

Asset pricing in the frequency domain: theory and empirics

Ian Dew-Becker and Stefano Giglio*

2nd April 2013

Abstract

In many affine asset pricing models, the innovation to the pricing kernel is a function of innovations to current and expected future values of an economic state variable, often consumption growth, aggregate market returns, or short-term interest rates. The impulse response of the priced state variable to various shocks has a frequency (Fourier) decomposition, and we show that the price of risk for a given shock can be represented as a weighted integral over that spectral decomposition. In terms of consumption growth, Epstein–Zin preferences imply that the weight of the pricing kernel lies almost entirely at low frequencies, while internal habit-formation models imply that the weight is shifted to high frequencies. We estimate spectral weighting functions for the equity market semi-parametrically and find that they place most of their weight at low frequencies, consistent with Epstein–Zin preferences. For Treasuries, we find that investors view increases in interest rates at low frequencies and decreases at business-cycle frequencies negatively.

*Federal Reserve Bank of San Francisco and Booth School of Business, University of Chicago. The views in this paper are those of the authors and do not represent the views of the Federal Reserve System or the Board of Governors. We appreciate helpful comments from Rhys Bidder, John Campbell, Nikola Mirkov, Marius Rodriguez, Eric Swanson, and seminar participants at the San Francisco Fed and the University of Bergen.

1 Introduction

This paper studies how risk prices for shocks depend on their dynamic effects on the economy. Theoretical asset pricing models have strong implications for how short- and long-term shocks should be priced, and we empirically estimate how the power of a shock at different frequencies determines its risk price.

Affine models are a workhorse of both theoretical and empirical asset pricing. In these models, innovations to the pricing kernel are linearly related to innovations in economic state variables. This paper shows that many widely used affine frameworks can be written, estimated, and interpreted in the frequency domain. The frequency-domain decompositions give a clear and compact characterization of exactly how the dynamics of the economy affect risk prices and provide sharp tests of competing asset pricing models.

The dynamic behavior of the economy is a key input to asset pricing models. In a representative-agent model with Epstein–Zin (1989) preferences, the risk premium of an asset depends on the covariance of its return with current and expected future consumption growth. For the intertemporal CAPM (Merton, 1973; Campbell, 1993), risk premia depend on covariances with shocks to both current market returns and also future expected returns. And in affine term structure models, we show that risk premia depend on covariances with innovations to current and future short-term interest rates.

In dynamic asset pricing models, then, the price of risk for a shock depends on how it affects the state of the economy in the current period and in the future. The dynamic response of the economy to a shock is represented in the time domain with an impulse response function (IRF). Long-run shocks to consumption growth that have large risk prices under Epstein–Zin preferences, for example, (Bansal and Yaron, 2004) have IRFs that decay slowly.

In this paper we propose and derive a new frequency-domain representation of risk prices. First, we map the IRF of a shock into the frequency domain. A shock that has strong long-run effects has high power at low frequencies, whereas shocks that dissipate rapidly have more power at high frequencies. We refer to the frequency-domain version of the IRF as the impulse transfer function.

Our key result is that the price of risk for a shock depends on the integral of the impulse transfer function weighted by a function $Z(\omega)$ over the set of all frequencies ω . The weighting function, Z , determines how shocks are priced depending on how they affect the economy at different frequencies. In other words, $Z(\omega)$ represents the price of exposure to shocks with frequency ω . In this paper we derive $Z(\omega)$ for various theoretical models and estimate it empirically in both equity and debt markets.

The advantage of studying risk prices in the frequency domain is that Z gives a compact and intuitive measure of how different shocks affect the pricing kernel. For example, under power utility,

the only thing that determines the price of risk for a shock is how it affects consumption today. So Z is perfectly flat across frequencies because cycles of all frequencies receive identical weight in the pricing kernel. Under Epstein–Zin preferences, long-run risks matter, and Z places much more weight at low than high frequencies; in fact, the weight is focused only at the very lowest frequencies. Conversely, for an agent with internal habit formation most of the weight of Z is located at high frequencies.

The weighting function representation we derive is useful for a number of reasons. First, structural models place strong and testable restrictions on the spectral weighting function Z . For example, Epstein–Zin preferences imply that Z is monotone and that the majority of its weight lies near frequency zero. A simple question, then, is whether the weighting function is actually monotone or whether, for example, business-cycle frequencies receive particularly high weight.

Second, the decompositions give a novel description of the determinants of the price of risk. While the literature has long known that non-separabilities in preferences can induce different loadings on consumption cycles of different frequencies, this is the first paper to give a fully analytic characterization of those loadings.

Third, the weighting function shows which aspects of the consumption process one needs to focus on most when calibrating models. Under Epstein–Zin preferences, we find that the key statistics are the unconditional standard deviation and the long-run standard deviation of consumption growth; other aspects of consumption’s dynamic behavior are unimportant. So when considering calibrations, the long-run standard deviation should always be reported; in most recent analyses, it is not.

Fourth, looking at risk prices in the frequency domain allows us to study how the prices vary across different ranges of frequencies in a broader way than what models literally imply. When we look at prices in the frequency domain, we may want to test whether, say, frequencies lower than the business cycle (which can be interpreted as long-run fluctuations in a broad sense) are significantly priced, rather than forcing the model to assume that only the very long-run is priced, as in the standard calibrations of Epstein–Zin models. Similarly, all the theoretical models we study have the feature that Z is monotone, so an obvious empirical question is whether that is the case or whether, for example, business-cycle frequencies carry particularly high weight. Our methodology makes these sorts of tests simple and straightforward.

After characterizing Z theoretically for various consumption-based models, we proceed to estimate it semi-parametrically, making minimal assumptions about its shape. Sections 3 and 4 estimate weighting functions for the equity market and find that it is the low-frequency dynamics of the economy that are relevant for determining risk prices, as predicted by Epstein–Zin preferences. In section 3, we study a range of measures of real activity, including consumption, output, and differ-

ent types of investment. While not all the results are statistically significant, we consistently find that low frequency shocks to the state of the economy drive the pricing kernel. Section 4 prices equity portfolios based on their covariance with short- and long-run shocks to equity market returns. Again, we find (consistent with Campbell and Vuolteenaho, 2004), that it is low-frequency shocks to equity market returns that drive the pricing kernel.

Next, we estimate weighting functions for affine term structure models. In this case we show that the price of risk for a shock depends on how it affects the dynamics of short-term interest rates in the future. Standard term structure models have the problem that the fundamental shocks are only identified up to a rotation, making interpretation of estimated risk prices difficult. Our frequency-domain estimates of Z , on the other hand, are invariant to a rotation of the shocks. So instead of interpreting risk prices for the usual “level”, “slope”, and “curvature” factors, risk prices are interpreted in terms of how investors price shocks to interest rates at low frequencies, business-cycle frequencies, and high frequencies.

We find that the low-frequency shocks to short-term interest rates have a significantly positive price of risk, consistent with the idea that investors want to hedge against persistent increases in interest rates (and, presumably, inflation). At business-cycle frequencies and higher, shocks to interest rates have a negative price of risk, as we would expect given that short-term interest rates are procyclical.

There is very little extant analysis of preference-based asset pricing in the frequency domain. Otrok, Ravikumar, and Whiteman (2002) is the most prominent example.¹ While their paper also presents a spectral decomposition, the object of the decomposition is different from ours. Instead of studying risk prices in the frequency domain, they ask how welfare depends on the spectral density of consumption.²

Empirically, a number of papers study the relationship between asset returns and consumption growth at long horizons as methods of testing the implications of Epstein–Zin preferences.³ These papers essentially assume that changes in expected consumption growth at any date in the future carry the same weight in the pricing kernel since they assume that the pricing kernel is driven by changes in the long-run expectation of consumption (which is what Epstein–Zin preferences imply). Our empirical estimates allow for a much more general specification where shocks to consumption at different horizons may have different risk prices.

The paper is also, in some respects, related to Hansen, Heaton, and Li (HHL; 2008). HHL decompose risk premia on assets in terms of their cash flows at different horizons, essentially deriving

¹See also Yu (2012).

²Moreover, unlike this paper, they do not obtain analytic relationships between the spectrum and asset prices; their results are all generated numerically.

³Parker and Julliard, ; Malloy, Moskowitz, and Vissing-Jorgensen, 2009; Bansal, Dittmar, and Lundblad, 2005; Yu, 2012; among others.

term structures for various types of zero-coupon claims (e.g. consumption claims, as in Lettau and Wachter, 2007). We, on the other hand, decompose risk prices for shocks based on how they affect some economic state variable at different horizons. Whereas HHL (and, relatedly, Alvarez and Jermann, 2005, and Otrok, Ravikumar, and Whiteman, 2007) study the present and future dynamics of the pricing kernel itself, we study how the dynamics of various shocks affect the pricing kernel only in the current period. In their empirical analysis, HHL look at how the dynamics of dividends of value and growth stocks relate to the dynamics of the pricing kernel over time. We have nothing to say about the term structure of discount rates on dividends. Rather, we ask how the dynamics of various shocks to consumption growth (i.e. short and long-run shocks) affect the pricing kernel today.

2 Spectral Decomposition and the Weighting Function

We derive our spectral decomposition of the pricing kernel under two main assumptions. First, the pricing kernel, m_t , depends on the current and future values of a state variable, x (perhaps consumption growth or market returns). Second, the dynamics of the economy are described by a vector moving average process X_t which includes x_t .

Assumption 1: Structure of the SDF.

Denote the log pricing kernel (or stochastic discount factor, SDF) m_t .⁴ We assume that m_t depends on current and future values of some state variable in the economy x_t :

$$m_{t+1} = F(x_t) - \Delta E_{t+1} \sum_{k=0}^{\infty} z_k x_{t+1+k} \quad (1)$$

where x_t is the (scalar) *priced variable*, F is some unspecified function, $\Delta E_{t+1} \equiv E_{t+1} - E_t$ denotes the innovation in expectations, and E_t is the expectation operator conditional on information available on date t . Note that this specification is sufficiently flexible to match standard empirical applications of power utility, habit formation, and Epstein–Zin preferences. It can also accommodate the CAPM and ICAPM. Equation (1) implies that risk prices are constant, but we relax that assumption below.

Assumption 2: Dynamics of the economy.

⁴We do not take a position on whether m_t is the pricing kernel for all markets or whether there is some sort of market segmentation. We also do not assume at this point that there is a representative investor.

x_t is driven by an N -dimensional vector moving average process

$$x_t = B_1 X_t \quad (2)$$

$$X_t = \Gamma(L) \varepsilon_t \quad (3)$$

where X_t has dimension $N \times 1$, L is the lag operator, $\Gamma(L)$ is an $N \times N$ matrix lag polynomial,

$$\Gamma(L) = \sum_{k=0}^{\infty} \Gamma_k L^k \quad (4)$$

and ε_t is an $N \times 1$ vector of (potentially correlated) martingale difference sequences. We refer to ε_t as the fundamental shocks to the economy. Throughout the paper B_j denotes a conformable (here, $1 \times N$) vector equal to 1 in element j and zero elsewhere. We assume without loss of generality that x_t is the first element of X_t . Furthermore, we require $\Gamma(L)$ has properties sufficient to ensure that x_t is covariance stationary with a finite and continuous spectrum.

Putting together the assumptions about m with those about the dynamics of the economy, we can write the innovations to the pricing kernel as function of the impulse-response functions (IRFs) of x_t to each of the fundamental shocks. In particular, for the j th fundamental shock, $\varepsilon_{j,t}$, the IRF of x_t is the set of $g_{j,k}$ for all horizons k defined as:

$$g_{j,k} \equiv \begin{cases} B_1 \Gamma_k B_j' & \text{for } k \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

We can then rewrite the innovation to the SDF as:

$$\Delta E_{t+1} m_{t+1} = - \sum_j \left(\sum_{k=0}^{\infty} z_k g_{j,k} \right) \varepsilon_{j,t+1} \quad (6)$$

and we refer to $(\sum_{k=0}^{\infty} z_k g_{j,k})$ as the price of risk for shock j . In this representation, the effect of a fundamental shock $\varepsilon_{j,t+1}$ on the pricing kernel is decomposed by horizon: for every horizon k , the effect of the shock on m_{t+1} depends on the response of x at that horizon (captured by $g_{j,k}$) and on the horizon-specific price of risk z_k .

Our main result is a spectral decomposition in which the price of risk of a shock depends on the response of x to that shock at each frequency ω and on a frequency-specific price of risk, $Z(\omega)$.

Result 1. *Under Assumptions 1 and 2, we can write the innovations to the SDF as,*

$$\Delta E_{t+1} M_{t+1} = - \sum_j \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} Z(\omega) G_j(\omega) d\omega \right) \varepsilon_{j,t+1} \quad (7)$$

where $Z(\omega)$ is a weighting function depending on the risk prices $\{z_k\}$ and $G_j(\omega)$ measures the dynamic effects of $\varepsilon_{j,t}$ on x in the frequency domain,

$$Z(\omega) \equiv z_0 + 2 \sum_{k=1}^{\infty} z_k \cos(\omega k) \quad (8)$$

$$G_j(\omega) \equiv \sum_{k=0}^{\infty} \cos(\omega k) g_{j,k} \quad (9)$$

Equivalently, the price of risk for a shock can be written as

$$\sum_{k=0}^{\infty} z_k g_{j,k} = \frac{1}{2\pi} \int_{-\pi}^{\pi} Z(\omega) G_j(\omega) d\omega \quad (10)$$

Derivation and discussion

For each shock $\varepsilon_{j,t}$, the set of coefficients $\{g_{j,k}\}$ is the impulse-response function of x_t at different horizons k . Moving into the frequency domain, the first step is to decompose the effects of each shock $\varepsilon_{j,t}$ on the future values of x_t into cycles of different frequencies. To do this, we use the discrete Fourier transform, and define

$$\tilde{G}_j(\omega) \equiv \sum_{k=0}^{\infty} e^{-i\omega k} g_{j,k} \quad (11)$$

If $\varepsilon_{j,t}$ has very long-lasting effects on x , then most of the mass of $\tilde{G}_j(\omega)$ will lie at low frequencies, while if $\varepsilon_{j,t}$ induces mainly transitory dynamics in x , then $\tilde{G}_j(\omega)$ will isolate high frequencies.⁵ We

⁵To be more rigorous about the sense in which \tilde{G} gives weights in terms of cycles of different frequencies, we refer to the spectral representation theorem. Specifically, denote $\bar{x}_{k,t}$ the process induced in x_t if the only shock realizations were for ε_k . That is,

$$\bar{x}_{k,t} = \sum_{j=0}^{\infty} g_{k,j} \varepsilon_{k,t-j}$$

$\varepsilon_{k,t}$ has a spectral representation

$$\varepsilon_{k,t} = \int_{-\pi}^{\pi} e^{it\omega} dZ(\omega)$$

where $dZ(\omega)$ is an orthogonal increment process with constant variance (see, e.g., Priestley, 1981, for a textbook statement and proof of the spectral representation theorem). The spectral representation of $\bar{x}_{k,t}$ is then

$$\bar{x}_{k,t} = \int_{-\pi}^{\pi} e^{it\omega} \sum_{j=0}^{\infty} g_{k,j} e^{-ij\omega} dZ(\omega) = \int_{-\pi}^{\pi} e^{it\omega} \tilde{G}_k(\omega) dZ(\omega)$$

\tilde{G}_k thus determines the magnitude of fluctuations in $\bar{x}_{k,t}$ at frequency ω .

refer to \tilde{G}_j as the *impulse transfer function* of shock j since it is the transfer function associated with the filter $\sum_{k=0}^{\infty} g_{j,k} L^k$.

