Price Dividend Ratio and Long-Run Stock Returns: a Score Driven State Space Model

Davide Delle Monache
Bank of Italy
davide.dellemonache@bancaditalia.it

Ivan Petrella
University of Warwick and CEPR
Ivan.Petrella@wbs.ac.uk

Fabrizio Venditti
European Central Bank
fabrizio.venditti@ecb.int

Sunday 26th April, 2020

Abstract

In this paper we develop a general framework to analyze state space models with time-varying system matrices, where time variation is driven by the score of the conditional likelihood. We derive a new filter that allows for the simultaneous estimation of the state vector and of the time-varying matrices. We use this method to study the time-varying relationship between the price dividend ratio, expected stock returns and expected dividend growth in the US since 1880. We find a significant increase in the long-run equilibrium value of the price dividend ratio over time, associated with a fall in the long-run expected rate of return on stocks. The latter can be attributed mainly to a decrease in the natural rate of interest, as the long-run risk premium has only slightly fallen.

JEL codes: C32, C51, C53, E44, G12.
Keywords: State space models, time-varying parameters, score-driven models, equity premium, present-value models.

 The authors would like to thank Juan Antolin-Diaz, Gino Cenedese, Laura Coroneo, Ana Galvao, Domenico Giannone, Arie Gozluklu, Andrew Harvey, Gary Koop, Siem Jan Koopman, André Lucas, Massimiliano Marcellino, Daniele Massacci, Pavol Povala, Barbara Rossi, Bernd Schwaab, Ron Smith and Shaun Vahey for their useful suggestions.

1 Introduction

A decade after the Great Recession the global economy is mired in an environment of low real interest rates, low growth and high stock valuations. Whether this is a "new normal" is a question that is at the center of a new research agenda on macro-financial trends in a changing environment (Caballero et al., 2017). The issue of structural breaks is back in the spotlight, as it is the use of time series models that allow for parameter instability. In this paper we contribute to this debate by developing a general method to analyze state space models where parameters change over time and by applying this method to study the evolving relationship between stock valuations, stock returns and dividend growth.

The econometric method that we propose posits a law of motion for the parameters that is a linear function of the score of the conditional likelihood, following Creal et al. (2008) and Harvey (2013). We derive the analytical expressions for a new set of recursions that, running in parallel with the Kalman Filter (KF), update at each point in time both the vector of time-varying parameters (TVP) and the latent states. Within this framework, the likelihood of any Gaussian state space model with TVP is available in closed form and the model can be estimated by maximum likelihood (ML). A unique feature of our method is that it can easily handle parameter constraints, which might arise from economic theory or from a deliberate choice of the econometrician. Constraints are taken into account directly in the estimation algorithm through a Jacobian function. Finally, missing observations, mixed frequencies and the shrinkage of the TVP towards desired values are easily dealt with. A Monte Carlo exercise shows that the method can replicate the salient features of various data generating processes. In particular, our method delivers constant coefficients when the data are simulated from a fixed coefficient model, and tracks time variation in parameters when this is present in the data.

We are not the first ones to tackle the issue of parameter variation in state space models. Harvey et al. (1992), Koopman et al. (2010), Eickmeier et al. (2015) and Koop and Korobilis

(2013) estimate state space models with TVP using either Maximum Likelihood or forgetting factors. All these papers focus on specific models for which they develop ad-hoc estimation methods. The algorithm that we develop in this paper is instead general and can be used to analyze any model that can be cast in state space form. A second strand of the literature analyzes models with TVP using Bayesian simulation techniques (see e.g., Cogley and Sargent, 2005; Stock and Watson, 2007; Durbin and Koopman, 2012). In some settings, our method presents a distinctive advantage with respect to these methods, as the likelihood of the model can always be computed using the standard Kalman filter, even in presence of non-linear restrictions in the time varying parameters. Restrictions on the system matrices of the state space are embedded in the updating steps of the unobserved components and of the model parameters with no additional computational costs.

We use the methodology developed in the first part of the paper to revisit the relationship between the price dividend ratio, the return on stocks and dividend growth in present value models. State space models are particularly attractive to study present value relationships because they handle efficiently complex dynamics, while avoiding over-parameterization (Binsbergen and Koijen, 2010). We estimate a generalization of the Campbell and Shiller decomposition (Campbell and Shiller, 1988), which allows for time variation in the steady state of expected dividend growth and of the expected return on stocks (and therefore the steady state price dividend ratio as in Lettau and Nieuwerburgh, 2008) as well as in the conditional variance of the shocks (as in Piatti and Trojani, 2017). We show that expected returns and, to a lesser extent, expected dividend growth, exhibit slow moving steady states. Expected returns, in particular, have experienced a continuous decline in their long-run equilibrium value. This decline accelerated in the 1960s and in the 1990s, prior to the stock market crashes of the early 1970s and 2000s, and is reflected in an upward trend of long-run price dividend ratio, which so far the literature failed to explain. An intuition for this result is sketched by Fama and French (2002) in their seminal paper on the equity risk premium. They use long historical

data on dividend growth and dividend yields to obtain an estimate of expected stock returns and of the implied risk premium. Upon noticing that dividend growth is broadly stationary but that the price dividend ratio has instead risen considerably, they conclude that the only logical explanation is a decline in expected returns. We provide an econometric method and a model specification that formalizes their intuition and that confirms its empirical validity. We also add to a growing literature on the relative role played by the natural rate of interest (r-star) and the risk-premium in explaining long-term trends in the return on stocks (Greenwald et al., 2019; Farhi and Gourio, 2018). We decompose long-run expected returns into a riskless component and a risk premium and find that the former has remained relatively stable until the beginning of the 1960s, to decrease rapidly thereafter. The long-run equity risk premium has only slightly fallen, from 4 to 3 percent.

Structure of the paper. Section 2 constitutes the methodological body of the paper. Section 3 discusses present value models with shifting steady states and provides details on the model specification. Sections 4 and 5 present the empirical analysis. Section 6 concludes.

2 Score driven state space models

Let us assume that a given time series model has the following state space representation:

$$\begin{aligned} y_t &= Z_t\alpha_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, H_t), \\ \alpha_t &= T_t\alpha_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, Q_t), \quad t = 1, ..., n, \end{aligned} \tag{1}$$

where y_t is the $N \times 1$ vector of observed variables, ϵ_t is the $N \times 1$ vector of measurement errors, α_t is the $m \times 1$ vector of state variables and η_t is the corresponding $m \times 1$ vector of disturbances. The two disturbances are assumed to be Gaussian distributed and uncorrelated for all time periods, that is $E(\epsilon_t\eta_s') = 0$ for $\forall t, s$. This assumption can be relaxed at the cost of a complication in the filtering formulae. The initial value of the state vector is also assumed to be Gaussian, $\alpha_0 \sim \mathcal{N}(a_0, P_0)$ and uncorrelated $\forall t$ with ε_t and η_t.

Following Harvey (1989, sec. 3.1) it is usually assumed that the *system matrices* Z_t, H_t, T_t

and Q_t are non-stochastic. As a result the system (1) is *linear* with respect to the state vector. Conditional on the information set $Y_{t-1} = \{y_{t-1}, ..., y_1\}$ and on the vector of parameters θ, the state vector and the observations are both Gaussian distributed; i.e. $y_t|Y_{t-1}; \theta \sim \mathcal{N}(Z_t a_t, F_t)$ and $\alpha_t|Y_{t-1}; \theta \sim \mathcal{N}(a_t, P_t)$, and the log-likelihood function at time t is:

$$\ell_t = \log p(y_t|Y_{t-1}, \theta) \propto -\frac{1}{2}\left(\log|F_t| + v_t' F_t^{-1} v_t\right). \tag{2}$$

The prediction error v_t, its covariance matrix F_t, the conditional mean of the state vector a_t, and its mean square error (MSE) matrix P_t, are recursively estimated by means of the KF:

$$\begin{aligned} v_t &= y_t - Z_t a_t, & F_t &= Z_t P_t Z_t' + H_t, \\ a_{t|t} &= a_t + P_t Z_t' F_t^{-1} v_t, & P_{t|t} &= P_t - P_t Z_t' F_t^{-1} Z_t P_t, \\ a_{t+1} &= T_{t+1} a_{t|t} & P_{t+1} &= T_{t+1} P_{t|t} T_{t+1}' + Q_{t+1}, \quad t = 1, ..., n. \end{aligned} \tag{3}$$

Specifically, we have that $a_t = \mathrm{E}(\alpha_t|Y_{t-1}, \theta)$ is the so-called *predictive* filter with its MSE matrix being $P_t = \mathrm{E}[(a_t - \alpha_t)(a_t - \alpha_t)'|Y_{t-1}, \theta]$, while $a_{t|t} = \mathrm{E}(\alpha_t|Y_t, \theta)$ is the so-called *real time* filter with MSE equal to $P_{t|t} = \mathrm{E}[(a_{t|t} - \alpha_t)(a_{t|t} - \alpha_t)'|Y_{t-1}, \theta]$. The state space model in (1) is the so-called *contemporaneous* form used in Harvey (1989). Durbin and Koopman (2012) use instead the so-called *forward* form. In this paper we prefer to use the former so that the system matrices have all the same time dependency with respect to the vector of TVP. Using the forward form, the time dependency of matrices T and Q needs to be adapted; for details see Appendix A.

We assume that the changes in the system matrices over time are endogenous and depend on past observations. Thus, although stochastic, the system matrices are *predetermined*, meaning that conditional on past observation they can be regarded ad being fixed. As a result, the model is still conditionally Gaussian like the one introduced by Harvey (1989, sec. 3.7.1).[1] This setup has three attractive features. First, both the state vector and the observations

[1]The KF generates the conditional Gaussian distribution of the state vector where the mean is no longer linear in the observations and the MSE (conditional error covariance) depends upon the particular realization of the observations in the sample. Chen et al. (1989) state the necessary conditions in order for the KF to generate the conditional mean and covariance of the Gaussian distributed state vector.

are conditionally Gaussian. Second, the likelihood function can be written in the form of the prediction error decomposition (2) and computed by means of the KF (3). Third, although the model is not linear in the observations, the KF delivers the minimum mean square estimates of the state vector (see Harvey, 1989, p. 342). The key analytic challenge here is represented by the joint updating of both the system matrices and of the state vector. To solve this problem, we propose a new set of recursions that run in parallel with the KF.

2.1 Score driven system matrices

The time-varying elements of the system matrices in (1) are collected in the vector f_t also known as the TVP vector. As in Creal et al. (2008) and Harvey (2013), we posit the following score driven law of motion for such vector:

$$f_{t+1} = \mathrm{c} + \mathrm{A} f_t + \mathrm{B} s_t, \quad s_t = \mathcal{S}_t \nabla_t, \quad t = 1, ..., n, \tag{4}$$

with

$$\nabla_t = \frac{\partial \ell_t}{\partial f_t}, \quad \mathcal{S}_t = -\mathrm{E}_t \left(\frac{\partial \ell_t^2}{\partial f_t \partial f_t'} \right)^{-1}, \tag{5}$$

where ℓ_t is the conditional log-likelihood function of the model (1), ∇_t is the score (gradient) with respect to f_t and the scaling matrix, $\mathcal{S}_t$, is the inverse of the information matrix $\mathcal{I}_t$. In this case, s_t has zero conditional mean and conditional variance equal to the inverse of the information matrix.[2] The parameters in B determine the sensitivity of the time-varying parameters to the score of the conditional likelihood, and therefore to the information contained in the prediction error. The special case of constant system matrices is obtained by setting this matrix to 0. The system matrices may contain both time-varying and constant elements. We collect the latter in the vector θ_m. Thus, at each point in time, the system matrices depend upon f_t and θ_m, namely $Z_t = Z(f_t, \theta_m)$, $T_t = T(f_t, \theta_m)$, $H_t = H(f_t, \theta_m)$, and $Q_t = Q(f_t, \theta_m)$.

[2] Alternatively, one could choose $\mathcal{S}_t = \mathcal{I}_t^{-1/2}$, in which case the conditional variance of the score is equal to the identity matrix. One can also set $s_t = \nabla_t$, in which case the score has conditional variance equal to the information matrix. In general, to avoid numerical instability in the scaling matrix we replace $\mathcal{S}_t$ with its smoothed counterpart $\tilde{\mathcal{S}}_t = (1 - \kappa_h)\mathcal{S}_t + \kappa_h \tilde{\mathcal{S}}_{t-1}$.

The score vector s_t is computed conditional on the information up to time t, thus the vector f_t is entirely determined by past observations and by the vector of static parameters $\theta_f = \{\mathrm{c}, \mathrm{A}, \mathrm{B}\}$. Since the dynamic of the system matrices is observation-driven, i.e. entirely determined by past observations and by the vector $\theta = (\theta_f', \theta_m')'$, the model is conditional Gaussian and the log-likelihood (2) can be evaluated recursively through the KF.

The gradient and the information matrix can be analytically compute by the expressions presented below.

Result 1 *Given the model (1)-(2), the score and the information matrix are:*

$$\begin{aligned} \nabla_t &= \tfrac{1}{2}\left[\dot{F}_t'(F_t \otimes F_t)^{-1} vec(v_t v_t' - F_t) - 2\dot{V}_t' F_t^{-1} v_t\right] \\ \mathcal{I}_t &= \tfrac{1}{2}\left[\dot{F}_t'(F_t \otimes F_t)^{-1}\dot{F}_t + 2\dot{V}_t' F_t^{-1}\dot{V}_t\right], \quad t = 1, ..., n, \end{aligned} \tag{6}$$

where $\dot{V}_t = \partial v_t / \partial f_t'$ and $\dot{F}_t = \partial vec(F_t)/\partial f_t'$ measure the sensitivity of the prediction error v_t and its variance F_t with respect to f_t. Proofs in Appendix A.1.

Notice that all the elements of the information matrix $\mathcal{I}_t$ are computed using information up to time $t-1$. On the other hand, the gradient ∇_t contains the current observation y_t via the prediction error v_t. The terms $\dot{V}_t$ and $\dot{F}_t$ play a key role in the gradient ∇_t. They measure the sensitivity of, respectively, the first and second moment of the state vector with respect to f_t. Together with the variance of the prediction error (F_t) and with the curvature of the conditional likelihood (proportional to $\mathcal{I}_t$), they determine the impact that new information, summarized in the prediction error v_t, has on the TVP vector. Notice that v_t and F_t are recursively computed by means of the KF (3), while their Jacobian counterparts, $\dot{V}_t$ and $\dot{F}_t$, are recursively computed through the new filter presented below.

Result 2 *The Jacobian counterpart of the KF leads to the following set of expressions:*

$$\begin{aligned} \dot{V}_t &= -[(a_t' \otimes I_N)\dot{Z}_t + (a_{t-1|t-1}' \otimes Z_t)\dot{T}_t], \qquad t = 1, ..., n, \\ \dot{F}_t &= 2N_N(Z_t P_t \otimes I_N)\dot{Z}_t + 2(Z_t \otimes Z_t)N_m(T_t P_{t-1|t-1} \otimes I_m)\dot{T}_t + \dot{H}_t + (Z_t \otimes Z_t)\dot{Q}_t, \end{aligned} \tag{7}$$

where $\dot{Z}_t = \partial vec(Z_t)/\partial f_t'$, $\dot{H}_t = \partial vec(H_t)/\partial f_t'$, $\dot{T}_t = \partial vec(T_t)/\partial f_t'$ and $\dot{Q}_t = \partial vec(Q_t)/\partial f_t'$ are the Jacobians of the system matrices with respect to f_t, and N_m is a symmetrizer matrix (i.e., for any $n \times n$ matrix, S, $N_n vec(S) = vec[\frac{1}{2}(S + S')]$). Proofs in Appendix A.2.

Putting together Results 1 and 2, we can compute the scaled score $s_t = \mathcal{S}_t \nabla_t$ and therefore recursively estimate the vector f_t using the the score-driven filter (4). Such *auxiliary filter* runs in parallel with the standard KF (3) as explained in the Algorithm described below.

Algorithm for the score driven state space model
Initialize $a_{0\|0}$, a_1, $P_{0\|0}$, P_1, f_1.
For $t = 1, ..., n$:
i. evaluate Z_t, T_t, H_t, Q_t, $\dot{Z}_t$, $\dot{T}_t$, $\dot{H}_t$, $\dot{Q}_t$
ii. compute v_t, F_t, $\dot{V}_t$, $\dot{F}_t$
iii. compute $a_{t\|t}$, $P_{t\|t}$, ∇_t $\mathcal{I}_t$, s_t
iv. compute f_{t+1}
v. evaluate Z_{t+1}, T_{t+1}, H_{t+1}, Q_{t+1}
vi. compute a_{t+1}, P_{t+1}

The vector of parameters θ can be estimated by ML, that is $\widehat{\theta} = \arg\max \sum_{t=1}^{n} \ell_t(\theta)$. Given the above algorithm, the evaluation of the log-likelihood function is straightforward and the maximization can be obtained numerically.[3]

Appendix B provides examples of popular models with TVP that have been used in the literature and that can be analyzed within our framework. In particular we discuss the local-level model with time-varying volatilities used by Stock and Watson (2007); the autoregressive models and the vector autoregression with time-varying parameters (as in Koop and Korobilis, 2013). Appendix C shows, through a detailed Monte Carlo analysis, that our score driven method successfully tracks parameters time variation for a variety of different Data Generating Processes (DGPs). Importantly, the Monte Carlo exercise also highlights that, when the true parameters do not change over time, the model correctly estimates them as constant.

2.2 Non-linearity in the system matrices

In many applications it is desirable to constrain the TVP at each point in time. Steady state relationships, for instance, often imply non-linear parameters constraints that must be taken into account when estimating the model. The score driven approach provides a general framework to deal with such constraints. In this respect it presents a clear advantage over

[3] As in Creal et al. (2013, sec. 2.3) one can conjecture that the usual ML results hold.

alternative methods, including Bayesian ones. When the system matrices (Z_t, T_t, H_t, Q_t) are a non-linear function of f_t, the effect of these parameters on the score is mediated via the Jacobians ($\dot{Z}_t$, $\dot{T}_t$, $\dot{H}_t$, $\dot{Q}_t$). Although the exact expression of the Jacobians is model specific, we offer a flexible expression for dealing with them. Let M_t be a generic system matrix of dimension $r \times c$, and decompose this matrix as follows:

$$\text{vec}\,(M_t) = S_0 + S_1\psi(S_2 f_t), \tag{8}$$

where S_0 is a $rc \times 1$ vector containing all the time-invariant elements in each of the entries of the matrix M_t, S_1 and S_2 are selection matrices that select respectively the time-varying elements of M_t and the sub-vector of f_t belonging to M_t. Finally, $\psi(\cdot)$ denotes the mapping function, also known as the link function, between f_t and the corresponding elements in M_t embedding parameter restrictions. Such function is assumed to be time-invariant, continuous, invertible, twice differentiable, and its Jacobian matrix is denoted by Ψ_t. Given the representation (8), the generic Jacobian matrix $\dot{M}_t$ can be computed as follows:

$$\dot{M}_t = \frac{\partial \text{vec}(M_t)}{\partial f_t'} = S_1\Psi_t S_2. \tag{9}$$

While equations (8)-(9) can be directly used to deal with Z_t and T_t, when modelling a generic covariance, Ω_t, it is often useful to decompose this in terms of volatilities and correlations, i.e. $\Omega_t = D_t R_t D_t$, where D_t is the diagonal matrix containing the standard deviations and R_t is the correlation matrix. Now R_t and D_t can be expressed in the form of (8), $\dot{R}_t$ and $\dot{D}_t$ can be computed as in (9) and $\dot{\Omega}_t$ is computed using standard rules of matrix differentiation.

2.3 Shrinkage, mixed frequency data, correlated disturbances

As the dimension of the system grows, it could be desirable to impose some shrinkage on the model parameters to avoid an increase in the estimation variance. In a Bayesian framework this is achieved through the prior distribution. In a classical setting, like the one hereby adopted, shrinkage can be achieved by means of stochastic constraints that lead to a mixed estimator

(Theil and Goldberger, 1961). In the appendix D we show how to modify the score-driven algorithm in order to take into account shrinkage on model parameters.

The presence of missing observations and data sampled at mixed frequencies does not present a particular challenge, as it only involves a re-weighing of the observations and temporal aggregation. A detailed treatment of these issues is discussed in Appendix E. Finally, Appendix F shows how to modify our algorithm in the presence of correlation between the innovations of the measurement and transition equations.

3 Price dividend ratio, expected returns and expected dividend growth

We use the method laid out above to study the relationship between the price dividend ratio, the return on stocks and dividend growth. This topic provides the ideal setting for score driven state space models. First, it involves present value relationships that can be conveniently analyzed via state space models (Binsbergen and Koijen, 2010). Second, there is significant evidence of instability in regressions of stock returns on the price dividend ratio (Paye and Timmermann, 2006; Lettau and Nieuwerburgh, 2008), suggesting that a state space model with TVP is the appropriate modelling choice. Third, the parameters that link these three objects are subject to a set of non-linear restrictions that pose a non-trivial challenge for other methods but can be easily dealt with in our framework. We start by recalling the steady state relationship between the return on stocks, the price dividend ratio and dividend growth.

3.1 Stock return and time-varying steady states

Let P_t and D_t denote stock prices and dividends. From the simple definition of gross return of an asset it follows that:

$$R_{t+1} \equiv \frac{P_{t+1} + D_{t+1}}{P_t} = \frac{D_{t+1}}{D_t} \frac{P_{t+1}/D_{t+1} + 1}{P_t/D_t}. \tag{10}$$

Lettau and Nieuwerburgh (2008) show that this implies the following relationship in logs:

$$\overline{\mathtt{pd}}_t = \overline{g}_t - \log(\exp \overline{\mu}_t - \exp \overline{g}_t), \tag{11}$$

where $\overline{\mathtt{pd}}_t$, $\overline{\mu}_t$ and $\overline{g}_t$ denote the steady state level of the price dividend ratio, of the return on stocks and of dividend growth, respectively.[4] Equation (11) has two important implications. First, changes in the steady state of the price dividend ratio reflect either changes in the steady state of the return on stocks or in the steady state of dividend growth or in both. Second, small changes in long-run growth (reflected in the steady state of dividend growth) and/or in the steady state of returns have large effects on the steady state of the price dividend ratio.

3.2 Preliminary evidence on parameter instability

Lettau and Nieuwerburgh (2008) report evidence in support of the hypothesis that the long-run mean of the price dividend ratio is subject to two structural breaks. In this section, we extend their analysis by considering not only the price dividend ratio but also the return on stocks and dividend growth and by testing for instability in the variance of these variables. Our analysis is based on annual data between 1873 and 2018. Annual data for the Standard and Poor Composite Stock Price Index and associated dividends are sourced from Robert Shiller's website.[5] We deflate total returns and dividends using data on US CPI.

