Article

Cash Flows Discounted Using a Model-Free SDF Extracted under a Yield Curve Prior

A. Ronald Gallant [1,*] and George Tauchen [2]

[1] Department of Economics, Penn State University, University Park, State College, PA 16802, USA
[2] Department of Economics, Duke University, Durham, NC 27708, USA; george.tauchen@duke.edu
* Correspondence: aronldg@gmail.com or arg21@psu.edu; Tel.: +1-919-428-1130

Abstract: We developed a model-free Bayesian extraction procedure for the stochastic discount factor under a yield curve prior. Previous methods in the literature directly or indirectly use some particular parametric asset-pricing models such as with long-run risks or habits as the prior. Here, in contrast, we used no such model, but rather, we adopted a prior that enforces external information about the historically very low levels of U.S. short- and long-term interest rates. For clarity and simplicity, our data were annual time series. We used the extracted stochastic discount factor to determine the stripped cash flow risk premiums on a panel of industrial profits and consumption. Interestingly, the results align very closely with recent limited information (bounded rationality) models of the term structure of equity risk premiums, although nowhere did we use any theory on the discount factor other than its implied moment restrictions.

Keywords: stochastic discount factor; discounting; cash flows; yield curve; moment functions; bounded rationality; Bayesian

JEL Classification: C32; C36; E27

1. Introduction

Under very mild conditions, there exists a scalar stochastic discount factor (SDF) process that generates moment restrictions on the returns (or cash flows) of traded securities. Knowledge of the SDF process allows one to check if a particular security is priced consistently with other traded assets, and it allows the valuation of uncertain future cash flows regarding nontraded assets as far as one is confident the pricing implications extend appropriately. In the literature, there are a plethora of methods either to nonparametrically extract the SDF process from historical data or to evaluate particular theories on the SDF process such as long-run risk models or habit models.

One issue associated with such approaches concerns the ex post implied level of real interest rates. Extant procedures use moment conditions based on asset returns, and the return horizons typically range from monthly to annual. The returns series are like first differenced asset prices and thus nearly white-noise processes; information on the levels of asset prices, and bond prices in particular, is negligible, leaving interest rates ill determined. Left on its own, an extracted SDF can give rise to somewhat implausible levels of real interest rates. As an example, for a recent long-run risk application, Christensen (2017) reports the long-term interest rate as a rather high at 7 percent per year; in additional computations using this author's code, we found that the entire yield curve from one year on out is essentially flat at just over 7 percent. Poorly determined real yield curves were encountered using the extraction procedure of this paper, where absent prior knowledge implied that yield curves shifted and bent in implausible configurations. As just noted, the moment conditions contain little, if any, level information, and external information from other sources needs to be imposed to discipline the SDF extraction.

An agreed upon fact is that U.S. real interest rates are very low. According to Campbell (2003, p. 812), the average short-term U.S. real rate was 0.896 percent over the period 1947–1998, and few would have argued for higher real short-term rates since then. As for longer-term real rates, Figure 2 of Tesar and Obstfeld (2015, p. 8) indicates that the 10-year real rate of interest over the period 1930–2014 was often negative, generally fluctuated between 0 and 2.5 percent per year, and only briefly bumped 5 percent during the interwar era and again during the disinflation period of the early 1980s. Additional information from Treasury Inflation Protected Securities (TIPS) real yields is seen in Table 1, which are remarkably low.

Table 1. Treasury Inflation Protected Securities (TIPS) (real) yields.

Year	5-Year	7-Year	10-Year	20-Year
2004	1.02	1.39	1.76	2.13
2005	1.50	1.63	1.81	1.97
2006	2.28	2.30	2.31	2.31
2007	2.15	2.25	2.29	2.36
2008	1.30	1.63	1.77	2.18
2009	1.06	1.32	1.66	2.21
2010	0.26	0.68	1.15	1.73
2011	−0.41	0.10	0.55	1.20
2012	−1.20	−0.88	−0.48	0.21
2013	−0.76	−0.30	0.07	0.75
2014	−0.09	0.32	0.44	0.86
2015	0.15	0.36	0.45	0.78

Displayed are annual averages of daily fixed-term Treasury Inflation Protected Securities (TIPS) yields. See Appendix A for details.

In what follows, we implemented a Bayesian SDF extraction procedure subject to a prior that enforces these known low values for U.S. real interest rates. The method's mathematical foundation requires a prior to ensure all the random variables are actually defined on a proper probability space, and here, we elected to use the yield curve prior in place of a specific model of the SDF. Specifically, the prior centers the one-year yield at 0.896 percent with a standard deviation of 1.00 percent, and it centers the 30-year yield at 2.00 percent with a standard deviation of 1.00 percent. The prior generally accommodates both the levels and fluctuations in real rates suggested in the above discussion and by Table 1. This prior was maintained throughout the entire sample period, although it would be relatively easy to impose a time-varying prior with possibly higher yields in the earlier parts of the sample if reliable real rates were available to inform the development of such a prior.

Post-extraction, we used the dynamics of the SDF and related variables in a standard log-Gaussian pricing framework to value various cash flows, with a focus on the risk premiums on stripped cash flows, often termed dividend strips in the literature. The concept is simple: if, from the perspective of period 0, an investment pays off the uncertain stream $\{CF_t\}_{t=1}^{\infty}$ into the indefinite future, then the stripped cash flow is the asset that pays just CF_t in period $t > 0$ and zero in all other periods. Recently, researchers have been investigating the term structure of the equity risk premiums on stripped cash flows to better understand the relationship between risk and reward at short- and long-term horizons. Asset-pricing models (van Binsbergen et al. 2012; Giglio et al. 2015) suggest that the term structure of risk premiums is upward sloping, with more distant cash flows earning higher risk premiums due to a long-run risk (Bansal and Yaron 2004) mechanism that makes investors fearful of volatility in the distant future. This prediction seems at odds with common sense intuition, but theory alone is not powerful enough to make an unambiguous prediction on the slope of the equity risk premium term structure. Backus et al. (2016) show how a wide range of levels and shapes of the term structures of claims

can be achieved by modifying the dynamics of the pricing kernel, the cash flow growth, and their interaction. Empirically, discussions about the true average slope of the equity return term structure have not yet been settled (Cochrane 2017; Bansal et al. 2017; van Binsbergen et al. 2017), and reconciling asset-pricing models with the possible slopes of the term structure of equity returns has recently become a very active area of research. Of particular interest here is Croce et al. (2015), who developed a bounded rationality model with long-run risk that appears to explain our findings below.

2. Ex Post Stochastic Discount Factor

The ex post-realized values of $SDF_{t-1,t}$ were extracted annually for 1930–2015 using the methodology developed in Gallant and Hong (2007). The differences are threefold: the dataset is longer due to the passage of time, all of the Fama–French portfolios (Fama and French 1992, 1993) can be used because a missing data problem has been resolved, and the prior tilts values toward a specified yield curve instead of toward long-run risk dynamics (Bansal and Yaron 2004). In brief, the ideas are as follows.