Using the inverse Fourier transform, the price of risk for shock j is

$$\sum_{k=0}^{\infty} z_k g_{j,k} = \sum_{k=0}^{\infty} \left(z_k \frac{1}{2\pi} \int_{-\pi}^{\pi} \tilde{G}_j(\omega) e^{i\omega k} d\omega \right) \quad (12)$$

Now note that $g_{j,k} = 0$ for all $k < 0$. Therefore, for any $k > 0$ we can rewrite (12) as:

$$\sum_{k=0}^{\infty} z_k g_{j,k} = \frac{1}{2\pi} \int_{-\pi}^{\pi} G_j(\omega) \left(z_0 + 2 \sum_{k=1}^{\infty} z_k \cos(\omega k) \right) d\omega \quad (13)$$

where $G_j(\omega)$ is the real part of $\tilde{G}_j(\omega)$,

$$G(\omega) = \text{re} \left(\tilde{G}_j(\omega) \right) = \sum_{k=0}^{\infty} \cos(\omega k) g_{j,k} \quad (14)$$

In other words, the price of risk for any shock depends on the integral of its response in the frequency domain, $G_j(\omega)$, weighted by a real-valued function $Z(\omega)$, where

$$Z(\omega) \equiv z_0 + 2 \sum_{k=1}^{\infty} z_k \cos(\omega k) \quad (15)$$

We thus have

$$\sum_{k=0}^{\infty} z_k g_{j,k} = \frac{1}{2\pi} \int_{-\pi}^{\pi} G_j(\omega) Z(\omega) d\omega \quad (16)$$

This equation maps element-by-element product of the infinite collections $\{g_{j,k}\}$ and $\{z_k\}$ into a simple integral over a finite range in the frequency domain. This result is closely related to Plancherel and Parseval's theorems, but is not identical because we take advantage of the fact that $g_{j,k} = 0$ for $k < 0$ to ensure that $Z(\omega)$ is real-valued.

The price of risk for shock ε_j thus depends on an integral over the function $G_j(\omega)$, with weights $Z(\omega)$. Recall that for each frequency ω , $G_j(\omega)$ tells us the effect of ε_j on x at frequency ω . $Z(\omega)$ therefore determines the price of risk for any shock to the variable x at frequency ω .

2.1 Examples of impulse transfer functions $G_j(\omega)$

Before proceeding further, it is helpful to see examples of what the impulse transfer looks like for some simple impulse response functions. For the sake of concreteness, suppose for the moment that the priced variable x_t is log consumption growth, Δc_t .

Figure 1 plots the impulse response (IRF) and impulse transfer functions for four different hypothetical shocks. Note that while we are ultimately interested in the effects of the shocks on consumption growth, Δc_t , we plot the IRF in terms of consumption levels, c_t , as they are the more natural way to think about consumption.

The first shock is a simple one-time increase in the level of consumption. This shock has a flat impulse transfer function on consumption growth, indicating it has power at all horizons. The second shock is a long-run-risk type shock, inducing persistently positive consumption growth, with the level of consumption ultimately reaching the same level as that induced by the first shock. In this case, there is much less power at high frequencies, but the power at frequency zero is identical, since $G(0)$ depends only on the long-run effect of the shock on the level of consumption ($G_j(0) = \sum_{k=0}^{\infty} g_{j,k}$).

The next two shocks have purely transitory effects. The third shock raises consumption for just a single period, and we see now zero power at frequency zero and positive power at high frequencies. The fourth shock is more interesting. Consumption rises initially, turns negative in the second period, and returns to its initial level in the third period. The transfer function is again equal to zero at $\omega = 0$, but it now actually has negative power at low and middle frequencies. This is a result of the fact that the impulse response of consumption is actually negative in some periods. The sign of G reflects the direction in which the shock drives consumption. If we had reversed the signs of the impulse responses for the first three shocks, their transfer functions would all have been negative.

3 Weighting functions in consumption-based models

This section applies the analysis above to a range of standard utility functions for which m can be written as a function of innovations to consumption growth. We analyze power utility, models of internal and external habit formation, and Epstein–Zin preferences.⁶ We then estimate weighting functions empirically using data on equity returns.

⁶While these models of preferences are often applied under the assumption of the existence of a representative agent, note that that assumption is not strictly necessary. In particular, the pricing kernel generated by an agent's Euler equation will hold for any market in which he participates. We thus do not concern ourselves, for now, with issues of market completeness or the existence of a representative agent. Of course, we *will* assume a representative agent when testing the model using data on aggregate consumption.

3.1 Weighting functions in theoretical models

3.1.1 Power utility

Under power utility, the log pricing kernel is

$$m_{t+1} = \log \beta - \alpha \Delta c_{t+1} \quad (17)$$

where c_t denotes the log of an agent's consumption, α is the coefficient of relative risk aversion, and $-\log \beta$ is the rate of pure time preference. (17) implies that $z_0 = \alpha$ and $z_k = 0$ for all $k > 0$, and thus the weighting function under power utility is simply

$$Z^{power}(\omega) = \alpha \quad (18)$$

Z^{power} is flat and exactly equal to the coefficient of relative risk aversion. Z^{power} is constant because the only determinant of the innovation to the SDF is the innovation to consumption on date $t + 1$. A shock to consumption growth has the same effect on the pricing kernel regardless of how long the innovation is expected to last.

3.1.2 Habits

Adding an internal habit to the preferences yields the lifetime utility function

$$V_t = \sum_{j=0}^{\infty} \beta^j \frac{(C_{t+j} - bC_{t+j-1})^{1-\alpha}}{1-\alpha} \quad (19)$$

where $C_t = \exp(c_t)$ is the level of consumption and $0 \leq b < 1$ is a parameter determining the importance of the habit. The pricing kernel is

$$\exp(m_{t+1}) = \beta \frac{(C_{t+1} - bC_t)^{-\alpha} - E_{t+1}b(C_{t+2} - bC_{t+1})^{-\alpha}}{(C_t - bC_{t-1})^{-\alpha} - E_t b(C_{t+1} - bC_t)^{-\alpha}} \quad (20)$$

If we log-linearize the pricing kernel in terms of Δc_{t+1} and Δc_{t+2} around a zero-growth steady-state, we obtain

$$\Delta E_{t+1} m_{t+1} \approx -\alpha (b(1-b)^{-2} + 1) \Delta E_{t+1} \Delta c_{t+1} + \alpha b(1-b)^{-2} \Delta E_{t+1} \Delta c_{t+2} \quad (21)$$

With internal habits the pricing kernel depends on both the innovation to current consumption growth and also the change in consumption growth between dates $t + 1$ and $t + 2$. The spectral

weighting function under habit formation is

$$Z^{internal}(\omega) = \alpha(1 + b(1 - b)^{-2}) - \alpha b(1 - b)^{-2} 2 \cos(\omega) \quad (22)$$

The weighting function with habits is equal to a constant plus a negative multiple of $\cos(\omega)$. As we would expect, $Z^{internal}(\omega) = Z^{power}(\omega)$ when $b = 0$.

The left panel of Figure 2 plots $Z^{internal}(\omega)$ for various values of b . Here and in all cases below we only plot Z between 0 and π as is standard, since Z is even across 0 and π . The x-axis lists the wavelength of the cycles, as opposed to the frequency ω . Given a frequency of ω , the corresponding cycle has length $2\pi/\omega$ periods (the smallest cycle we can discern lasts two periods).

As b rises, there are two effects. First, the integral over Z gets larger, and second, its mass shifts to higher frequencies. The latter effect is consistent with the usual intuition about internal habit formation that households prefer to smooth consumption growth and avoid high-frequency fluctuations to a greater extent than they would under power utility.⁷

One lesson from the equation for $Z^{internal}$ is that as long as b is the only parameter we can vary, there is little flexibility in controlling preferences over different frequencies. $\cos(\omega)$ always crosses zero at $\pi/2$, so the pricing kernel will always place higher weight on cycles of frequency greater than $\pi/2$ and relatively less weight on cycles with frequency less than $\pi/2$. Furthermore, $Z^{internal}$ is monotone, regardless of the value of b .⁸

Under external habit formation, the SDF is

$$\exp(m_{t+1}) = \beta \frac{(C_{t+1} - b\bar{C}_t)^{-\alpha}}{(C_t - b\bar{C}_{t-1})^{-\alpha}} \quad (23)$$

where \bar{C} denotes some external measure of consumption (e.g. aggregate consumption or that of an agent's neighbors). In this case, the innovation to the SDF depends only on the innovation to C_{t+1} . So the weighting function with an external habit will be completely flat. Otrok, Ravikumar, and Whiteman (2002) show that the external habit has a strong effect on what weights *utility* places on consumption cycles of different frequencies, but what we show here is the SDF is driven entirely by one-period innovations, so all cycles receive the same weight. The pricing kernel in models with external habit formation, e.g. Campbell and Cochrane (1999), places equal weight on all frequencies.

⁷Otrok, Ravikumar, and Whiteman (2002) obtain a similar result, but in a different manner. Rather than characterize the volatility of the pricing kernel, they characterize the price of a Lucas tree, which is equivalent to simply characterizing lifetime utility as a function of the spectral density of consumption growth. While lifetime utility is important, it is not the same as the price of risk in the economy. Our results are therefore complements rather than substitutes.

⁸One potential way to enrich preferences to allow preferences to isolate smaller ranges of the spectrum may be to allow for more lags of consumption to enter the utility function.

On the other hand, the internal habit models of Constantinides (1990) and Abel (1990) are heavily weighted towards high-frequency fluctuations.

3.1.3 Epstein–Zin preferences

An alternative way of incorporating non-separabilities over time to habits is Epstein and Zin’s (1991) formulation of recursive preferences. In general, under recursive preferences, anything that affects an agent’s welfare affects the pricing kernel. So not only shocks to current and future consumption growth, but also innovations to higher moments will be priced. We begin by focusing on the case where consumption growth is log-normal and homoskedastic. Subsequent sections consider models with stochastic volatility.

Suppose an agent has lifetime utility

$$V_t = \left\{ (1 - \beta) C_t^{1-\rho} + \beta (E_t [V_{t+1}^{1-\rho}])^{1-\rho} \right\}^{\frac{1}{1-\rho}} \quad (24)$$

Campbell (1993) and Restoy and Weil (1998) show that if consumption growth is log-normal and homoskedastic, the stochastic discount factor for these preferences can be log-linearized as

$$\Delta E_{t+1} m_{t+1} \approx - \left(\rho \Delta E_{t+1} \Delta c_{t+1} + (\alpha - \rho) \Delta E_{t+1} \sum_{j=0}^{\infty} \theta^j \Delta c_{t+1+j} \right) \quad (25)$$

where ρ is the inverse elasticity of intertemporal substitution (EIS), and α is the coefficient of relative risk aversion. θ is a parameter (generally close to 1) that comes from the log-linearization of the return on the agent’s wealth portfolio (Campbell and Shiller, 1988).⁹ θ is a measure of impatience: if the agent is highly impatient, then he consumes a large fraction of his wealth in each period and θ is small. In the case where $\rho = 1$, (25) is exact.

For the case of (25), the weighting function is

$$Z^{EZ}(\omega) \equiv \alpha + (\alpha - \rho) \sum_{j=1}^{\infty} \theta^j 2 \cos(\omega j) \quad (26)$$

which can be further simplified using Euler’s formula as

$$Z^{EZ}(\omega) = \rho + (\alpha - \rho) \left(\frac{1 - \theta^2}{1 - 2\theta \cos(\omega) + \theta^2} \right) \quad (27)$$

⁹Specifically, $\theta = (1 + \overline{DP})^{-1}$, where \overline{DP} is the dividend-price ratio for the wealth portfolio (i.e. the consumption-wealth ratio) around which we approximate. θ generalizes the rate of pure time preference somewhat because it also depends on discounting due to uncertainty about future consumption.

Under power utility, $\alpha = \rho$ and $Z^{EZ}(\omega) = \alpha$ is flat, so all frequencies receive equal weight, as discussed above. On the other hand, if $\alpha \neq \rho$, then weights can vary across frequencies due to the second term.

Z^{EZ} is much richer than what we obtain in the case of power utility and it has a number of interesting properties. First, as with power utility, its average value is exactly equal to the coefficient of relative risk aversion,

$$\frac{1}{\pi} \int_0^\pi Z^{EZ}(\omega) d\omega = \alpha \quad (28)$$

So the total weight placed on the spectrum depends only on risk aversion. To the extent that the volatility of the pricing kernel depends on the EIS, it is due only to how $\alpha - \rho$ affects which frequencies receive weight. Moreover, in the special case where consumption follows a random walk with standard deviation σ , $G(\omega) = \sigma$ and the standard deviation of the log pricing kernel is simply $\alpha\sigma$.

Looking at Z^{EZ} , we have

$$Z^{EZ}(0) = (\alpha - \rho) \frac{1 + \theta}{1 - \theta} + \rho \quad (29)$$

$$Z^{EZ}(\pi) = (\alpha - \rho) \frac{1 - \theta}{1 + \theta} + \rho \quad (30)$$

$$\frac{dZ^{EZ}(\omega)}{d\omega} \propto \rho - \alpha \text{ for } \omega \in [0, \pi] \quad (31)$$

For θ near 1, $Z^{EZ}(0)$ is driven by the term $\alpha - \rho$ since $\frac{1+\theta}{1-\theta}$ is large (approaching ∞ as $\theta \rightarrow 1$). Since the integral of Z is always equal to α , the term $\alpha - \rho$ determines the relative weight of Z^{EZ} near frequency zero.

The sign of $dZ^{EZ}/d\omega$ depends only on $\alpha - \rho$. If risk aversion is higher than the inverse EIS – agents prefer an early resolution of uncertainty – then $Z^{EZ}(0)$ is high and Z^{EZ} is decreasing on $(0, \pi)$. If $\alpha < \rho$, then $Z^{EZ}(0)$ is low (or negative) and Z^{EZ} is increasing. So, essentially, $Z^{EZ}(0)$ depends on the magnitude of $(\alpha - \rho)$, and Z^{EZ} then monotonically moves towards ρ as ω increases.