We start by testing the null hypothesis of no breaks against the alternative hypothesis of one break (Bai and Perron, 2003). The results, reported in the top panel of Table 1, convey two clear messages. First, the hypothesis of no breaks can not be rejected for returns and dividend growth. The second result is that there is strong evidence of structural breaks in the mean of the price dividend ratio. Two of the dates for which the Bai and Perron procedure detects a break (1954 and 1995) are consistent with the findings in Lettau and Nieuwerburgh

[4] Lettau and Nieuwerburgh (2008) assume that, at the steady state, the level of the price dividend ratio is constant (i.e. $\overline{P/D}_{t+1} \approx \overline{P/D}_t$). Moreover, denoting with $\overline{DY}_t$ the steady state dividend yield, equation (11) implies that $\overline{DY}_t = \overline{R}_t - \overline{G}_t$, consistent with the traditional Gordon model.

[5] http://www.econ.yale.edu/shiller/ (see Shiller, 1989, Ch.26 for a discussion).

(2008). Evidence of a third break in 1913 is somewhat weaker. In fact, the null hypothesis of two breaks against the alternative of a third one cannot be rejected on the basis of sequential break tests (Table 1, central panel). These sequential tests also cannot reject the null that the mean of the return on stocks and of dividend growth has remained stable over time.

Last, we employ Nyblom (1989) test of the null hypothesis of constant parameters (both mean and variance) against the alternative that the parameters follow a martingale. The test detects significant shifts in the mean and volatility of the price dividend ratio as well as in the volatility of dividend growth. This result casts doubts on the hypothesis of constant variances maintained by Lettau and Nieuwerburgh (2008) and confirms that shifts in the steady state of the price dividend ratio are a robust feature of the data.

3.3 A score driven present value model with drifting steady states

The structural breaks analysis leaves one question open, that is how to reconcile the evidence of breaks in the price dividend ratio with the apparent stability of returns and dividend growth, given that the former is a function of the latter two. A plausible explanation is that changes in the long-run mean of returns and dividend growth are overshadowed by the presence of a very volatile transitory component. In such an environment state of the art break tests have low power against the alternative of no breaks (Cogley and Sargent, 2005). The changing relationship between the price dividend ratio, returns and dividend growth could be better captured by a flexible model that allows for gradual shifts in their long-run mean as well as in their volatility, which we now specify.

Lettau and Nieuwerburgh (2008) generalize the Campbell and Shiller decomposition to allow for time-varying steady states in expected returns and in expected dividend growth. They show that a first order approximation of (10) around a time-varying steady state yields:[6]

$$\mathtt{pd}_t - \overline{\mathtt{pd}}_t \simeq \sum_{j=1}^{\infty} \rho_t^j (\Delta d_{t+j} - \overline{g}_t) - \sum_{j=1}^{\infty} \rho_t^j (r_{t+j} - \overline{\mu}_t), \tag{12}$$

[6] Lettau and Nieuwerburgh (2008) assume that steady state log returns and dividend growth are martingales

where $\rho_t = \exp \overline{\mathtt{pd}}_t/(1+\exp \overline{\mathtt{pd}}_t)$ and, through $\overline{\mathtt{pd}}_t$, depends on $\overline{\mu}_t$ and $\overline{g}_t$ according to (11). Equation (12) shows that stock prices fall relative to dividends when there is bad news about future cash flows, or when discount rates (i.e. expected returns) rise. Therefore, the price dividend ratio should forecast returns and dividend growth (Cochrane, 2008). Most importantly, the presence of shifts in long-run expected returns and expected dividend growth implies that the sensitivity of $\mathtt{pd}_t - \overline{\mathtt{pd}}_t$ to news about cash flows and discount rates also changes over time (i.e. the higher the equilibrium level of the price dividend ratio, the higher ρ_t).

To take the model to the data one needs to make some additional assumptions on expected returns and expected dividend growth. Binsbergen and Koijen (2010) assume that these expected values can be modeled as AR(1) processes. Here we depart from their specification and assume that expected returns and expected dividend growth are the sum of a transitory and a persistent component. The latter, which constitutes a shift in long-run expected returns and expected dividend growth, is the key novelty of our model and we discuss it more in detail in the next sub-section. Specifically, we assume that:

$$\mathrm{E}_t(\Delta d_{t+1}) = \overline{g}_{t+1|t} + \tilde{g}_{t+1|t} \tag{13}$$

$$\mathrm{E}_t(r_{t+1}) = \overline{\mu}_{t+1|t} + \tilde{\mu}_{t+1|t}, \tag{14}$$

where the equilibrium levels of expected returns and expected dividend growth are defined as $\lim_{h\to\infty} \mathrm{E}_t(r_{t+h}) = \overline{\mu}_{t+1|t}$ and $\lim_{h\to\infty} \mathrm{E}_t(\Delta d_{t+h}) = \overline{g}_{t+1|t}$. Using the notation in Binsbergen and Koijen (2010) we denote the transitory component of the expectations ($\tilde{g}_{t+1|t}$ and $\tilde{\mu}_{t+1|t}$) with $\tilde{g}_t$ and $\tilde{\mu}_t$ and assume that they can be characterized by simple AR(1) models:

$$\tilde{g}_{t+1} = \phi_g \tilde{g}_t + \varepsilon_{g,t+1} \tag{15}$$

$$\tilde{\mu}_{t+1} = \phi_\mu \tilde{\mu}_t + \varepsilon_{\mu,t+1}. \tag{16}$$

The present value relationship (12) implies that the transitory component of the price dividend

(i.e. $\mathrm{E}_t(\overline{\mu}_{t+j}) = \overline{\mu}_t$ and $\mathrm{E}_t(\overline{g}_{t+j}) = \overline{g}_t$), as well as $\mathrm{E}_t(\overline{\mathtt{pd}}_{t+j}) = \overline{\mathtt{pd}}_t$ and $\mathrm{E}_t(\rho_{t+j}) = \rho_t$, and that that deviations from the mean price-dividend ratio are uncorrelated with ρ_t (i.e. $\mathrm{E}_t[\rho_{t+j}(\mathtt{pd}_{t+j} - \overline{\mathtt{pd}}_t)] = 0$). The specification of the model for $\overline{\mu}_t$ and $\overline{g}_t$, that we describe below, is consistent with the martingale assumption. As for $\overline{\mathtt{pd}}_t$ and ρ_t, Lettau and Nieuwerburgh (2008) show that the martingale assumption is satisfied to a very good approximation for reasonable break processes for $\overline{\mu}_t$ and $\overline{g}_t$.

ratio is related to $\tilde{\mu}_t$ and $\tilde{g}_t$:

$$\mathtt{pd}_t - \overline{\mathtt{pd}}_{t|t-1} \simeq -b_{1,t|t-1}\tilde{\mu}_t + b_{2,t|t-1}\tilde{g}_t, \tag{17}$$

with the following constraints on the parameters:

$$b_{1,t|t-1} = \frac{1}{1-\rho_{t|t-1}\phi_\mu}, \qquad b_{2,t|t-1} = \frac{1}{1-\rho_{t|t-1}\phi_g}. \tag{18}$$

Hence the loadings of the transitory component of the price dividend ratio are not only time-varying, but also a non-linear function of the steady state level of the price dividend ratio. These restrictions must be imposed exactly when estimating the model, a challenge that our score driven modeling approach can easily overcome.

The decomposition of dividend growth into the expected dividend growth plus an unexpected shock, $\varepsilon_{d,t+1}$, provides the first measurement equation:

$$\Delta d_{t+1} - \mathtt{E}_t(\Delta d_{t+1}) = \varepsilon_{d,t+1}. \tag{19}$$

The second measurement equation is (17), which we augment with a measurement error $\nu_t \sim \mathcal{N}(0, \sigma_\nu^2)$ to take into account the approximation error associated with the solution of the present value model. Moreover, we collect the three shocks in the vector $\varepsilon_t = (\varepsilon_{d,t}, \varepsilon_{g,t}, \varepsilon_{\mu,t})'$, respectively the shock to dividend growth, the shock to expected dividend growth and the shock to expected returns. It is easy to see that they also map into unexpected changes in returns, specifically:

$$r_{t+1} - \mathrm{E}_t(r_{t+1}) = -\rho_{t+1|t}b_{1,t+1|t}\varepsilon_{\mu,t+1} + \rho_{t+1|t}b_{2,t+1|t}\varepsilon_{g,t+1} + \varepsilon_{d,t+1}. \tag{20}$$

3.3.1 Slow-moving trends and time-varying risk

We assume that long-run expected dividend growth and long-run expected returns are martingales driven by the score of the conditional likelihood:

$$\overline{\mu}_{t+1|t} = \overline{\mu}_{t|t-1} + b_\mu s_{\mu,t} \tag{21}$$

$$\overline{g}_{t+1|t} = \overline{g}_{t|t-1} + b_g s_{g,t}, \tag{22}$$

where $s_{\mu,t}$ $s_{g,t}$ are the appropriate elements of the score vector. The steady states, $\overline{\mu}_{t|t-1}$ and $\overline{g}_{t|t-1}$, are therefore updated (through the scaled score) using information on dividend growth

and on the price dividend ratio, as highlighted by equation (11), as well as the present value restrictions embodied in equations (17)-(18), which imply that long-run returns and dividend growth can also change, through $\rho_{t|t-1}$, the sensitivity of the transitory components of price dividend ratio ($\mathtt{pd}_t - \overline{\mathtt{pd}}_{t|t-1}$) to expected returns and dividend growth.

Last, we assume that the vector of shock ε_t is normally distributed with time-varying covariances that are themselves driven by the score. Formally, $\varepsilon_t \sim \mathcal{N}(0, \Omega_t)$, and we decompose the covariance matrix $\Omega_t = D_t R_t D_t$, where R_t denotes the time-varying correlation matrix and $D_t = \text{diag}([\sigma_{d,t}, \sigma_{g,t}, \sigma_{\mu,t}])$ contains the standard deviations. Whereas the three shocks could be all correlated, not all elements of the covariance matrix can be separately identified. In Appendix G.1 we show that, in order to identify the model, one restriction needs to be imposed on the correlations among the innovations of the model. Below, we follow Binsbergen and Koijen (2010) and assume that the measurement error in dividend growth ($\varepsilon_{d,t}$) and the stochastic disturbance in expected dividend growth ($\varepsilon_{g,t}$) are uncorrelated, i.e. $\text{Corr}(\varepsilon_{g,t}, \varepsilon_{d,t}) = 0$.

Inspired by Joe (2006), we take advantage of the mapping between the correlations and the partial correlations. This allows us to constrain the partial correlations to lie in the unit circle, and yet guarantees a well-defined correlation matrix in every period. Modeling partial correlations has the additional advantage that we can easily impose that one of the correlation coefficients is always zero: ordering the innovations so that the restriction is placed on the first column of the correlation matrix, one has that $\varrho_{dg,t} = \pi_{dg,t}$ (where ϱ_{ij} and π_{ij} are, respectively, the generic ij element of the correlation and partial correlation matrix) and the identification restriction simply requires that $\pi_{dg,t} = 0, \forall t$. Therefore, through the Jacobian, the algorithm translates the score of the likelihood with respect to the correlation matrix into the appropriate updating of the unrestricted partial correlations. An alternative parametrization that guarantees a positive-definite correlation matrix uses *hyperspherical coordinates* (Creal et al., 2011; Buccheri et al., 2020). This implies constraining the partial correlations by the cosine function, hence it is a special case of our approach. For full details on modelling the correlation matrix

through partial correlations see Appendix G.2.

3.3.2 State space representation and estimation

Let $y_t = (\Delta d_t, \mathtt{pd}_t)'$ be the vector of observed variables and $\alpha_t = (1, \tilde{g}_t, \tilde{\mu}_t, \tilde{g}_{t-1}, \varepsilon_{d,t}, \varepsilon_{g,t}, \varepsilon_{\mu,t})'$ be the state vector. The measurement equations of the model are (17) and (19). The law of motion of the transitory components (15)-(16) and the definition of the innovations of the model constitute the transition equations.[7] The model can be set in the state space representation (1), with the following specification of the system matrices:[8]

$$Z_t = \begin{bmatrix} \bar{g}_t & 0 & 0 & 1 & 1 & 0 & 0 \\ \overline{\mathtt{pd}}_t & \frac{1}{1-\rho_t\phi_g} & -\frac{1}{1-\rho_t\phi_\mu} & 0 & 0 & 0 & 0 \end{bmatrix}, \tag{23}$$

$$T = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \phi_g & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \phi_\mu & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \tag{24}$$

$\epsilon_t = (0, \nu_t)'$ and $H = \mathtt{diag}([0, \sigma_\nu^2])$. The innovations to the state vector are $\eta_t = S_\eta \varepsilon_t$, where S_η is a selection matrix appropriately defined, and $Q_t = S_\eta \Omega_t S_\eta'$. The time-varying elements of the system matrices Z_t and Q_t are collected in the vector:

$$f_t = (\overline{\mu}_t, \overline{g}_t, \log\sigma_{d,t}, \log\sigma_{g,t}, \log\sigma_{\mu,t}, \mathrm{atanh}\pi_t^{d\mu}, \mathrm{atanh}\pi_t^{g\mu})' \tag{25}$$

where $\mathrm{atanh}(\cdot)$ denotes the inverse hyperbolic tangent so that the partial correlations $\pi_{ij,t} \in (-1, 1)$, $\forall t$. The law of motion of f_t is a restricted version of (4). Specifically, we assume A and B to be diagonal and we further restrict some of the elements of c and A in accordance with the random walk specification for the slow moving trends discussed in section 3.3.1. The model

[7] The measurement equation for the price dividend ratio and for dividend growth imply a measurement equation for returns (20), which is therefore redundant for the estimation of the model.

[8] To simplify the notation, in this section any generic score driven time-varying parameter, $x_{t|t-1}$, is simply

likelihood is computed using the KF, coupled with the updating algorithm for the score driven parameters as discussed in section 2 and parameters are easily estimated by maximising this likelihood. Confidence intervals are obtained following Blasques et al. (2016). Appendix G.3 provides details on the state space representation, the non-linear mapping between the score driven time-varying parameters and the system matrices, as well as the associated Jacobians that are required in the score driven filter.

4 Results

Parameter estimates. The estimation results are shown in Table 2. Expected returns have a root of 0.829, which implies a half life of 4.7 years. Expected dividend growth is less persistent, with an autoregressive root of 0.345, implying a half life of 1.7. These numbers are, as expected, lower than those estimated by Binsbergen and Koijen (2010) and Piatti and Trojani (2017), as part of the persistence is captured in our case by the shifts in the steady states. The volatility and correlation for the innovations of the model map into an average volatility of the shocks to returns of roughly 0.18 and correlation between the innovations to the return on stocks and expected returns in a range between -0.6 and -0.7, in line with the estimates in Carvalho et al. (2018). The variance of the measurement error of the price dividend ratio is negligible. The second column reports the autoregressive roots of the five time-varying parameters that are modelled as stationary processes, namely the three volatilities and the partial correlations. All of them are far from having a unit root, justifying our modeling choice. The third column presents the seven loadings on the score and the smoothing coefficient of the Hessian term. The long-run mean of expected returns has a loading of 0.151 on the likelihood score, three times as large as that of dividend growth (0.052). This is a signal that the low frequency component of expected returns exhibits much more time variation than that of dividend growth. In sum, expected returns are more predictable than expected dividend growth, they are characterized by a slowly changing steady state and by a very persistent transitory component.

Filtered steady states. Figure 1 shows returns and dividend growth together with the expected components. For both series, movements in the long-run steady states are completely dominated by very volatile transitory shocks, and seem relatively flat. As we shall see below, this is in fact not the case. Consistently with the parameter estimates in Table 2, expected dividend growth is more volatile and falls considerably in recessions. It is worth stressing that its fall in the 2008 recession is not exceptional by historical standards and is in fact significantly milder than those observed in 1929 and during WWII. Expected returns, on the other hand, are more persistent, with periods of high and low valuations spanning multiple recessions and expansions of the business cycle.

Figure 2 shows the estimated trend components. These are the central results of our empirical analysis. In the left panel we plot $\overline{\mu}_{t|t-1}$ and $\overline{g}_{t|t-1}$, which are estimated in our model as score driven martingales. It is clear that what seemed to be a flat line, i.e. the permanent component of expected returns, is in fact a downward sloping trend. It starts off at around 9 percent at the end of the 19^{th} century and fluctuates between 7 and 8 percent until the 50s. Thereafter it experiences two sharp falls: one between the 50s and the 70s, and one in the 90s, the former being interrupted by the stock market crash in 1973 and the latter by the burst of the dotcom bubble in 2000. At the end of the sample, the long-run expected return on stocks stands at around 4 percent, less than half of its initial value. The steady state of expected dividend growth has also fallen slightly over the sample, from around 2 percent to 1.3 percent at the turn of the century, to then rebound to 1.5 percent. This minor fall in long-run dividend growth is quantitatively consistent with the rise of alternative forms of payout (such as shares repurchases and issuance) which have become quantitatively more important since the 90s (see e.g., Boudoukh et al., 2007). The right hand side panel shows the implications of these results for the steady state of the price dividend ratio. Our estimate of $\overline{\mathtt{pd}}_t$ has an upward trend over the whole sample, with three large changes, one early in the sample, one after WWII and one in the Nineties. These are indeed the dates for which the Bai and Perron

(2003)'s test detects significant structural breaks. However, structural break tests are unable to offer a structural explanation for such an upward trend in valuations, nor can they provide a narrative behind these episodes. Our model reveals that these long waves of strong valuations have largely built on falling expected returns. The rally of stock prices in the 90s, in particular, was associated with a sharp fall in discount rates. These results formalize and provide evidence for the intuition put forward by Fama and French (2002). In their seminal paper on the equity premium they notice that, between 1950 and 2000, realized average stock returns in excess of the riskless interest rate have been substantially higher than the equity risk premium. Given the relative stationarity of dividend growth, they argue that realized capital gains must have come from a fall in discount rates. Our empirical analysis quantifies such shifts in long-run discount rates and identifies their timing within a formal model.

Dynamic price dividend ratio decomposition. Figure 3 decomposes the cyclical component of the price dividend ratio in a linear combination of the transitory components of expected returns and of expected dividend growth, based on equation (17). The results are quite striking. First, most of the cyclical variation in valuations is due to changes in expected returns. In particular, high stock prices in the 60s and in the 90s were due to a slow fall (of transitory nature) in the discount rate, which eventually reverted back towards its mean in the bear markets of the early 70s and early 2000s. This finding is consistent with the relatively higher persistence of $\tilde{\mu}_t$ compared to $\tilde{g}_t$ (AR root of 0.829 versus 0.345) as well as with the fact that changes in expected dividends affect both prices as well as actual dividends, with relative little impact on valuations (Cochrane, 2011; Campbell and Ammer, 1993). Second, the role of cash flows is episodic but non-negligible. A fall in cash flows expectations, for instance, contributed as much as discount rates to the stock market crash in 1929 and in WWII. It had an even larger role in explaining the 2008 crash.

denoted with x_t.

Term Structure of expected returns and dividend growth. Iterating forward equations (15)-(16) we can obtain an expression for the average expected returns and expected dividend growth over a given horizon n:

$$\mu_t^{(n)} = \frac{1}{n}\mathtt{E}_t\left[\sum_{j=1}^{n} r_{t+j}\right] = \overline{\mu}_{t+1|t} + \frac{1}{n}\frac{1-\phi_\mu^n}{1-\phi_\mu}\tilde{\mu}_t \tag{26}$$

$$g_t^{(n)} = \frac{1}{n}\mathtt{E}_t\left[\sum_{j=1}^{n} \Delta d_{t+j}\right] = \overline{g}_{t+1|t} + \frac{1}{n}\frac{1-\phi_g^n}{1-\phi_g}\tilde{g}_t. \tag{27}$$

In Figure 4, left hand side panel, we plot expected returns at the ten years horizon (blue dashed line) as well as the slope of the term structure of expected returns (red solid line), that is the difference between average expected returns ten and two years out. At low frequencies, expected returns inherit the properties of the shifting steady state $\overline{\mu}_{t|t-1}$: they fall from an average of around 7 percent in the first part of the sample to around 2 percent at the end of the sample. The slope, despite some short run movements, fluctuates around zero, suggesting that over this period the whole term structure of discount rates has shifted downwards. The level and the slope are clearly negatively correlated. This negative correlation is driven by a number of a recessionary episodes in which long term discount rates rise but the slope falls, i.e. discount rates rise more sharply at short than at long maturities. This evidence, consistent with findings in Gormsen (2020), is directly connected to the predictability of returns: high expected returns today predict low expected returns tomorrow.[9] The right hand side panel shows the level and the slope of the term structure of expected dividend growth. Long-term expectations are rather smooth (much more than those of expected returns) and have remained relatively stable at 2 percent,[10] consistently with the relative stability of $\overline{g}_{t|t-1}$. The slope, on the other hand, is countercyclical and quite volatile fluctuating between minus 10 to 10 percent, a range that is almost five times as wide as that in which the slope of expected returns moves.

[9]These findings stand partially in contrast with those in Piatti and Trojani (2017) who keep the steady state of expected returns constant. Keeping the steady state of expected returns constant has two consequences. First, it constraints long-term expectations to remain relatively stable, the more the longer the horizon that one considers. This is hard to reconcile with the secular upward trend in the price dividend ratio. Second, in a model with constant steady states the terminal point of the term structure is fixed. This implies that a model with constant steady states ends up overstating the predictability of returns, especially at very long horizons.

[10]Few rare exceptions when pessimistic expectations on cash flows persisted for more than 10 years include

Indeed, as shown in table 2, the volatility of shocks to $\tilde{g}_t$ is on average almost four times as high as that of $\tilde{\mu}_t$ (0.083 versus 0.024) and this is reflected in the volatility of expected annual cash flows. This slope has some notable peaks, corresponding to sharp recessions, such as the post WWI recession, the Great Depression, the 1973 recession and more recently the Great Financial Crisis. Therefore, the term structure of dividend growth highlights that short-term cash flows expectations fall substantially during recessions, but are anticipated to mean revert relatively quickly. The countercyclicality of the slope of the term structure of dividend growth expectations is in line with the finding of Binsbergen et al. (2013) who document similar properties from dividend derivatives for the post 2003 sample. Appendix H zooms in on some large stock market corrections and shows that discount rates shocks contributed greatly to the severity of the recessions in 1929 and 2008.