For time $t = 1, \ldots, n = 86$, where $t = 1$ corresponds to 1930 and $t = n$ corresponds to 2015, denote the real gross returns on the 25 Fama–French 5×5 size and value portfolios with the vector R_{st}, denote the real gross annual return on the thirty-day T-bill with R_{bt}, denote the real per capita consumption growth with $\frac{C_t}{C_{t-1}}$, and denote the per capita labor income growth with $\frac{L_t}{L_{t-1}}$. Let x_t be a vector of length 28 containing these variables. Let x be an array with x_t as columns; x has the dimensions 28 by $n = 86$. Let θ_t denote the stochastic discount factor $SDF_{t-1,t}$ and set $\theta = (\theta_1, \ldots, \theta_{86})$.

The vector θ is random and endogenous, so determining a likelihood $p^*(x \mid \theta)$ for Bayesian analysis requires some care. The likelihood construction proceeds as follows.

We presume the existence of, but not knowledge of, a general equilibrium model with a financial sector. The general equilibrium model determines a joint probability space on which all the random variables that enter the model live and, hence, a marginal probability space on which the random variables $(x, \theta) \in \mathcal{X} \times \Theta$ live. The marginal probability space determines a conditional distribution of x given θ. We presume that this conditional distribution has the density $p^o(x \mid \theta)$. Ours is a partial equilibrium analysis, so any general equilibrium parameters that might be involved in an expression for $p^o(x \mid \theta)$, were they known, do not affect our analysis and can be ignored or, to be pedantic, are fixed and calibrated by nature.

Let $p(\theta)$ denote the prior we intend to use for a partial equilibrium Bayesian analysis. The analysis will be with reference to the probability space $(\mathcal{X} \times \Theta, \mathcal{C}^o, P^o)$, where P^o has the density $p^o(x \mid \theta) \, p(\theta)$.

Denote the vector of the (conditional) moment-equation errors with

$$e_{t,t-1}(\theta_t) = 1 - \theta_t \begin{pmatrix} R_{st} \\ R_{bt} \end{pmatrix}, \tag{1}$$

where 1 denotes a vector of 1's of length twenty-six. Define the instruments

$$V_t = \begin{pmatrix} R_{st} - 1 \\ R_{bt} - 1 \\ C_t/C_{t-1} - 1 \\ L_t/L_{t-1} - 1 \\ 1 \end{pmatrix},$$

where $R_{st} - 1$ denotes 1 subtracted from each element of R_{st}. Consider the moment conditions

$$m(x_t, x_{t-1}, \theta_t) = V_{t-1} \otimes e_{t,t-1}(\theta), \tag{2}$$

where $\otimes$ denotes the Kronecker product, and their sample average

$$\bar{m}(x,\theta) = \frac{1}{n}\sum_{t=2}^{n} m(x_t, x_{t-1}, \theta_t). \tag{3}$$

The length of the vector $m(x_t, x_{t-1}, \theta_t)$ is $K = 754$, so the number of overidentifying restrictions on $\theta_2, \ldots, \theta_{86}$ is 669. Note that θ_1 is not yet identified because θ_1 does not appear in (3); it is identified by the prior as discussed later in this subsection.

Following Gallant and Hong (2007), we assume that $e_{t,t-1}(\theta)$ has a factor structure. There is one error common to all the elements of $\theta_t R_{st}$ and twenty-six idiosyncratic errors, one for each element of $(\theta_t R_{st}, \theta_t R_{bt})$. Denote this matrix with Σ_e (or with $\Sigma_{e,t}$ if one wants to allow for heterogeneity, which makes no difference in what follows). A set of orthogonal eigenvectors U_e for Σ_e is easy to construct (Gallant and Hong 2007, p. 535) and can be used to diagonalize Σ_e. Similarly, U_v and Σ_v for V_t can be determined. Let $H_t(\theta) = (U_v \otimes U_e)'m(x_t, x_{t-1}\theta)$ with elements $h_{i,t}(\theta)$. Diagonalization implies that we can estimate the variance of $H_t(\theta)$ by a diagonal matrix $S_n(\theta)$ with the elements

$$s_i(\theta) = \frac{1}{n}\sum_{t=2}^{n}\left[h_{t,i}(\theta) - \frac{1}{n}\sum_{t=2}^{n} h_{t,i}(\theta)\right]^2.$$

Let $S_n^{-1/2}(\theta)$ denote this matrix with the diagonal elements replaced by their square roots.

The extraction of the ex post realization of the SDF is based on the random variable

$$Z(x,\theta) = \sqrt{n}\, S_n^{-1/2}(\theta)(U_z \otimes U_e)'\bar{m}_n(x,\theta) \tag{4}$$

with a range $\mathcal{Z}$ defined on the aforementioned probability space $(\mathcal{X} \times \Theta, \mathcal{C}^o, P^o)$. $Z(x,\theta)$ is the normalized sum of transformed draws (x_t, θ_t) and is asymptotically multivariate normal with a zero mean and identity variance under plausible regularity conditions on $(\mathcal{X} \times \Theta, \mathcal{C}^o, P^o)$. Note, specifically, that θ_t is random and jointly distributed with x_t, so issues of uniformity in θ do not arise. Thus, it is reasonable to assume that Z follows the standard normal distribution $\Phi(z)$ with a density of $\phi(z)$.

The assumption that $Z(x,\theta)$ has a density of $\phi(z)$ induces a probability space $(\mathcal{X} \times \Theta, \mathcal{C}, P)$, where $\mathcal{C}$ is the σ-algebra of preimages $\mathcal{C} = \{C = Z^{-1}(B), B \subset \mathcal{Z}, B \text{ Borel}\}$ and $P[C = Z^{-1}(B)] = \int_B d\Phi(z)$. Define $\mathcal{C}^*$ to be the smallest σ-algebra that contains all the sets in $\mathcal{C}$ plus all the sets of the form $R_B = (\mathcal{X} \times B)$, where B is a Borel subset of Θ. Under a semipivotal assumption for (4), (Which is that $\{x : Z(x,\theta) = z\}$ is not empty for any choice of (z,θ) in the parameter space Θ and range space $\mathcal{Z}$; a sufficient condition is that each element of Z is continuous in μ when some element x_{it} of x_t is replaced by $x_{it} + \mu$ for all t and is unbounded from above and below in μ.) there is an extension of $(\mathcal{X} \times \Theta, \mathcal{C}, P)$ to a space $(\mathcal{X} \times \Theta, \mathcal{C}^*, P^*)$ on which the conditional density of x given θ is

$$f(x \mid \theta) = \phi[Z(x,\theta)] \tag{5}$$

(Gallant 2016a). This density is termed the "method of moments representation" of the likelihood and may be used for Bayesian inference in connection with the prior $p(\theta)$ (Gallant 2016a, 2016b). (Missing in (5) is a multiplicative Jacobian term that experience indicates has a negligible impact on computations when omitted (Gallant 2020)).

In short, the Bayesian method used in Gallant and Hong (2007) and here uses the moment conditions $z = Z(x,\theta)$ given by (4), takes

$$p(x \mid \theta) = \phi(z) \tag{6}$$

as the likelihood, and proceeds directly to Bayesian inference using a prior $p(\theta)$.

Next, we describe the prior.