An obvious question is how rapidly Z^{EZ} falls as ω rises above zero. That is, how much of the mass of Z^{EZ} is concentrated at very low frequencies? In the limit as $\theta \rightarrow 1$, i.e. the case where households are indifferent about when consumption occurs, $Z^{EZ}(\omega)$ approaches

$$Z^{EZ}(\omega) = (\alpha - \rho) D_\infty(\omega) + \rho \quad (32)$$

where D_∞ is the limit of the Dirichlet kernel (closely related to the Dirac delta function), with the

key properties

$$D_\infty(\omega) = 0 \text{ for } \omega \neq 0 \tag{33}$$

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} D_\infty(\omega) d\omega = 1 \tag{34}$$

for ω in the interval $[-\pi, \pi]$. For an agent who is effectively infinitely patient, then, two features of the consumption process matter: the permanent innovations at $\omega = 0$ ($\lim_{j \rightarrow \infty} \Delta E_{t+1} c_{t+j}$), which are weighted by α , and all other innovations, which have no effect on $\lim_{j \rightarrow \infty} \Delta E_{t+1} c_{t+j}$, and are weighted by ρ .

Moving away from the limiting case, the right-hand panel of figure 2 plots Z^{EZ} for a variety of parameterizations. The parameterizations are meant to correspond to annual data, so we take $\theta = 0.975$ as our benchmark, which corresponds to a 2.5 percent annual dividend yield. For $\alpha = 5$ and $\rho = 0.5$, we see a large peak near frequency zero, with little weight elsewhere. In fact, half the mass of Z^{EZ} in this case lies on cycles with length of 230 *years* or more, and 75 percent lies on cycles with length 72 years or more. In this parameterization, it is effectively only the very longest cycles in consumption (up to permanent shocks) that carry any substantial weight in the pricing kernel. Purely temporary shocks to the level of consumption (which is what are induced by shocks to monetary policy in standard models, for example) are essentially unpriced.

The line that is highly negative near $\omega = 0$ is for $\alpha = 0.5$ and $\rho = 5$, where households prefer a late resolution of uncertainty. In this case, the mass of Z^{EZ} is still effectively isolated near zero, but because households now prefer an early resolution of uncertainty, Z^{EZ} is negative at that point. The integral of Z^{EZ} is still equal to α , though, so it turns positive at higher frequencies.¹⁰

The final two lines in the right-hand panel of figure 2 plot Z^{EZ} for $\alpha = 5$ and $\rho = 0.5$ with lower values of θ , 0.9 and 0.5. These are values that imply a decidedly unrealistic discount rate, but they help show what is necessary for Epstein–Zin to place any meaningful weight on higher frequencies. Even with $\theta = 0.9$, half the weight of Z^{EZ} is on cycles with length 50 years or more. When we push θ all the way to 0.5, the median cycle finally has length 9 years, roughly corresponding to a long business cycle (whereas under power utility, half the weight is on cycles with length 4 years or more).

3.1.4 Ambiguity averse interpretation

As usual, the analysis of Epstein–Zin preferences naturally also applies to the preferences of an ambiguity averse agent (e.g. Hansen and Sargent, 2001; Barillas, Hansen, and Sargent, 2009).

¹⁰Note, though, that the case where $\rho > \alpha$ is not taken as a benchmark and is not widely viewed as empirically relevant (see, e.g., Bansal and Yaron, 2004)

When the agent has a preference for robustness, he can be viewed as having a reference distribution (the true distribution) and a worst-case distribution, which is what he uses to actually price assets. Under the reference distribution, the agent simply has power utility, so his weighting function would be flat. Under the worst-case distribution, though, he places relatively more weight on certain "bad" states of the world (based on a joint entropy condition on the two distributions). Our weighting function shows the effect of that reweighing in the frequency domain. Agents essentially place more weight on the possibility of the occurrence of low-frequency fluctuations, which gives them a relatively high weight in the function Z^{EZ} .¹¹

3.2 Estimates of weighting functions

We now estimate the weighting function $Z(\omega)$ for consumption growth using the cross-section of equity prices. Estimating Z involves three main steps. First, we need to estimate the dynamics of the economy and identify the fundamental shocks ε_{t+1} and the dynamic response of consumption growth to these shocks. To do this, we represent the dynamics of the economy with a VAR and we estimate it; second, we parametrize the function $Z(\omega)$; third, we estimate the parameters of Z using the cross-section of equity returns.

3.2.1 Step 1: Estimation of the dynamics

We assume the process driving the priced variable, x_t , follows a finite-order VAR,

$$\bar{X}_t = \bar{\Phi}(L) \bar{X}_{t-1} + \varepsilon_t \quad (35)$$

where $x_t = B_1 \bar{X}_t$ is the first element of \bar{X}_t and \bar{X}_t has dimension $N \times 1$. In our benchmark results, x_t is log consumption growth, Δc_t . If the lag polynomial $\bar{\Phi}(L)$ has order K , then we can stack K consecutive observations of \bar{X}_t so that $X_t \equiv [\bar{X}'_t, \bar{X}'_{t-1}, \dots]'$ follows a VAR(1)

$$X_t = \Phi X_{t-1} + \varepsilon_t \quad (36)$$

and $x_t = B_1 X_t$. We estimate this VAR using OLS yielding estimates of Φ and ε_t .

3.2.2 Step 2: Parametrization of the Spectral Weighting Function

The weighting function that we want to estimate, $Z(\omega)$, is potentially infinite-dimensional. However, we only have a finite number of risk prices (one for each estimated shock in ε_t) with which to

¹¹Hansen and Sargent (2007) provide a similar interpretation of their multiplier preferences in the frequency domain for a linear-quadratic control problem.

estimate it. We therefore need to choose a functional form for Z with a finite number of parameters. We consider two specifications, a flexible function motivated by the utility functions discussed above, and a step function.

The utility basis The analysis of the utility functions in the previous sections suggests modeling Z as:

$$Z^U(\omega) = q_1 \sum_{j=1}^{\infty} \theta^j \cos(\omega j) + q_2 + q_3 \cos(\omega) \quad (37)$$

where q_1 , q_2 , and q_3 are unknown coefficients. We call (37) the utility basis because it nests the weighting functions derived from utility-based models. If $q_3 = 0$ (37) matches the weighting function for Epstein–Zin preferences in (26). If $q_1 = 0$, the long-run component that is crucial in the Epstein–Zin case is shut off, and we obtain the weighting function for internal habit formation in (22). Finally, if both $q_1 = 0$ and $q_3 = 0$, then we have the weighting function for power utility.

Note that we have an extra parameter θ here. In the empirical analysis, we experiment with different calibrations for it.¹² If the utility basis is motivated by Epstein–Zin preferences, then a natural choice of θ is one that matches the average dividend/price ratio on the aggregate equity market. We therefore calibrate $\theta = 0.975^{1/4}$ for quarterly data, corresponding to a 2.5 percent annual dividend/price ratio. That choice has the drawback that it isolates only the very lowest frequencies. We therefore also study $\theta = 0.6^{1/4}$, which places more weight on higher frequencies and helps improve the identification.

Because the utility basis is so closely related to the weighting functions we derived under various preference specifications, the constituent functions are already plotted in Figure 2. In particular, the lines in the right-hand panel represent the second function, $\sum_{j=1}^{\infty} \theta^j \cos(\omega j)$, shifted upward by a constant. This function clearly isolates very low frequencies, and the extent to which the lowest frequencies are isolated depends on the parameter θ .

The bandpass basis One advantage of working in the frequency domain is that it is straightforward to estimate risk prices for ranges of frequencies of interest. We simply break the interval $[0, \pi]$ into three intervals, corresponding to business-cycle length fluctuations with wavelength between 6 and 32 quarters (as is standard in the macro literature, e.g. Christiano and Fitzgerald, 2003), and frequencies above and below that window. Under Epstein–Zin preferences, we would expect most of the weight of Z to lie in the range of frequencies below the business cycle, while habit formation implies that the mass should lie at higher frequencies.

We refer to the set of three step functions as the bandpass basis, since $Z(\omega)$ is composed of the

¹²We tried to estimate θ and found it was very poorly identified.

sum of three bandpass filters. Specifically, we define

$$Z^{(a,b)}(\omega) \equiv \begin{cases} 1 & \text{if } a < \omega \leq b \\ 0 & \text{otherwise} \end{cases} \quad (38)$$

For quarterly data, our three basis functions are then $Z^{(0,2\pi/32)}(\omega)$, $Z^{(2\pi/32,2\pi/6)}(\omega)$, and $Z^{(2\pi/6,\pi)}(\omega)$. We therefore estimate the function

$$Z^{BP}(\omega) = q_1 Z^{(0,2\pi/32)}(\omega) + q_2 Z^{(2\pi/32,2\pi/6)}(\omega) + q_3 Z^{(2\pi/6,\pi)}(\omega) \quad (39)$$

3.2.3 Step 3: Estimation of the spectral weighting function

Result 1 and the estimated VAR imply that the innovations to the SDF are:

$$\Delta E_{t+1} M_{t+1} = -W(\bar{q}) \varepsilon_{t+1} \quad (40)$$

for a $1 \times N$ vector W that depends on the parameters $\bar{q} \equiv [q_1 \ q_2 \ q_3]'$. We then estimate the vector \bar{q} using the cross-section of asset prices.

To find $W(\bar{q})$ for a given basis, use back to the VAR representation to write:

$$\Delta E_{t+1} M_{t+1} = - \sum_{k=0}^{\infty} z_k B_1 \Phi^k \varepsilon_{t+1} \quad (41)$$

According to Result 1, the time-domain weights $\{z_k\}$ are transformations of the weighting function,

$$z_k = \begin{cases} \frac{1}{2\pi} \int_{-\pi}^{\pi} Z(\omega) d\omega & \text{for } k = 0 \\ \frac{1}{\pi} \int_{-\pi}^{\pi} Z(\omega) \cos(\omega k) d\omega & \text{for } k > 0 \end{cases} \quad (42)$$

For both the utility and bandpass basis, $Z(\omega)$ is linear in the coefficients \bar{q} , which implies that z_k is also linear in \bar{q} . Specifically,

$$z_k = \bar{q}' H_k \quad (43)$$

where H_k contains the integrals of the basis functions function for Z . For the utility basis,

$$H_0 = \begin{bmatrix} \frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^{\infty} \theta^i \cos(\omega i) d\omega \\ \frac{1}{2\pi} \int_{-\pi}^{\pi} 1 d\omega \\ \frac{1}{2\pi} \int_{-\pi}^{\pi} \cos(\omega) d\omega \end{bmatrix} \quad H_{k>0} = \begin{bmatrix} \frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^{\infty} \theta^i \cos(\omega i) \cos(\omega k) d\omega \\ \frac{1}{2\pi} \int_{-\pi}^{\pi} \cos(\omega k) d\omega \\ \frac{1}{2\pi} \int_{-\pi}^{\pi} \cos(\omega) \cos(\omega k) d\omega \end{bmatrix} \quad (44)$$

For the bandpass basis, we obtain:

$$H_0 = \begin{bmatrix} \frac{1}{2\pi} \int_{-\pi}^{\pi} Z^{(0,2\pi/32)}(\omega) d\omega \\ \frac{1}{2\pi} \int_{-\pi}^{\pi} Z^{(2\pi/32,2\pi/6)}(\omega) d\omega \\ \frac{1}{2\pi} \int_{-\pi}^{\pi} Z^{(2\pi/6,\pi)}(\omega) d\omega \end{bmatrix} \quad H_{k>0} = \begin{bmatrix} \frac{1}{\pi} \int_{-\pi}^{\pi} Z^{(0,2\pi/32)}(\omega) \cos(\omega k) d\omega \\ \frac{1}{\pi} \int_{-\pi}^{\pi} Z^{(2\pi/32,2\pi/6)}(\omega) \cos(\omega k) d\omega \\ \frac{1}{\pi} \int_{-\pi}^{\pi} Z^{(2\pi/6,\pi)}(\omega) \cos(\omega k) d\omega \end{bmatrix} \quad (45)$$

which can be further simplified as a function of sines (without integrals) and then computed numerically. The set of vectors $\{H_j\}$ is thus determined exogenously by the choice of the basis.

Given that $z_k = \bar{q}' H_k$, (41) becomes

$$\Delta E_{t+1} M_{t+1} = -\bar{q}' \left(\sum_{j=0}^{\infty} H_j B_1 \Phi^j \right) \varepsilon_{t+1} \quad (46)$$

$\Delta E_{t+1} M_{t+1}$ is thus a function of the VAR parameters Φ , the innovations $\varepsilon_t = X_t - \Phi X_{t-1}$, and the three parameters q_1, q_2 and q_3 .

The risk prices are then estimated from the asset pricing condition,

$$E[r_{it} - r^f] = -Cov(m_{t+1}, r_{it+1}) \quad (47)$$

$$= E[\bar{q}' u_{t+1} r_{it+1}] \quad (48)$$

Our full set of moment conditions identifying the parameters of the model is

$$G_{t+1}(\Phi, \bar{q}) = \left[\underbrace{(X_{t+1} - \Phi X_t) \otimes X_t}_{\text{VAR moments}}, \underbrace{\mathbf{r}_{t+1} - r_{t+1}^f - \bar{q}' \left(\sum_{j=0}^{\infty} H_j B_1 \Phi^j \right) (X_{t+1} - \Phi X_t) \mathbf{r}_{t+1}}_{\text{Asset pricing moments}} \right] \quad (49)$$

where \mathbf{r}_t is the vector of test asset returns and r_t^f is the risk-free rate.

While we could in principle minimize the GMM objective function for all the parameters simultaneously, that method has the drawbacks that the optimization is difficult to perform (due to the large number of parameters) and that it allows errors in the asset pricing model to affect the VAR estimates. We therefore construct estimates of Φ and \bar{q} by minimizing the two moment conditions separately. That is, Φ is simply estimated through OLS and then \bar{q} is estimated taking Φ as given, using GMM.¹³ Given estimates $\hat{\Phi}$ and $\hat{\bar{q}}$, we construct standard errors using the full set of moments, $G_t(\hat{\Phi}, \hat{\bar{q}})$. The standard errors we report for the risk prices \bar{q} therefore always incorporate uncertainty about the dynamics of the economy through Φ .

¹³Optimizing the full GMM objective function (or even using two-stage GMM) would be more efficient, so our standard errors will in general be larger than if we used a fully efficient method.

