q}_t(v_{t+1}^1 + c_{t+1} - c_t) = U_v \cdot x_{t+1} - \frac{1}{\beta}U_v \cdot x_t + \left(1 - \frac{1}{\beta}\right)\mu_v$$

and

$$\begin{aligned} Dv_{t+1}^1 - \tilde{E}(Dv_{t+1}^1 | \mathcal{F}_t) &= -\frac{1}{2}(Hw_{t+1})'\Upsilon_dHw_{t+1} - (Hw_{t+1})'[\Upsilon_dGx_t - U_d] \\ &\quad + \frac{1}{2}(1 - \theta)^2\gamma(\beta)H'\Upsilon_dH\gamma(\beta)' + (1 - \theta)\gamma(\beta)H'[\Upsilon_dGx_t - U_d] \\ &\quad + \frac{1}{2}\text{trace}(H'\Upsilon_dH) \end{aligned}$$

Recall that a precision matrix is the inverse of a corresponding covariance matrix. The precision matrix for w_{t+1} is given by:

$$I + (\rho - 1)H'\Upsilon_dH$$

provided that this matrix is positive definite. Since the matrix Υ_d is positive semidefinite, it suffices to check if the matrix $I - H'\Upsilon_dH$ is positive semidefinite because ρ is restricted to

be positive. Under the risk-neutral distribution for prices, precision increases in ρ at least locally around unity.

The mean under the risk neutral measure for w_{t+1} is

$$[I + (\rho - 1)H'\Upsilon_d H]^{-1} [-\gamma(0) - (\theta - 1)\gamma(\beta) + (\rho - 1)H'U_v - (\rho - 1)(\theta - 1)(H'U_d - H'\Upsilon_d Gx_t)].$$

This mean can be interpreted as the negative of a risk premia. A component of this mean is the undiscounted (by the risk free rate) price an investor is willing to pay for contingent claim to the corresponding component of the shock w_{t+1} . This mean is state dependent when both ρ and θ are different from unity. Typically, this price will be negative because the investor is risk averse. Risk premium for nonlinear claims to w_{t+1} can be also be computed using both the mean adjustment, the precision adjustment and the corresponding normal adjustment.

B Calculating Eigenvalues for Example 2.1

Consider the first-order autoregressive specification in example 2.1 except that we exclude the state that remains one over time from x_t . Let \hat{x}_t be same as x_t except that it does not include a state that remains one over time, then it still follows a first-order Markov process model

$$\hat{x}_{t+1} = \hat{A}\hat{x}_t + \hat{B}w_{t+1}.$$

where we restrict \hat{A} to have eigenvalues with absolute values that are strictly less than one. We now presume that

$$c_{t+1} - c_t = \hat{U}_c \hat{x}_t + \mu_c.$$

Write:

$$s_{t+1,t} + \lambda(c_{t+1} - c_t) = -\frac{1}{2}(\hat{x}_{t+1})'\Xi_1\hat{x}_{t+1} - \xi_1 \cdot \hat{x}_{t+1} - \frac{1}{2}(\hat{x}_t)'\Xi_2\hat{x}_t - \xi_2 \cdot \hat{x}_t - \xi_0$$

We seek an eigenfunction that is log-quadratic:

$$\log \psi(x) = -\frac{1}{2}x'\Omega x - \omega \cdot x.$$

When $\rho = 1$, the matrices Ξ_1 and Ξ_2 are all zero and the eigenfunction will be log-linear ($\Omega = 0$).

Let $\hat{\Omega} = \Omega + \Xi_1$. Then $\hat{\Omega}$ satisfies the Riccati equation:

$$\Xi_1 + \Xi_2 + \hat{A}'\hat{\Omega}\hat{A} - \hat{A}'\hat{\Omega}\hat{B}(I + \hat{B}'\hat{\Omega}\hat{B})^{-1}\hat{B}'\hat{\Omega}\hat{A} = \hat{\Omega}$$

which is easily solved. When $\rho = 1$, the matrices Ξ_1 and Ξ_2 are all zero, $\hat{\Omega} = 0$ and the eigenfunction will be log-linear ($\Omega = 0$).