Let $w_t = \left(\log(\theta_t), \log\left(\frac{GDP_t}{GDP_{t-1}}\right)\right)'$, where GDP_t growth is observed for $t = 1, 2, \ldots,$ $n = 86$. GDP is not involved in the SDF extraction up to this point. It is now included as prior information regarding past business-cycle conditions. Consider the recursion

$$w_t = d_0 + D w_{t-1} + u_t, \tag{7}$$

where the u_t are independent and bivariate normal with a zero mean and variance Σ_d. Markov chain Monte Carlo (MCMC) (Gamerman and Lopes 2006) is used in the Gallant and Hong (2007) method which means that the proposed θ_t is available to compute w_t before the prior and likelihood need to be computed. From the w_t the parameters of (7) can be determined by the least squares. With the least-squares values replacing parameters in (7), a yield curve for maturities one year through thirty years can be computed analytically from (7) conditional on a specified initial condition w_0; see Equations (12), (13) and (15) of Section 3. In particular, the one-year and 30-year yields, $Y^*_{1,t}$ and $Y^*_{30,t}$, can be computed successively for $w_0 = w_t$, $t = 1, \ldots, n = 86$. Our prior is

$$p(\theta) = \prod_{t=1}^{n} \phi[(Y^*_{1,t} - 0.00896)/0.01]\phi[(Y^*_{30,t} - 0.02)/0.01] \tag{8}$$

Note, in particular, that the prior identifies θ_1.

With likelihood (5) and prior (8) in hand, Bayesian inference can be carried out using MCMC in the usual way; see, for exaxmple, Gamerman and Lopes (2006). After the transients died out, we ran an MCMC chain of length 8,000,000. The θ in the chain with the highest value of the posterior was selected as the estimate $\hat{\theta}$ of the ex post SDFs for the years 1930 through 2015. The estimate is plotted as Figure 1. The shaded areas are NBE R recessions.
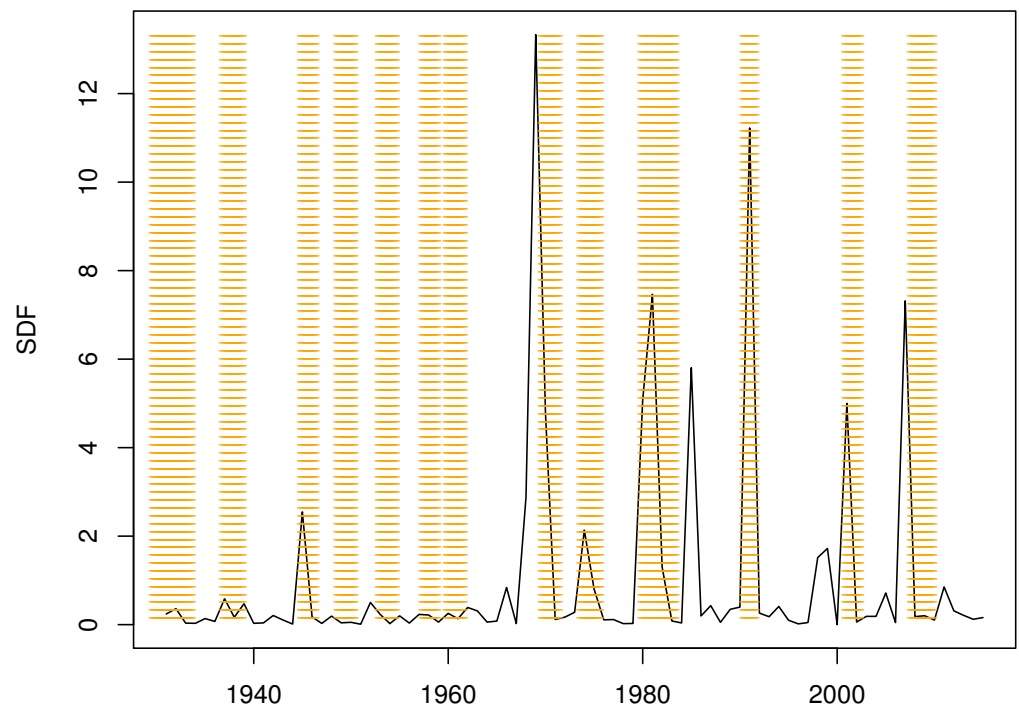

Figure 1. Posterior valuation. Plotted is the posterior mode of $\theta_1 = \text{SDF}_{1930,1931}$ through $\theta_{86} = \text{SDF}_{2014,2015}$.

3. Discounted Cash Flow Estimation

We now used the extracted SDF series to value cash flows for assets outside the span of returns used in the extraction step. For this part, we used annual data on corporate profits from various large sectors of the U.S. economy. We assembled annual data for seven sectors

but for the shorter period 1959–2015, as data limitations precluded going any farther back. We also treated consumption as a cash flow, making a total of eight under consideration. These data were concatenated with the extracted SDF data and various macroaggregates for this shorter time span.

For the valuation step, consider the trivariate series

$$
y_t = \begin{pmatrix} \log(SDF_{t-1,t}) \\ \log(GDP_t) - \log(CF_t) \\ \log(CF_t) - \log(CF_{t-1}) \end{pmatrix} = \begin{pmatrix} sdf_{t-1,t} \\ gdp_t - cf_t \\ \Delta cf_{t-1,t} \end{pmatrix} \tag{9}
$$

where CF_t denotes a cash flow payoff at time t, such as the annual corporate profits in year t; GDP_t denotes the gross domestic product; and $SDF_{t-1,t}$ denotes the extracted stochastic discount factor $\hat{\theta}_t$ of Section 2. Note that the second variable in the autoregression, $\log(GDP_t) - \log(CF_t)$, plays no direct role in subsequent pricing, but it is included because it conveys information on future cash flows. The specification presumes co-integration between GDP and CF, which is discussed more fully in Section 4 below.

The time zero present value of the cash flow CF_t is

$$
PV_{0,t}(CF) = \mathcal{E}(SDF_{0,t}\, CF_t \,|\, \mathcal{F}_0) = \mathcal{E}\left[\exp\left(\sum_{s=1}^{t} \Delta cf_{s-1,s} + \sum_{s=1}^{t} sdf_{s-1,s}\right) \,\Big|\, \mathcal{F}_0\right] \tag{10}
$$

where $\mathcal{F}_0$ denotes the time 0 information set. (Note that in (10), the time zero value of the SDF must be unity; the time zero value of CF is irrelevant because we work in terms of ratios—see (19) below—and therefore, we will normalize it to be unity.) (In (10), the expectation $\mathcal{E}$ refers to the probability space $(\mathcal{X} \times \Theta, \mathcal{C}^o, P^o)$ defined in Section 2. In (14)–(16) and thereafter, $\mathcal{E}$ refers to the VAR (11).) For a risk-free payoff of one real dollar at time t, the time zero present value is $PV_{0,t}(1) = \mathcal{E}(SDF_{0,t} \,|\, \mathcal{F}_0)$, where $SDF_{0,t} = \exp(\sum_{s=1}^{t} sdf_{s-1,s})$. The corresponding yield is $Y_t^* = -\log[PV_{0,t}(1)]/t$.