5 Expected excess returns and the equilibrium real rate

The central result of our empirical analysis is that the long-run expected return on equity has fallen over time. In this section, we elaborate further on this finding by splitting long-run expected returns in a riskless component and a risk premium. A wealth of research has shown that the natural rate has fallen dramatically not only in the US but also in a set of advanced countries (Holston et al., 2017, and Del Negro et al., 2019).[11]

We follow Cochrane (2008) and rewrite the present value decomposition of the price dividend ratio in equation (12) in terms of excess returns, by simply subtracting a measure of the safe real rate of interest (let us call it r^f) from both expected discounted returns and expected discounted dividend growth.[12] Our baseline model can then be re-estimated using as observable variables for the measurement equations the price dividend ratio and the difference between dividend

the height of the Great Depression, the oil price shock of 1973 and the Great Financial Crisis.

[11] Three main competing explanations have been put forward to rationalize the fall in equilibrium real rate: a permanent decline in the rate of growth of the economy (secular stagnation), an increase in desired savings due to aging population (saving glut) and a rise in economic risk (or a fall in its tolerance) that has raised the liquidity and safety premium of safe assets like US Treasuries, see Del Negro et al. (2017) for an extensive discussion.

[12] Returns and dividend growth enter the decomposition with opposite signs, so that r_t^f cancels out leaving unaffected the price dividend ratio.

growth and r_t^f.[13] Besides this simple modification, the model is essentially unchanged, apart from the fact that, following the suggestion in Campbell and Thompson (2008), we impose the restriction that the long run equity premium needs to be always positive. This alternative model provides us with three time series, $\overline{\texttt{pd}}_{t|t-1}$, $\overline{g}^{ex}_{t|t-1}$, $\overline{\mu}^{ex}_{t|t-1}$ that measure, respectively, the time-varying equilibrium price to dividend ratio for this alternative model specification, the long-run expected *excess* dividend growth (i.e. the long-run expected dividend growth minus the equilibrium riskless real rate) and the steady state equity risk premium or expected *excess* return. A comforting result of this exercise is that the estimate of the equilibrium price to dividend ratio that we retrieve using this alternative specification is very close to the one obtained in the baseline model in Section 3 (see Appendix I for a comparison). Yet, we now have a measure of the equilibrium excess returns $\overline{\mu}^{ex}_{t|t-1}$. Subtracting this from $\overline{\mu}_{t|t-1}$ we can then obtain a measure of the riskless long-run real rate, $\overline{r}_{t|t-1} = \overline{\mu}_{t|t-1} - \overline{\mu}^{ex}_{t|t-1}$.[14]

The left panel of figure 5 shows that the long-run expected excess return, $\overline{\mu}^{ex}_{t|t-1}$, has only slightly fallen, from around 4 percent at the beginning of the 20th century, to reach a minimum of 3 percent in 2000 and to rebound thereafter to around 3.5 percent. These numbers are in the ballpark of the estimates provided by Avdis and Wachter (2017) who report an annual equity premium of 3.86 per cent in the post WWII period and of around 4.5 in a longer sample (1927-2011) as well as with the estimate of the equity risk premium (between 2.6 percent and 4.3 percent) for the period 1950-2000 in Fama and French (2002). These results are also in line with Greenwald et al. (2019) who find that the fall in the equity premium played a limited role in the overall fall in the equilibrium rate of return on stocks after the 80s. We add to their analysis by showing that the trend in valuations is almost entirely driven by the fall in the long-run riskless rate.

The implied long-run natural rate of interest $\overline{r}_t$, our measure of r-star, has remained stable (between 3 and 4 percent) for about a century until 1960. It has then fallen by 1 full percentage

[13] A long series of the risk free rate is taken from Amit Goyal's website http://www.hec.unil.ch/agoyal/.

[14] Alternatively one could compute the implied riskless real rate as $\overline{r}_{t|t-1} = \overline{g}_{t|t-1} - \overline{g}^{ex}_{t|t-1}$. The two offer a

point in the 60s, and by 2 percentage points in the 90s. Its terminal estimate is only slightly above zero. The sudden fall in the equilibrium rate of interest after the 60s is also documented, with a completely different method, by Del Negro et al. (2019). The behaviour of our measure of r-star is qualitatively in line with that of Laubach and Williams (2003) and Holston et al. (2017), also plotted in the chart.[15]

Concluding, our results confirm that the natural rate of interest has fallen substantially over the last three decades, but add to previous studies by drawing a clear link between this secular trend and the increase in stock valuations (Caballero et al., 2017).

6 Conclusions

State space models with time-varying parameters can help us better understanding the co-movement in macro financial aggregates, in a world in which returns, long-run growth and asset valuations appear to have undergone long-lasting shifts. These models present computational as well as analytical challenges. In this paper we propose a method for analyzing state space models with time-varying system matrices where the parameters are driven by the score of the conditional likelihood. We derive a new set of recursions that, running in parallel with the KF, update at each point in time both the vector of TVP and the latent states. A unique feature of our method is that it can easily incorporate in the estimation a broad class of parameter constraints. These are taken into account directly in the estimation algorithm through a Jacobian function, without the need for rejection sampling or complicated filtering techniques. A Monte Carlo analysis provides support for the proposed method. Recent empirical analyses that use our method also testify its usefulness (Delle Monache et al., 2016; Buccheri et al., 2020; Gorgi et al., 2019).

We have then used this framework to fill a gap in the literature that studies the relationship

similar picture, as we show in Appendix I. Figure 5 reports the average between these two alternative measures.

[15] Quantitatively, our measures are lower in the 60s and 70s. In defence of our results we point out that, given our estimate of the expected return on equity, the level of r-star implied by the Laubach and Williams (2003) model prior to the 90s would require a level of the equity risk premium substantially lower than what is typically found in the literature.

between the price dividend ratio, the expected return on stocks and expected dividend growth in present value models. Our estimates reveal that the secular upward trend in the price dividend ratio, so far unexplained in the literature, is associated with a persistent decline (from 9 to 4 percent) of the long-run expected return on stocks. A decomposition into a riskless component and a risk premium further reveals that most of this decline (four percentage points) is accounted for by the riskless component (i.e. the natural rate of interest, r-star), that is virtually zero at the end of our sample. The long-term equity risk premium has instead remained relatively stable over the past 150 years.

In sum, our work substantiates, and provides quantitative evidence for, the argument that in a world of permanently low rates valuation rates will be persistently higher than in the past. Falling productivity growth, scarcity of safe assets and population aging depress r-star and risk-free rates, inducing investors to reach for yield and permanently lifting stock valuations (Campbell and Sigalov, 2020).

Table 1: Structural Break Tests

	r_t	Δd_t	pd_t
$SupF_T(k)$			
k=1	0.954	0.458	48.836***
	[1921]	[1996]	[1991]
k=2	1.078	0.365	49.804***
	[1921; 1962]	[1918; 1944]	[1954; 1995]
k=3	1.619	0.488	35.862***
	[1921; 1960; 1982]	[1896; 1918; 1944]	[1913; 1954; 1995]
$SupF_T(k+1\|k)$			
k=2	1.676	0.233	14.896***
k=3	2.229	0.846	2.227
$U_d$max	1.619	0.488	49.804***
Nyblom Test			
μ	0.023	0.057	7.401***
σ^2	0.319	0.693**	2.608***
Joint L_c	0.348	0.750	7.873***

Note. $SupF_T(k)$ denotes the Bai and Perron (2003) test where the null hypothesis of no breaks is tested against the alternatives of $k = 1, 2$, or 3 breaks. Dates in square brackets are the most likely break date(s) for each of the specifications. $SupF_T(k+1|k)$ denotes the test of k breaks against the alternative of $k+1$ breaks. The U_d max statistics is the result of testing the null hypothesis of absence of breaks against the alternative of an unknown number of breaks. The bottom panel reports the Nyblom (1989) test that in the model: $y_t = \mu + \sigma\epsilon_t$, either μ or σ are constant against the alternative that they evolve as random walks. The symbols */**/*** indicate significance at the 10/5/1% level.

Table 2: Model Estimation Results

ϕ_μ	0.829			b_μ	0.151
	[0.010]				[0.010]
ϕ_g	0.345			b_g	0.052
	[0.010]				[0.010]
$\bar{\sigma}_d$	0.075	a_{σ_d}	0.881	b_{σ_d}	0.015
	[0.074; 0.083]		[0.026]		[0.002]
$\bar{\sigma}_g$	0.083	a_{σ_g}	0.899	b_{σ_g}	0.012
	[0.082; 0.127]		[0.050]		[0.004]
$\bar{\sigma}_\mu$	0.024	a_{σ_μ}	0.902	b_{σ_μ}	0.014
	[0.023; 0.040]		[0.040]		[0.004]
$\bar{\rho}_{d,\mu}$	0.339	$a_{\pi_{d,\mu}}$	0.820	$b_{\pi_{d,\mu}}$	0.013
	[0.280; 0.371]		[0.068]		[0.003]
$\bar{\rho}_{g,\mu}$	-0.232	$a_{\pi_{g,\mu}}$	0.844	$b_{\pi_{g,\mu}}$	0.017
	[-0.253; -0.130]		[0.040]		[0.003]
σ^2_ν	0.001			κ_h	0.020
	[0.0005]				[0.001]
Log Lik.	311.567				

Note. First column: autoregressive coefficients of expected returns and expected dividend growth (ϕ_μ and ϕ_g) and average (over the whole sample) estimates of the volatilities ($\bar{\sigma}_d$, $\bar{\sigma}_g$ and $\bar{\sigma}_\mu$) and correlations ($\bar{\rho}_{d,\mu}$ and $\bar{\rho}_{g,\mu}$) that form the matrix Q_t. σ^2_ν is the volatility of the measurement error for the price dividend ratio. The second and third columns show the estimates of the coefficients that enter the law of motion of the score driven time-varying processes (4) where A and B are diagonal matrices, and the smoothing coefficient applied to the Hessian term (κ_h). For each coefficient we report in square brackets the associated standard error. For the average volatilities and correlations in the first column we report the 68% confidence interval from 1000 simulations of the model (calculated as in Blasques et al., 2016).

Figure 1: Expected Returns and Expected Dividend Growth

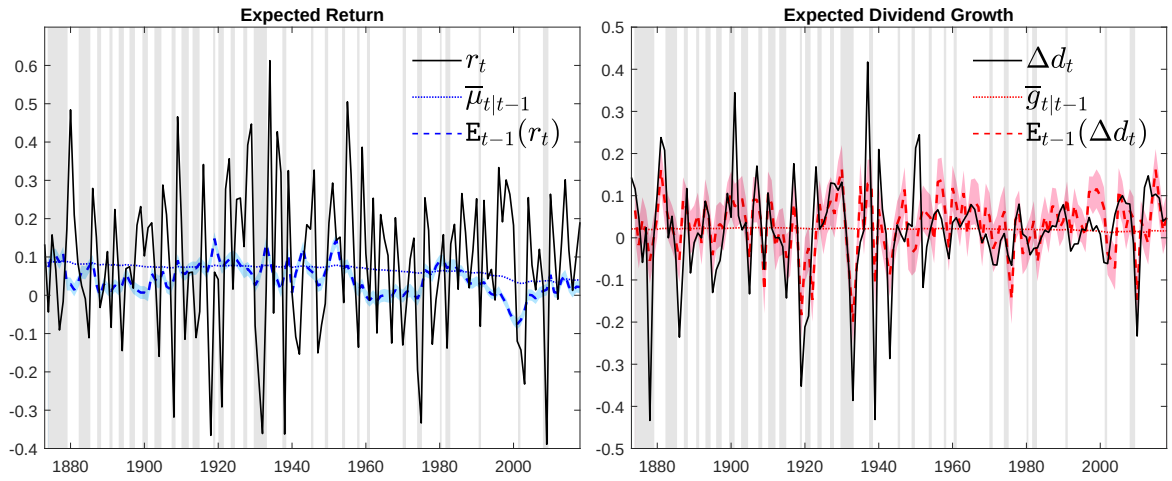

Note. The panel on the left of Figure 1 shows the return on stocks (r_t), expected returns ($\mathrm{E}_{t-1}(r_t)$), and the long-run component of expected returns ($\bar{\mu}_{t|t-1}$). The panel on the right reports dividend growth (Δd_t), together with expected dividend growth ($\mathrm{E}_{t-1}(\Delta d_t)$) and the long-run component of expected dividend growth ($\bar{g}_{t|t-1}$). In both panels the colored bands denote the 68% confidence interval. Vertical shadows indicate recessions as identified by the National Bureau of Economic Research (NBER).

Figure 2: Long-run Expected Returns, Expected Dividend Growth and Price-Dividend Ratio

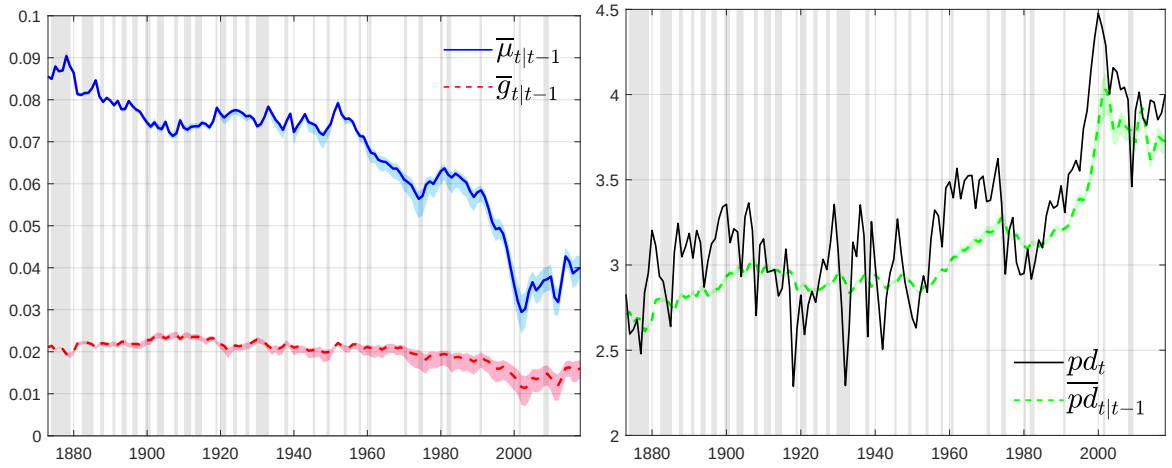

Note. The panel on the left of Figure 2 shows long-run expected returns ($\overline{\mu}_{t|t-1}$, continuous blue line) and long-run expected dividend growth ($\overline{g}_{t|t-1}$, broken red line). The implied long-run price dividend ratio, obtained from $\overline{\mu}_{t|t-1}$ and $\overline{g}_{t|t-1}$ on the basis of equation (11), is shown in the panel on the right (green broken line) together with the actual level of the (log) price dividend ratio (black solid line). Bands around the estimates correspond to the 68% confidence interval obtained through simulation, as discussed in Blasques et al. (2016). Vertical shadows indicate recessions as identified by the National Bureau of Economic Research (NBER).

Figure 3: Price Dividend Ratio Decomposition

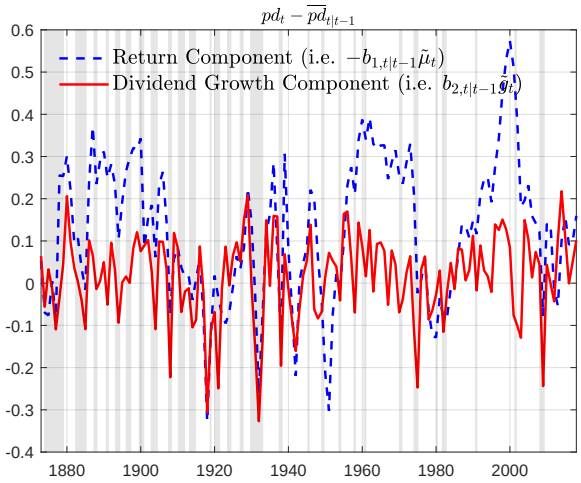

Note. Figure 3 reports the decomposition of the transitory component of the price dividend ratio into the contribution of the transitory components of, respectively, expected returns and expected dividend growth according to equation (17). Vertical shadows indicate recessions as identified by the National Bureau of Economic Research (NBER).

Figure 4: Term Structure of Expected Returns and Expected Dividend Growth

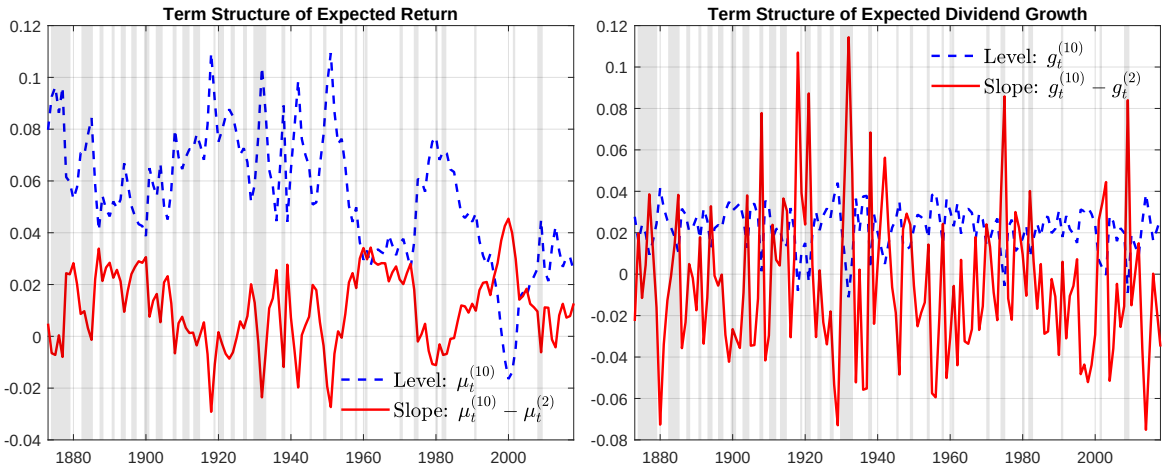

Note. Figure 4 reports with a broken blue line the level (i.e. 10 years) and with a continuous red line the slope (i.e 10 year minus 2 years) of conditional expected returns (panel on the left) and of conditional expected dividend growth (panel on the right). Vertical shadows indicate recessions as identified by the National Bureau of Economic Research (NBER).

Figure 5: Long-Run Excess Return on Stocks and Long-Run Riskless Rate

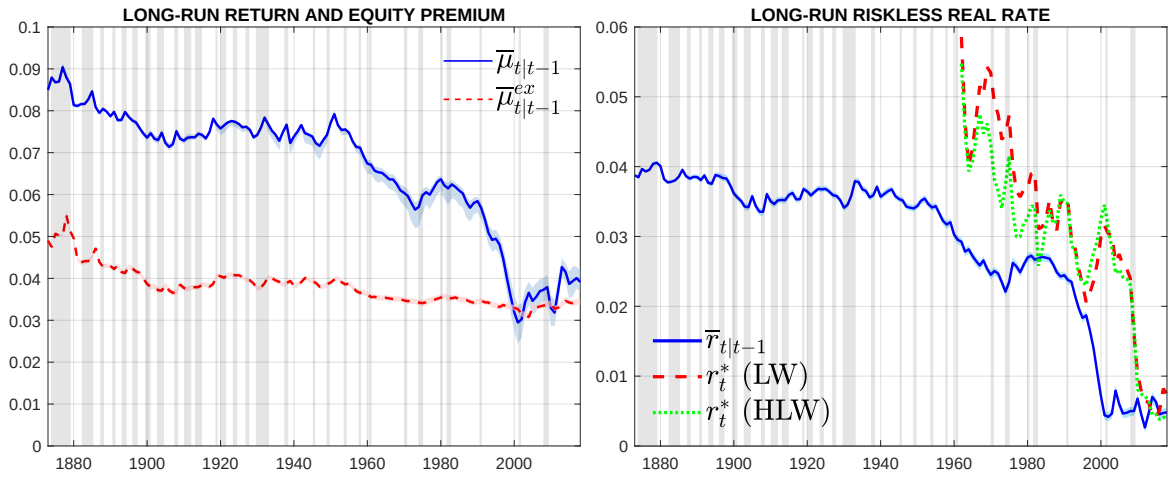

Note. The panel on the left of Figure 5 reports in blue the long-run expected return on stocks ($\overline{\mu}_{t|t-1}$) estimated using the baseline model specification described in Section 3 and in red the long run equity premium ($\overline{\mu}^{ex}_{t|t-1}$) estimated as described in Section 5. The panel on the right shows (blue line) the long-run real riskless rate constructed as $\overline{r}_{t|t-1} = 0.5(\overline{\mu}_{t|t-1} - \overline{\mu}^{ex}_{t|t-1}) + 0.5(\overline{g}_{t|t-1} - \overline{g}^{ex}_{t|t-1})$. The panel on the right also shows two alternative measures of r-star for the US, as estimated in Laubach and Williams (2003) denoted with "LW" and in Holston et al. (2017) denoted with "HLW". Vertical shadows indicate recessions as identified by the National Bureau of Economic Research (NBER).

References

Avdis, E. and Wachter, J. A. (2017). Maximum likelihood estimation of the equity premium. *Journal of Financial Economics*, 125(3):589–609.

Bai, J. and Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1):1–22.

Binsbergen, J. H. V., Hueskes, W., Koijen, R., and Vrugt, E. (2013). Equity yields. *Journal of Financial Economics*, 110(3):503–519.

Binsbergen, J. H. V. and Koijen, R. S. J. (2010). Predictive Regressions: A Present Value Approach. *Journal of Finance*, 65(4):1439–1471.

Blasques, F., Koopman, S. J., Lasak, K., and Lucas, A. (2016). In-sample confidence bands and out-of-sample forecast bands for time-varying parameters in observation-driven models. *International Journal of Forecasting*, 32(3):875–887.

Boudoukh, J., Michaely, R., Richardson, M., and Roberts, M. R. (2007). On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing. *Journal of Finance*, 62(2):877–915.

Buccheri, G., Bormetti, G., Corsi, F., and Lillo, F. (2020). A Score-Driven Conditional Correlation Model for Noisy and Asynchronous Data: an Application to High-Frequency Covariance Dynamics. *Journal of Business Economic Statistics*, forthcoming.

Caballero, R. J., Farhi, E., and Gourinchas, P.-O. (2017). Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares. *American Economic Review*, 107(5):614–620.

Campbell, J. Y. and Ammer, J. (1993). What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns. *Journal of Finance*, 48(1):3–37.

Campbell, J. Y. and Shiller, R. J. (1988). The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies*, 1(3):195–228.

Campbell, J. Y. and Sigalov, R. (2020). Portfolio choice with sustainable spending: A model of reaching for yield. unpublished manuscript.

Campbell, J. Y. and Thompson, S. B. (2008). Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average? *Review of Financial Studies*, 21(4):1509–1531.

Carvalho, C. M., Lopes, H. F., and McCulloch, R. E. (2018). On the long-run volatility of stocks. *Journal of the American Statistical Association*, 113(523):1050–1069.