We perform the GMM estimation of \bar{q} , taking Φ as given, using either one-step GMM (using the identity matrix to weight the asset pricing moments) or two-step GMM (using the estimated variance-covariance matrix of the moment residuals to construct the weighting matrix for the second step).¹⁴

3.3 Empirical results

3.3.1 Data

The most natural choice for the priced variable, x_t , is consumption growth, but we also explore using other variables: GDP, durable consumption, and investment growth. The rationale for using variables other than consumption, even though we are motivated by consumption-based models, is that to the extent that the pricing kernel is driven by permanent shocks to consumption, permanent shocks to any variable that is cointegrated with consumption should also proxy for the pricing kernel, since the permanent shocks to consumption and any variable it is cointegrated with must be perfectly correlated. That said, households want to smooth consumption compared to income, so we cannot view estimates of the spectral weighting function for aggregates other than consumption as yielding direct tests comparing utility functions. Rather, we interpret them as simply illustrating how the dynamics of the economy are priced. Furthermore, Cochrane (1996) argues that investment growth should price the cross-section of asset returns. Our results on investment are a generalization of his analysis that asks whether and how future dynamics of investment growth are priced.

For the vector of state variables \bar{X}_t , we want variables that are both priced and can forecast our priced variable x_t . Since the number of parameters of the VAR increases quadratically with the dimension of \bar{X}_t , we look for a parsimonious representation. We include in \bar{X}_t the lagged values of the priced variable, and we add the first two principal components of a set of 13 standard forecasting variables: the aggregate price/earnings and price/dividend ratios (with and without an inflation adjustment); the 10 year/3 month term spread; the AAA–Baa corporate yield spread (default spread); the small-stock value spread; the unemployment rate minus its 8-year moving average; RREL, the detrended version of the short-term interest rate that Campbell (1991) finds forecasts market returns; the three-month Treasury yield rate; NTIS, a measure of net aggregate equity issuance (from Goyal and Welch, 2008, as updated on Amit Goyal’s website); the investment-

¹⁴When computing the standard errors incorporating the full estimation uncertainty (according to eq. 49), we take into account the weighting matrix we have used to estimate \bar{q} . We construct the full weighting matrix in the following way. We assign equal weight to all VAR moment conditions (i.e. we use the identity matrix for the block of the weighting matrix that corresponds to the VAR moment conditions). For the block that corresponds to the asset pricing moments, we use the same weighting matrix we used in the estimation of \bar{q} . We set to zero the weight on the cross-product between VAR and asset pricing moment conditions. Finally, we scale the VAR moment conditions by a (common) constant such that on average the block of VAR moments and the block of asset pricing moments get the same weight.

capital ratio, IK; and Lettau and Ludvigson’s (2001) *cay*. Because many of the variables used are only available after 1952, in the analysis that follows we use the quarterly data over the period 1952–2011. Finally, in the analysis that follows we use 3 lags of quarterly data, but results are robust to the choice of the number of lags. Table 1 reports the VAR coefficients using consumption growth as a priced variable.

3.3.2 Parameter estimates

Table 2 reports the estimation results using two-step GMM to estimate the risk prices. The left-hand side uses the set of 25 size and book/market-sorted portfolios, while the right-hand side adds in a set of 49 industry portfolios (both sets of portfolios are obtained from Ken French’s website). For both portfolio sets we estimate both the bandpass basis and the utility basis with $\theta = 0.975^{1/4}$ and $\theta = 0.6^{1/4}$. For the bandpass basis, q_1 corresponds to the price of lower-than-business cycle risks, q_2 to business cycle risks, and q_3 to higher-than-business cycle risks. For the utility basis, q_1 is price the long-run component, q_2 is the constant, and q_3 is the high-frequency component (coefficient on $\cos(\omega)$).

The first set of rows in table 2 reports results obtained using consumption growth as priced variable in the SDF. As is common in the literature, we do not find any statistically significant results, regardless of the set of basis functions or test assets.

However, looking across the other priced variables using the 25 Fama–French portfolios, some consistent results appear. In the bandpass basis, q_1 is statistically significantly positive for durable consumption and the three measures of investment that we examine, implying that investors view positive long-term shocks positively. For the utility basis with $\theta = 0.6^{1/4}$, we find that q_1 is significantly positive for the same set of variables, which makes sense since that function isolates similar frequencies to the ones that are isolated by the low-frequency component. The bandpass and utility bases thus lead to similar qualitative conclusions about what features of economic dynamics are priced in equity markets.

All of that said, $\theta = 0.975^{1/4}$ is the more natural choice since that corresponds to an annual 2.5 percent dividend yield on a consumption claim (Bansal and Yaron, 2004, and Campbell and Vuolteenaho, 2004, use similar values). Reassuringly the signs of the parameter estimates for $\theta = 0.975^{1/4}$ are generally the same as those for $\theta = 0.6^{1/4}$, we just see substantially larger standard errors. Intuitively, this is because $\theta = 0.975^{1/4}$ isolates frequencies for which the impulse transfer functions are particularly poorly identified, making it difficult to obtain significant results.

The right-hand set of columns adds the set of 49 industry portfolios to the sample. All the main results hold in this case. Table 3 shows that the main results hold also when using one-step GMM (that weighs all moments equally), though they appear to be statistically weaker.

3.3.3 Impulse transfer and weighting functions

We now plot the impulse transfer functions, G_j , and estimated weighting functions for two priced variables – aggregate consumption growth and the growth in residential investment. We choose the second variable because it delivers the most highly significant estimates in tables 2 and 3.

Figure 3 plots the impulse transfer functions for the three shocks with the two priced variables. The shaded regions in each figure are 95-percent confidence intervals. There are two key features of the transfer functions to note. First, there are meaningful qualitative differences across the functions in how power is distributed, which helps identify the underlying risk prices. If the transfer functions were all highly similar, then we would not expect to be able to distinguish risk prices across frequencies very well. Looking at the confidence bands, though, it is clear that the transfer functions are poorly estimated near frequency zero. $\omega = 0$ corresponds to the the long-horizon response to each shock, so it is not surprising that it is most difficult to estimate. Nevertheless, the fact that the uncertainty rises so much at very low frequencies helps explain why we have trouble estimating the coefficient on the low-frequency component of the utility basis, especially with $\theta = 0.975^{1/4}$.

Figure 4 plots the estimated spectral weighting functions for consumption and residential investment, estimated with the 25 Fama–French portfolios, using the bandpass basis and the utility basis with $\theta = 0.6^{1/4}$. The weighting functions clearly focus on the low-frequency features of the data, and the confidence bands make it clear that the weight at low frequencies is both statistically and economically significantly higher than at other frequencies for residential investment. For consumption, the point estimates place high weight at low frequencies, but the confidence intervals are wide enough that we cannot draw any strong conclusions. In both panels, the estimates with the bandpass and utility bases are highly similar, illustrating that the results are not sensitive to the precise specification that we use.

4 Weighting functions in returns-based models

4.1 Weighting functions in theoretical models

4.1.1 The CAPM

Under the CAPM, innovations to the SDF are proportional to innovations to the market return,

$$m_{t+1} - E_t m_{t+1} = -\frac{E[r_{m,t+1} - r_{f,t+1}]}{Var(r_{m,t+1} - r_{f,t+1})} (r_{m,t+1} - Er_{m,t+1}) \quad (50)$$

where $r_{m,t+1}$ is the market return. The weighting function under the CAPM is thus simply

$$Z^{CAPM}(\omega) = \frac{E[r_{m,t+1} - r_{f,t+1}]}{Var(r_{m,t+1} - r_{f,t+1})} \quad (51)$$

4.1.2 Epstein–Zin and power utility

In a model with a representative agent with Epstein–Zin preferences (with power utility as a special case) and where consumption growth is log-normal and homoskedastic, Campbell (1993) shows that innovations to the pricing kernel can be written purely in terms of returns on the representative agent’s wealth portfolio,

$$m_{t+1} - E_t m_{t+1} = -\alpha \Delta E_{t+1} r_{w,t+1} - (\alpha - 1) \Delta E_{t+1} \sum_{j=1}^{\infty} \theta^j r_{w,t+1+j} \quad (52)$$

where r_w is the log return of the wealth portfolio of the representative agent. θ is the same log-linearization parameter as in the previous section. Campbell (1993) interprets (52) as a version of Merton’s (1973) intertemporal CAPM because both current returns and changes in the investment opportunity set are priced risk factors.

The weighting function for (52) is

$$Z^{EZ-returns}(\omega) = \alpha + (\alpha - 1) \sum_{j=1}^{\infty} \theta^j 2\cos(\omega j) \quad (53)$$

As $\theta \rightarrow 1$ we obtain the limit

$$Z(\omega) = (\alpha - 1) D_{\infty}(\omega) + 1 \quad (54)$$

with D_{∞} the limit of the Dirichlet kernel, that essentially corresponds to a point mass at 0. All agents, then, regardless of ρ (i.e., regardless of whether they have power utility or more general recursive preferences) place high weight on low-frequency fluctuations in equity returns.

4.1.3 Returns-based asset pricing when we can forecast returns but not consumption

Campbell’s (1993) analysis, and that used in Campbell and Vuolteenaho (2004, 2009), assumes that risk premia are constant and that consumption growth is potentially forecastable. Suppose, alternatively, that we cannot forecast consumption growth at all, and that when we forecast asset returns we are simply forecasting risk premia. For example, return predictability might arise from stochastic volatility (as in Bansal and Yaron, 2004 and Campbell, Giglio, Polk and Turley, 2012)

or time-varying risk aversion (Campbell and Cochrane, 1999; Dew-Becker, 2012). The Campbell–Shiller approximation when consumption is unpredictable reduces to

$$\Delta E_{t+1} r_{w,t+1} = \Delta E_{t+1} \Delta c_{t+1} - \Delta E_{t+1} \sum_{j=1}^{\infty} \theta^j r_{w,t+j+1} \quad (55)$$

and the pricing kernel is

$$\Delta E_{t+1} m_{t+1} = -\alpha \Delta E_{t+1} r_{w,t+1} - \rho \frac{1-\alpha}{1-\rho} \Delta E_{t+1} \sum_{j=1}^{\infty} \theta^j r_{w,t+j+1} \quad (56)$$

This result is notably different from that of Campbell (1993) and equation (52) above, which are derived assuming risk premia are constant. Specifically, if the EIS is greater than 1 ($\rho < 1$), then the coefficient on expected future returns becomes proportional to $-(1-\alpha)$: it has the *opposite* sign than in Campbell (1993) and equation (52). The intuition for this result is as follows. In Campbell (1993), news about high future returns corresponds to an improvement in future expected consumption growth (or, in the ICAPM, the investment opportunity set), which is unambiguously good.¹⁵ If, however, high expected returns are due to high future risk aversion or volatility, then there is only a discounting effect: agents dislike news about high future expected returns because it is associated with low lifetime utility. An increase in risk aversion or volatility is purely bad news.

4.2 Estimation of the weighting function

4.2.1 Methods

Given that the weighting function presented above can be decomposed in two of the three constituents functions that we saw for the case of consumption (and that are plotted in Figure 2), the utility basis representation in the case of returns will simply be:

$$Z^U(\omega) = q_1 \sum_{j=1}^{\infty} \theta^j \cos(\omega j) + q_2 \quad (57)$$

Since we are mostly interested in estimating the pricing of long-run discount rate news, we parametrize the bandpass basis to only include a constant and a long-run component,

$$Z^{BP}(\omega) = q_1 Z^{(0,2\pi/32)}(\omega) + q_2 \quad (58)$$

¹⁵In the terminology of the ICAPM, the cash flow effect from higher expected future consumption growth outweighs the discounting effect that comes from higher discount rates in equation (55).

Like in Campbell and Vuolteenaho (2004), we use a VAR(1) with state vector composed of log excess returns, the price/earnings ratio, term spread and default spread. We use quarterly data from 1926q3 to 2011q2. We estimate the VAR using OLS, and set $\theta = 0.95$ per year.

We then use GMM as above to estimate the two parameters q_1 and q_2 using the cross-section of 25 Fama-French assets or the combination of those assets and the 49 industry portfolios. As before, we estimate Φ and $\bar{q} \equiv [q_1, q_2]$ separately. Again, we report results using both one-step and the two-step GMM to estimate \bar{q} , and compute standard errors for \bar{q} taking into account the uncertainty related to the estimation of the VAR parameters as explained in Section 3. For robustness, we also compute the results using the three-parameters bandpass and utility basis we presented in Section 3.

4.2.2 Results

Table 4 shows the results using only the 25 FF assets (left columns) or adding the 49 industry portfolios (right columns). The top panel reports the version with two parameters (where the first one captures the long-run risks) discussed in the previous section. For both the bandpass basis and the utility basis, we find evidence that the long-run component of discount rate news is priced, at least when using only two parameters and using the efficient matrix to estimate q . Consistent with equation (53), when we use the utility basis we find that both the constant and the discount-rate news (long-run component) are priced, and that q_1 is approximately equal to $q_2 - 1$. Similarly, for the bandpass basis, the price for frequencies below the business cycle is positive and significant. The bottom panel of Table 4 reports estimates of the three-parameter version described in Section 3. Here we find much weaker evidence that long-run innovations in returns are priced. Overall, then, looking at returns we find mild evidence that news about the long-term expected returns carry a positive risk price, though the results seem to be more sensitive to the specification used than in the previous section.

5 Multiple priced variables and stochastic volatility

So far the analysis has focused only on the case where there is a single priced variable. In some models, though, the dynamics of multiple variables matter for asset pricing. For example, in many applications with Epstein–Zin preferences, both consumption growth and variation in volatility or disaster risk are priced (e.g. Bansal and Yaron, 2004; Campbell et al., 2012; Gourio, 2012). It turns out that the results above map easily into a multivariate setting.

Assumption 1a: Structure of the SDF

Instead of there being a single priced variable x_t , suppose there is an $M \times 1$ vector of priced

variables, \vec{x}_t , with

$$m_{t+1} = F(\vec{x}_t) - \Delta E_{t+1} \sum_{k=0}^{\infty} Z_k \vec{x}_{t+1+k} \quad (59)$$

where Z_k is a $1 \times M$ vector of weights and $F(\vec{x}_t) : \mathbb{R}^M \rightarrow \mathbb{R}$ is a scalar valued function.

Assumption 2a: Dynamics of the economy

We assume that \vec{x}_t is driven by a vector moving average process as before,

$$\vec{x}_t = JX_t \quad (60)$$

$$X_t = \Gamma(L) \varepsilon_t \quad (61)$$

for some matrix J of dimension $M \times N$.