Let

$$
y_t = b_0 + B y_{t-1} + e_t, \tag{11}
$$

where the e_t are independent and trivariate normal with a zero mean and variance Σ_b. The sum

$$
\sum_{s=1}^{t} y_s = \left(\sum_{s=1}^{t} (t+1-s) B^{s-1}\right) b_0 + \left(\sum_{s=1}^{t} B^s\right) y_0 + \sum_{s=1}^{t} \left(\sum_{u=1}^{s} B^{u-1}\right) e_{t+1-s}
$$

where $B^0 = I$, $B^1 = B$, $B^2 = B B$, etc., has moments

$$
\mathcal{E} \sum_{s=1}^{t} y_s = \left(\sum_{s=1}^{t} (t+1-s) B^{s-1}\right) b_0 + \left(\sum_{s=1}^{t} B^s\right) y_0 \tag{12}
$$

$$
Var \sum_{s=1}^{t} y_s = \sum_{s=1}^{t} \left(\sum_{u=1}^{s} B^{u-1}\right) \Sigma_b \left(\sum_{u=1}^{s} B^{u-1}\right)' \tag{13}
$$

One can use (12) and (13) to evaluate

$$
\begin{aligned}
PV_{0,t}(CF) \; = \; & \exp\left[\mathcal{E}\left(\sum_{s=1}^{t}\Delta cf_{s-1,s}\right) + \mathcal{E}\left(\sum_{s=1}^{t}sdf_{s-1,s}\right)\right] \\
& + \exp\left[\frac{1}{2}Var\left(\sum_{s=1}^{t}\Delta cf_{s-1,s}\right) + \frac{1}{2}Var\left(\sum_{s=1}^{t}sdf_{s-1,s}\right)\right] \\
& + \exp\left[Cov\left(\sum_{s=1}^{t}\Delta cf_{t-1,t'}\sum_{s=1}^{t}sdf_{s-1,s}\right)\right]
\end{aligned} \tag{14}
$$

$$
PV_{0,t}(1) \; = \; \exp\left[\mathcal{E}\left(\sum_{s=1}^{t}sdf_{s-1,s}\right) + \frac{1}{2}Var\left(\sum_{s=1}^{t}sdf_{s-1,s}\right)\right] \tag{15}
$$

$$
EV_{0,t}(CF) \; = \; \exp\left[\mathcal{E}\left(\sum_{s=1}^{t}\Delta cf_{s-1,s}\right) + \frac{1}{2}Var\left(\sum_{s=1}^{t}\Delta cf_{s-1,s}\right)\right] \tag{16}
$$

where $\mathcal{E}$ refers to expectation with respect to the VAR (11). Note that in the above, the time zero value of the SDF must be unity; the time zero value of CF is irrelevant because we work in terms of ratios—see, for example, (19)—therefore, we normalize it to be unity.

We now describe the imposition of a yield curve prior on the estimation of the b_0, B, and Σ_b that appear in the VAR (11). Consider a state-space representation of VAR (11)

$$
v_t = \begin{pmatrix} sdf_{t-2,t-1} \\ gdp_{t-1} - cf_{t-1} \\ \Delta cf_{t-2,t-1} \\ sdf_{t-1,t} \\ gdp_t - cf_t \\ \Delta cf_{t-1,t} \end{pmatrix} = \begin{pmatrix} 0 \\ b_0 \end{pmatrix} + \begin{pmatrix} 0 & I \\ 0 & B \end{pmatrix}v_{t-1} + \begin{pmatrix} 0 \\ e_t \end{pmatrix} \tag{17}
$$

The estimation of (17) subject to the indicated parameter restrictions gives the same estimates $\hat{b}_0$, and $\hat{B}$, $\hat{\Sigma}_b$ as does the unconstrained estimation of (11). Let

$$
A = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}
$$

and note that the fourth and fifth elements of Ax_t are $sdf_{t-1,t}$ and $\Delta gdp_{t-1,t} = \log(GDP_t) - \log(GDP_{t-1})$. An implication is that we can insert the parameters b_0^*, B^*, and Σ_b^* of

$$
Av_t = b_0^* + B^* Av_{t-1} + e_t^* = A\begin{pmatrix} 0 \\ b_0 \end{pmatrix} + A\begin{pmatrix} 0 & I \\ 0 & B \end{pmatrix}A^{-1}Av_{t-1} + A\begin{pmatrix} 0 \\ e_t \end{pmatrix} \tag{18}
$$

into Equations (12), (13), and (15) to compute the one-year yield, $Y_{1,t}^*$, and 30-year yield, $Y_{30,t}^*$, with y_0 set to v_t successively for $t = 1, ..., n = 86$, and impose the prior (8).

The computational procedure is, within an MCMC loop, for the proposed b_0, B, and Σ_b, to evaluate the likelihood implied by VAR (11); compute b_0^*, B^*, and Σ_b^* as indicated by expression (18) from the proposed b_0, B, and Σ_b; evaluate the prior (8); and use the likelihood and prior so computed to make the accept/reject decision of the MCMC loop.

Serendipitously, the state-space complications can be avoided because it turns out that the yields $Y_{1,t}^*$ and $Y_{30,t}^*$ obtained by applying Equations (12), (13) and (15) directly to the b_0, B, and Σ_b of VAR (11) are identical to those computed from the b_0^*, B^*, and Σ_b^* of VAR (18). Apparently, the reason is that the only difference between the distributions of Av_t and v_t is the location parameter of their fifth element, and the location parameter of

the fifth element is not involved in Equation (12), (13) or (15). For ourselves, we are more comfortable relying on having performed the computations using both (11) and (18) and obtaining identical results than relying on the distributional argument.

4. Empirical Implementation

4.1. Valuation, Expectation, and Risk Premiums

Display (14) shows the valuation operator $PV_{0,t}(CF)$ as the economic value in period 0 of the cash flow CF_t received in t; evidently, $PV_{0,t}(\bullet)$ is a linear operation on random variables realized at time t. Likewise, (16) shows the conventional statistically expected value operator $EV_{0,t}(CF)$, also a linear operator on the same space of time t random variables where $PV_{0,t}(\bullet)$ operates. As usual, two linear operators on a space are connected via a Randon–Nikodym-style change in measure/density, which is the usual risk-neutral change of measure, noted but not used here.

For the risk-free asset,

$$\text{bond price}_{0,t} = PV_{0,t}(1).$$

The (geometric) risk-free yield to maturity $r^f_{0,t}$ at time 0 of the t-maturity zero-coupon bond is defined via the relationship

$$e^{t\,r^f_{0,t}} \times PV_{0,t}(1) = 1,$$

because the invested amount $PV_{0,t}(1)$ grows at the continuously compounded rate $r^f_{0,t}$ up to \$1 at time t. Equivalently, the rate is defined by

$$r^f_{0,t} = -\frac{\log(PV_{0,t}(1))}{t}.$$

Just as a coupon-bearing bond can be thought of as a portfolio of stripped payments valued as described immediately above, finance economists have become interested in "dividend strips", where the dividend asset that pays the owner the infinite stream $\{CF_s\}_{s=1}^{\infty}$ is viewed as a portfolio of stripped payments. The value of each stripped payment is given by the pricing operation worked out above as $PV_{0,t}(CF)$. Only if agents are neutral to risk would it be the case that $PV_{0,t}(\bullet) = e^{-r_{f,0,t}t} EV_{0,t}(\bullet)$.