Chen, H., Kumar, P., and van Schuppen, J. (1989). On Kalman filtering for conditionally Gaussian systems with random matrices. *Systems & Control Letters*, 13(5):397–404.

Cochrane, J. H. (2008). The Dog That Did Not Bark: A Defense of Return Predictability. *Review of Financial Studies*, 21(4):1533–1575.

Cochrane, J. H. (2011). Discount Rates. *Journal of Finance*, 66(4):1047–1108.

Cogley, T. and Sargent, T. J. (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics*, 8(2):262–302.

Creal, D., Koopman, S. J., and Lucas, A. (2008). A General Framework for Observation Driven Time-Varying Parameter Models. Tinbergen Institute Discussion Papers 08-108/4, Tinbergen Institute.

Creal, D., Koopman, S. J., and Lucas, A. (2011). A Dynamic Multivariate Heavy-Tailed Model for Time-Varying Volatilities and Correlations. *Journal of Business & Economic Statistics*, 29(4):552–563.

Creal, D., Koopman, S. J., and Lucas, A. (2013). Generalized Autoregressive Score Models With Applications. *Journal of Applied Econometrics*, 28(5):777–795.

Del Negro, M., Giannone, D., Giannoni, M. P., and Tambalotti, A. (2017). Safety, Liquidity, and the Natural Rate of Interest. *Brookings Papers on Economic Activity*, 48(1 (Spring):235–316.

Del Negro, M., Giannone, D., Giannoni, M. P., and Tambalotti, A. (2019). Global trends in interest rates. *Journal of International Economics*, 118(C):248–262.

Delle Monache, D., Petrella, I., and Venditti, F. (2016). Common Faith or Parting Ways? A Time Varying Parameters Factor Analysis of Euro-Area Inflation. In *Dynamic Factor Models*, volume 35 of *Advances in Econometrics*, pages 539–565. Emerald Publishing Ltd.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*. Oxford University Press.

Eickmeier, S., Lemke, W., and Marcellino, M. (2015). Classical time varying factor-augmented vector auto-regressive models. *Journal of the Royal Statistical Society Series A*, 178:493–533.

Fama, E. F. and French, K. R. (2002). The Equity Premium. *Journal of Finance*, 57:637–659.

Farhi, E. and Gourio, F. (2018). Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia. *Brookings Papers on Economic Activity*, 49(2 (Fall)):147–250.

Gorgi, P., Koopman, S. J. S., and Schaumburg, J. (2019). Time-Varying Vector Autoregressive Models with Structural Dynamic Factors, mimeo. Technical report.

Gormsen, N. J. (2020). Time Variation of the Equity Term Structure. *Journal of Finance*, forthcoming.

Greenwald, D. L., Lettau, M., and Ludvigson, S. C. (2019). How the Wealth Was Won: Factors Shares as Market Fundamentals. NBER Working Papers 25769, National Bureau of Economic Research, Inc.

Harvey, A., Ruiz, E., and Sentana, E. (1992). Unobserved component time series models with Arch disturbances. *Journal of Econometrics*, 52(1-2):129–157.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails*. Cambridge University Press.

Holston, K., Laubach, T., and Williams, J. C. (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, 108(S1):59–75.

Joe, H. (2006). Generating random correlation matrices based on partial correlations. *Journal of Multivariate Analysis*, 97(10):2177 – 2189.

Koop, G. and Korobilis, D. (2013). Large time-varying parameter VARs. *Journal of Econometrics*, 177(2):185–198.

Koopman, S. J., Mallee, M. I. P., and Van der Wel, M. (2010). Analyzing the Term Structure of Interest Rates Using the Dynamic Nelson-Siegel Model With Time-Varying Parameters. *Journal of Business & Economic Statistics*, 28(3):329–343.

Laubach, T. and Williams, J. C. (2003). Measuring the Natural Rate of Interest. *The Review of Economics and Statistics*, 85(4):1063–1070.

Lettau, M. and Nieuwerburgh, S. V. (2008). Reconciling the Return Predictability Evidence. *Review of Financial Studies*, 21(4):1607–1652.

Nyblom, J. (1989). Testing for the constancy of parameters over time. *Journal of the American Statistical Association*, 84(405):223–230.

Paye, B. S. and Timmermann, A. (2006). Instability of return prediction models. *Journal of Empirical Finance*, 13(3):274–315.

Piatti, I. and Trojani, F. (2017). Predictable Risks and Predictive Regression in Present-Value Models. Working paper, Said Business School.

Stock, J. H. and Watson, M. W. (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33.

Theil, H. and Goldberger, A. S. (1961). On pure and mixed statistical estimation in economics. *International Economic Review*, 2(1):65–78.

Price Dividend Ratio and Long-Run Stock Returns: a Score Driven State Space Model

Supplementary Material

Davide Delle Monache
Bank of Italy
davide.dellemonache@bancaditalia.it

Ivan Petrella
University of Warwick and CEPR
Ivan.Petrella@wbs.ac.uk

Fabrizio Venditti
European Central Bank
fabrizio.venditti@ecb.int

A Proofs

We follow the notation and the results in Abadir and Magnus (2005, ch. 13). Given a $N \times m$ matrix X, $\text{vec}(X)$ is the vector obtained by stacking the columns of X one underneath the other. The $Nm \times Nm$ commutation matrix $C_{N,m}$ is such that $C_{N,m}\text{vec}(X) = \text{vec}(X')$. For $N = m$ the $m^2 \times m^2$ commutation matrix is denoted by C_m. The identity matrix of order k is denoted by I_k, and '$\otimes$' is the Kronecker product. Given a square matrix U, the symmetrizer matrix is $N_n = \frac{1}{2}(I_{n^2} + C_n)$ and $N_n\text{vec}(U) = \text{vec}\left[\frac{1}{2}(U + U')\right]$.

A.1 Gradient and information matrix

The gradient vector is

$$\begin{aligned}
\nabla_t = \left(\frac{\partial \ell_t}{\partial f_t'}\right)' &= -\frac{1}{2}\left[\frac{\partial \log|F_t|}{\partial f_t'} + \frac{\partial v_t' F_t^{-1} v_t}{\partial f_t'}\right]' \\
&= -\frac{1}{2}\left[\frac{1}{|F_t|}\frac{\partial |F_t|}{\partial \text{vec}(F_t)'}\frac{\partial \text{vec}(F_t)}{\partial f_t'} + \frac{\partial v_t' F_t^{-1} v_t}{\partial v_t}\frac{\partial v_t}{\partial f_t'} + \frac{\partial v_t' F_t^{-1} v_t}{\partial \text{vec}(F_t^{-1})'}\frac{\partial \text{vec}(F_t^{-1})}{\partial \text{vec}(F_t)'}\frac{\partial \text{vec}(F_t)}{\partial f_t'}\right]' \\
&= -\frac{1}{2}\left[\text{vec}(F_t^{-1})'\dot{F}_t + 2v_t' F_t^{-1}\dot{V}_t - (v_t' \otimes v_t')(F_t^{-1} \otimes F_t^{-1})\dot{F}_t\right]' \\
&= \frac{1}{2}\left[\dot{F}_t'(F_t^{-1} \otimes F_t^{-1})(v_t \otimes v_t) - \dot{F}_t'(F_t^{-1} \otimes F_t^{-1})\text{vec}(F_t) - 2\dot{V}_t' F_t^{-1} v_t\right] \\
&= \frac{1}{2}\left[\dot{F}_t'(F_t^{-1} \otimes F_t^{-1})\text{vec}(v_t v_t' - F_t) - 2\dot{V}_t' F_t^{-1} v_t\right]. \qquad \text{(A.1)}
\end{aligned}$$

We now compute the information matrix as the expected value of the Hessian matrix

$$\mathcal{I}_t = -\text{E}_t\left[\frac{\partial^2 \ell_t}{\partial f_t \partial f_t'}\right]. \qquad \text{(A.2)}$$

From the third line of (A.1) we can write that

$$\begin{aligned}
\nabla_t &= -\frac{1}{2}\left[\dot{F}_t'[\text{vec}(F_t^{-1}) - \text{vec}(F_t^{-1} v_t v_t' F_t^{-1})] + 2\dot{V}_t' F_t^{-1} v_t\right] \\
&= -\frac{1}{2}\left[\dot{F}_t'\text{vec}(F_t^{-1} - F_t^{-1} v_t v_t' F_t^{-1}) + 2\dot{V}_t' F_t^{-1} v_t\right] \\
&= -\frac{1}{2}\left[\dot{F}_t'\text{vec}[F_t^{-1}(I_N - v_t v_t' F_t^{-1})] + 2\dot{V}_t' F_t^{-1} v_t\right] \\
&= -\frac{1}{2}\left[\dot{F}_t'(I_N \otimes F_t^{-1})\text{vec}(I_N - v_t v_t' F_t^{-1}) + 2\dot{V}_t' F_t^{-1} v_t\right]. \qquad \text{(A.3)}
\end{aligned}$$

The negative Hessian is equal to:

$$
\begin{aligned}
-\frac{\partial^2 \ell_t}{\partial f_t \partial f_t'} &= \frac{1}{2}\frac{\partial \Phi_t}{\partial f_t'} + \frac{\partial \text{vec}(\dot{V}_t' F_t^{-1} v_t)}{\partial \text{vec}(\dot{V}_t')'}\frac{\partial \text{vec}(\dot{V}_t')}{\partial f_t'} \qquad (A.4)\\
&\quad + \frac{\partial \text{vec}(\dot{V}_t' F_t^{-1} v_t)}{\partial \text{vec}(F_t^{-1})}\frac{\partial \text{vec}(F_t^{-1})}{\partial \text{vec}(F_t)'}\frac{\partial \text{vec}(F_t)}{\partial f_t'} + \frac{\partial \text{vec}(\dot{V}_t' F_t^{-1} v_t)}{\partial v_t}\frac{\partial v_t}{\partial f_t'}\\
&= \frac{1}{2}\frac{\partial \Phi_t}{\partial f_t'} + (v_t' F_t^{-1} \otimes I_N)\ddot{V}_t - (v_t \otimes \dot{V}_t)'(F_t^{-1} \otimes F_t^{-1})\dot{F}_t + \dot{V}_t' F_t^{-1}\dot{V}_t, \qquad (A.5)
\end{aligned}
$$

where $\Phi_t = \dot{F}_t'(I_N \otimes F_t^{-1})\text{vec}(I_N - v_t v_t' F_t^{-1})$. Let us now compute the following Jacobian:

$$
\begin{aligned}
\frac{\partial \Phi_t}{\partial f_t'} &= \frac{\partial \Phi_t}{\partial \text{vec}(\dot{F}_t')'}\frac{\partial \text{vec}(\dot{F}_t')}{\partial f_t'} + \frac{\partial \Phi_t}{\partial \text{vec}(I_N \otimes F_t^{-1})'}\frac{\partial \text{vec}(I_N \otimes F_t^{-1})}{\partial \text{vec}(F_t^{-1})'}\frac{\partial \text{vec}(F_t^{-1})}{\partial \text{vec}(F_t)'}\frac{\partial \text{vec}(F_t)}{\partial f_t'}\\
&\quad - \frac{\partial \Phi_t}{\partial \text{vec}(I_N - v_t v_t' F_t^{-1})'}\frac{\partial \text{vec}(v_t v_t' F_t^{-1})'}{\partial f_t'}\\
&= \left[\text{vec}(I_N - v_t v_t' F_t^{-1})'(I \otimes F_t^{-1}) \otimes I\right]\ddot{F}_t - \left[\text{vec}(I_N - v_t v_t' F_t^{-1})' \otimes \dot{F}_t'\right](F_t^{-1} \otimes F_t^{-1})\dot{F}_t\\
&\quad - \dot{F}_t'(I \otimes F_t^{-1})\left[\frac{\partial \text{vec}(v_t v_t' F_t^{-1})}{\partial \text{vec}(v_t v_t')'}\frac{\partial \text{vec}(v_t v_t')}{\partial v_t'}\frac{\partial v_t}{\partial f_t'} + \frac{\partial \text{vec}(v_t v_t' F_t^{-1})}{\partial \text{vec}(F_t^{-1})'}\frac{\partial \text{vec}(F_t^{-1})}{\partial \text{vec}(F_t)'}\frac{\partial \text{vec}(F_t)}{\partial f_t'}\right]\\
&= [\text{vec}(I_N - v_t v_t' F_t^{-1})'(I_N \otimes F_t^{-1}) \otimes I_N]\ddot{F}_t - \left[\text{vec}(I_N - v_t v_t' F_t^{-1})' \otimes \dot{F}_t'\right](F_t^{-1} \otimes F_t^{-1})\dot{F}_t\\
&\quad - \dot{F}_t'(F_t^{-1} \otimes F_t^{-1})(v_t \otimes I_N + I_N \otimes v_t)\dot{V}_t + \dot{F}_t'(F_t^{-1} \otimes F_t^{-1} v_t v_t' F_t^{-1})\dot{F}_t. \qquad (A.6)
\end{aligned}
$$

where $\ddot{F}_t = \frac{\partial \text{vec}(\dot{F}_t')}{\partial f_t'}$. Putting together (A.5) and (A.6) we obtain the following expression:

$$
\begin{aligned}
-\frac{\partial^2 \ell_t}{\partial f_t f_t'} &= \frac{1}{2}[\text{vec}(I_N - v_t v_t' F_t^{-1})'(I_N \otimes F_t^{-1}) \otimes I_N]\ddot{F}_t - \frac{1}{2}\left[\text{vec}(I_N - v_t v_t' F_t^{-1})' \otimes \dot{F}_t'\right](F_t^{-1} \otimes F_t^{-1})\dot{F}_t\\
&\quad - \frac{1}{2}\dot{F}_t'(F_t^{-1} \otimes F_t^{-1})(v_t \otimes I_N + I_N \otimes v_t)\dot{V}_t + \frac{1}{2}\dot{F}_t'\left(F_t^{-1} \otimes F_t^{-1} v_t v_t' F_t^{-1}\right)\dot{F}_t\\
&\quad + (v_t' F_t^{-1} \otimes I_N)\ddot{V}_t - (v_t \otimes \dot{V}_t)'(F_t^{-1} \otimes F_t^{-1})\dot{F}_t + \dot{V}_t' F_t^{-1}\dot{V}_t, \qquad (A.7)
\end{aligned}
$$

where $\ddot{V}_t = \frac{\partial \text{vec}(\dot{V}_t')}{\partial f_t'}$. Following Harvey (1989, p.141), taking the conditional expectation of (A.7) the fourth and the seventh term in (A.7) are the only nonzero elements and the information matrix is equal to

$$
\mathcal{I}_t = \frac{1}{2}\dot{F}_t'\left(F_t^{-1} \otimes F_t^{-1}\right)\dot{F}_t + \dot{V}_t' F_t^{-1}\dot{V}_t. \qquad (A.8)
$$

A.2 Jacobians of the Kalman filter

Let us write the KF recursions (3) in terms of the predictive filter:

$$
\begin{aligned}
&v_t = y_t - Z_t a_t, && F_t = Z_t P_t Z_t' + H_t,\\
&K_t = T_{t+1} P_t Z_t' F_t^{-1}, && L_t = T_{t+1} - K_t Z_t, \qquad (A.9)\\
&a_{t+1} = T_{t+1} a_t + K_t v_t && P_{t+1} = T_{t+1} P_t L_t' + Q_{t+1}, \quad t = 1, ..., n.
\end{aligned}
$$

Given the model specific Jacobian matrices:

$$\dot{Z}_t = \frac{\partial \text{vec}(Z_t)}{\partial f_t'}, \quad \dot{H}_t = \frac{\partial \text{vec}(H_t)}{\partial f_t'}, \quad \dot{T}_t = \frac{\partial \text{vec}(T_t)}{\partial f_t'}, \quad \dot{Q}_t = \frac{\partial \text{vec}(Q_t)}{\partial f_t'},$$

we compute the following Jacobian matrices:

$$\dot{V}_t = \frac{\partial v_t}{\partial f_t'} = \left[\frac{\partial v_t}{\partial \text{vec}(Z_t)'}\frac{\partial \text{vec}(Z_t)}{\partial f_t'} + \frac{\partial v_t}{\partial a_t'}\frac{\partial a_t}{\partial f_t'}\right] = -[(a_t' \otimes I_N)\dot{Z}_t + Z_t\dot{A}_t]. \tag{A.10}$$

$$\begin{aligned}\dot{F}_t = \frac{\partial \text{vec}(F_t)}{\partial f_t'} &= \frac{\partial \text{vec}(F_t)}{\partial \text{vec}(Z_t)'}\frac{\partial \text{vec}(Z_t)}{\partial f_t'} + \frac{\partial \text{vec}(F_t)}{\partial \text{vec}(P_t)'}\frac{\partial \text{vec}(P_t)}{\partial f_t'} + \dot{H}_t \\ &= 2N_N(Z_tP_t \otimes I_N)\dot{Z}_t + (Z_t \otimes Z_t)\dot{P}_t + \dot{H}_t. \end{aligned} \tag{A.11}$$

$$\dot{K}_t = \frac{\partial \text{vec}(K_t)}{\partial f_{t+1}'} = \frac{\partial \text{vec}(K_t)}{\partial \text{vec}(T_{t+1})'}\frac{\partial \text{vec}(T_{t+1})}{\partial f_{t+1}'} = (F_t^{-1}Z_tP_t \otimes I_m)\dot{T}_{t+1}. \tag{A.12}$$

$$\begin{aligned}\dot{L}_t = \frac{\partial \text{vec}(L_t)}{\partial f_{t+1}'} &= \frac{\partial \text{vec}(L_t)}{\partial \text{vec}(T_{t+1})'}\frac{\partial \text{vec}(T_{t+1})}{\partial f_{t+1}'} + \frac{\partial \text{vec}(L_t)}{\partial \text{vec}(K_t)'}\frac{\partial \text{vec}(K_t)}{\partial f_{t+1}'} \\ &= \dot{T}_{t+1} - (Z_t' \otimes I_m)\dot{K}_t. \end{aligned} \tag{A.13}$$

$$\begin{aligned}\dot{A}_{t+1} = \frac{\partial a_{t+1}}{\partial f_{t+1}'} &= \frac{\partial T_{t+1}a_t}{\partial \text{vec}(T_{t+1})'}\frac{\partial \text{vec}(T_{t+1})}{\partial f_{t+1}'} + \frac{\partial K_tv_t}{\partial \text{vec}(K_t)'}\frac{\partial \text{vec}(K_t)}{\partial f_{t+1}'} \\ &= (a_t' \otimes I_m)\dot{T}_{t+1} + (v_t' \otimes I_m)\dot{K}_t. \end{aligned} \tag{A.14}$$

$$\begin{aligned}\dot{P}_{t+1} = \frac{\partial \text{vec}(P_{t+1})}{\partial f_{t+1}'} &= \frac{\partial \text{vec}(T_{t+1}P_tL_t')}{\partial \text{vec}(T_{t+1})'}\frac{\partial \text{vec}(T_{t+1})}{\partial f_{t+1}'} + \frac{\partial \text{vec}(T_{t+1}P_tL_t')}{\partial \text{vec}(L_t')'}\frac{\partial \text{vec}(L_t')}{\partial \text{vec}(L_t)'}\frac{\partial \text{vec}(L_t)}{\partial f_{t+1}'} + \dot{Q}_{t+1} \\ &= (L_tP_t \otimes I_m)\dot{T}_{t+1} + (I_m \otimes T_{t+1}P_t)C_m\dot{L}_t + \dot{Q}_{t+1}. \end{aligned} \tag{A.15}$$

Plugging (A.12) in (A.14) we obtain

$$\begin{aligned}\dot{A}_{t+1} &= [(a_t' \otimes I_m) + (v_t'F_t^{-1}Z_tP_t \otimes I_m)]\dot{T}_{t+1} \\ &= [(a_t' + v_t'F_t^{-1}Z_tP_t) \otimes I_m]\dot{T}_{t+1} \\ &= (a_{t|t}' \otimes I_m)\dot{T}_{t+1}. \end{aligned} \tag{A.16}$$

Plugging (A.12) and (A.13) in (A.15) we obtain

$$\begin{aligned}
\dot{P}_{t+1} &= (L_t P_t \otimes I_m)\dot{T}_{t+1} + (I_m \otimes T_{t+1}P_t)C_m[I_{m^2} - (Z_t' F_t^{-1} Z_t P_t \otimes I_m)]\dot{T}_{t+1} + \dot{Q}_{t+1} \\
&= (L_t P_t \otimes I_m)\dot{T}_{t+1} + C_m[(T_{t+1}P_t \otimes I_m) - (T_{t+1}P_t Z_t' F_t^{-1} Z_t P_t \otimes I_m)]\dot{T}_{t+1} + \dot{Q}_{t+1} \\
&= (L_t P_t \otimes I_m)\dot{T}_{t+1} + C_m[(T_{t+1}P_t - T_{t+1}P_t Z_t' F_t^{-1} Z_t P_t) \otimes I_m]\dot{T}_{t+1} + \dot{Q}_{t+1} \\
&= (T_{t+1}P_{t|t} \otimes I_m)\dot{T}_{t+1} + C_m(T_{t+1}P_{t|t} \otimes I_m)\dot{T}_{t+1} + \dot{Q}_{t+1} \\
&= 2N_m(T_{t+1}P_{t|t} \otimes I_m)\dot{T}_{t+1} + \dot{Q}_{t+1}. \qquad \text{(A.17)}
\end{aligned}$$

Note that expressions (A.10), (A.11), (A.16) and (A.17) can be be also obtained by differentiating the recursions in (3), therefore avoiding the computation of (A.12)-(A.15). Shifting one period backward (A.16) and substituting into (A.10) we obtain:

$$\dot{V}_t = -[(a_t' \otimes I_N)\dot{Z}_t + (a_{t-1|t-1}' \otimes Z_t)\dot{T}_t]. \qquad \text{(A.18)}$$

Shifting one period backward (A.17) and substituting into (A.11) we obtain

$$\dot{F}_t = 2N_N(Z_t P_t \otimes I_N)\dot{Z}_t + 2(Z_t \otimes Z_t)N_m(T_t P_{t-1|t-1} \otimes I_m)\dot{T}_t + \dot{H}_t + (Z_t \otimes Z_t)\dot{Q}_t. \qquad \text{(A.19)}$$

A.3 State space model in forward form

Let us consider the state space model in the so-called *forward form*:

$$\begin{aligned}
y_t &= Z_t\alpha_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, H_t), \\
\alpha_{t+1} &= T_t\alpha_t + \eta_t, \quad \eta_t \sim \mathcal{N}(0, Q_t), \quad \alpha_1 \sim \mathcal{N}(a_1, P_1), \quad t = 1, ..., n.
\end{aligned} \qquad \text{(A.20)}$$

The log-likelihood function is the same as in (2), thus ∇_t and $\mathcal{I}_t$ are the same as in (6). The recursions in (3) are replaced by the following ones:

$$\begin{aligned}
&v_t = y_t - Z_t a_t, && F_t = Z_t P_t Z_t' + H_t, \\
&K_t = T_t P_t Z_t' F_t^{-1}, && L_t = T_t - K_t Z_t, \\
&a_{t+1} = T_t a_t + K_t v_t && P_{t+1} = T_t P_t L_t' + Q_t, \quad t = 1, ..., n.
\end{aligned} \qquad \text{(A.21)}$$

Here we assume the following time dependency in the system matrices: $Z_t = Z(f_t, \theta_m)$, $H_t = H(f_t, \theta_m)$, but $T_t = T(f_{t+1}, \theta_m)$ and $Q_t = Q(f_{t+1}, \theta_m)$. The formulae (A.10)-(A.19) remain unchanged by simply replacing T_{t+1}, $\dot{T}_{t+1}$, Q_{t+1}, and $\dot{Q}_{t+1}$ with T_t, $\dot{T}_t$, Q_t, and $\dot{Q}_t$.