The appendix derives the following extension to Result 1,

Result 2. *Under Assumptions 1a and 2a, we can write the innovations to the SDF as,*

$$\Delta E_{t+1} M_{t+1} = - \sum_j \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \vec{Z}(\omega) G(\omega) d\omega \right) \varepsilon_{j,t+1} \quad (62)$$

where (ω) is a vector-valued weighting function and $G(\omega)$ measures the dynamic effects of $\varepsilon_{j,t}$ on x in the frequency domain,

$$\vec{Z}(\omega) \equiv Z_0 + 2 \sum_{k=1}^{\infty} Z_k \cos(\omega k) \quad (63)$$

$$G(\omega) \equiv \sum_{k=0}^{\infty} \cos(\omega k) \bar{g}_k \quad (64)$$

and \bar{g}_k is the impulse response function,

$$\bar{g}_k \equiv J\Gamma_k$$

In this case, then, we have multiple variables whose impulse responses we track in G , and each of the priced variables has its own weighting function, represented as one of the elements of $\vec{Z}(\omega)$. We can thus also write the price of risk for a shock as

$$\sum_{k=0}^{\infty} Z_k \bar{g}_k B_j = \frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_m \vec{Z}_m(\omega) G_{m,j}(\omega) d\omega$$

where $\vec{Z}_m(\omega)$ denotes the m th element of $\vec{Z}(\omega)$ and $G_{m,j}(\omega)$ denotes the m, j th element of $G(\omega)$.

The M weighting functions each multiply N different impulse transfer functions, $G_{m,j}(\omega)$. The price of risk for shock j depends on how it affects the various priced variables at all horizons.

5.1 Epstein–Zin with stochastic volatility

Using Result 2, we now extend the results on Epstein–Zin preferences to also allow for stochastic volatility. We use the same log-normal and log-linear framework as above. The log stochastic discount factor under Epstein–Zin preferences is,

$$m_{t+1} = -\rho \frac{1-\alpha}{1-\rho} \Delta c_{t+1} + \frac{\rho-\alpha}{1-\rho} r_{w,t+1} \quad (65)$$

where $r_{w,t+1}$ is the return on a consumption claim on date $t+1$. Whereas we previously assumed that consumption growth was log-normal and homoskedastic, we now allow for time-varying volatility driven by a variable σ_t^2 . We assume that σ_t^2 follows a linear, homoskedastic, and stationary process.

The volatility term may affect consumption growth in myriad ways. For example, consumption growth might be a simple ARMA process with σ_t^2 driving the volatility of the innovations. Alternatively, consumption growth might follow the long-run risk process from Bansal and Yaron (2004), and σ_t^2 could affect either the persistent or transitory consumption growth. Regardless, though, it is straightforward to show that we will have

$$E_t r_{w,t+1} = k_0 + \rho E_t \Delta c_{t+1} + k_1 \sigma_t^2 \quad (66)$$

where k_0 and k_1 are constants that depend on the underlying process driving consumption growth. Using the Campbell–Shiller approximation, we can then write the innovation to the SDF as

$$\Delta E_{t+1} m_{t+1} = -\alpha \Delta c_{t+1} - (\alpha - \rho) \Delta E_{t+1} \sum_{j=1}^{\infty} \theta^j \Delta c_{t+1+j} \quad (67)$$

$$-\frac{\rho-\alpha}{1-\rho} \Delta E_{t+1} \theta k_1 \sigma_{t+1}^2 - \frac{\rho-\alpha}{1-\rho} \Delta E_{t+1} \sum_{j=1}^{\infty} \theta^j \theta k_1 \sigma_{t+j+1}^2 \quad (68)$$

The weighting functions for consumption growth and volatility are now

$$Z_C^{EZ-SV}(\omega) = \alpha + (\alpha - \rho) \sum_{j=1}^{\infty} \theta^j 2 \cos(\omega j) \quad (69)$$

$$Z_{\sigma^2}^{EZ-SV}(\omega) = \theta k_1 \frac{\rho-\alpha}{1-\rho} \left(1 + \sum_{j=1}^{\infty} \theta^j 2 \cos(\omega j) \right) \quad (70)$$

In the case where $\rho = 0$, Z_C^{EZ-SV} is exactly proportional to $Z_{\sigma^2}^{EZ-SV}$. In any case, even for $\rho > 1$ they are highly similar. Z_C^{EZ-SV} is in fact the same we obtained in the homoskedastic case. Both weighting functions have a constant and also allow for a point mass near zero. $Z_{\sigma^2}^{EZ-SV}$ always has the same basic shape regardless of the value of ρ : unless we are in the particular case $\rho = \alpha$ in which $Z_{\sigma^2}^{EZ-SV}(\omega) = 0$, agents always place high weight on the low-frequency features of volatility.

Alternatively, the weighting functions can be written in terms of returns and their volatility,

$$Z_R^{EZ-SV-R}(\omega) = \alpha - (1 - \alpha) \sum_{j=1}^{\infty} \theta^j 2 \cos(\omega j) \quad (71)$$

$$Z_{\sigma^2}^{EZ-SV-R}(\omega) = \theta k_1 \frac{1 - \alpha}{1 - \rho} \left(1 + \sum_{j=1}^{\infty} \theta^j 2 \cos(\omega j) \right) \quad (72)$$

which yields conceptually very similar results.

6 Spectral weighting functions in affine term structure models

The spectral analysis laid out in Section 2 applies to affine asset pricing models, so it can also be used to understand models of the term structure. This section shows that standard essentially affine asset pricing models can be recast in the frequency domain. The risk prices are reinterpreted as weighting functions in terms of shocks to short-term interest rates. The usual “level factor”, for example, corresponds to low-frequency shocks to short-term interest rates.

Our method applies to affine and essentially affine (Duffee, 2002) models. It can accommodate standard yields-only models, models with macro factors (Ang and Piazzesi, 2003), and models with hidden factors that are not reflected in the yield curve (Duffee, 2011). In all these cases, we show that the risk premium for a given shock (e.g. a shock to the level factor) depends on its dynamic effects on the short-term interest rate.

It is in general difficult to interpret risk prices in term structure models because they are only identified up to a rotation (at least in yields-only models). That is, the underlying factors, and consequently their risk prices, can be rotated without having any observable effect on the dynamics of bond prices. Some papers identify the shocks by assuming, for example, that the coefficient matrix in the VAR driving the factors is lower triangular, while others assume the factors are principal components or that the factors are fixed functions of the yields (Hamilton and Wu, 2012, Joslin, Singleton, and Zhu, 2011, and Christensen, Diebold, and Rudebusch, 2011, respectively). These restrictions are not necessarily economically motivated, however, and there is no guarantee that

they will deliver economically interpretable factors (we are fortunate that the principal components in bond yields seem to have a “level” and “slope” factors that we can tell intuitive stories about). Moreover, it is not obvious how the estimated risk prices can be compared across samples, even if the identifying assumptions are the same (the three principal components in one sample will not be identical to those in another sample). Our spectral decomposition, on the other hand, has a clear and stable interpretation and the risk prices can be easily compared across various sample periods or datasets.

6.1 A canonical term structure model

We consider here a standard yields-only model in which the factors follow a homoskedastic VAR(1) because it is widely studied in the literature. The analysis here can be easily generalized to other settings.

Suppose the state of the economy is summarized by a vector X_t (with dimension $N \times 1$) that follows a VAR(1),

$$X_t = \rho X_{t-1} + \varepsilon_t \quad (73)$$

where ε_t is a vector of mean-zero normally distributed random variables. The SDF takes the essentially affine form,

$$m_{t+1} = -\delta_0 - \delta'_1 X_t - \frac{1}{2} \lambda'_t \lambda_t - \lambda'_t \varepsilon_{t+1} \quad (74)$$

$$\lambda_t \equiv \lambda + \Lambda X_t \quad (75)$$

where the short-term interest rate, r_t , follows

$$r_t = \delta_0 + \delta'_1 X_t \quad (76)$$

We now show that for a model of this form, the innovation to the SDF can be written as

$$\Delta E_{t+1} m_{t+1} = - \sum_{k=0}^{\infty} (z_{0,k} + z_{1,k} X_t) \Delta E_{t+1} r_{t+k+1} \quad (77)$$

where $z_{0,k}$ is a scalar and $z_{1,k}$ is a $1 \times N$ vector of coefficients. First, suppose that (77) is the true model of the SDF. The innovation to future expected values of the short rate is

$$\Delta E_{t+1} r_{t+k+1} = \delta'_1 \rho^k \varepsilon_{t+1} \quad (78)$$

and (77) can then be rewritten as

$$\Delta E_{t+1} m_{t+1} = - \left(\sum_{k=0}^{\infty} (z_{0,k} + z_{1,k} X_t)' \delta'_1 \rho^k \right) \varepsilon_{t+1} \quad (79)$$

So in order for (74) and (77) to be equivalent representations of the pricing kernel, we simply need to be able to solve the pair of vector equations

$$\sum_{k=0}^{\infty} z'_{0,k} \delta'_1 \rho^k = \lambda' \quad (80)$$

$$\sum_{k=0}^{\infty} z'_{1,k} \delta'_1 \rho^k = \Lambda' \quad (81)$$

Equations (80) and (81) give us a way to directly link a vector of risk prices λ_t into the frequency domain in terms of sets of coefficients $\{z_{0,k}\}$ and $\{z_{1,k}\}$. As above, we can map the coefficients $z_{k,t} = z_{0,k} + z_{1,k} X_t$ into the frequency domain through the cosine transform,

$$Z_t(\omega) = z_{0,t} + 2 \sum_{k=1}^{\infty} z_{k,t} \cos(\omega j) \quad (82)$$

where the subscript on $Z_t(\omega)$ denotes the dependence of the weighting function on time (which follows from the fact that the risk prices vary over time through the term ΛX_t).

In standard term structure models (where the factors can be recovered from bond yields alone), the shocks are not uniquely identified – any full-rank rotation produces an observationally equivalent model, which means that the risk prices are also not uniquely identified. It is thus not obvious how to interpret the risk prices. The spectral representation of the risk prices, $Z_t(\omega)$, is invariant to a rotation of the shocks and thus can be interpreted without any ambiguity. For example, standard intuition tells us that highly persistent increases in nominal interest rates are associated with persistent increases in inflation and are generally viewed negatively. So we would expect Z_t to generally take on negative values for low values of ω . An increase in interest rates just at business-cycle frequencies, though, tends to represent good news as short-term interest rates are empirically procyclical, so we would expect Z_t to be positive at business cycle frequencies.

6.2 Empirics

We start our empirical analysis by estimating a standard three-factor model using the method of Hamilton and Wu (2012). We assume that three yields are observed perfectly and one is measured with error (as in, e.g., Duffee, 2002; Kim and Wright, 2005; Joslin, Singleton, and Zhu, 2011). We

use data on 1, 12, 36, and 60-month yields from 1980–2003 (a period with stable monetary policy where the zero lower bound does not bind) and assume that the 36-month yield is measured with error.

We only report a single set of estimates here. In experiments with various specifications, we have found that the results are sensitive to choices about the yields used, the sample period, and details of the estimation (this is true of both the reduced-form and spectral risk prices). For now, then, the results here are meant to be illustrative of the method, rather than definitive.

We use the same bandpass basis for pricing bonds as for equities. Specifically, we solve the equations,

$$\sum_{k=0}^{\infty} K_0 \begin{bmatrix} F_k^{(0,2\pi/32)} \\ F_k^{(2\pi/32,2\pi/6)} \\ F_k^{(2\pi/6,\pi)} \end{bmatrix} \delta'_1 \rho^k = \lambda' \quad (83)$$

$$\sum_{k=0}^{\infty} K_1 \begin{bmatrix} F_k^{(0,2\pi/32)} \\ F_k^{(2\pi/32,2\pi/6)} \\ F_k^{(2\pi/6,\pi)} \end{bmatrix} \delta'_1 \rho^k = \Lambda' \quad (84)$$

for K_0 and K_1 . The vector K_0 gives the set of steady-state frequency-domain risk prices, and the matrix K_1 determines how those risk prices respond to the factors X_t .

6.3 Results

We begin by reporting and interpreting the risk prices in the usual way. First, the left-hand panel of figure 5 plots the loadings of bond yields with maturities from 1 to 60 months on the three factors. The factors are identified as principal components – that is, we rotate them so that they are independent and have unit variance. The first factor can be interpreted as a level factor since it affects all yields roughly equally. The second factor is generally thought of as affecting the slope, and the third factor is a curvature factor.

The center panel of figure 5 plots the response of the 1-month interest rate to shocks to the three factors. As we would expect, the level shock has persistent effects, and the curvature shock has a hump-shaped response. Interestingly, the shock to the slope factor has strongly persistent effects on interest rates. Even though the slope of the yield curve rises, the short-term interest rate is driven persistently lower, which implies that the slope factor mainly represents a shift in the term premium.

The right-hand panel of figure 5 plots the corresponding impulse transfer functions. For the sake of readability, the horizontal axis only covers cycles with length greater than 12 months (i.e.

$\omega < 2\pi/12$). The key result in this plot is that the three shocks have noticeably different impacts in the frequency domain. For example, all three have positive effects at business-cycle frequencies, while the only the level factor has positive effects at low frequencies. This result implies that the transfer functions of the three factors are not highly collinear, which should help identification of the spectral weighting functions.

Table 5 reports the estimated risk prices for the three reduced-form shocks. λ represents the steady-state reduced-form risk prices. The risk prices for all three shocks are statistically significant. The shock to the level factor has a negative price, implying that, conditional on slope and curvature, periods when the level of the term structure is high are viewed as bad times – investors want to purchase assets that insure them against the possibility of a decline in the level factor. This is consistent with the view that long-term bonds are risky because they are exposed to persistent shifts in inflation (e.g. Bekaert, Cho, and Moreno, 2010).

The slope factor, on the other hand, has a positive price, implying that periods when the term structure is upward sloping are viewed as good times (conditional on the level and curvature), a result that runs against our priors. Periods in which the slope of the term structure is high are usually times when the economy is depressed (the term spread is highly correlated with the unemployment rate, for example), so periods of a high slope should be “bad times”.¹⁶

Table 5 also reports time parameters determining time-variation in risk prices, Λ . The nine parameters in Λ are somewhat more difficult to interpret. Only three of them are statistically significant. The price of risk for shocks to the level factor depends significantly on the level and slope. It is higher when either of those factors is higher, consistent with previous findings that the term spread forecasts returns on long-term bonds (and risky assets more generally; Campbell and Shiller, 1991; Fama and French, 1989)

The right-hand side of table 5 reports estimates of K_0 and K_1 , where K_0 collects the steady-state prices of low, business cycle, and high frequencies, and K_1 reports the 9 coefficients of the function that links the variation of the risk prices on the three groups of frequencies on the three state variables. All three elements of K_0 are significantly different from zero. The low-frequency shocks have a negative price of risk, implying that a highly persistent increase in nominal interest rates is viewed as a bad shock. This finding fits with our priors that market participants view permanent increases in interest rates (presumably from increases in long-run inflation expectations) negatively. Unlike the interpretation of λ , though, the interpretation (and estimation) of K_0 is completely independent of any rotation of the factors. Regardless of how we may define “level”, “slope”, and “curvature”, highly persistent shocks to nominal interest rates carry a positive price of risk.