By analogy with the pure discount bond, we can define the geometric rate $r_{0,t}$ at which the amount $PV_{0,t}(CF)$ invested at time 0 grows continuously compounded to its statistically expected value at time t by way of

$$e^{t\,r_{0,t}} \times PV_{0,t}(CF) = EV_{0,t}(CF) \Rightarrow$$
$$r_{0,t} = \frac{1}{t} \log\left(\frac{EV_{0,t}(CF)}{PV_{0,t}(CF)}\right) \tag{19}$$

Note that $r_{0,t}$ is a number known at time 0 that pertains to a cash flow received at time t. The quantity

$$\text{risk premium: } r_{0,t} - r^f_{0,t} \tag{20}$$

is the excess (geometric) return over cash of the investment (stripped cash flow) that pays CF_t. The amount (20) represents the required rate of return above cash necessary to compensate for the economic risk embedded in the investment. To aid interpretation below, we recall from elementary asset pricing that in the iid log-Gaussian case, the one-period relationship is

$$EV_{0,1}(r_{0,1}) = r^f_{0,1} - \text{Cov}(\Delta cf_{0,1}, sdf_{0,1}), \tag{21}$$

using notation defined in (9). The risk premium in (21) is $-\text{Cov}(\Delta cf_{0,1}, sdf_{0,1})$, the fundamental notion in the finance of reward for bearing covariance risk. Expression (21) extends to cover (20) in the obvious way for the t horizon case.

4.2. Cash Flow Data

We now apply the preceding to the valuation of eight cash flows. There are seven industrial profit series for real per capita income and real per capita consumption of non-durables and services, annually, for 1959–2015. (For the cash flow sources and construction, see Appendix A.2 of the data appendix, Appendix A.) The eighth cash flow is obtained by treating measured consumption as a cash flow; i.e., we compute the risk premium on the (endowment) asset that pays out annual consumption. The basic statistics for the cash flows, labeled 1–8 are shown in Table 2.

Table 2. Summary statistics.

	Industry	$gdp_t - cf_t$		$\Delta cf_{t-1,t}$	
		Mean	**Std. Dev.**	**Mean**	**Std. Dev.**
1	Total Corporate	2.572	0.227	0.019	0.131
2	Federal Reserve Banks	5.834	0.343	0.045	0.131
3	Other Financial	4.366	0.468	0.028	0.360
4	Total Nonfinancial	2.828	0.272	0.016	0.146
5	Manufacturing	3.626	0.473	0.004	0.239
6	Wholesale Trade	5.194	0.260	0.027	0.137
7	Retail Trade	5.141	0.332	0.028	0.252
8	Consumption (NDS)	0.597	0.063	0.022	0.013
		$sdf_{t-1,t}$		$gdp_t - gdp_{t-1}$	
	log-MRS and GDP growth	−1.2231	1.8595	0.0196	0.0202

Displayed are sample means and standard deviations. See display (9) of the text for notation and the Appendix A sources. Note that Fed Reserve Banks means commercial banks in the Federal Reserve System.

Some of the industrial cash flows are aggregates. but none is a complete aggregate of any of the others. For example, cash flow 1, Total Corporate profits, includes items such as transportation and utilities, which are not among the other categories because of a lack of consistent data over the entire sample period. The bottom section of Table 2 also shows the basic statistics for the extracted log-SDF and GDP growth processes.

These cash flows do not correspond to the payoffs of traded securities, but using the above methods, we can compute the risk premiums on the stripped cash flows using (19) and (20). Among other things, we can then examine issues such as the reasonableness of the risk premiums relative to the characteristics of the industries and their term structure at short- and long-term horizons.

The specification of (9) as a VAR presumes that log-cash flow and log-GDP processes are co-integrated with $gdp_t - cf_t$ as the stationary error correction process. As a check, the first seven panels (excluding cash flow 8) of Figure 2 show time-series plots of the $gdp_t - cf_t$ process for the seven industrial profit series, each of which appears to be reasonably treated as realizations of a stationary process.

We estimated autoregressions of the form

$$gdp_t - cf_t = a + \rho(gdp_{t-1} - cf_{t-1}) + e_t, \tag{22}$$

and for the seven industrial profit series, the estimates of ρ ranged between 0.71 and 0.93, with a median of 0.86, and generally, we rejected, quite strongly, $H_0 : \rho = 1$ in favor of $H_1 : \rho < 1$. For the consumption cash flow series, the results are different, as the consumption of nondurables and services grew steadily relative to GDP during the transition from a production to a service economy over our particular sample period. Thus, for the consumption cash flow, we cannot invoke co-integration, and we simply use $gdp_t - gdp_{t-1}$ in place of $gdp_t - cf_t$ as the predictor variable in the vector autoregression (11). The log consumption growth series is displayed in the lower-right corner of Figure 2.
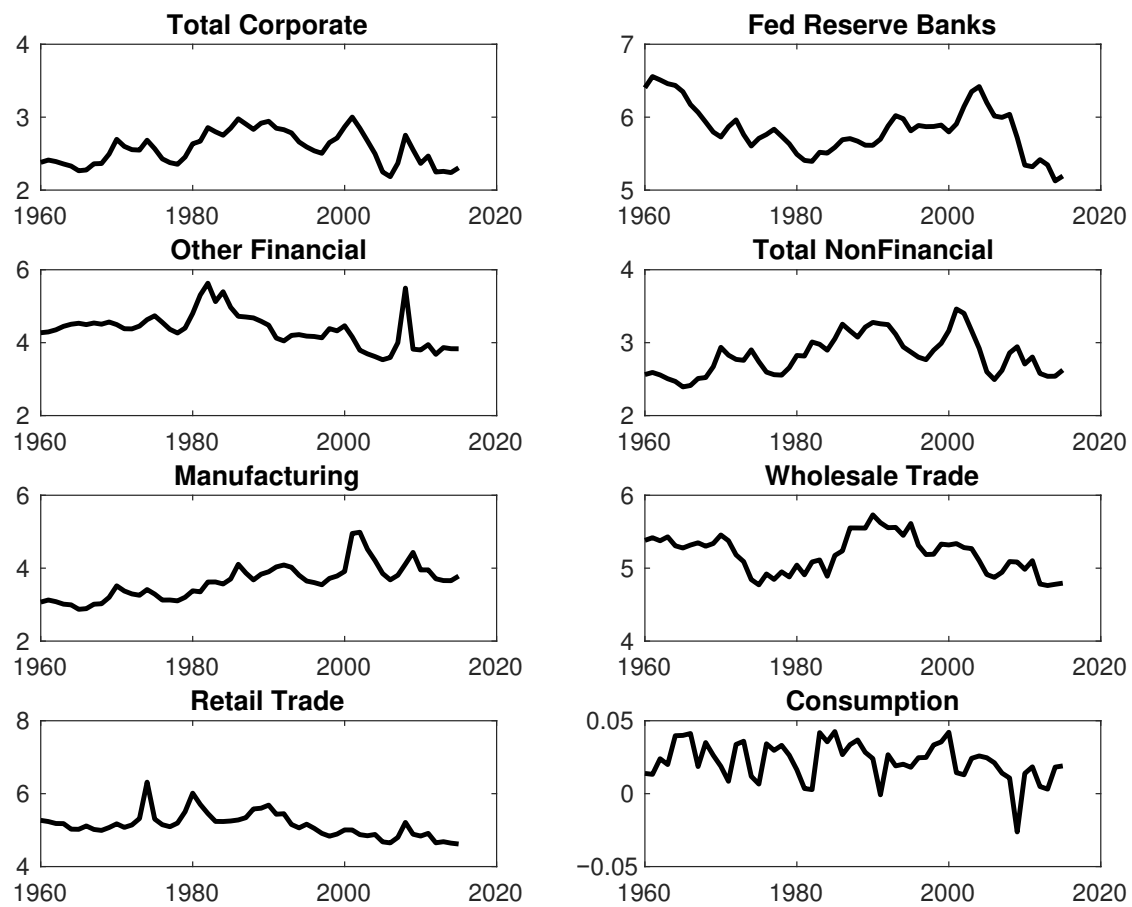
Figure 2. Plotted are $gdp_t - cf_t$ for the seven industrial cash flows 1–7 and $\Delta cf_{t-1,t}$ for cash flow 8 (consumption), where the variables are defined in display (9) of the text.