B Examples

In this section, we look at some examples of time-varying state space models that have been considered in the literature and show how they can be analyzed with our score driven algorithm. In particular, in section B.1 we consider the local level model with time-varying volatility, a

model that has been popularized by Stock and Watson (2007) to study inflation dynamics. In section, B.2 we consider autoregressive processes with time-varying parameters.

B.1 Local level model

Let us consider the local level model with time-varying volatilities:

$$\begin{aligned} y_t &= \mu_t + \varepsilon_t, \qquad & \varepsilon_t &\sim \mathcal{N}(0, \sigma^2_{\varepsilon,t}), \\ \mu_t &= \mu_{t-1} + \eta_t, \qquad & \eta_t &\sim \mathcal{N}(0, \sigma^2_{\eta,t}). \end{aligned} \tag{B.1}$$

This model has been proposed by Stock and Watson (2007) to model US inflation. The estimation of (B.1) using the score-driven approach was initially proposed by Creal et al. (2008, sec. 4.4). Their algorithm, however, contains some inconsistencies and below we show how it should be modified. First, the state vector and the system matrices are equal to $\alpha_t = \mu_t$, $Z_t = T_t = 1$, $H_t = \sigma^2_{\varepsilon,t}$, $Q_t = \sigma^2_{\eta,t}$. Thus, the application of the Kalman Filter leads to the following recursions:

$$\begin{aligned} v_t &= y_t - a_t, \qquad & a_{t+1} &= a_t + k_t v_t, \\ d_t &= p_t + \sigma^2_{\varepsilon,t}, \qquad & p_{t+1} &= (1 - k_t)p_t + \sigma^2_{\eta,t+1}, \\ k_t &= p_t / d_t, \qquad & t &= 1, ..., n. \end{aligned} \tag{B.2}$$

Second, the vector of time-varying parameters, which is recursively estimated using the score-driven filter (4), is equal to $f_t = (\log \sigma_{\varepsilon,t}, \log \sigma_{\eta,t})'$. Third, the corresponding Jacobian matrices are $\dot{H}_t = \left(2\sigma^2_{\varepsilon,t}, 0\right)$, $\dot{Q}_t = \left(0, 2\sigma^2_{\eta,t}\right)$, $\dot{Z}_t = \dot{T}_t = (0, 0)'$. Finally, the conditional log-likelihood is equal to $\ell_t \propto -\frac{1}{2}(\log d_t + v_t^2/d_t)$, and the gradient vector and information matrix in (6) are:[1]

$$\nabla_t = \frac{1}{2d_t^2}\dot{d}_t'(v_t^2 - d_t), \quad \mathcal{I}_t = \frac{1}{2d_t^2}\left[\dot{d}_t'\dot{d}_t\right], \tag{B.3}$$

where $\dot{d}_t = \left(2\sigma^2_{\varepsilon,t}, 2\sigma^2_{\eta,t}\right)$. The score's driving mechanism is represented by $(v_t^2 - d_t)$, that is the deviation of the current estimate of the prediction error variance (by means of v_t) from its past estimate (d_t). Such term is weighted by $\dot{d}_t$, which determines the different impact on the two time-varying volatilities.[2] A multivariate extension of the score driven model considered in this section has been used by Buccheri et al. (2020) to model high-frequency multivariate financial time-series.

[1] Note that the information matrix is singular. Therefore, it needs to be smoothed to be used as scaling matrix.

[2] Note that the resulting algorithm is different from the one derived in Creal et al. (2008). In fact, the two volatilities are only updated using information in the second moments of the data and the level of the prediction error, v_t, does not enter directly the filter.

B.2 Autoregressive models

Here, we consider models that are perfectly observable. In this case, our algorithm collapses to the score-driven filter proposed in the literature by Blasques et al. (2014) and Delle Monache and Petrella (2017). Let us consider the following autoregressive model with time-varying parameters:

$$y_t = \phi_t y_{t-1} + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \sigma_t^2), \tag{B.4}$$

the model can be easily cast in state space form by setting $\alpha_t = y_t$, $Z_t = 1$, $\epsilon_t = 0$, $H_t = 0$, $T_t = \phi_t$, $\eta_t = \xi_t$ and $Q_t = \sigma_t^2$. The vector of interest is $f_t = (\phi_t, \sigma_t^2)'$, and the corresponding Jacobians are $\dot{T}_t = (1, 0)$, $\dot{Q}_t = (0, 1)$, $\dot{Z}_t = \dot{H}_t = 0$. For simplicity we do not impose any restrictions on f_t.[3] Combining the KF (3) with the new filter (6)-(7) leads to the following expression for the scaled-score vector:

$$s_t = \mathcal{I}_t^{-1} \nabla_t = \begin{bmatrix} \frac{1}{y_{t-1}^2}(y_{t-1}\xi_t) \\ (\xi_t^2 - \sigma_t^2) \end{bmatrix}. \tag{B.5}$$

The driving process for the coefficient ϕ_t is the prediction error scaled by the regressor, while the volatility σ_t^2 is driven by the squared prediction error. These match exactly those in Blasques et al. (2014) and Delle Monache and Petrella (2017).

We can easily extend to the case of vector autoregressive model of order p:

$$y_t = \Phi_{1,t} y_{t-1} + \cdots + \Phi_{p,t} y_{t-p} + c_t + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \Gamma_t). \tag{B.6}$$

The SSF representation can be obtain by setting:

$$\alpha_t = (y_t', \ldots, y_{t-p}', 1)', \quad Z_t = I, \quad T_t = \begin{bmatrix} \Phi_{1,t} & \ldots & \Phi_{p,t} & c_t \\ I & & & \\ & \ddots & & \\ & & I & \\ 0 & & & 0 \end{bmatrix}, \quad Q_t = \begin{bmatrix} \Gamma_t & & & \\ & 0 & & \\ & & \ddots & \\ & & & 0 \end{bmatrix},$$

where $f_t = (\text{vec}(\Phi_t)', \text{vec}(\Gamma_t)')'$, with $\Phi_t = [\Phi_{1,t}, \ldots, \Phi_{p,t}, c_t]$. We therefore have that $\dot{Z}_t = \dot{H}_t = 0$, while $\dot{T}_t$ and $\dot{Q}_t$, are selection matrices. After some algebra, the scaled-score can be specialized in two sub-vectors driving the coefficients and the volatilities:

$$s_t = \mathcal{I}_t^{-1} \nabla_t = \begin{bmatrix} (X_t' \Gamma_t^{-1} X_t)^{-1} X_t' \Gamma_t^{-1} \xi_t \\ \text{vec}(\xi_t \xi_t') - \text{vec}(\Gamma_t) \end{bmatrix}, \tag{B.7}$$

where $X_t = (\alpha_{t-1}' \otimes I)$. Interestingly, the algorithm proposed by Koop and Korobilis (2013)

[3] Delle Monache and Petrella (2017) show how to impose stable roots.

can be obtained as a special case of ours by imposing some restrictions on the model. First, set the law of motion (4) c = 0, A = I and let B depend on two scalar parameters (one driving the coefficients and one the volatility). Second, replace the information matrix for the time-varying coefficients by its smoothed estimator: $\tilde{\mathcal{I}}_{c,t} = (1-\kappa)\tilde{\mathcal{I}}_{c,t-1} + \kappa(X_t'\Gamma_t^{-1}X_t)$.

It is well known that VAR models tend to suffer from the 'curse of dimensionality' and to overfit the data (see, e.g., Litterman, 1979; Doan et al., 1986). In the context of fixed coefficients, Bayesian techniques are used to shrink the parameters, therefore reducing estimation variance. Our algorithm can easily accommodate shrinkage, as detailed in Appendix D. A regularized version of the model proposed by Koop and Korobilis (2013) can then be easily obtained, where the parameters can be shrunk towards any type of prior that can be reformulated in the form of stochastic constraints. These include the most popular priors typically considered in the Bayesian VAR literature, including the Minnesota prior, the sum of coefficients prior and the long-run prior (see, e.g., Del Negro, 2012; Kapetanios et al., 2019).

C Monte Carlo exercise

The Monte Carlo exercise is based on a battery of simple bivariate models that share a common component. We simulate two time series ($y_{1,t}$ and $y_{2,t}$) that load (with parameters $\lambda_{1,t}$ and $\lambda_{2,t}$) on a common factor μ_t and are each affected by idiosyncratic noise. The common factor evolves as an AR(1), with coefficient ρ_t. In this model, we look at the time variation of a subset of parameters one at the time. Specifically, DGP1 lets the loading on the common factor vary over time and keeps all the remaining parameters fixed. In DGP2 we keep both factor loadings constant ($\lambda_{1,t} = 1$ and $\lambda_{2,t} = 1$) while allowing for time variation in the AR coefficient of the common factor, ρ_t. In DGP3 and DGP4 we experiment with time-varying volatility, either in the measurement or in the transition equations, keeping everything else fixed. In these latter cases, the simulated model is a univariate signal plus noise model. In all cases, we consider two different sample sizes, $n = 250$ and $n = 500$. As for the laws of motions for the TVPs entering the various DGPs, we experiment with 6 different possibilities:

Case 1: CONSTANT $f_t = a_1, \forall t;$
Case 2: SINE $f_t = a_2 + b_2 \sin\left(\frac{2\pi t}{T/2}\right);$
Case 3: SINGLE STEP $f_t = a_3 + b_3 (t \geq \tau);$
Case 4: DOUBLE STEP $f_t = a_4 + b_4 I (t \geq \tau_1) + c_4 I (t \geq \tau_2);$
Case 5: RAMP $f_t = a_5 + \left(\frac{b_5}{T/c_5}\right) \mod (t);$
Case 6: AR(1) MODEL $f_t = a_6(1-b_6) + b_6 f_{t-1} + \xi_t, \quad \xi_t \sim \mathcal{N}(0, c_6);$

where $f_t = \lambda_t$ in DGP1, $f_t = \rho_t$ in DGP2, $f_t = \sigma^2_{\varepsilon,t}$ in DGP3, and $f_t = \sigma^2_{u,t}$ in DGP4.[4] We

[4]Moreover, f_t in the AR(1) model is transformed to be within the unit circle for DGP2, and to be positive

start with a baseline case in which we keep the parameters constant over time. We then move to four cases where the parameters change according to a deterministic process. In case 2 the parameters follow a cyclical pattern determined by a sine function. In cases 3 and 4 we let the parameters break at discrete points in time, allowing for either one or two breaks. We set the location of the discrete breaks at given point in the sample. In the case of a single break, this occurs in the middle of the sample. When we consider two breaks, they are located at 1/3 and 2/3 of the sample. Case 5 (RAMP) is a rather challenging case, whereby the parameters increase for some time before returning abruptly to their starting levels. Finally, case 6 is the only one in which we let the parameters vary stochastically, following a persistent AR(1) model. We consider two cases, one with a near unit root process (i.e. with an AR root of 0.99) and a low variance, one with lower persistence (AR root of 0.97) but substantially higher variance. We obtain very similar results in these two specifications. The DGPs that we design are simple, in that time variation is introduced in all the channels in which it can manifest itself, but only one at the time. By focusing on a single channel at the time, we can better identify the situations in which our model either succeeds or fails at tracking parameter variation. Further details on the calibration of the parameters are provided in the next subsections.

The results of the Monte Carlo exercise are shown in Table C.1. The table is organized in four panels, corresponding to each of the four DGPs. On the left hand side of the table we show the results for a sample size of $n = 250$, on the right hand side those obtained when setting $n = 500$. For each DGP the analysis is based on 300 simulations. In each panel we report the results obtained for the six alternative laws of motion described above. We base our assessment on five different statistics, namely the Root Mean Squared Error (RMSEs), the Mean Absolute Error (MAE), the correlation between actual and estimated coefficients, the Coverage (i.e. percentage of times that the actual parameters fall in a given estimated confidence interval) and the number of cases in which a pile-up occurs (#Pile-up). The last statistics consists of the number of simulations in which the static coefficients that pre-multiply the score end up being lower than 10^{-6}, which we take as sufficient evidence that the estimated parameters are effectively zero, i.e. that the model does not detect any time variation.

We highlight five results. First, for all the DGPs in which the true parameters are constant the model performs extremely well. This means that the adaptive filter correctly collapses the parameters to a constant. As a result, RMSEs and MAEs are virtually nil, the actual coverage extremely precise and a pile-up at zero occurs in about 75 percent of the cases for the volatility models and more than half of the cases for the loadings and AR coefficients.[5] This result implies that our estimation method does not generate spurious time variation in the coefficients when

for DGP3 and DGP4.

[5]For the latter two cases, in an additional 20% of the simulations the estimated parameters are virtually constant, despite not being classified as a pile-up according to the criterion we have set above.

this is not present in the data generating process. Second, the pile-up problem is not of primary concern for our estimator. The number of instances in which our method (incorrectly) concludes that there is no time variation is basically zero in most cases. Third, for all the DGPs and across all the specifications for the parameters we obtain extremely good coverage. Coverage is almost perfect when time variation involves the autoregressive coefficients, somewhat lower when it affects the volatility of the measurement and of the transition equation, in particular when parameters break at discrete points in the sample. Fourth, across all DGPs the RAMP specification is the one that the model finds more challenging to estimate. This specification generally leads to low correlation between actual and estimated parameters, higher RMSEs and MAEs and lower coverage. This is not surprising, since our model is, by construction, designed to detect smooth changes, whereas in the RAMP model periods of small, continuous changes are suddenly interrupted by large breaks. Fifth, the adaptive filter is very effective in estimating time-varying loadings and autoregressive coefficients, but it is rather conservative in the estimation of the time-varying variances, especially when these are driven by a near unit root process. For this DGP, in one third of the cases the filter ends in a pile-up when the sample is relatively small (T=250). However, when time variation *is* detected, the algorithm yields relatively low RMSEs and MAEs and a satisfactory coverage. We take these results as evidence that, in the case of time-varying variances, the algorithm needs relatively more evidence of breaks in the parameters to move away from zero. A larger sample size (of the type encountered in financial applications that use high frequency data) basically eliminates the problem.

C.1 Specification of the DGPs

DGP1 - Time-Varying loadings

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} = \begin{bmatrix} 1 \\ \lambda_t \end{bmatrix} \mu_t + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix}, \quad \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix} \sim \mathcal{N}(0, I),$$
$$\mu_t = 0.8\mu_{t-1} + u_t \qquad u_t, \sim \mathcal{N}(0, 1).$$

DGP2 - Time-Varying AR coefficient

$$\begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \mu_t + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix}, \quad \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix} \sim \mathcal{N}(0, I),$$
$$\mu_t = \rho_t \mu_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0, 1).$$

DGP3 - Time-Varying Volatility in the measurement equation

$$\begin{aligned} y_t &= \mu_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2_{\varepsilon,t}), \\ \mu_{t+1} &= 0.8\mu_t + u_t, \quad u_t \sim \mathcal{N}(0, 1). \end{aligned}$$

DGP4 - Time-Varying Volatility in the transition equation

$$\begin{aligned} y_t &= \mu_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1), \\ \mu_{t+1} &= 0.8\mu_t + u_t, \quad u_t \sim \mathcal{N}(0, \sigma^2_{\eta,t}). \end{aligned}$$

C.2 Calibration

DGP1:Time-varying loadings

CONSTANT: $a_1 = 1$;

SINE: $a_2 = 2$, $b_2 = 1.5$;

SINGLE STEP: $a_3 = 1$, $b_3 = 2$, $\tau = (2/5)n$;

DOUBLE STEP: $a_4 = 1$, $b_4 = c_4 = 1.5$, $\tau_1 = (1/5)n$, $\tau_2 = (3/5)n$;

RAMP: $a_5 = 0.5$, $b_5 = 4$, $c_5 = 2$;

AR(1) $[b_6 = 0.99]$: $a_6 = 1$, $b_6 = 0.99$, $c_6 = 0.08^2$.

AR(1) $[b_6 = 0.97]$: $a_6 = 1$, $b_6 = 0.97$, $c_6 = 30.24^2$.

DGP2: Time-varying autoregressive coefficient

CONSTANT: $a_1 = 0.7$;

SINE: $a_2 = 0$, $b_2 = 0.7$;

SINGLE STEP $a_3 = 0.8$, $b_3 = -0.6$, $\tau = (2/5)n$;

DOUBLE STEP:$a_4 = 0.8$, $b_4 = c_4 = -0.5$, $\tau_1 = (1/5)n$, $\tau_2 = (3/5)n$;

RAMP: $a_5 = 0.3$, $b_5 = -0.9$. $c_5 = 2$;

AR(1) $[b_6 = 0.99]$: $a_6 = 0.2$, $b_6 = 0.99$, $c_6 = 0.08^2$;

AR(1) $[b_6 = 0.97]$: $a_6 = 0.2$, $b_6 = 0.97$, $c_6 = 0.24^2$;

and in the latter two cases we also impose the restriction that $|\rho_t| < 1$.

DGP3 and DGP4: Time-varying volatilities

CONSTANT: $a_1 = 1$;

SINE: $a_2 = 1$, $b_2 = 0.9$;

SINGLE STEP: $a_3 = 1$, $b_3 = 4$, $\tau = (2/5)n$;

DOUBLE STEP:$a_4 = 1$, $b_4 = c_4 = 3$, $\tau_1 = (1/5)n$, $\tau_2 = (3/5)n$;

RAMP: $a_5 = 0.5$, $b_5 = 8$, $c_5 = 2$;

AR(1) $[b_6 = 0.99]$: $a_6 = 0$, $b_6 = 0.99$, $c_6 = 0.08^2$;

AR(1) $[b_6 = 0.97]$: $a_6 = 0$, $b_6 = 0.97$, $c_6 = 0.24^2$;

In DGP3 and DGP4, after having simulated the dynamic of the volatility the time-varying volatilities are rescaled so as to have a fixed ratio between the measurement and transition error variances equal to 1.

For each DGP we target 300 simulations. However, the actual number of samples changes

depending on the specifications. In the case of constant coefficients, where we would like to see our estimator to end up in a pile-up situation as often as possible, we perform 300 simulations and compute all the statistics on these samples. For the remaining specifications, on the other hand, we keep on drawing artificial samples until we obtain 300 simulations in which the estimated parameters are different from zero and compute RMSEs, MAEs, correlations and coverage ratios on these 300 artificial samples. At the same time, we also keep track of the number of times in which the pile-up problem arises. To better understand how we proceed, let us take a concrete example, that is the bottom-left panel of Table 1 (DGP4, i.e. the model with time-varying volatility in the transition equation, $n = 250$). In the first row we report the results for the constant coefficient case. As explained, for this case we simulate 300 artificial samples and estimate the model using our algorithm. It turns out that in 236 out of 300 simulations our estimation method ends up in a pile-up. The RMSEs, MAEs, Correlations and Coverages, are estimated on all the 300 simulations. Now let us take in the same panel the last line, referring to one of the AR(1) specifications. In this case we need to draw up to 314 samples to obtain 300 simulations in which the estimation algorithm does not end being stuck in a region of the likelihood where the model loading is zero. Now, in this case all the remaining statistics are computed on the 300 'good' samples. We proceed in this way because we want to appraise two different issues. The former is the percentage of cases in which the algorithm ends up in the pile-up even if the true DGP implies time variation. The second is how well it estimates the parameters *conditional on the model correctly detecting time variation*. The two points are of independent interest. If we were to find that the model often ends up in the pile-up but it is very precise when it does not, one could decide to force the algorithm to stay away from zero, for instance by using a grid-based estimation method. This is the choice made, for instance, by Koop and Korobilis (2013). Similarly, in their Monte Carlo Markov Chain estimation, Stock and Watson (2007) reject draws in which the variances are very close to zero.

In figures C.1-C.8 we report the simulated true process for the time-varying parameters (red line), and the fan chart associated to the $5^{th}, 10^{th}, 20^{th}, 30^{th}, 40^{th}, 60^{th}, 70^{th}, 80^{th}, 90^{th}$ and 95^{th} quantile of the filtered parameters. In the case of the AR(1) specification we focus on the more persistent AR(1) DGP and report the difference between actual and estimated parameters. The figures are based on 300 replications.