¹⁶This counterintuitive result is also present in Hamilton and Wu’s (2012) estimates.

The fact that the second element of K_0 is positive means that positive shocks to interest rates at business-cycle frequencies are viewed positively, which is consistent with the idea that short-term interest rates have a procyclical component. For K_1 , none of the parameter estimates are significant.

In the end, even though estimating the spectral risk prices is more difficult than the reduced-form risk prices since they are dependent on estimates of ρ and δ_1 , we still obtain highly significant estimates for K_0 , and the results fit well with standard intuition about the prices of risk for different fluctuations in interest rates.

7 Conclusion

This paper studies risk prices in the frequency domain. The impulse response of consumption growth to a given shock to the economy can be decomposed into components of varying frequencies. In a model where innovations to current and expected future consumption growth drive the pricing kernel, the price of risk for a given shock then depends on a weighted integral over the frequency-domain representation of the impulse response function. We study this weighting function both theoretically and empirically. Theoretically, we find that the weighting function helps us gain a deeper understanding of the behavior of asset pricing models. Empirically, our estimates of the weighting function are consistent with the standard version of Epstein–Zin preferences, with a high coefficient of relative risk aversion and elasticity of intertemporal substitution. When we use variables other than consumption growth as the priced risk factor, we obtain stronger results and the long-run components still come out as the most strongly priced.

The method of analysis used here is generally applicable in asset pricing models where the pricing kernel is a linear (or log-linear) function of some state variable. Frequency domain analysis is useful for showing what aspects of the dynamics of the data are important to focus on in studying asset pricing models. For example, under Epstein–Zin preferences, we show that the variance of the pricing kernel is driven essentially by the long-run standard deviation of consumption growth, i.e. its spectral density at frequency zero. Whereas calibrations in the consumption-based asset pricing literature tend to focus on matching the unconditional standard deviation of consumption growth, our results show that they should match the long-run standard deviation.

Our spectral method is also useful for studying the term structure. We show that the price of risk for a shock to the term structure can be written in terms of its dynamic effects on short-term interest rates. This formulation has the advantage that it expresses risk prices in a way that is invariant to any rotation of the underlying factors. As we would have expected, we find that investors want to hedge increases in interest rates at very low frequencies and decreases at business-cycle frequencies.

A Derivation of weighting function with multiple priced variables

The impulse response function is denoted

$$\bar{g}_k \equiv J\Gamma_k \quad (85)$$

where \bar{g}_k is an $M \times N$ matrix whose $\{m, n\}$ element determines the effect of a shock to the n th element of ε_t on the m th element of \bar{X}_{t+k} . The innovation to the SDF is then

$$\Delta E_{t+1} m_{t+1} = - \left(\sum_{k=0}^{\infty} Z_k \bar{g}_k \right) \varepsilon_{t+1} \quad (86)$$

The price of risk for the j th element of ε is simply the j th element of $\sum_{k=0}^{\infty} Z_k \bar{g}_k$.

As before, we take the discrete Fourier transform of $\{\bar{g}_k\}$, defining

$$\tilde{G}(\omega) \equiv \sum_{k=0}^{\infty} e^{-i\omega k} \bar{g}_k \quad (87)$$

Following the same steps as in section 2 and defining $G(\omega) \equiv re\left(\tilde{G}(\omega)\right)$, we arrive at

$$\sum_{k=0}^{\infty} Z_k \bar{g}_k B_j = \frac{1}{2\pi} \int_{-\pi}^{\pi} \vec{Z}(\omega) G(\omega) B_j d\omega \quad (88)$$

$$= \frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_m \vec{Z}_m(\omega) G_{m,j}(\omega) d\omega \quad (89)$$

where

$$\vec{Z}(\omega) \equiv Z_0 + 2 \sum_{k=1}^{\infty} Z_k \cos(\omega k) \quad (90)$$

and where $\vec{Z}_m(\omega)$ denotes the m th element of $\vec{Z}(\omega)$ and $G_{m,j}(\omega)$ denotes the m, j th element of $G(\omega)$. We thus have M different weighting functions, one for each of the priced variables. The M weighting functions each multiply N different impulse transfer functions, $G_{m,j}(\omega)$. The price of risk for shock j depends on how it affects the various priced variables at all horizons.

References

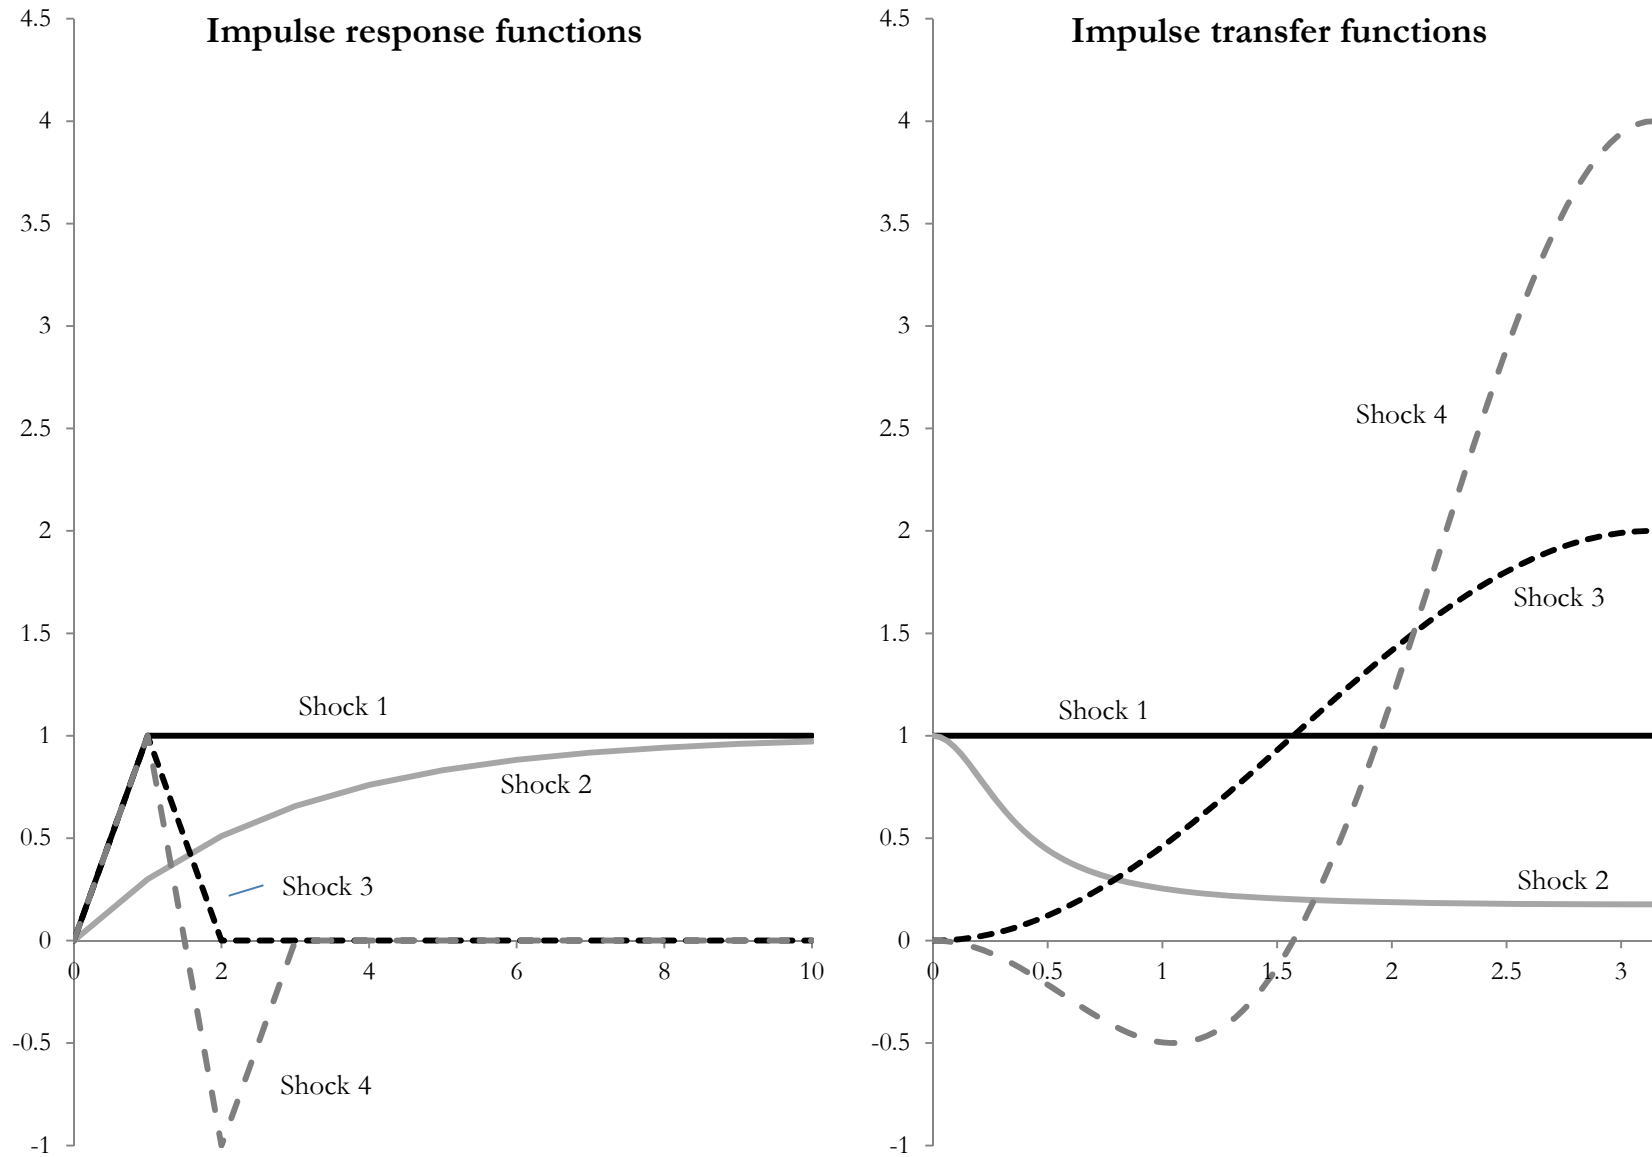
- Abel, Andrew B.**, “Asset Prices under Habit Formation and Catching up with the Joneses,” *The American Economic Review, Papers and Proceedings*, 1990, *80(2)*, 38–42.
- Alvarez, Fernando and Urban J. Jermann**, “Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth,” *Econometrica*, 2005, *73(6)*, 1977–2016.
- Ang, Andrew and Monika Piazzesi**, “No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables,” *Journal of Monetary Economics*, 2003, *50*, 745–787.
- Bansal, Ravi and Amir Yaron**, “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 2004, *59(4)*, 1481–1509.
- , **Robert F. Dittmar, and Christian T. Lundblad**, “Consumption, Dividends, and the Cross Section of Equity Returns,” *Journal of Finance*, 2005, *60(4)*, 1639–1672.
- Barillas, Francisco, Lars P. Hansen, and Thomas J. Sargent**, “Doubts or Variability?,” *Journal of Economic Theory*, 2009, *144(6)*, 2388–2418.
- Bekaert, Geert, Seonghoon Cho, and Antonio Moreno**, “new Keynesian Macroeconomics and the Term Structure,” *Journal of Money, Credit and Banking*, 2010, *42(1)*, 33–62.
- Campbell, John Y.**, “A Variance Decomposition for Stock Returns,” *The Economic Journal*, 1991, *101(405)*, 157–179.
- , “Intertemporal Asset Pricing Without Consumption Data,” *American Economic Review*, 1993, *83(3)*, 487–512.
- **and John H. Cochrane**, “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 1999, *107(2)*, 205–251.
- **and Robert J. Shiller**, “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1988, *1(3)* (3), 195–228.
- **and** —, “Yield Spreads and Interest Rate Movements: A Bird’s Eye View,” *Review of Economic Studies*, 1991, *58(3)*, 495–514.
- **and Tuomo Vuolteenaho**, “Bad Beta, Good Beta,” *American Economic Review*, 2004, *94(5)*, 1249–1275.

- and — , “Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns,” *Journal of Financial Economics*, Forthcoming.
- Christensen, Jens H.E., Francis X. Diebold, and Glenn D. Rudebusch**, “The Affine Arbitrage-Free Class of Nelson–Siegel Term Structure Models,” *Journal of Econometrics*, 2011, *164*(1), 4–20.
- Christiano, Lawrence J. and Terry J. Fitzgerald**, “The Band Pass Filter,” *International Economic Review*, 2003, *44*(2), 435–465.
- Cochrane, John H.**, “A Cross-Sectional Test of an Investment-Based Asset Pricing Model,” *Journal of Political Economy*, 1996, *104* (3), 572–621.
- Constantinides, George M.**, “Habit Formation: A Resolution of the Equity Premium Puzzle,” *The Journal of Political Economy*, 1990, *98*(3), 519–543.
- Dew-Becker, Ian**, “A Model of Time-Varying Risk Premia with Habits and Production,” 2012. Working paper.
- Duffee, Gregory R.**, “Term Premia and Interest Rate Forecasts in Affine Models,” *Journal of Finance*, 2002, *57*(1), 405–443.
- , “Information in (and not in) the Term Structure,” *Review of Financial Studies*, 2011, *24*(9), 2895–2934.
- Epstein, Larry G. and Stanley E. Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 1989, *57*(4), 937–969.
- and — , “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *The Journal of Political Economy*, 1991, *99*(2), 263–286.
- Fama, Eugene F. and Kenneth R. French**, “Business Conditions and Expected Returns on Stocks and Bonds,” *Journal of Financial Economics*, 1989, *25*(1), 23–49.
- Goyal, Amit and Ivo Welch**, “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction,” *Review of Financial Studies*, 2008, *21*(4), 1455–1508.
- Hamilton, James D. and Cynthia Wu**, “Identification and Estimation of Gaussian Affine-Term-Structure Models,” *Journal of Econometrics*, 2012. Working Paper.