Consider next the equation for ω . Let $\hat{\omega} = \xi_1 + \omega$ and

$$A^* = \hat{A} - \hat{B}(I + \hat{B}'\hat{\Omega}\hat{B})^{-1}\hat{B}'\hat{\Omega}\hat{A}.$$

Then

$$(A^*)'\hat{\omega} + \xi_1 + \xi_2 = \hat{\omega}$$

implying that

$$\hat{\omega} = (I - (A^*)')^{-1}(\xi_1 + \xi_2).$$

Finally the equation for eigenvalue is given by

$$\nu_\lambda = -\frac{1}{2}\hat{\omega}'\hat{B}(I + \hat{B}'\hat{\Omega}\hat{B})^{-1}\hat{B}'\hat{\omega} + \xi_0 + \frac{1}{2}\log \det (I + \hat{B}'\hat{\Omega}\hat{B})$$

For this example it is also straightforward to compute the decomposition of the price-dividend ratio (12) provided that the dividend process is can be expressed as:

$$d_t - \lambda c_t - \zeta t = -x_t\Phi_d x_t - \phi_d \cdot x_t - \mu_d.$$

This entails iterating on the Riccati equation for matrices in the quadratic forms, along with the coefficients of the linear and constant terms.

References

- Alvarez, F. and U. Jermann (2001). The size of the permanent component of asset pricing kernels. University of Chicago.
- Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). A quartet of semi-groups for model specification, detection, robustness and the price of risk. *Journal of European Economic Association* 1, 66–123.
- Bansal, R., R. Dittmar, and C. Lundblad (2002). Consumption, dividends, and the cross-section of equity returns. Duke University.
- Bansal, R. and B. Lehmann (1997). Growth optimal portfolio restrictions on asset pricing models. *Macroeconomic Dynamics* 1, 1135–1171.
- Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* forthcoming.
- Blanchard, O. J. and D. Quah (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review* 79, 655–673.
- Breeden, D. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics* 7, 265–296.
- Campbell, J. (1996). Understanding risk and return. *Journal of Political Economy* 104, 298–345.
- Campbell, J. and J. Cochrane (1999). By force of habit. *Journal of Political Economy* 107, 205–251.
- Campbell, J., C. Polk, and T. Vuolteenaho (2003). Growth or glamour? Harvard University.
- Campbell, J. and R. Shiller (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1, 195–227.
- Campbell, J. and T. Vuolteenaho (2003). Bad beta, good beta. Harvard University.
- Campbell, J. Y. (1994). Intertemporal asset pricing without consumption data. *American Economic Review* 83, 487–512.
- Campbell, J. Y. and L. M. Viceira (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford, U. K.: Oxford University Press.
- Chacko, G. and L. M. Viceira (2003). Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. unpublished, Harvard University.
- Constantinides, G. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy* 104, 519–543.
- 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. 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.
- Gibbons, M. and W. Ferson (1985). Testing asset pricing models with changing expectations and an unobservable market portfolio. *Journal of Financial Economics* 14, 217–236.
- Hansen, L. P., J. Heaton, and N. Li (2004). Intangible risk. University of Chicago.
- Hansen, L. P. and S. J. Richard (1987). The role of conditioning information in deducing testable implications of asset pricing models. *Econometrica* 55(3), 587–613.
- Hansen, L. P. and T. Sargent (1995, May). Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control* 40(5), 968–971.
- Hansen, L. P. and J. Scheinkman (2003). Semigroup pricing. Princeton University.
- 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.
- 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 S. Ludvigson (2001a). Consumption, aggregate wealth and expected stock returns. *Journal of Finance* 56, 815–849.
- Lettau, M. and S. Ludvigson (2001b). Resurrecting the (c)capm: A cross-sectional test when risk premia are time-varying. *Journal of Political Economy* 109, 1238–1287.
- 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.
- Parker, J. A. and C. Julliard (2004). Consumption risk and cross-sectional returns. Princeton University.

- Restoy, F. and P. Weil (1998, June). Approximate equilibrium asset prices. NBER Working Paper 6611.
- 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.