Table C.1: Monte Carlo Exercise

	DGP 1: time-varying LOADINGS COEFFICIENT											
	T = 250						T = 500					
	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up
CONSTANT	0.003	0.003		0.678	0.900	165	0.000	0.000		0.680	0.900	180
SINE	0.473	0.380	0.909	0.636	0.852	0	0.386	0.305	0.940	0.648	0.868	0
SINGLE STEP	0.406	0.280	0.927	0.656	0.876	0	0.335	0.229	0.951	0.660	0.882	0
DOUBLE STEP	0.462	0.339	0.936	0.640	0.860	0	0.390	0.277	0.953	0.652	0.874	0
RAMP	0.695	0.461	0.723	0.648	0.856	0	0.575	0.367	0.817	0.658	0.870	0
AR(1) [$b_6 = 0.99$]	0.265	0.213	0.727	0.676	0.892	0	0.274	0.217	0.807	0.676	0.892	0
AR(1) [$b_6 = 0.97$]	0.523	0.415	0.803	0.660	0.872	0	0.527	0.413	0.828	0.662	0.872	0
	DGP 2: time-varying AUTOREGRESSIVE COEFFICIENT											
	T = 250						T = 500					
	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up
CONSTANT	0.006	0.006		0.676	0.900	147	0.005	0.004		0.680	0.900	156
SINE	0.330	0.267	0.780	0.684	0.900	0	0.268	0.212	0.866	0.682	0.900	0
SINGLE STEP	0.228	0.166	0.769	0.684	0.900	0	0.203	0.140	0.811	0.686	0.902	0
DOUBLE STEP	0.240	0.185	0.872	0.684	0.900	0	0.209	0.160	0.892	0.686	0.900	0
RAMP	0.341	0.261	0.392	0.682	0.900	0	0.299	0.221	0.548	0.683	0.900	0
AR(1) [$b_6 = 0.99$]	0.297	0.237	0.608	0.684	0.900	4	0.301	0.241	0.695	0.686	0.902	1
AR(1) [$b_6 = 0.97$]	0.477	0.377	0.575	0.686	0.900	0	0.478	0.375	0.593	0.685	0.900	0
	DGP 3: TIME-VARYING VOLATILITY - MEASUREMENT EQUATION ERROR											
	T = 250						T = 500					
	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up
CONSTANT	0.000	0.000		0.676	0.896	231	0.000	0.000		0.682	0.899	232
SINE	0.981	0.768	0.747	0.672	0.876	1	0.829	0.637	0.813	0.678	0.882	0
SINGLE STEP	0.808	0.605	0.843	0.618	0.848	0	0.659	0.477	0.883	0.652	0.870	0
DOUBLE STEP	0.702	0.551	0.856	0.628	0.848	0	0.595	0.460	0.889	0.648	0.870	0
RAMP	0.960	0.764	0.498	0.640	0.860	20	0.803	0.599	0.646	0.656	0.874	1
AR(1) [$b_6 = 0.99$]	0.717	0.571	0.568	0.664	0.880	93	0.748	0.578	0.608	0.668	0.886	24
AR(1) [$b_6 = 0.97$]	1.446	0.998	0.600	0.664	0.868	22	1.489	1.005	0.626	0.674	0.878	3
	DGP 4: TIME-VARYING VOLATILITY - TRANSITION EQUATION ERROR											
	T = 250						T = 500					
	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up	RMSE	MAE	Corr.	68% Cov.	90% Cov.	# Pile up
CONSTANT	0.000	0.000		0.676	0.896	236	0.000	0.000		0.680	0.898	241
SINE	1.052	0.832	0.714	0.672	0.880	1	0.871	0.673	0.794	0.672	0.886	0
SINGLE STEP	0.829	0.614	0.834	0.644	0.864	0	0.680	0.485	0.874	0.656	0.878	0
DOUBLE STEP	0.754	0.592	0.849	0.644	0.868	0	0.620	0.481	0.885	0.656	0.876	0
RAMP	1.015	0.822	0.468	0.644	0.868	1	0.829	0.640	0.615	0.659	0.881	0
AR(1) [$b_6 = 0.99$]	0.768	0.623	0.622	0.668	0.888	95	0.776	0.607	0.613	0.668	0.887	34
AR(1) [$b_6 = 0.97$]	1.533	1.069	0.590	0.664	0.876	14	1.523	1.042	0.619	0.664	0.880	5

Note. The results shown in the first and in the second panel (DGP1 and DGP2) refer to a bivariate factor model in which two variables are driven by a single common factor that evolves as an autoregressive process of order 1. In the first case (DGP1) the loading of the second variable on the common factor varies over time and all the other parameters are kept constant. In the second case (DGP2) the autoregressive component of the common factor varies over time and all the other parameters are kept constant. The results shown in the third and in the fourth panel (DGP3 and DGP4) refer to ARMA(1,1) models that are cast in state space and feature time-varying variances of the random disturbance in, respectively, the measurement and the transition equation. The abbreviations *Corr.* and *Cov.* stand, respectively for Correlation and Coverage, while # Pile-up denotes the number of simulations in which the algorithm delivers constant parameters. The different laws of motion of the parameters in the first column (Constant, Sine, Single Step, Double Step, Ramp and AR(1) are described in Section 4).