- Hansen, Lars P. and Thomas J. Sargent**, “Robust Control and Model Uncertainty,” *American Economic Review*, 2001, *91*(2), 60–66.
- , **John C. Heaton**, and **Nan Li**, “Consumption Strikes Back? Measuring Long-Run Risk,” *Journal of Political Economy*, 2008, *116*(2), 260–302.
- Hansen, Lars Peter and Thomas J. Sargent**, *Robustness*, Princeton University Press, 2007.
- Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu**, “A New Perspective on Gaussian Dynamic Term Structure Models,” *Review of Financial Studies*, 2011, *24*(3), 926–970.
- Kim, Don H. and Jonathan H. Wright**, “An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates,” Federal Reserve Board, Finance and Economics Discussion Series 2005-33 2005.
- Lettau, Martin and Jessica A. Wachter**, “Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium,” *Journal of Finance*, 2007, *62* (1), 55–92.
- and **Sydney Ludvigson**, “Consumption, Aggregate Wealth, and Expected Stock Returns,” *Journal of Finance*, 2001, *56* (3), 815–849.
- Malloy, Christopher J., Tobias J. Moskowitz, and Annette Vissing-Jørgensen**, “Long-Run Stockholder Consumption Risk and Asset Returns,” *Journal of Finance*, 2009, *64*(6), 2427–2479.
- Merton, Robert C.**, “An Intertemporal Capital Asset Pricing Model,” *Econometrica*, 1973, *41* (5), 867–887.
- Otrok, Christopher, B. Ravikumar, and Charles H. Whiteman**, “Habit Formation: A Resolution of the Equity Premium Puzzle?,” *Journal of Monetary Economics*, 2002, *49*(6), 1261–1288.
- , —, and —, “A generalized volatility bound for dynamic economies,” *Journal of Monetary Economics*, 2007, *54*(8), 2269–2290.
- Parker, Jonathan A. and Christian Julliard**, “Consumption Risk and the Cross Section of Expected Returns,” *Journal of Political Economy*, 2005, *113*(1), 185–222.
- Restoy, Fernando and Philippe Weil**, “Approximate Equilibrium Asset Prices,” 1998. NBER Working Paper.

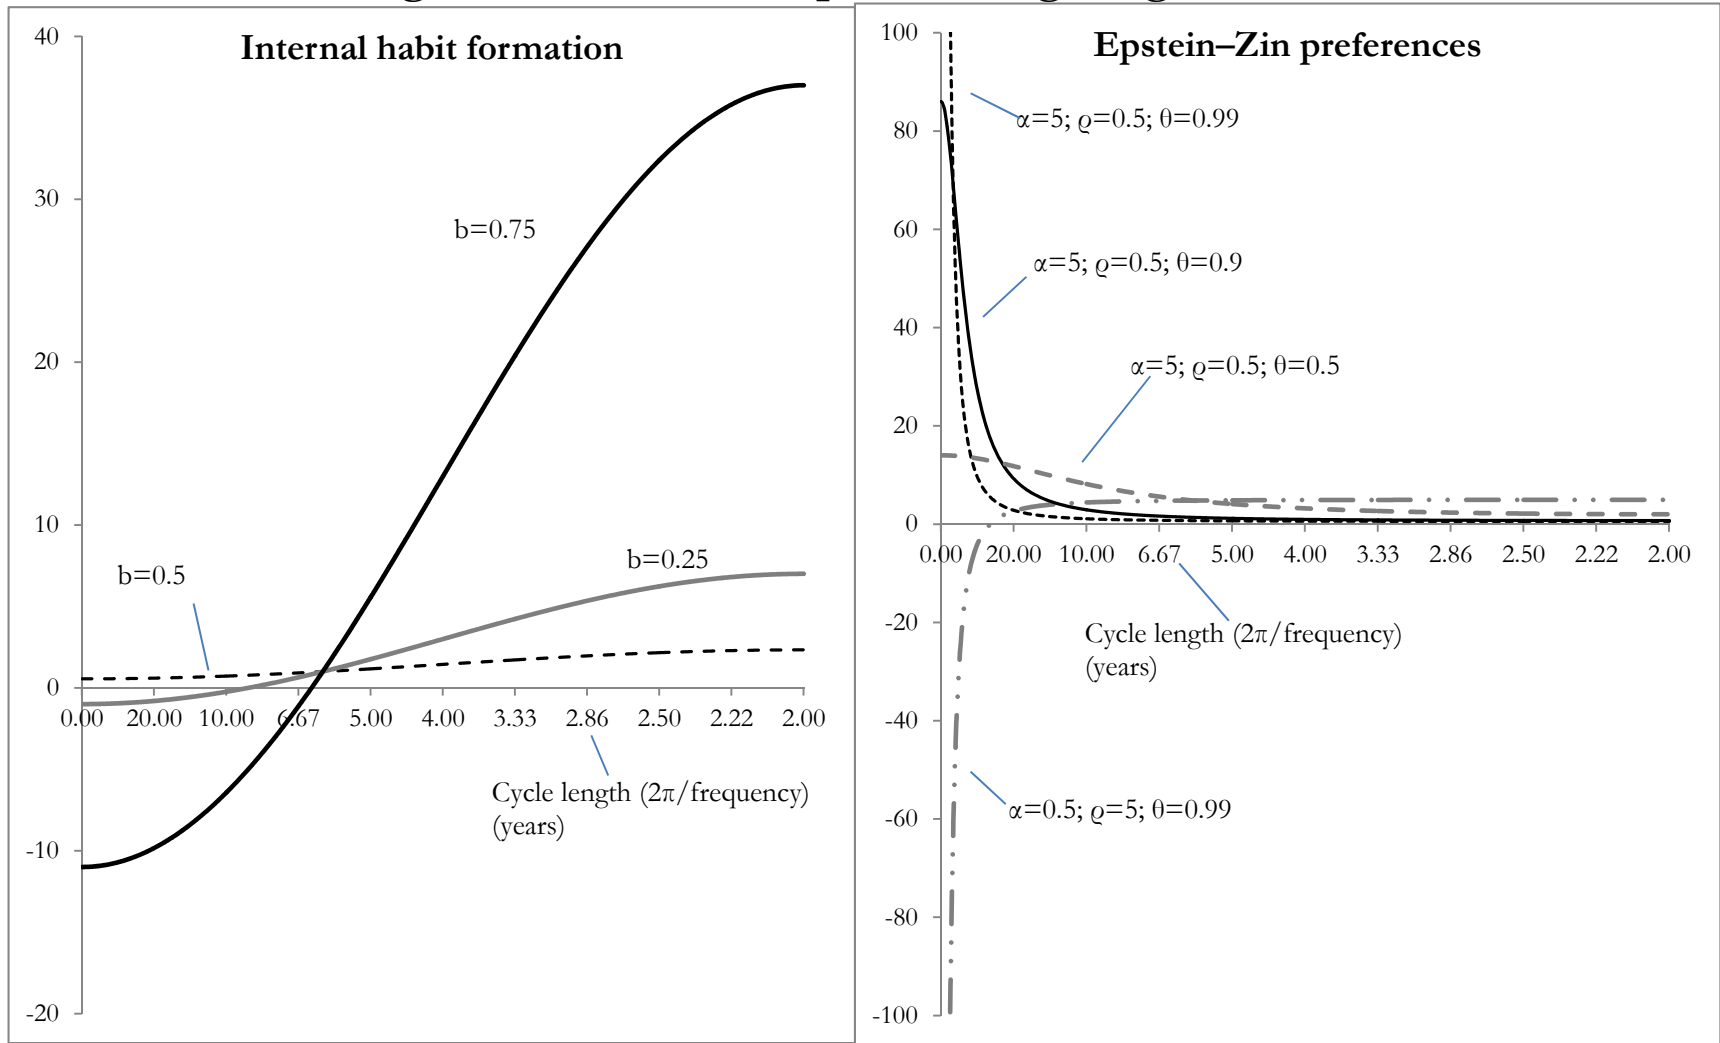
Yu, Jianfeng, “Using Long-Run Consumption-Return Correlations to Test Asset Pricing Models,”
2012. Working paper.

Figure 1. Impulse response functions and impulse transfer functions



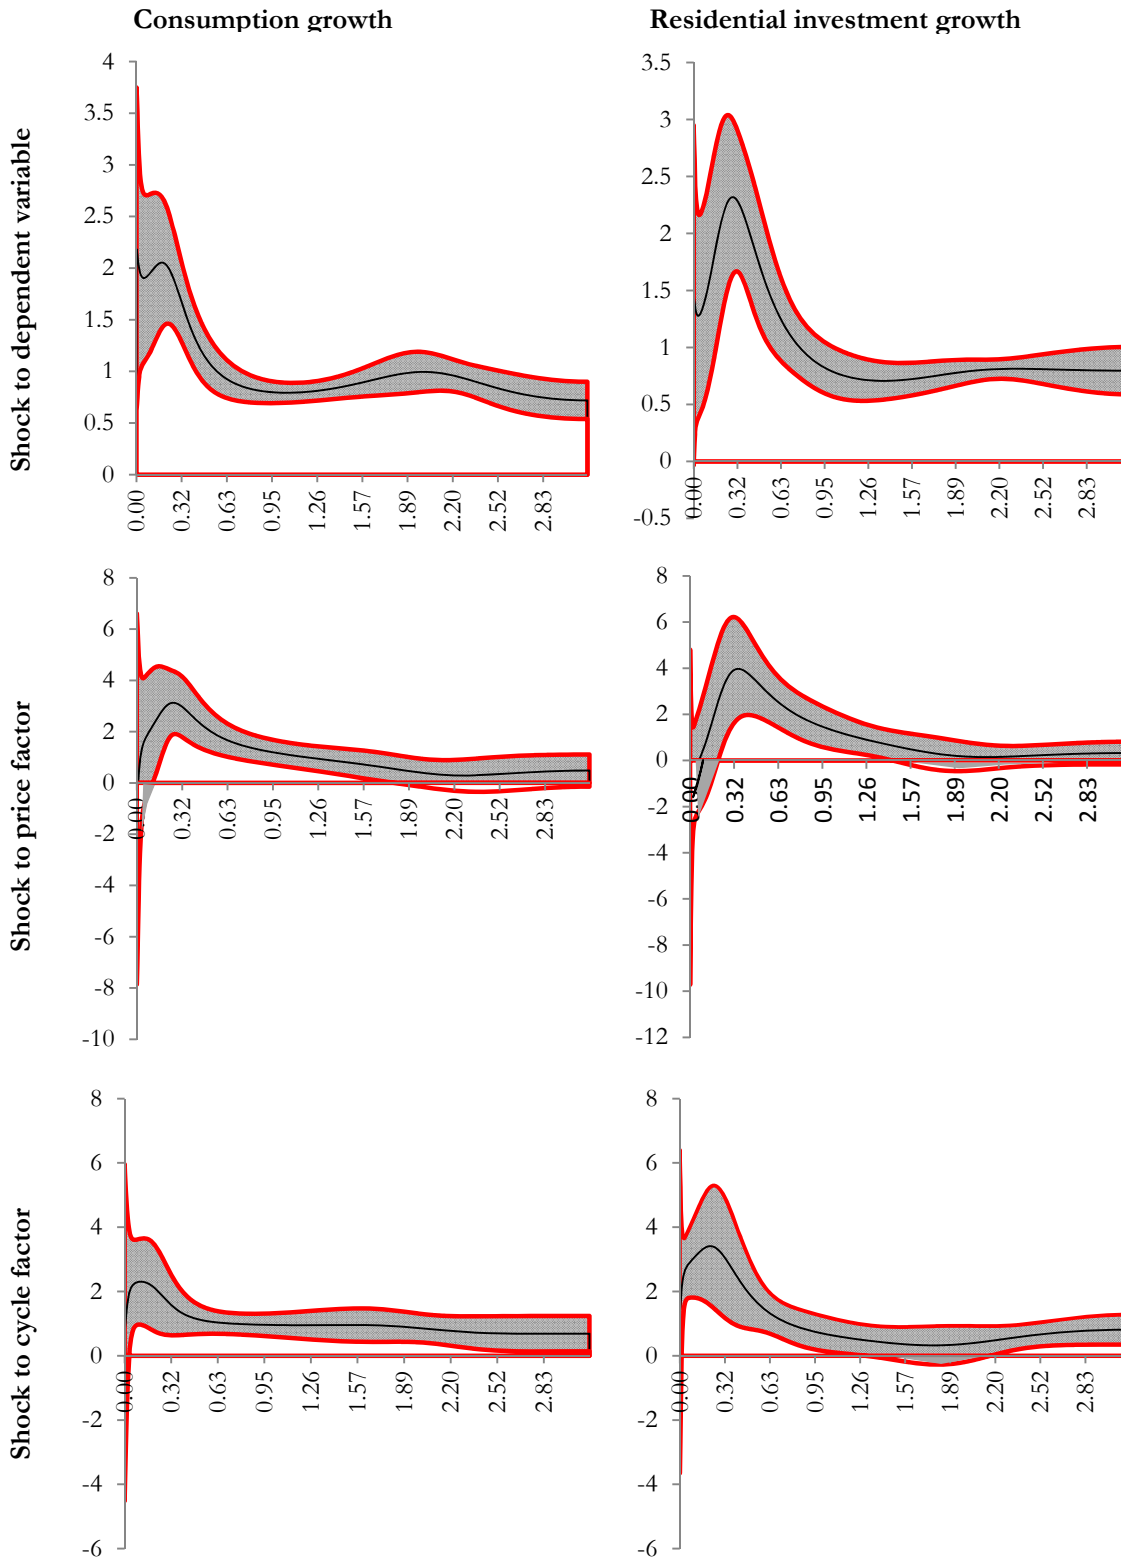
Notes: The left panel plots responses of the level of consumption to four hypothetical shocks. The right-hand panel plots the fourier transforms of the shocks to consumption *growth*, which we refer to as the impulse transfer functions