To verify the predictability, we estimated forecasting regressions of the form

$$\Delta cf_{t,t+1} = a + b_1 \Delta cf_{t-1,t} + b_2(gdp_t - cf_t) + u_t, \tag{23}$$

for each of the seven industrial cash flows, and we found very strong evidence for additional predictability coming from the error correction variable $gdp_t - cf_t$. For consumption (the eighth cash flow), we performed the regression as

$$\Delta cf_{t,t+1} = a + b_1 \Delta cf_{t-1,t} + b_2(gdp_t - gdp_{t-1}) + u_t, \tag{24}$$

and found only mild evidence for additional predictability coming from the second right-hand variable.

4.3. Risk Premiums

Using the methods described in Sections 3 and 4.1 above, we computed for each of the eight cash flows the implied risk premiums (20) at horizons $t = 1, 2, \ldots, 50$ for each available year (1960–2015) after lags. Since the risk premiums show negligible temporal variation, we onlyt report and discuss the full sample averages. Table 3 shows the one-period computations, where the average $\mathrm{Cov}(\Delta cf_{t-1,t}, sdf_{t-1,t})$ is the logarithmic SDF exposure of the cash flow. To ease interpretation, the table shows the average correlations as well. The negative of the covariance is the one-period risk premium calculated by way of (21) in this log-Gaussian framework. As seen from the table, all the cash flows carry

positive risk premiums except for Retail Trade, which is plausibly seen to be a hedging cash flow. For easier interpretation, we convert the covariance exposure to beta exposure,

$$\beta = -\overline{\text{Cov}(\Delta c f_{t-1,t}, s d f_{t-1,t})/\text{Var}(s d f_{t-1,t})} \tag{25}$$

and plot the relationship between the risk premiums and the (negative) log-SDF beta exposure in Figure 3. The essentially exact linear relationship seen in the figure is a mechanical consequence of the computations, but nonetheless, it is interesting to note that the price of risk is very close to 0.05, meaning an increase in the average return of 5 percent per year per unit exposure of the cash flow to moves in $-s d f_{t-1,t}$.

Table 3. Cash flow exposures.

	Industry	Exposure	
		Cov.	**Corr.**
1	Total Corporate	−0.083	−0.37
2	Fed Reserve Banks	−0.042	−0.21
3	Other Financial	−0.138	−0.23
4	Total Nonfinancial	−0.087	−0.36
5	Manufacturing	−0.161	−0.41
6	Wholesale Trade	−0.082	−0.34
7	Retail Trade	0.044	0.11
8	Consumption	−0.010	−0.35

Displayed are the implied average $\text{Cov}(\Delta c f_{t-1,t}, s d f_{t-1,t})$ and, similarly, the correlation.

Far more interesting are the paths of the risk premiums in (20) for t forward into the future, which are shown in Figure 4 for $t = 1, 2, \ldots, 50$. For the industrial cash flows 1–6, i.e., those other than the hedging cash flow, the risk premiums show some increase at short horizons but then decrease with horizons, and they decrease by a factor of one half from 1 to 50 years out; by contrast, for the consumption cash flow, the risk premium increases from about 1 percent per year to 4 percent per year when moving from 1 to 50 years out.

These results are exactly in line with what one would expect from Croce et al. (2015, p. 723). In their limited information (bounded rationality) model, there is long-run risk embedded in consumption, which under usual parameterizations thereby carry an increasing risk premium at longer horizons, just as seen in Table 2. There is also long-run risk embedded in individual cash flows, but it is obscured by a high level of cash flow noise that can be correlated with short-term consumption risk. Agents following optimal filtering rules thereby view the assets as much more covariance-risky in the short run than in the long run, and so we expect to observe higher risk premiums for the short run than the long run. It bears noting these findings are model free in that no a priori economic theory of the discount factor was imposed in the estimation.

A final matter is the summability of the stripped cash flows, which relates to the issue of whether the asset that pays the entire stream $\{CF_s\}_{s=1}^{\infty}$ is even sensible. Using basic computations such as those of Burnside (1998) for log-Gaussian models, in the iid case, the convergence of $\sum_{t=1}^{\infty} PV_{0,t}(CF)$ is assured if

$$\mathcal{E}(\Delta c f_{t-1,t}) - r^f + \text{Cov}(\Delta c f_{t-1,t}, s d f_{t-1,t}) < 0. \tag{26}$$

In a general case, such as that considered by De Groot (2015), the summability conditions are more involved, as they involve interactions between conditional means and variances, but the basic intuition of (26) remains: the covariance between the cash flow growth has to be sufficiently negative to overcome any excess of the expected cash flows over the risk-free rate. In our case, the cash flows grow at implied rates around 2 percent per year, while the long-term interest rate prior centers on 2 percent, but the covariances in Table 3 are negative, so the sums converge numerically, albeit very slowly. An exception is