Figure C.1: TIME-VARYING LOADINGS, N=250

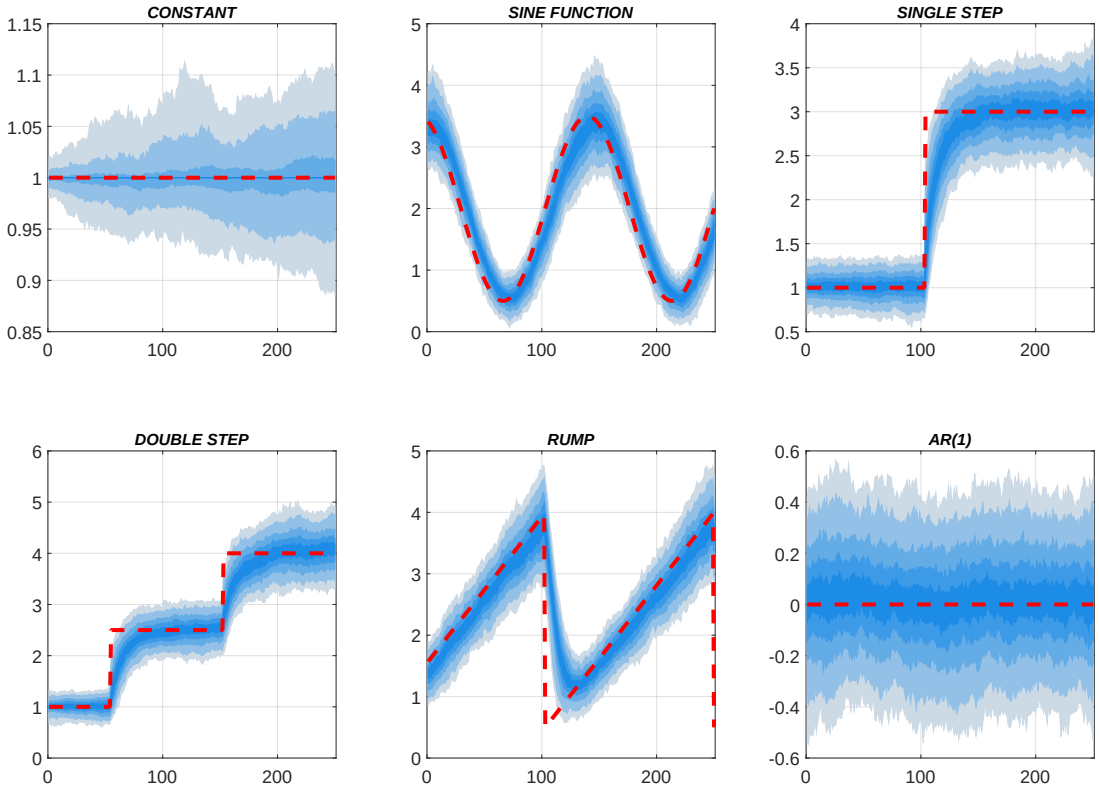

Figure C.2: TIME-VARYING LOADINGS, N=500

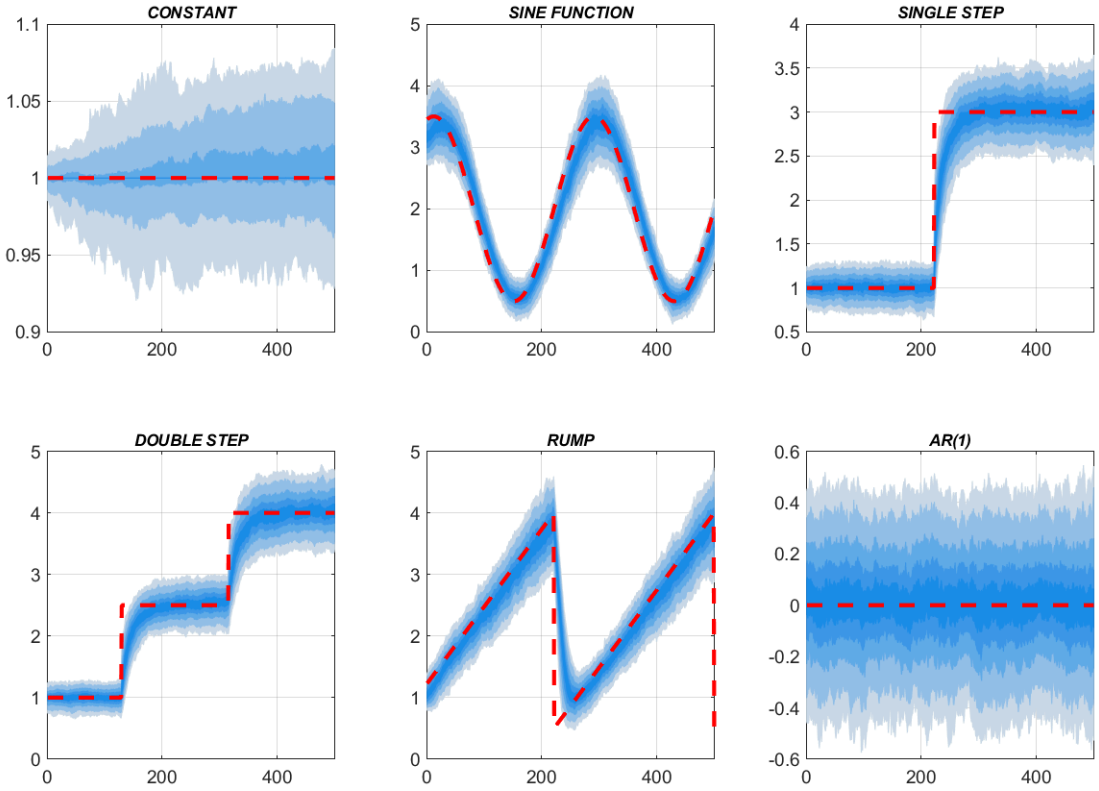

Figure C.3: TIME-VARYING AUTOREGRESSIVE COEFFICIENTS, n=250

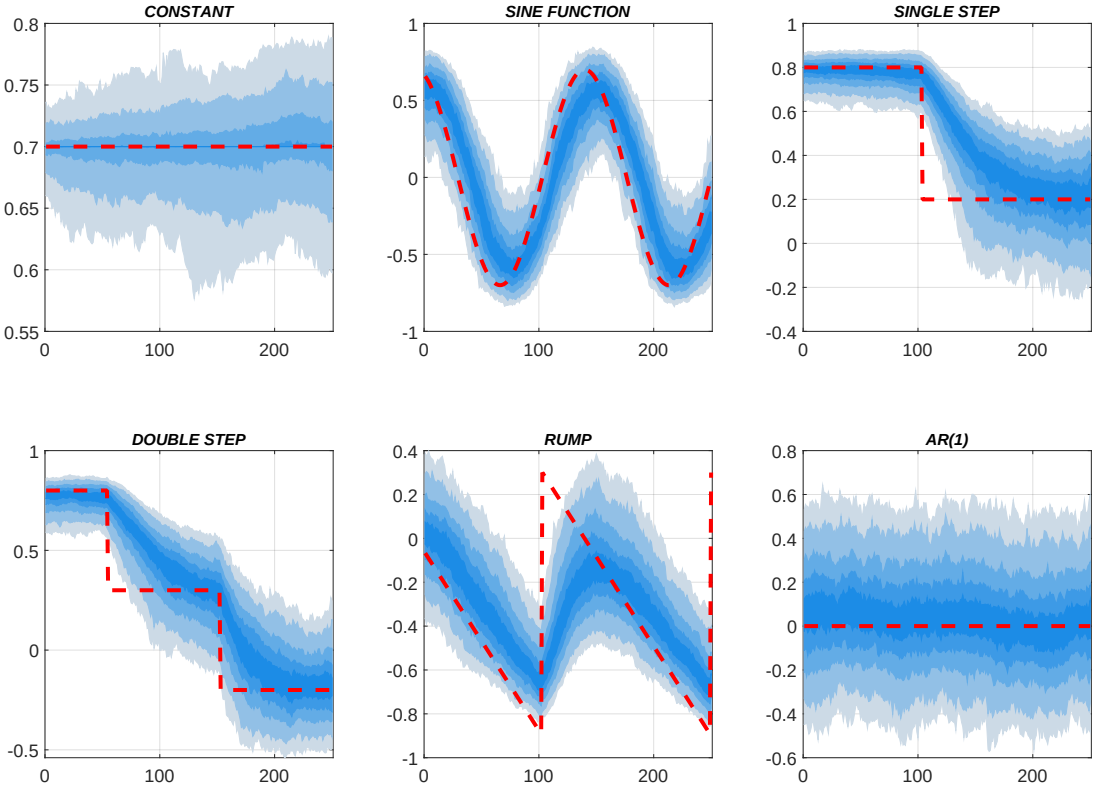

Figure C.4: TIME-VARYING AUTOREGRESSIVE COEFFICIENTS, n=500

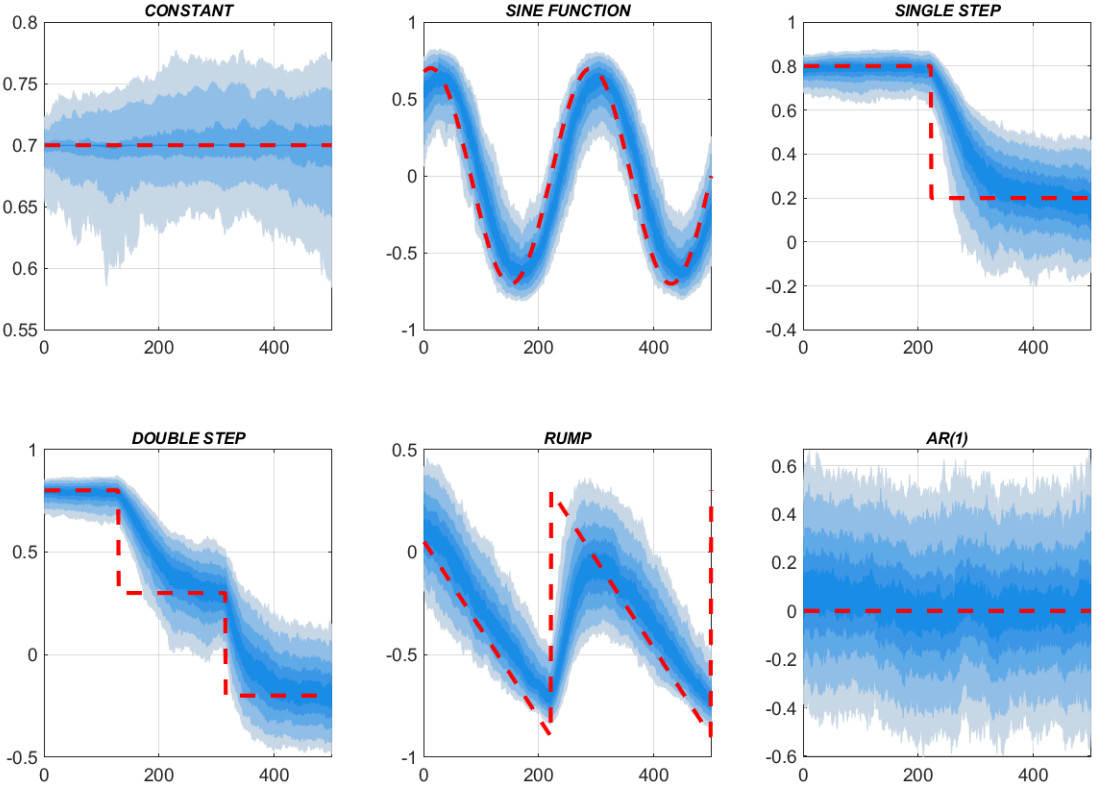

Figure C.5: TIME-VARYING MEASUREMENT EQUATION ERROR VARIANCE, n=250

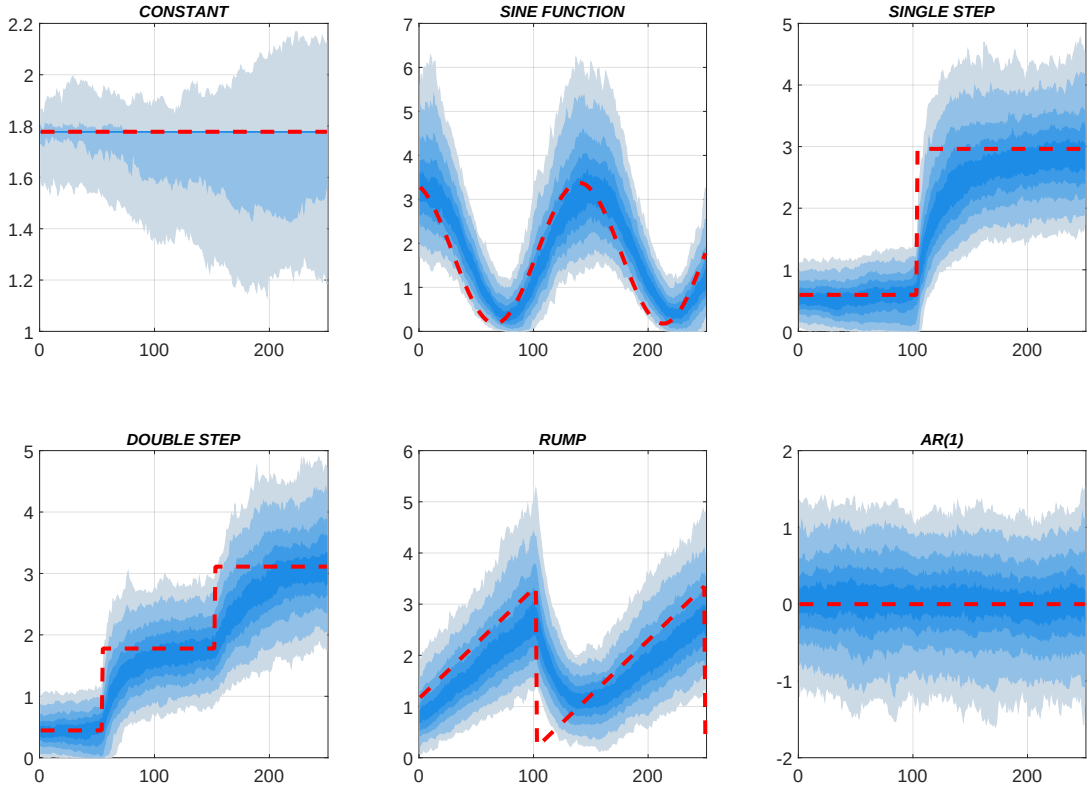

Figure C.6: TIME-VARYING MEASUREMENT EQUATION ERROR VARIANCE, n=500

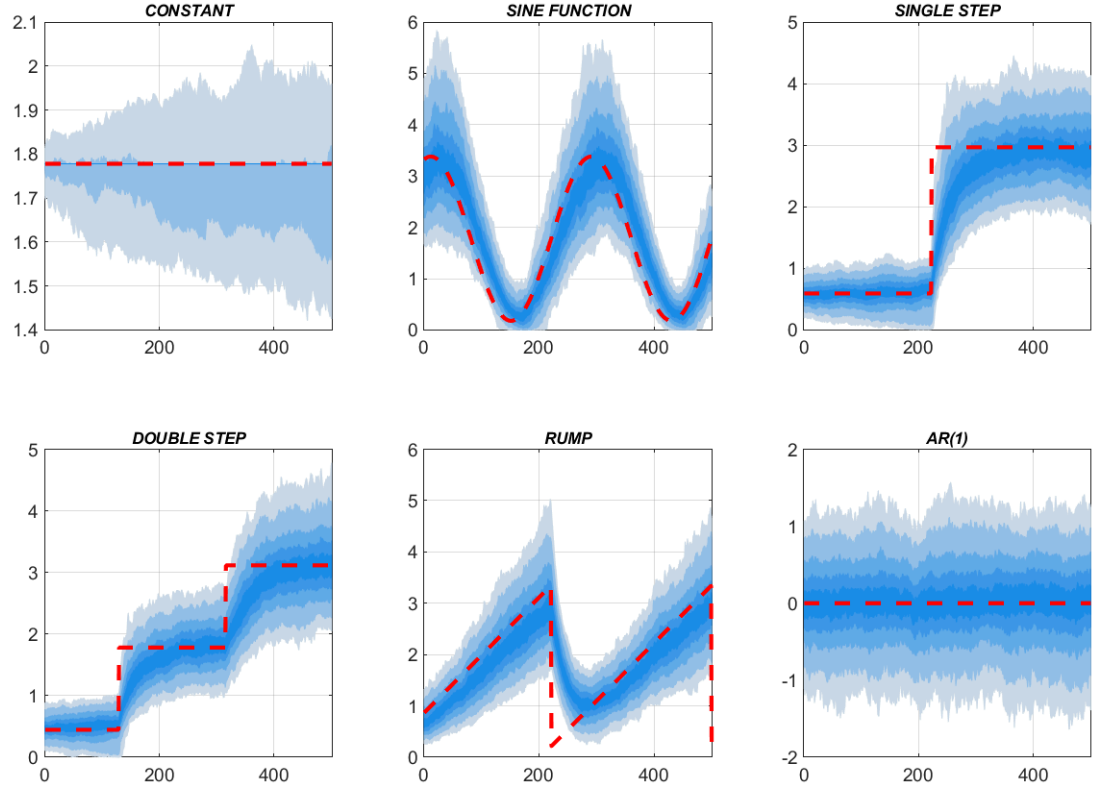

Figure C.7: TIME-VARYING TRANSITION EQUATION ERROR VARIANCE, n=250

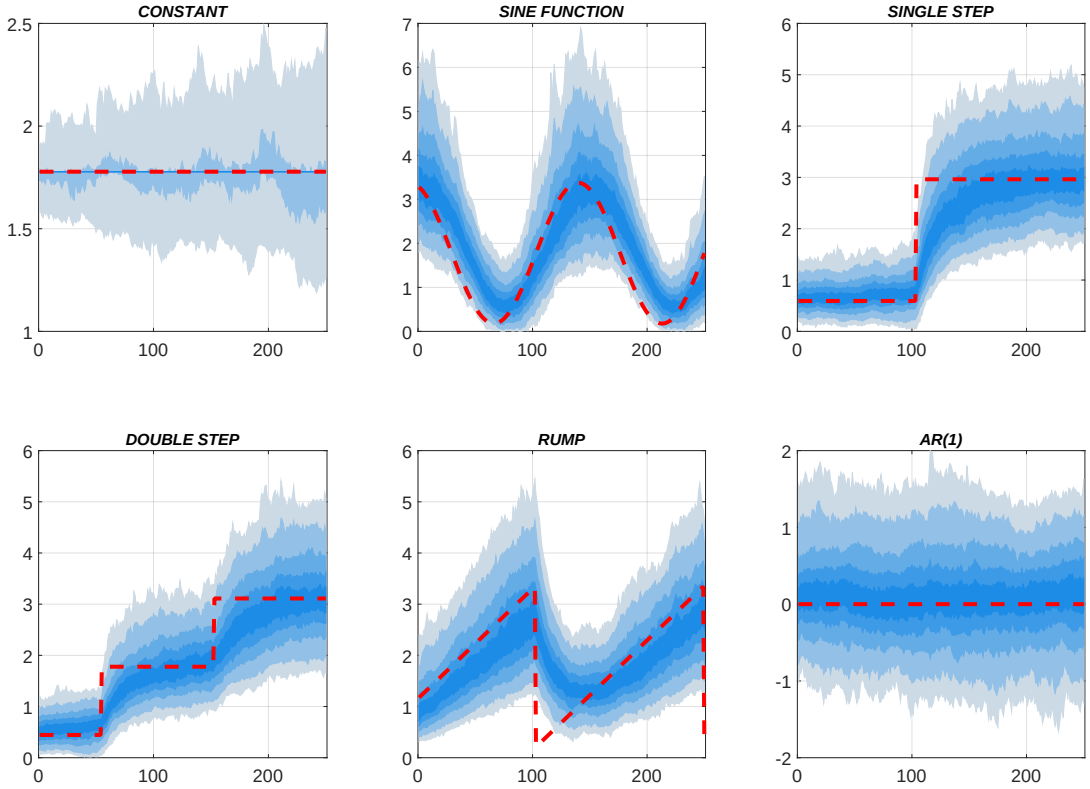

Figure C.8: TIME-VARYING TRANSITION EQUATION ERROR VARIANCE, n=500

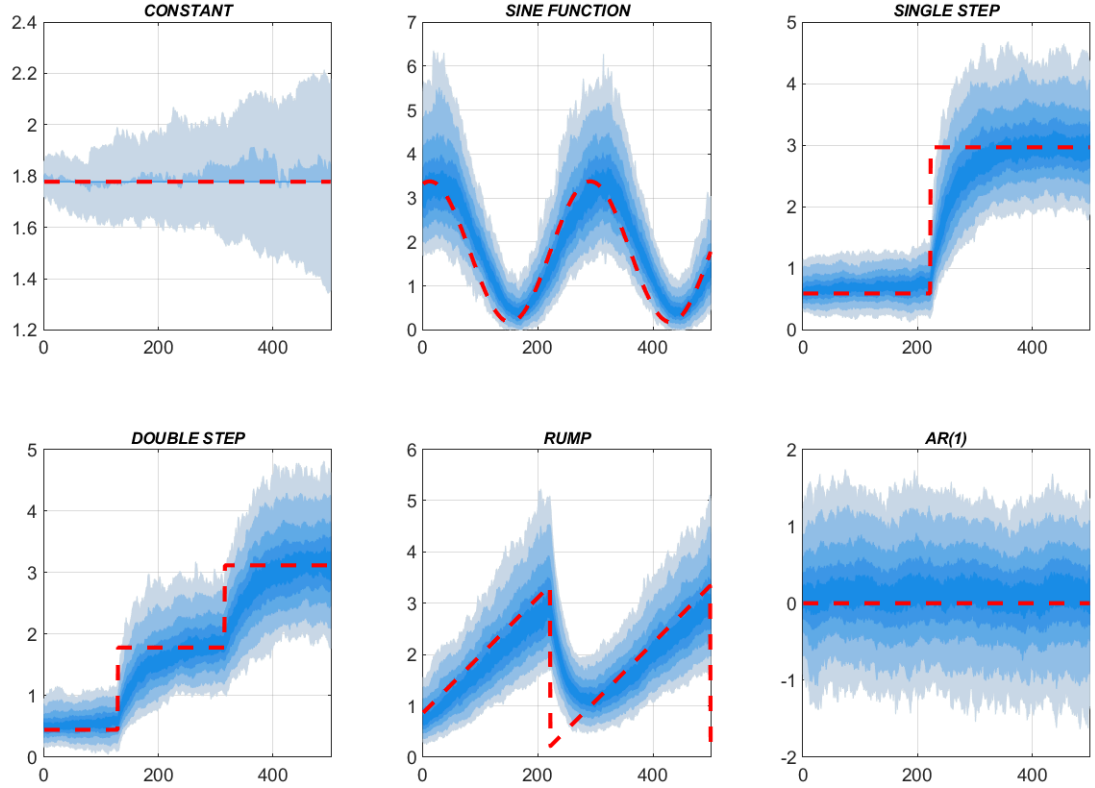

D Shrinking the vector of parameters by the L2 penalty

As the dimension of the system grows, it could be desirable to impose some shrinkage on the model parameters to avoid an increase in the estimation variance (Hastie et al., 2001). Following (Theil and Goldberger, 1961), priors can be incorporated into the state space model (1), with the score driven system matrices described by (4), augmenting the model using the following Gaussian constraints:

$$r_t = R_t f_t + u_t \quad u_t \sim \mathcal{N}(0, \Sigma_t), \tag{D.1}$$

where r_t, Σ_t and R_t are known and the random disturbance u_t is Gaussian and uncorrelated with the other two disturbances ε_t and η_t. Since u_t is Gaussian, the resulting penalty will be quadratic with the matrix Σ_t regulating the degree of shrinkage; e.g for $\Sigma_t \to \infty$ the constraints vanish, while for $\Sigma_t \to 0$ the constraints hold exactly. The vector r_t can be considered a vector of artificial observations. The likelihood function for the vector of 'augmented data', $\tilde{y}_t = (y_t', r_t')'$ is equal to:

$$\ell_t^p = \log p(\tilde{y}_t | Y_{t-1}, \theta) \propto -\frac{1}{2}\left(\log|F_t| + v_t' F_t^{-1} v_t\right) - \frac{1}{2}\left(\log|\Sigma_t| + u_t' \Sigma_t^{-1} u_t\right), \tag{D.2}$$

which can be interpreted as a 'penalized likelihood' with a quadratic penalty function. The resulting 'penalized' score is:

$$s_t^p = (\mathcal{I}_t^p)^{-1} \nabla_t^p = (\mathcal{I}_t + R_t' \Sigma_t^{-1} R_t)^{-1} (\nabla_t + R_t' \Sigma_t^{-1} u_t). \tag{D.3}$$

Using the Woodbury identity we obtain[6]:

$$(\mathcal{I}_t^p)^{-1} = (\mathcal{I}_t + R_t' \Sigma_t^{-1} R_t)^{-1} = \mathcal{I}_t^{-1} - \mathcal{I}_t^{-1} R_t' (R_t \mathcal{I}_t^{-1} R_t' + \Sigma_t)^{-1} R_t \mathcal{I}_t^{-1} = (I - \Upsilon_t R_t) \mathcal{I}_t^{-1}, \tag{D.4}$$

where $\Upsilon_t = \mathcal{I}_t^{-1} R_t' (R_t \mathcal{I}_t^{-1} R_t' + \Sigma_t)^{-1}$. Finally, we can express the penalized (regularized) score as follows:

$$s_t^p = (I - \Upsilon_t R_t) s_t + \Upsilon_t u_t, \tag{D.5}$$

From the last expression it is straightforward to obtain the two polar cases for $\Sigma_t \to 0$ and $\Sigma_t \to \infty$. To illustrate with a simple example how such 'penalized score' works let us define $R_t = I$, $r_t = \bar{f}$, and $\Sigma_t = \frac{1}{\lambda} I$ so that the constraints reduce to $f_t \sim \mathcal{N}(\bar{f}, \frac{1}{\lambda} I)$. This corresponds to a Ridge-type penalty with λ regulating the degree of penalization. The penalized score will be $s_t^p = (I - \Lambda_t) s_t + \Lambda_t (\bar{f} - f_t)$, where $\Lambda_t = \lambda (\mathcal{I}_t + \lambda I)^{-1}$. If we assume a simple random-walk law of motion for f_t, that is $f_{t+1} = f_t + B s_t^p$, and the information matrix equals to identity

[6] In case $\mathcal{I}_t$ is not invertible we use its smoothed estimator $\tilde{\mathcal{I}}_t = (1-\kappa)\tilde{\mathcal{I}}_{t-1} + \kappa \mathcal{I}_t$ which is invertible.

matrix, the resulting regularized score-driven filter is:[7]

$$f_{t+1} = \frac{\mathrm{B}\lambda}{1+\lambda}\bar{f} + \left(I - \frac{\mathrm{B}\lambda}{1+\lambda}\right) f_t + \frac{\mathrm{B}}{1+\lambda} s_t. \tag{D.6}$$

The law of motion is now mean reverting towards the desired value $\bar{f}$. Everything else equal, the larger is λ, the lower is the weight assigned to actual data and the more binding is the constraint. Notice that the strategy of stochastic constraints is very flexible, and generalizes a number of shrinkage methods, including Ridge regressions, Minnesota priors, sum of coefficients priors as well as the long-run prior in Giannone et al. (2019). Kapetanios et al. (2019) discuss in details these additional cases.

E Mixed frequencies and missing observations

Assume to have a data set containing missing observations. The observed vector is represented by $W_t y_t$, where W_t is an $N_t \times N$ selection matrix with $1 \leq N_t \leq N$, meaning that at least one observation is available at time t. Note that W_t is obtained by eliminating the $i-th$ row from I_N when the $i-th$ variable is missing. In this setting, at each time t the likelihood ℓ_t is computed using N_t observations; i.e. $\ell_t = \log p(W_t y_t | Y_{t-1}, \theta)$, that is the *marginal* likelihood. In practice, the score of the marginal likelihood is computed and the updating of f_t is based on the available information.[8] Given this reparameterization, the measurement equation in (1) is modified as follows:

$$W_t y_t = W_t Z_t \alpha_t + W_t \varepsilon_t, \quad W_t \epsilon_t \sim \mathcal{N}(0, W_t H_t W_t'), \tag{E.1}$$

and the first four expressions of the KF (3) are modified as follows:

$$\begin{aligned} v_t &= W_t(y_t - Z_t a_t), & F_t &= W_t(Z_t P_t Z_t' + H_t)W_t', \\ a_{t|t} &= a_t + P_t Z_t' W_t' F_t^{-1} v_t, & P_{t|t} &= P_t - P_t Z_t' W_t' F_t^{-1} W_t Z_t P_t. \end{aligned} \tag{E.2}$$

The expressions in (7) become

$$\begin{aligned} \dot{V}_t &= -[(a_t' \otimes W_t)\dot{Z}_t + (a_{t-1|t-1}' \otimes W_t Z_t)\dot{T}_t], \\ \dot{F}_t &= 2N_{N_t}(W_t Z_t P_t \otimes W_t)\dot{Z}_t + (W_t Z_t \otimes W_t Z_t)2N_m(T_t P_{t-1|t-1} \otimes I_m)\dot{T}_t \\ &\quad + (W_t \otimes W_t)\dot{H}_t + (W_t Z_t \otimes W_t Z_t)\dot{Q}_t. \end{aligned} \tag{E.3}$$

Mixed frequencies typically involve missing observations and temporal aggregation. The case of mixed frequencies is of particular interest for a number of applications, like for instance forecasting low frequency variables using higher frequency predictors (nowcasting). Indeed low frequency indicators can be modeled as a latent process that is observed at regular low

[7] The same regularized filter can be obtained by setting $f_t \sim \mathcal{N}(\bar{f}, \frac{1}{\lambda}\mathcal{I}_t)$, which is the Litterman-type of prior.

[8] A formal discussion of dealing with missing values in score-driven models can be found in Lucas et al. (2016).

frequency intervals and missing at higher frequency dates. The relation between the observed low frequency variable and the corresponding (latent) higher frequency indicator depends on whether the variable is a flow or a stock and on how the variable is transformed before entering the model. In all cases, the variable can be rewritten as a weighted average of the unobserved high frequency indicator (see e.g., Banbura et al., 2013).

F Correlated disturbances

Let consider the case in which the disturbances in (1) are correlated, that is $\mathrm{E}(\eta_t \epsilon_t') = G_t$. In this case the following expressions in (3) need to be modififed:

$$\begin{aligned} F_t &= Z_t P_t Z_t' + Z_t G_t + G_t' Z_t' + H_t, \\ a_{t|t} &= a_t + (P_t Z_t' + G_t) F_t^{-1} v_t, \\ P_{t|t} &= P_t - (P_t Z_t' + G_t) F_t^{-1} (P_t Z_t' + G_t)'. \end{aligned} \tag{F.1}$$

Therefore, the expression for $\dot{F}_t$ in (7) need to be modified as follows:

$$\begin{aligned} \dot{F}_t \;=\; & [2N_N(Z_t P_t \otimes I_N) + (G_t' \otimes I_N) + (I_N \otimes G_t') C_{Nm}] \dot{Z}_t \\ & + (Z_t \otimes Z_t) 2N_m (T_t P_{t-1|t-1} \otimes I_m) \dot{T}_t + \dot{H}_t + (Z_t \otimes Z_t) \dot{Q}_t \\ & + [(I_N \otimes Z_t) + (Z_t \otimes I_N) C_{mN}] \dot{G}_t. \end{aligned} \tag{F.2}$$

G Empirical application

G.1 Identification of the model

In this section, we consider the identification of the model presented in section 3. Let us start by looking at the constant parameter version of the model:

$$\begin{aligned} &\mathtt{pd}_{t+1} = \overline{\mathtt{pd}} - b_1 \tilde{\mu}_{t+1} + b_2 \tilde{g}_{t+1} + \nu_{t+1}, \qquad \nu_{t+1} \sim \mathcal{N}(0, \sigma_\nu^2), \\ &\Delta d_{t+1} = \overline{g} + \tilde{g}_t + \varepsilon_{d,t+1}, \\ &\tilde{\mu}_{t+1} = \phi_\mu \tilde{\mu}_t + \varepsilon_{\mu,t+1}, \\ &\tilde{g}_{t+1} = \phi_g \tilde{g}_t + \varepsilon_{g,t+1}. \end{aligned} \tag{G.1}$$

The measurement error ν_t is uncorrelated with the innovations of the model for which we assume a general a covariance structure, $\varepsilon_t = (\varepsilon_{d,t}, \varepsilon_{\mu,t}, \varepsilon_{g,t})' \sim \mathcal{N}(0, \Omega)$. Below we discuss the restrictions needed for this model to be identified.

Model (G.1) is equivalent to one estimated by Binsbergen and Koijen (2010), whose identification issue is discussed at lenght in Cochrane (2008), with the key difference being that we have added the measurement error in the $\mathtt{pd}_t$ equation. To be more precise, the specification for Δd_{t+1} implies an ARMA(1,1) process, while the model for $\mathtt{pd}_t$ without the measurement error ν_t is an ARMA(2,1). The resulting bivariate model is a restricted VARMA(2,1) with

five parameters[9] to identify the covariance Ω. Specifically, we set one correlation to zero as in Binsbergen and Koijen (2010) and Rytchkov (2012), i.e. $\text{Corr}(\varepsilon_{d,t}, \varepsilon_{g,t}) = 0$. Adding the measurement error ν_t in the $\texttt{pd}_t$ equation, the additional parameter σ_ν^2 is identified by the additional moving average coefficients.[10] By introducing time-varying long-run mean $\overline{\texttt{pd}}_t$, and $\overline{g}_t$, the implied reduced form models for $\texttt{pd}_t$ and Δd_t become ARIMA(2,1,3) and ARIMA(1,1,2), respectively. Thus, the two additional moving average coefficients are used to identify the two parameters, b_μ and b_g, scaling the score-driven filters for $\overline{\mu}_t$ and $\overline{g}_t$ in (21)-(22). Since our model features a time-varying Ω_t, at each point in time for a given covariance matrix Ω the model is identified; i.e. the model is locally identified.

G.2 Modelling the correlation matrix by partial correlations

Here we show how to model a time-varying correlation matrix by imposing bounds on the partial correlations. In order to save in notation we drop the subscript t. Let consider the following covariance matrix $\Omega = DRD'$, where $D = \text{diag}([\sigma_1, \sigma_2, \sigma_3)]$ and R is the correlation matrix:

$$R = \begin{bmatrix} 1 & \varrho_{12} & \varrho_{13} \\ \varrho_{12} & 1 & \varrho_{23} \\ \varrho_{13} & \varrho_{23} & 1 \end{bmatrix}.$$

To ensure positive standard deviations we model $\delta_i = \log \sigma_i$ so that $\sigma_i = \exp \delta_i$. For the correlations we model $\gamma = (\gamma_{12}, \gamma_{13}, \gamma_{23})'$, where $\gamma_{ij} = h(\varrho_{ij})$ and $h(\cdot)$ is the inverse function of the transformation $\varrho_{ij} = \psi_r(\gamma_{ij})$ that we describe below.

A well defined correlation matrix R must be positive semidefinite with ones on the main diagonal, this poses a non-trivial problem; see e.g. Budden et al. (2008). On the other hand, the one-to-one mapping between the correlations and the partial correlations allows us to impose simple constraints to the partial correlations. Inspired by Joe (2006), Daniels and Pourahmadi (2009) and Lewandowski et al. (2009), we re-parametrize the correlation matrix with respect to the partial correlation matrix. Specifically, R is positive semidefinite if the corresponding partial correlation matrix

$$\Pi = \begin{bmatrix} 1 & \pi_{12} & \pi_{13} \\ \pi_{12} & 1 & \pi_{23} \\ \pi_{13} & \pi_{23} & 1 \end{bmatrix}$$

has all the elements $\pi_{ij} \in (-1, 1)$, where π_{ij} are the partial correlations between variables i and j. To satisfy those bounds on the partial correlations π_{ij} we use the Fisher transformation, that is $\pi_{ij} = \tanh(\gamma_{ij})$, so that we model $\gamma_{ij} = \text{atanh}\pi_{ij}$. The function mapping the elements

[9] The tree autoregressive coefficients and the two constants are identified by construction. The two moving average coefficients and three parameters of the covariance matrix are used to identify the six elements of the matrix Ω.

[10] Adding the measurement error ν_t, the reduced form model for $\texttt{pd}_t$ becomes an ARMA(2,2).

of R into the elements of Π is:[11]

$$\varrho_{12} = \pi_{12}, \quad \varrho_{13} = \pi_{13}, \quad \varrho_{23} = \pi_{23}\sqrt{(1-\pi_{12}^2)(1-\pi_{13}^2)} + \pi_{12}\pi_{13}. \tag{G.2}$$

Thus, we perform two transformations:

$$\varrho_{ij} = \psi_r(\gamma_{ij}) = \psi_{r,2}(\psi_{r,1}(\gamma_{ij})), \tag{G.3}$$

where $\psi_{r,1}(\cdot) = \tanh(\cdot)$, $\psi_{r,2}(\cdot)$ is defined in (G.2). The resulting Jacobian is:

$$\frac{\partial \varrho}{\partial \gamma'} = \frac{\partial \varrho}{\partial \pi'}\frac{\partial \pi}{\partial \gamma'} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \varkappa_{12} & \varkappa_{13} & \varkappa_{23} \end{bmatrix} \begin{bmatrix} 1-\pi_{12}^2 & 0 & 0 \\ 0 & 1-\pi_{13}^2 & 0 \\ 0 & 0 & 1-\pi_{23}^2 \end{bmatrix},$$

where

$$\varkappa_{12} = \pi_{13} - \pi_{12}\pi_{23}\sqrt{\frac{1-\pi_{13}^2}{1-\pi_{12}^2}}, \quad \varkappa_{13} = \pi_{12} - \pi_{13}\pi_{23}\sqrt{\frac{1-\pi_{12}^2}{1-\pi_{13}^2}}, \quad \varkappa_{23} = \sqrt{(1-\pi_{12}^2)(1-\pi_{13}^2)}.$$

Remark: If the partial correlations π_{ij} are bounded using the cosine function, i.e. $\psi_{r,1}(\cdot) = \cos\gamma_{ij}$, the transformation (G.3) turns out to be the same as the *hyperspherical coordinates* used in, e.g., Creal et al. (2011) and Buccheri et al. (2020). This means that the use of hyperspherical coordinates implies modelling inverse consine of the partial correlations. The proof for a correlation matrix of general dimension is beyond the scope of this paper.