Figure 2. Theoretical spectral weighting functions



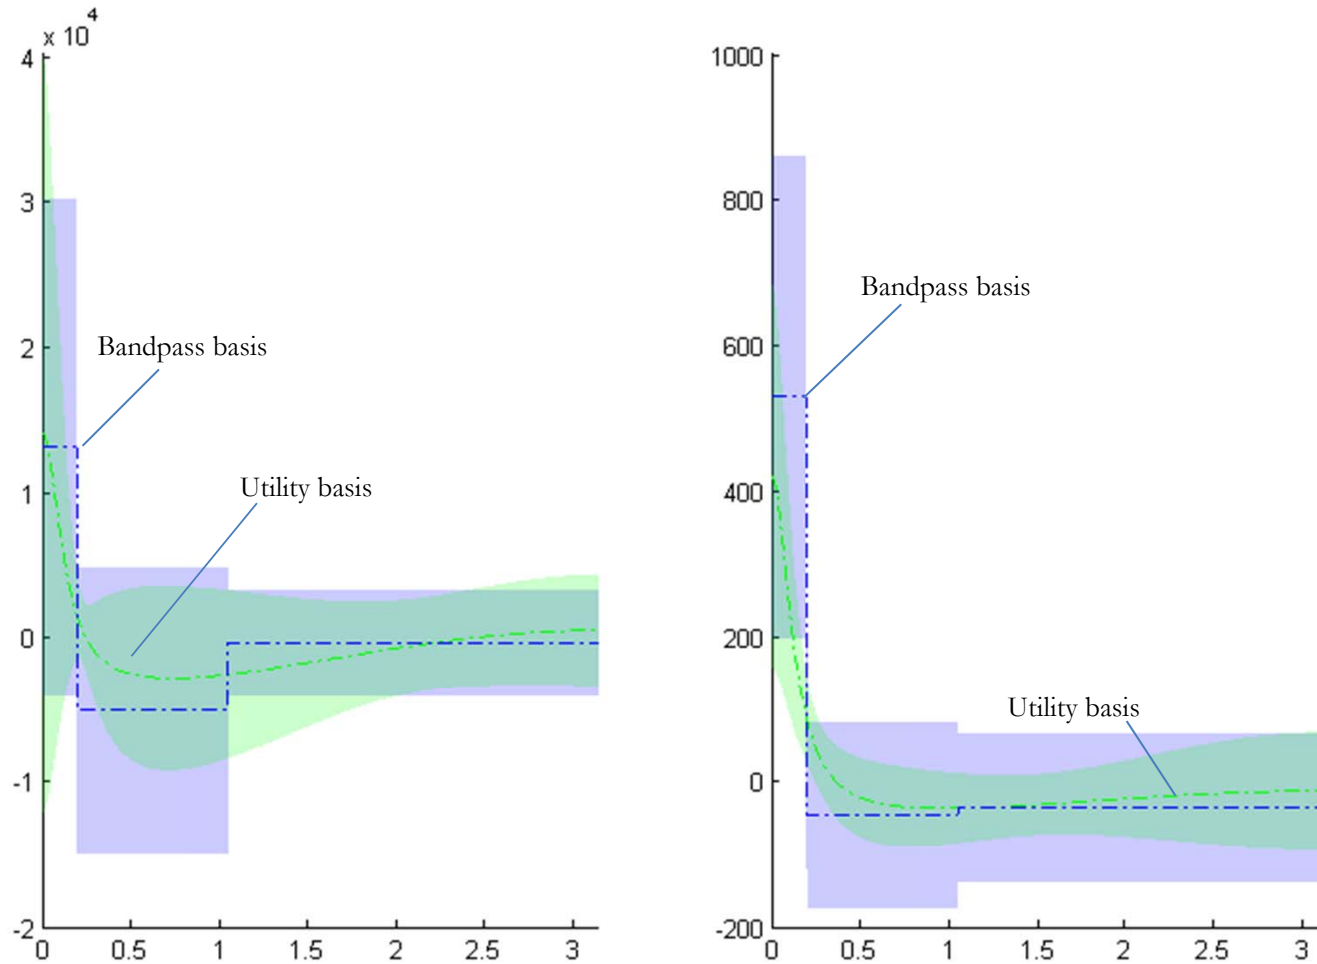
Notes: Plots of the spectral weighting function Z for various utility functions. The x-axis is the cycle length. In the left-hand panel, the parameter b determines the importance of the internal habit in the agent's utility function. In the right-hand panel, α is the coefficient of relative risk aversion; ρ is the inverse elasticity of intertemporal substitution; and θ is the discount factor.

Figure 3. Estimated impulse transfer functions for equity model




Notes: Impulse transfer functions estimated from separate VARs for consumption and residential investment growth and the two principal components. Shocks are not orthogonalized. Shaded regions represent 95-percent confidence intervals. The x-axis gives frequencies in quarters.

Figure 4. Estimated spectral weighting functions for equities
Consumption **Residential investment**



Notes: Estimated weighting functions. The left-hand plot uses (nondurables and services) consumption growth as the priced variable, while the right-hand plot uses residential investment growth. Risk prices are estimated using the 25 Fama–French portfolios. Shaded areas denote 95-percent confidence regions. The utility basis uses a discount factor of 0.60 at the annual horizon. The x-axis gives frequencies in quarters.

Figure 5. Bond pricing factor dynamics



Notes: Estimates from a term structure model in which the underlying factors are rotated so as to be orthogonal and have unit variance. The left-hand panel gives the loading of yields from 1 to 60 months on the three factors. The center panel gives the response of the 1-month interest rate to a unit increase in each of the three factors from 1 to 60 months. The right-hand panel plots the impulse transfer functions for frequencies between 0 and 0.52, where $\omega=0.52$ corresponds to a wavelength of 12 months. Units are annualized percentage points.

Table 1. Regression coefficients from VARs

	Lag 1			Lag 2			Lag 3		
	Cons.	Price	Cycle	Cons.	Price	Cycle	Cons.	Price	Cycle
Cons.	0.2909 ***	0.9253 ***	0.4035 **	0.1377 **	-0.839 **	-0.47 **	0.1698 **	-0.1298	0.1253
se	[0.08]	[0.28]	[0.17]	[0.07]	[0.39]	[0.22]	[0.08]	[0.21]	[0.17]

Notes: VAR results for consumption growth and the two principal components. The table reports the regression of consumption growth on its own lags and those of the two principal components. The sample is 1952:1–2011:2, quarterly. Standard errors are reported in brackets. *** indicates significance at the 5 percent level, ** at the 1% level.

Table 2. Parameter estimates for the spectral weighting function (efficient matrix for GMM)

Portfolios:		FF25						FF25+IND					
Basis:		Bandpass	t-stat	Utility (0.975)	t-stat	Utility (0.6)	t-stat	Bandpass	t-stat	E-Z (0.975)	t-stat	E-Z (0.6)	t-stat
Consumption growth	q1	816	1.49	1596.75	0.44	2017.54	1.03	264	1.31	468.72	0.44	592.24	0.92
	q2	-1360	-1.00	-1880.45	-0.27	-608.15	-0.49	-347	-0.75	-688.48	-0.33	-315.00	-0.79
	q3	-255	-0.20	1384.38	0.17	-2164.75	-0.87	-442	-1.06	661.24	0.27	-380.59	-0.50
GDP	q1	502	1.57	597.48	0.52	647.51	1.23	285	1.74 *	345.98	0.55	374.96	1.60
	q2	-534	-1.28	144.62	0.58	153.69	1.61	-272	-1.28	63.34	0.47	68.59	1.55
	q3	415	1.50	-627.97	-1.21	-999.42	-1.45	162	1.32	-282.64	-1.16	-497.74	-1.69 *
Durables	q1	130	2.31 **	98.79	1.43	236.27	1.65 *	40	2.27 **	28.02	1.36	67.01	1.90 *
	q2	-106	-1.17	33.48	1.31	12.46	0.43	-19	-0.82	8.87	1.50	2.91	0.41
	q3	48	1.00	-146.13	-0.91	-306.83	-1.31	-7	-0.54	-8.47	-0.22	-54.05	-1.00
Investment	q1	94	2.43 **	112.40	0.83	100.95	1.99 **	48	2.49 **	61.52	0.85	55.25	2.37 **
	q2	-53	-1.54	60.67	3.09 ***	43.67	3.57 ***	-26	-1.44	23.15	2.90 ***	13.85	3.35 ***
	q3	76	3.82 ***	-62.57	-1.60	-118.50	-2.55 **	19	2.43 **	-21.20	-1.06	-51.81	-2.32 **
Fixed Investment	q1	76	2.17 **	74.66	0.68	103.64	1.83 *	43	2.08 **	46.24	0.69	64.19	2.28 **
	q2	-100	-1.80 *	63.02	1.86 *	59.51	2.50 **	-61	-1.95 *	3.68	0.21	1.50	0.16
	q3	147	3.39 ***	-160.01	-1.68 *	-226.38	-2.37 **	23	1.37	-63.23	-1.15	-104.34	-2.74 ***
Residential Investment	q1	33	3.12 ***	30.99	2.75 ***	53.07	2.94 ***	5	1.95 *	4.63	1.62	7.92	1.84 *
	q2	-12	-0.68	-16.69	-0.40	-3.23	-0.17	2	0.55	-15.02	-2.05 **	-13.01	-3.16 ***
	q3	-23	-0.66	54.44	0.79	-19.69	-0.60	-32	-4.01 ***	28.33	2.46 **	17.27	2.26 **

Notes: Risk price estimates for the period 1952:1–2011:2 using quarterly data. The priced variable is listed in the left-hand column. The left-hand set of columns uses the Fama–French portfolios as the test assets; the right-hand columns add 49 industry portfolios from Ken French's website. For the bandpass basis, q1 is the price of low-frequency risk, q2 business-cycle frequency, and q3 high frequency. For the utility basis, q1 is the low-frequency component, q2 the constant, and q3 the coefficient on $\cos(w)$. The asset pricing moments are estimated using two-step GMM. The "t-stat" column gives the t statistics for the risk prices. * indicates significance at the 10-percent level, ** the 5-percent level, and *** the 1-percent level. t-stats take into account VAR estimation uncertainty, using GMM. The weighting matrix is constructed using the variance-covariance matrix of the asset pricing moment residuals.

Table 3. Parameter estimates for the spectral weighting function (identity matrix for GMM)

Portfolios:		FF25						FF25+IND					
Basis:		Bandpass	t-stat	Utility (0.975)	t-stat	Utility (0.6)	t-stat	Bandpass	t-stat	E-Z (0.975)	t-stat	E-Z (0.6)	t-stat
Consumption growth	q1	922	1.28	1822.56	0.45	2302.85	0.91	266	1.05	472.52	0.42	597.05	0.76
	q2	-1562	-0.88	-2051.88	-0.26	-599.66	-0.40	-352	-0.58	-690.39	-0.33	-313.88	-0.79
	q3	-75	-0.06	1455.88	0.16	-2595.13	-0.84	-435	-0.96	659.26	0.27	-391.02	-0.37
GDP	q1	600	1.35	718.35	0.51	778.51	1.19	287	1.50	349.25	0.51	378.50	1.33
	q2	-638	-1.13	153.35	0.53	164.25	1.58	-275	-1.05	63.17	0.43	68.47	1.25
	q3	455	1.41	-729.78	-1.12	-1176.38	-1.26	163	0.72	-286.20	-0.84	-503.33	-1.27
Durables	q1	152	1.66 *	115.69	1.16	276.69	1.42	41	1.35	28.51	0.88	68.18	1.16
	q2	-126	-1.04	38.77	1.29	14.16	0.43	-20	-0.50	8.94	0.84	2.88	0.26
	q3	59	0.96	-175.44	-0.80	-363.63	-1.11	-6	-0.24	-9.96	-0.17	-56.34	-0.56
Investment	q1	112	1.60	134.64	0.75	120.91	1.40	48	1.82 *	61.48	0.78	55.21	1.65 *
	q2	-62	-1.06	74.43	2.73 ***	54.07	3.00 ***	-26	-1.29	22.98	2.19 **	13.69	1.95 *
	q3	93	2.95 ***	-73.73	-1.19	-140.71	-1.45	19	1.09	-21.04	-0.85	-51.63	-1.60
Fixed Investment	q1	119	1.92 *	118.37	0.69	164.33	1.98 **	43	1.33	45.94	0.62	63.78	1.67 *
	q2	-160	-1.59	86.01	1.48	80.43	1.75 *	-61	-1.48	3.37	0.16	1.20	0.09
	q3	207	2.49 **	-241.43	-1.50	-346.66	-2.33 **	22	0.61	-63.11	-0.80	-103.96	-1.74 *
Residential Investment	q1	41	2.45 **	38.85	2.22 **	66.54	2.29 **	5	0.79	4.83	0.73	8.28	0.75
	q2	-15	-0.68	-19.52	-0.29	-2.64	-0.08	2	0.33	-14.98	-1.08	-12.88	-1.15
	q3	-25	-0.42	65.86	0.62	-27.08	-0.78	-32	-1.54	28.38	1.55	16.81	1.42

Notes: See table 2. These estimates differ only in that they use the identity matrix for the GMM weighting matrix.

Table 4. Parameter estimates with returns as priced variable

		FF25				FF25 + Industry			
		Weighting: S		Weighting: I		Weighting: S		Weighting: I	
		coeff	t-stat	coeff	t-stat	coeff	t-stat	coeff	t-stat
Utility basis	q1 Long-run	8.36	2.19 **	7.99	1.65 *	7.01	3.06 ***	7.07	1.57
	q2 Constant	9.99	3.34 ***	9.56	1.81 *	8.60	6.25 ***	8.63	2.09 **
Bandpass basis	q1 Long-run	26.20	2.36 **	27.32	1.53	19.28	2.61 ***	19.29	1.39
	q2 Constant	-4.74	-0.65	-5.54	-0.72	-2.62	-0.41	-2.72	-0.48

		FF25				FF25 + Industry			
		Weighting: S		Weighting: I		Weighting: S		Weighting: I	
		coeff	t-stat	coeff	t-stat	coeff	t-stat	coeff	t-stat
Utility basis	q1 Long-run	1.98	0.29	2.61	0.26	6.36	1.33	6.59	0.83
	q2 Constant	10.40	0.95	10.57	0.90	8.69	4.79 ***	8.75	2.09 **
	q3 High Freq	114.20	1.05	108.64	1.01	12.58	0.17	10.07	0.07
Bandpass basis	q1 Long-run	17.11	0.86	16.91	0.40	30.35	2.51 **	30.46	1.43
	q2 Business cycle	109.40	0.32	125.25	0.24	-143.13	-0.81	-144.29	-0.41
	q3 Short-run	-106.83	-0.33	-122.28	-0.25	131.46	0.81	132.46	0.40

Notes: Risk price estimates for the period 1926:3 - 2011:2, using quarterly data. The top panel uses two parameters for the weighting function (a long-run component and a constant), the bottom panel uses three parameters corresponding to the decomposition of Table 2. t-statistics take into account VAR estimation uncertainty, using GMM. The weighting matrix used is either the inverse of the variance-covariance matrix of the moment residuals (Weighting: S) or the identity matrix (Weighting: I).

Table 5. Estimates of bond pricing model

lam			K0		
Level	Slope	Curvature	Low	Business-cycle	High
2.59 ***	-0.91 **	0.77 **	-5409 **	14975 **	-7561 ***
[3.03]	[-2.42]	[2.22]	[-2.01]	[2.50]	[-2.77]

LAM			K1				
Level	Slope	Curvature	Level	Slope	Curvature		
Level	2.16 ***	1.26 ***	-0.46	Low	-2678	-1386	583
	[4.72]	[2.87]	[-0.94]		[-1.62]	[-0.86]	[0.35]
Slope	-0.27	-0.0019	0.23	Business-Cycle	-1447	-6411	-6291
	[-0.89]	[-0.01]	[0.77]		[-0.19]	[-0.84]	[-0.84]
Curvature	0.36 *	-0.046	-0.43	High	-1216	1697	3240
	[1.90]	[-0.34]	[-1.51]		[-0.47]	[0.69]	[1.29]

Notes: lam and LAM are the reduced-form risk prices for the three-factor term structure model. k0 gives the set of steady-state risk prices for low, business-cycle, and high frequencies. k1 determines how the risk prices interact with the underlying factors driving the model. t-statistics are reported in brackets. * denotes significance at the 10 percent level, ** the 5 percent level, and *** the 1 percent level.