Retail Trade, the hedging cash flow where the average covariance is positive. Some numerical instability of the partial sums for this cash flow was seen if the range was extended to 100+ years, which is not surprising for an extrapolation so far beyond the range of the data.
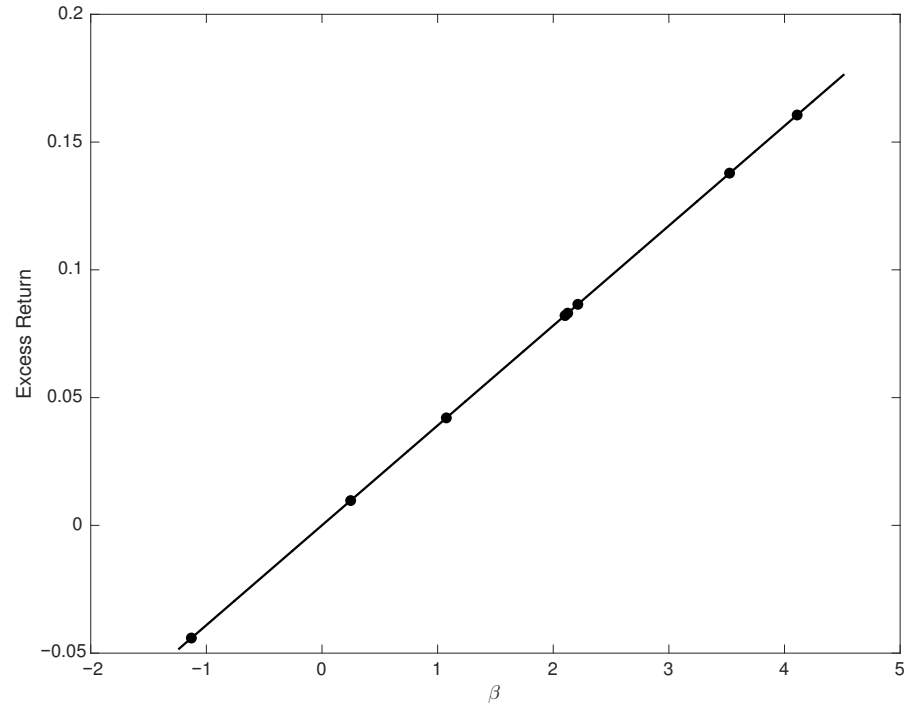
Figure 3. Plotted are the average 1-year risk premiums versus the cash flow beta exposures to sdf risk. Points for cash flows 6 and 8 nearly over plot, creating the illusion of seven points, while there are actually eight.
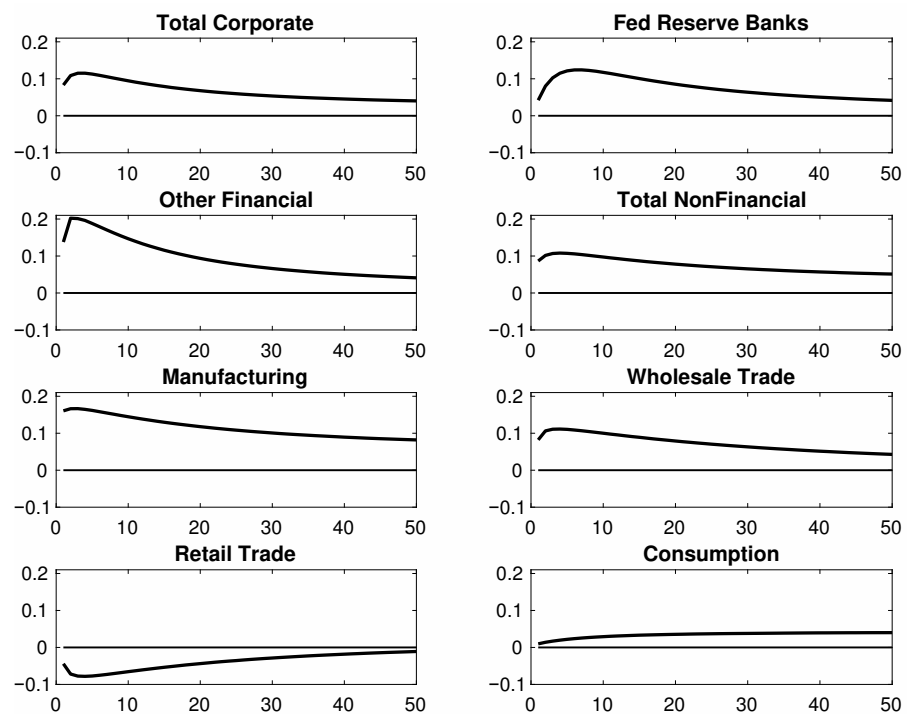
Figure 4. Plotted are the annualized risk premiums on cash flow strips for 1–50 years ahead.

5. Robustness

The objective of the paper was to adapt a Bayesian methodology (Gallant and Hong 2007) to a completely model-free data-only setting and then value important nontraded cash flows for the economic analysis of risk premiums. It only considers this specific objective and reaches interesting economic conclusions, but it is not exhaustive.

There are two issues regarding robustness that should be remarked upon: Is the methodology sensitive to the choice of assets? Is the methodology sensitive to the choice of prior?

As regards the assets, the methodology extracts the ex post SDF from the asset-pricing errors on 754 dynamic portfolio returns induced by instrumenting a smaller core set of asset returns. The core set, described in the Appendix A, is representative of those used for evaluating asset-pricing models; see, for example, the review by Bryzgalova et al. (2020, p. 3). The methods for the extraction of an SDF can be dependent on the assets used for the extraction, e.g., Nieto and Rubio (2014). However, our surmise is confidence that the large set of portfolio returns spans the factors on which all assets load, especially upon taking into account that the portfolios are dynamic and include instrumental variables such as consumption and labor income growth. Rather than the assets, it is the prior to which the Gallant and Hong (2007) methodology is sensitive. This issue we have examined, as reported above.

A referee suggests other references that a reader might consider. Lewellen et al. (2010) and Ghosh et al. (2017) discuss other methods for the extraction of the SDF, rather than the valuation of nontraded cash flows, though it would be interesting to extend these studies to our application. That step is beyond the scope of this paper. Additionally, Gormsen (2020); Bansal et al. (2019); and Giglio et al. (2020) further examine evidence on the slope of the term structure of equity returns without consensus. Our main finding regards the contrast between the term structure of the risk premiums on consumption versus that of those on the industrial cash flows we considered.

6. Conclusions

We developed a model-free Bayesian procedure to extract the SDF process using a yield curve prior that enforces the historically-very-low U.S. short- and long-term interest rates. The prior thereby enforces known information, but no particular theory on the SDF process itself. Using annual data for 1959–2015, we used the extracted SDF to compute the implied stripped cash flow risk premiums on panel corporate profits for eight major industrial sectors and consumption. The magnitudes of the risk premiums on the stripped cash flows are plausible, and, with one exception, the risk premiums show a decreasing term structure for 1–50-year horizons. The exception is Retail Trade, which is found to be a hedging asset in the short run but not the long run. By contrast, the risk premiums on the stripped consumption cash flow are found to be positive and rather low in the short term but increase with the horizon to about 4 percent per year 50 years out. The observed term structures of the equity risk premiums generally confirm the limited information (bounded rationality) model of Croce et al. (2015).

Appendix A. Data

Appendix A.1. SDF Extraction

In the extraction step, all the data were annual for the years 1930 through 2015. The raw data were converted from nominal to real using the annual consumer price index obtained from Table 2.3.4 on the Bureau of Economic Analysis web site. Conversions to per capita were performed by means of the mid-year population data from Table 7.1 on the Bureau of Economic Analysis web site.

The raw data for stock returns were value-weighted returns including dividends for the NYSE, AMEX, and NASDAQ from the Center for Research in Security Prices data on the Wharton Research Data Services web site (http://wrds.wharton.upenn.edu) (accessed on 23 February 2021). Likewise, the raw data for returns on U.S. Treasury 30-day debt were obtained from the Center for Research in Security Prices data on the Wharton Research Data Services web site.

Raw annual returns including dividends on the twenty-five Fama and French (1993) portfolios were obtained from Kenneth French's web site, http://mba.tuck.dartmouth.edu/pages{\protect\penalty\z@}/faculty/ken.french (accessed on 23 February 2021). The portfolios were the intersections of five portfolios formed on market equity and five portfolios formed on the ratio of book equity to market equity. The portfolios were for all NYSE, AMEX, and NASDAQ stocks for which equity data were not missing and book equity data were positive. The portfolios were constructed at the end of each June with breakpoints determined by the NYSE quintiles at the end of June. The complete details are on Kenneth French's web site. The advantage of the Fama–French portfolios here is that they appeared to isolate and exhaust the risk factors for holding equities (Fama and French 1992, 1993).

The raw labor income data were the "compensation of employees received" from Table 2.2 on the Bureau of Economic Analysis web site.

The Real Treasury Inflation Protected Securities (TIPS yields displayed in Table 1 are annual averages of daily values from https://www.treasury.gov/resource-center/data-chart-center/interest-rates/Pages/TextView.aspx?data=realyieldAll (accessed on 23 February 2021). According to the Treasury, "These rates are commonly referred to as "Real Constant Maturity Treasury" rates, or R-CMTs. Real yields on Treasury Inflation Protected Securities (TIPS) at "constant maturity" are interpolated by the U.S. Treasury from Treasury's daily real yield curve. These real market yields are calculated from composites of secondary market quotations obtained by the Federal Reserve Bank of New York. This method provides a real yield for a 10 year maturity, for example, even if no outstanding security has exactly 10 years remaining to maturity".

Appendix A.2. Valuation

The raw data for the cash flow valuation step were annual data for 1959–2015, for 56 observations, net, after a provision for the initial lag.

The industrial cash flow data were annual corporate profits with inventory valuation adjustments and without capital consumption allowances for major sectors. The data were spliced together for consistency from Table B-6 of the 2017 *Economic Report of the President* and Table B-91 of the 2004 *Economic Report of the President*. The GDP data were from NIPA Table 1.1.5; the consumption of nondurable goods and services data were from NIPA Table 2.3.5; nominal data were converted to real using the implicit GDP deflator (Price Index for Gross Domestic Product) from NIPA Table 1.1.4. All the data series were converted to per capita using the total U.S. resident population plus armed forces overseas (annual averages of monthly estimates) obtained from FRED, https://fred.stlouisfed.org/series/B230RC0A052NBEA (accessed on 23 February 2021).

References

Backus, David, Nina Boyarchenko, and Mikhail Chernov. 2016. Term structures of asset prices and returns. *Journal of Financial Economics* 129: 1–23. [CrossRef]

Bansal, Ravi, Shane Miller, and Amir Yaron. 2017. Is the Term Structure of Equity Risk Premia Upward Sloping? Working Paper. Available online: https://faculty.fuqua.duke.edu/~rb7/bio/BMY_dividend_strips.pdf (accessed on 22 March 2020).

Bansal, Ravi, and Amir Yaron. 2004. Risks For the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59: 1481–509. [CrossRef]

Bansal, Ravi, Shane Miller, Dongho Song, and Amir Yaron. 2019. The Term Structure of the Equity Risk Premia. NBER Working Paper No. 25690. Available online: https://www.nber.org/papers/w25690 (accessed on 22 March 2020).

Burnside, Craig. 1998. Solving asset pricing models with Gaussian shocks. *Journal of Economic Dynamics & Control* 22: 329–40.

Bryzgalova, Svetlana, Jiantao Huang, and Christian Julliard. 2020. Bayesian Solutions for the Factor Zoo: We Just Ran Two Quadrillion Models. Working Paper. Available online: https://apps.olin.wustl.edu/conf/SBIES/files/pdf/2020/102.pdf (accessed on 17 February 2021).

Campbell, John Y. 2003. Consumption-Based Asset Pricing. In *Handbook of the Economics of Finance*. Amsterdam, London and New York: Elsevier, vol. 1B.

Christensen, Timothy M. 2017. Nonparametric Stochastic Discount Factor Decomposition. *Econometrica* 85: 1501–36. [CrossRef]

Cochrane, John H. 2017. Macro-finance. *Review of Finance* 21: 945–85. [CrossRef]

Croce, Marino M., Martin Lettau, and Sydney C. Ludvigson. 2015. Investor Information, Long-Run Risk, and the Term Structure of Equity. *The Review of Financial Studies* 28: 706–42. [CrossRef]

De Groot, Oliver. 2015. Solving asset pricing models with stochastic volatility. *Journal of Economic Dynamics & Control* 52: 308–21.

Fama, Eugene F., and Kenneth R. French. 1992. The Cross-Section of Expected Stock Returns. *Journal of Finance* 59: 427–65. [CrossRef]

Fama, Eugene F., and Kenneth R. French. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33: 3–56. [CrossRef]

Gallant, A. Ronald. 2016a. Reflections on the Probability Space Induced by Moment Conditions with Implications for Bayesian Inference. *Journal of Financial Econometrics* 14: 227–28. [CrossRef]

Gallant, A. Ronald. 2016b. Reply to Comment on Reflections. *Journal of Financial Econometrics* 14: 284–94. [CrossRef]

Gallant, A. Ronald. 2020. Complementary Bayesian Method of Moments Strategies. *The Journal of Applied Econometrics* 35: 422–39. [CrossRef]

Gallant, A. Ronald, and Han Hong. 2007. A Statistical Inquiry into the Plausibility of Recursive Utility. *Journal of Financial Econometrics* 5: 523–59. [CrossRef]

Gamerman, Dani, and Hedibert F. Lopes. 2006. Markov Chain Monte Carlo: Stochastic Simulation. In *For Bayesian Inference*, 2nd ed. Boca Raton: Chapman and Hall.

Ghosh, Anisha, Christian Julliard, and Alex P. Taylor. 2017. What Is the Consumption CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models. *The Review of Financial Studies* 30: 442–504. [CrossRef]

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel. 2015. Very long-run discount rates. *The Quarterly Journal of Economics* 130: 1–53. [CrossRef]

Giglio, Stefano, Bryan T. Kelly, and Serhiy Kozak. 2020. Equity Term Structures without Dividend Strips Data. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3533486 (accessed on 17 February 2021).

Gormsen, Niels Joachim. 2020. Time variation of the equity term structure. *Journal of Finance* forthcoming. [CrossRef]

Lewellen, Jonathan, Stefan Nagelb, and Jay Shanken. 2010. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics* 96: 175–94. [CrossRef]

Nieto, Belen, and Gonzalo Rubio. 2014. Volatility Bounds, Size, and Real Activity Prediction. *Review of Finance* 18: 373–415. [CrossRef]

Tesar, Linda, and Maurice Obstfeld. 2015. Long-Term Interest Rates—A Survey. In *Council of Economic Advisers White Paper*. pp. 1–54. Available online: https://obamawhitehouse.archives.gov/sites/default/files/docs/interest_rate_report_final.pdf (accessed on 17 February 2021)

van Binsbergen, Jules, Michael Brandt, and Ralph Koijen. 2012. On the timing and pricing of dividends. *American Economic Review* 102: 1596–618. [CrossRef]

van Binsbergen, Jules, Michael Brandt, and Ralph Koijen. 2017. The term structure of returns: Facts and theory. *Journal of Financial Economics* 124: 1–21. [CrossRef]