In our application, the identification of the model requires to set to zero one of the correlations. Without loss of generality we set to zero the correlation between the first and second innovation. Eploiting the mapping between the correlations and partial correlations we have that $\pi_{12} = 0$ implies $\varrho_{12} = 0$. Therefore, we model the following vectors $\varrho = (\varrho_{13}, \varrho_{23})'$, $\pi = (\pi_{13}, \pi_{23})'$, $\gamma = (\gamma_{13}, \gamma_{23})'$ The mapping between correlations and partial correlation is $\varrho_{13} = \pi_{13}$, $\varrho_{23} = \pi_{23}\sqrt{1-\pi_{13}^2}$, and the Jacobian is

$$\frac{\partial \varrho}{\partial \gamma'} = \sqrt{1-\pi_{13}^2}\begin{bmatrix} \sqrt{1-\pi_{13}^2} & 0 \\ -\pi_{13}\pi_{23} & 1-\pi_{23}^2 \end{bmatrix}.$$

G.3 State space, score driven vector and jacobians

The model in section 3 can be cast easily in state space form:

$$\begin{aligned} y_t &= Z_t\alpha_t + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, H), \\ \alpha_t &= T\alpha_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, Q_t), \end{aligned}$$

[11] See also Yule and Kendall (1965, ch. 12) and Anderson (1984, p. 41).

where

$$y_t = \begin{bmatrix} \Delta d_t \\ \mathtt{pd}_t \end{bmatrix}, \quad Z_t = \begin{bmatrix} \bar{g}_t & 0 & 0 & 1 & 1 & 0 & 0 \\ \overline{\mathtt{pd}}_t & \frac{1}{1-\rho_t\phi_g} & -\frac{1}{1-\rho_t\phi_\mu} & 0 & 0 & 0 & 0 \end{bmatrix}, \quad H = \begin{bmatrix} 0 & 0 \\ 0 & \sigma_\nu^2 \end{bmatrix},$$

$$\alpha_t = \begin{bmatrix} 1 \\ \tilde{g}_t, \\ \tilde{\mu}_t, \\ \tilde{g}_{t-1} \\ \varepsilon_{d,t} \\ \varepsilon_{g,t} \\ \varepsilon_{\mu,t} \end{bmatrix}, \quad T = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \phi_g & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \phi_\mu & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad \eta_t = S_\eta \begin{bmatrix} \varepsilon_{d,t} \\ \varepsilon_{g,t} \\ \varepsilon_{\mu,t} \end{bmatrix}, \quad S_\eta = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

$Q_t = S_\eta \Omega_t S_\eta'$, $\Omega_t = D_t R_t D_t$, with D_t contains the standard deviations, and R_t denotes the correlation matrix. The zero correlation between the measurement error in dividend growth ($\varepsilon_{d,t}$) and the stochastic disturbance in expected dividend growth ($\varepsilon_{g,t}$), which is required for the identification of the model, is appropriately imposed. The resulting matrices are is:

$$D_t = \begin{bmatrix} \sigma_{d,t} & 0 & 0 \\ 0 & \sigma_{g,t} & 0 \\ 0 & 0 & \sigma_{\mu,t} \end{bmatrix}, \quad R_t = \begin{bmatrix} 1 & 0 & \varrho_{d\mu,t} \\ 0 & 1 & \varrho_{g\mu,t} \\ \varrho_{d\mu,t} & \varrho_{g\mu,t} & 1 \end{bmatrix}.$$

The vector of time-varying parameters is:

$$f_t = \begin{bmatrix} \varphi_t \\ \delta_t \\ \gamma_t \end{bmatrix}, \quad \varphi_t = \begin{bmatrix} \overline{\mu}_t, \\ \overline{g}_t \end{bmatrix}, \quad \delta_t = \begin{bmatrix} \log \sigma_{d,t} \\ \log \sigma_{g,t} \\ \log \sigma_{\mu,t} \end{bmatrix}, \quad \gamma_t = \begin{bmatrix} \operatorname{atanh} \pi_{d\mu,t} \\ \operatorname{atanh} \pi_{g\mu,t} \end{bmatrix}.$$

The vector f_t follows the score driven model discussed in section 2, with the following specification of the static parameters:

$$\begin{aligned} \mathrm{c} &= [0, 0, c_{\sigma_d}, c_{\sigma_g}, c_{\sigma_\mu}, c_{\pi_{d,\mu}}, c_{\pi_{g,\mu}}]', \\ \mathrm{A} &= \mathrm{diag}([1, 1, a_{\sigma_d}, a_{\sigma_g}, a_{\sigma_\mu}, a_{\pi_{d,\mu}}, a_{\pi_{g,\mu}}]), \\ \mathrm{B} &= \mathrm{diag}([b_\mu, b_g, b_{\sigma_d}, b_{\sigma_g}, b_{\sigma_\mu}, b_{\pi_{d,\mu}}, b_{\pi_{g,\mu}}]). \end{aligned}$$

Time variation in the Z matrix. Using the notation in section 2.2 we have that

$$\mathrm{vec}(Z_t) = S_{0,z} + S_{1,z}\psi_z\left(S_{2,z} f_t\right),$$

where

$$S_{0,z} = \begin{bmatrix} 0_{6\times 1} \\ 1 \\ 0 \\ 1 \\ 0_{5\times 1} \end{bmatrix}, \quad S_{1,z} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ & & 0_{8\times 4} & \end{bmatrix}, \quad S'_{2,z} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix},$$

$$\psi_z(\varphi_t) = \begin{bmatrix} \overline{g}_t \\ \overline{\mathtt{pd}}_t \\ \frac{1}{1-\rho_t\phi_g} \\ -\frac{1}{1-\rho_t\phi_\mu} \end{bmatrix}, \quad \overline{\mathtt{pd}}_t = \overline{g}_t - \log(\exp\overline{\mu}_t - \exp\overline{g}_t), \quad \rho_t = \frac{\exp\overline{\mathtt{pd}}_t}{1+\exp\overline{\mathtt{pd}}_t}.$$

The Jacobian matrix is:

$$\dot{Z}_t = S_{1,z}\Psi_{z,t}S_{2,z}, \quad \Psi_{z,t} = \begin{bmatrix} 0 & 1 \\ \frac{\partial\overline{\mathtt{pd}}_t}{\partial\overline{\mu}_t} & \frac{\partial\overline{\mathtt{pd}}_t}{\partial\overline{g}_t} \\ \frac{\phi_g}{(1-\phi_g\rho_t)^2}\frac{\partial\rho_t}{\partial\overline{\mu}_t} & \frac{\phi_g}{(1-\phi_g\rho_t)^2}\frac{\partial\rho_t}{\partial\overline{g}_t} \\ -\frac{\phi_\mu}{(1-\phi_\mu\rho_t)^2}\frac{\partial\rho_t}{\partial\overline{\mu}_t} & -\frac{\phi_\mu}{(1-\phi_\mu\rho_t)^2}\frac{\partial\rho_t}{\partial\overline{g}_t} \end{bmatrix}, \quad \begin{array}{l} \frac{\partial\overline{\mathtt{pd}}_t}{\partial\overline{\mu}_t} = -\frac{\exp\overline{\mu}_t}{\exp\overline{\mu}_t-\exp\overline{g}_t}, \\ \frac{\partial\overline{\mathtt{pd}}_t}{\partial\overline{g}_t} = -\frac{\partial\overline{\mathtt{pd}}_t}{\partial\overline{\mu}_t}, \\ \frac{\partial\rho_t}{\partial\overline{\mu}_t} = -\frac{\rho_t(1-\rho_t)\exp\overline{\mu}_t}{\exp\overline{\mu}_t-\exp\overline{g}_t}, \\ \frac{\partial\rho_t}{\partial\overline{g}_t} = -\frac{\partial\rho_t}{\partial\overline{\mu}_t}. \end{array}$$

Time variation in the Q matrix. Recall that the covariance matrix of the transition equation is $Q_t = S_\eta\Omega_t S'_\eta$ where $\Omega_t = D_tR_tD_t$. Using the notation in Section 2.2, and the standard rules of matrix differentiation, we have that:

$$\dot{Q}_t = (S_\eta \otimes S_\eta)\left[(D_tR_t \otimes I + I \otimes D_tR_t)\dot{D}_t + (D_t \otimes D_t)\dot{R}_t\right].$$

We now express the matrices of volatilities and correlations as follows:

$$\mathrm{vec}(D_t) = S_{1,d}\psi_d(S_{2,d}f_t), \quad \mathrm{vec}(R_t) = S_{0,r} + S_{1,r}\psi_r(S_{2,r}f_t),$$

where $S_{1,d}$ $S_{2,d}$, $S_{1,r}$, $S_{2,r}$ are selection matrices

$$S_{1,d} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad S'_{2,d} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad S_{0,r} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \quad S_{1,r} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad S'_{2,r} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

The functions $\psi_d(\delta_t)$ and $\psi_r(\gamma_t)$ and their Jacobians are described in section G.2. Specifically, we have that:

$$\dot{D}_t = S_{1,d}\Psi_{d,t}S_{2,d}, \quad \dot{R}_t = S_{1,r}\Psi_{r,t}S_{2,r},$$

where

$$\Psi_{d,t} = D_t, \quad \Psi_{r,t} = \sqrt{1-\pi^2_{d\mu,t}} \begin{bmatrix} \sqrt{1-\pi^2_{d\mu,t}} & 0 \\ -\pi_{d\mu,t}\pi_{g\mu,t} & 1-\pi^2_{g\mu,t} \end{bmatrix}.$$

H Term structure of expected returns and dividend growth in recessions

In Figure H.1 we plot the whole term structure of expected returns and expected dividend growth for three historical episodes. In particular, we look at the year before the recession, the peak of the recession and the year after the recession. We find that discount rates shocks, especially at the short end of the curve, contributed greatly to the severity of the recessions in 1929 and 2008, while they played a relatively minor role in the 2001 recession episode. These results are consistent with the narrative in Campbell et al. (2013).

Figure H.1: Expected Returns and Dividend Growth: Selected Episodes

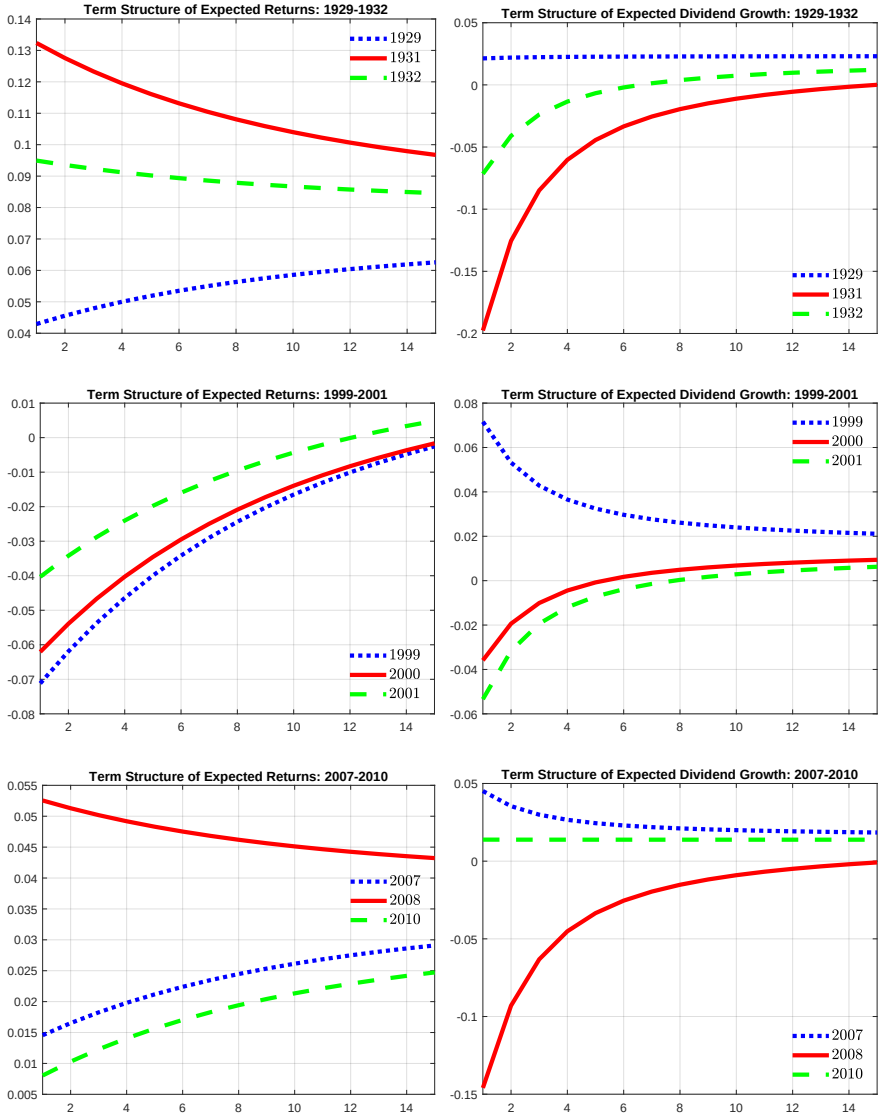

Note. Figure H.1 plots the term structure of expected return (left panel) and dividend growth (right panel) around some specific events. In particular, the upper panel looks at the Great Depression, the middle panel looks at the years around the 2001 recession and the bottom panel looks at the years of the Great Recession.

I Additional Results

Figure I.1: Term Structure of Expected Returns and Dividend Growth

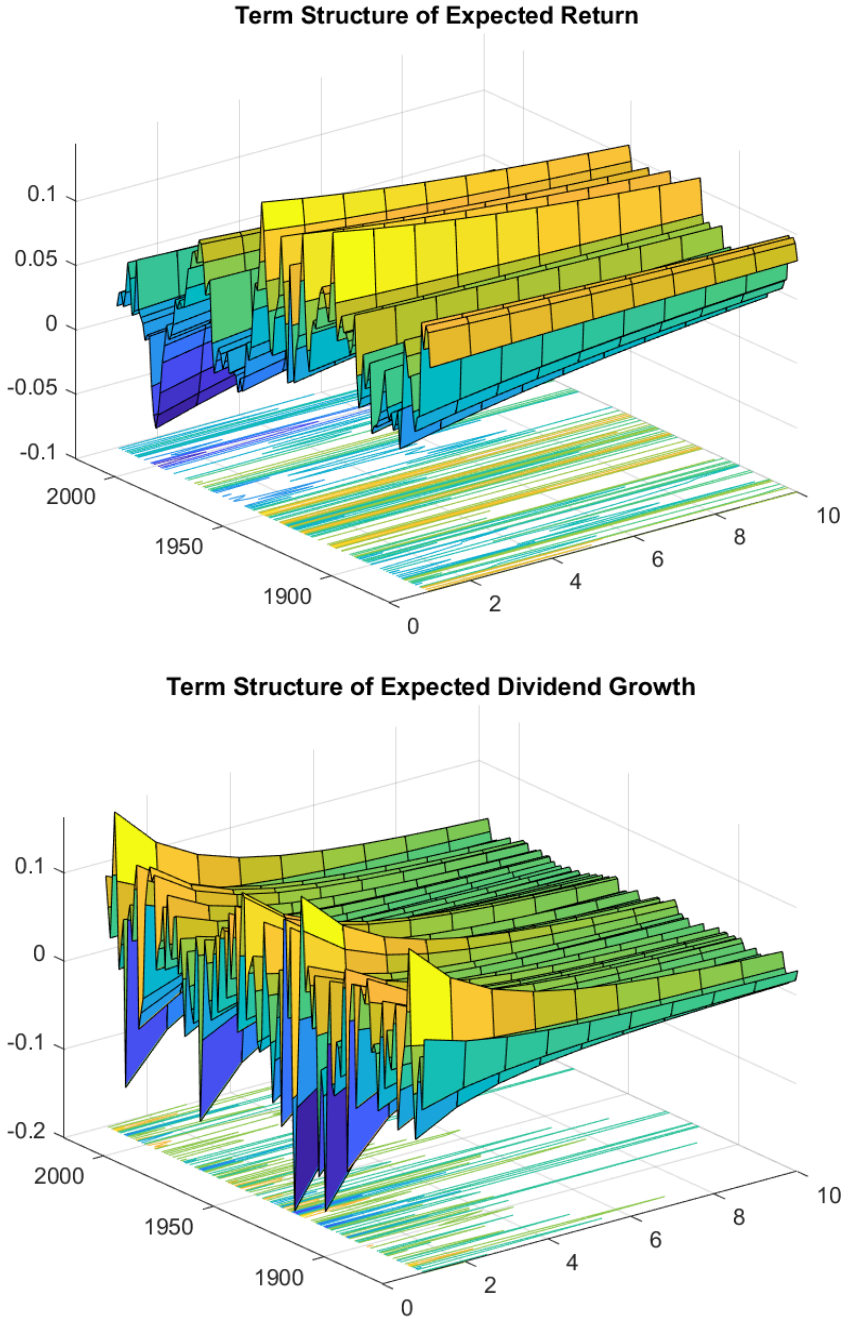

Note. Figure I.1 plots the term structure of expected return and dividend growth.

Figure I.2: Steady State Comparisons

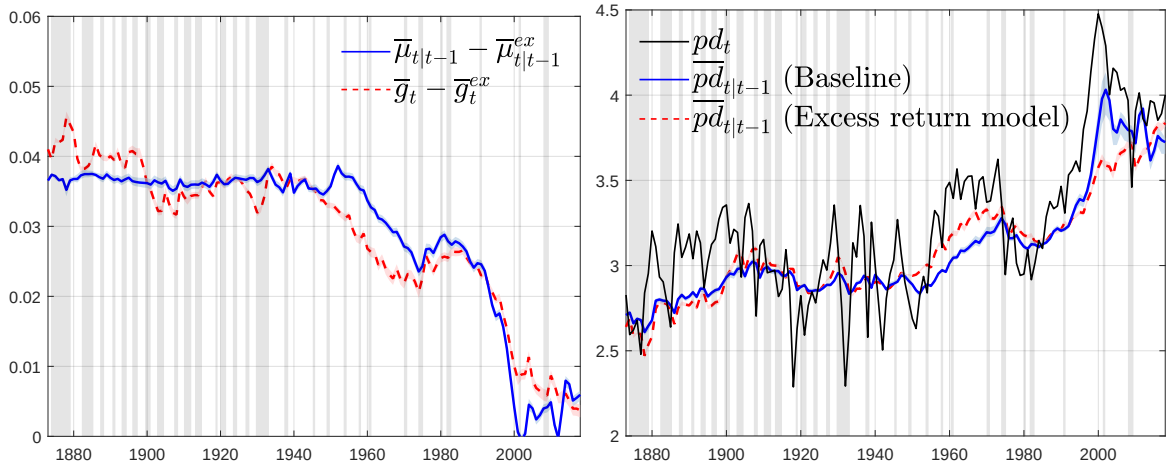

Note. The left panel of Figure I.2 reports two alternative measures of the long-run riskless real rate that we recover from the estimates in section 5. Specifically, $\overline{\mu}_{t|t-1} - \overline{\mu^{ex}}_{t|t-1}$ and $\overline{g}_{t|t-1} - \overline{g}^{ex}_{t|t-1}$. The estimates of $\overline{\mathtt{pd}}_{t|t-1}$ from the two models are reported in the panel on the right together with the (log) level of the price dividend ratio.

Table I.1: Excess Return Model: Estimation Results

ϕ_μ	0.863			b_μ	0.081
	[0.007]				[0.010]
ϕ_g	0.355			b_g	0.053
	[0.011]				[0.007]
$\bar{\sigma}_d$	0.062	a_{σ_d}	0.858	b_{σ_d}	0.016
	[0.061; 0.065]		[0.023]		[0.001]
$\bar{\sigma}_g$	0.097	a_{σ_g}	0.765	b_{σ_g}	0.012
	[0.096; 0.104]		[0.052]		[0.003]
$\bar{\sigma}_\mu$	0.019	a_{σ_μ}	0.847	b_{σ_μ}	0.015
	[0.018; 0.024]		[0.051]		[0.003]
$\bar{\rho}_{d,\mu}$	0.888	$a_{\pi_{d,\mu}}$	0.980	$b_{\pi_{d,\mu}}$	0.025
	[0.660; 0.892]		[0.010]		[0.010]
$\bar{\rho}_{g,\mu}$	-0.001	$a_{\pi_{g,\mu}}$	0.820	$b_{\pi_{g,\mu}}$	0.025
	[-0.026; -0.018]		[0.047]		[0.005]
σ^2_ν	0.008			κ_h	0.020
	[0.0002]				[0.0001]
Log Lik.	322.232				

Note. Table I.1 reports parameter estimates for the model estimated in section 5. First column: autoregressive coefficients of expected returns and expected dividend growth (ϕ_μ and ϕ_g) and average (over the whole sample) estimates of the volatilities ($\bar{\sigma}_d$, $\bar{\sigma}_g$ and $\bar{\sigma}_\mu$) and correlations ($\bar{\rho}_{d,\mu}$ and $\bar{\rho}_{g,\mu}$) that form the matrix Q_t. σ^2_ν is the volatility of the measurement error for the price dividend ratio. The second and third columns show the estimates of the coefficients that enter the law of motion of the score driven time-varying processes (4) where A and B are diagonal matrices, and the smoothing coefficient applied to the Hessian term (κ_h). For each coefficient we report in square brackets the associated standard error. For the average volatilities and correlations in the first column we report the 68% confidence interval from 1000 simulations of the model (calculated as in Blasques et al., 2016).

References

Abadir, K. and Magnus, J. (2005). *Matrix Algebra*. Cambridge University Press, Cambridge, UK.

Anderson, T. W. (1984). *An Introduction to Multivariate Statistical Analysis*. Wiley series in probability and mathematical statistic.

Banbura, M., Giannone, D., Modugno, M., and Reichlin, L. (2013). *Now-Casting and the Real-Time Data Flow*, volume 2 of *Handbook of Economic Forecasting*, pages 195–237. Elsevier.

Binsbergen, J. H. V. and Koijen, R. S. J. (2010). Predictive Regressions: A Present Value Approach. *Journal of Finance*, 65(4):1439–1471.

Blasques, F., Koopman, S. J., Lasak, K., and Lucas, A. (2016). In-sample confidence bands and out-of-sample forecast bands for time-varying parameters in observation-driven models. *International Journal of Forecasting*, 32(3):875–887.

Blasques, F., Koopman, S. J., and Lucas, A. (2014). Optimal Formulations for Nonlinear Autoregressive Processes. Tinbergen Institute Discussion Papers 14-103/III, Tinbergen Institute.

Buccheri, G., Bormetti, G., Corsi, F., and Lillo, F. (2020). A Score-Driven Conditional Correlation Model for Noisy and Asynchronous Data: an Application to High-Frequency Covariance Dynamics. *Journal of Business Economic Statistics*, forthcoming.

Budden, M., Hadavas, P., and Hoffman, L. (2008). On the generation of correlation matrices. *Applied Mathematics E-Notes*, 8:279–282.

Campbell, J. Y., Giglio, S., and Polk, C. (2013). Hard Times. *Review of Asset Pricing Studies*, 3(1):95–132.

Cochrane, J. H. (2008). State-Space vs. VAR Models for Stock Returns. Unpublished manuscript.

Creal, D., Koopman, S. J., and Lucas, A. (2008). A General Framework for Observation Driven Time-Varying Parameter Models. Tinbergen Institute Discussion Papers 08-108/4, Tinbergen Institute.

Creal, D., Koopman, S. J., and Lucas, A. (2011). A Dynamic Multivariate Heavy-Tailed Model for Time-Varying Volatilities and Correlations. *Journal of Business & Economic Statistics*, 29(4):552–563.

Daniels, M. and Pourahmadi, M. (2009). Modeling covariance matrices via partial autocorrelations. *Journal of Multivariate Analysis*, 100(10):2352–2363.

Del Negro, M. (2012). Bayesian Macroeconometrics. In *The Oxford Handbook of Bayesian Econometrics*. Oxford University Press.

Delle Monache, D. and Petrella, I. (2017). Adaptive models and heavy tails with an application to inflation forecasting. *International Journal of Forecasting*, 33(2):482–501.

Doan, T., Litterman, R. B., and Sims, C. A. (1986). Forecasting and conditional projection using realistic prior distribution. Staff Report 93, Federal Reserve Bank of Minneapolis.

Giannone, D., Lenza, M., and Primiceri, G. E. (2019). Priors for the Long Run. *Journal of the American Statistical Association*, 114(526):565–580.

Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA.

Joe, H. (2006). Generating random correlation matrices based on partial correlations. *Journal of Multivariate Analysis*, 97(10):2177 – 2189.

Kapetanios, G., Marcellino, M., and Venditti, F. (2019). Large timevarying parameter VARs: A nonparametric approach. *Journal of Applied Econometrics*, 34(7):1027–1049.

Koop, G. and Korobilis, D. (2013). Large time-varying parameter VARs. *Journal of Econometrics*, 177(2):185–198.

Lewandowski, D., Kurowicka, D., and Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9):1989 – 2001.

Litterman, R. B. (1979). Techniques of forecasting using vector autoregressions. Working Papers 115, Federal Reserve Bank of Minneapolis.

Lucas, A., Opschoor, A., and Schaumburg, J. (2016). Accounting for missing values in score-driven time-varying parameter models. *Economics Letters*, 148(C):96–98.

Piatti, I. and Trojani, F. (2017). Predictable Risks and Predictive Regression in Present-Value Models. Working paper, Said Business School.

Rytchkov, O. (2012). Filtering Out Expected Dividends and Expected Returns. *Quarterly Journal of Finance*, 2(03):1–56.

Shiller, R. J. (1989). *Market Volatility*. MIT Press, Cambridge, MA.

Stock, J. H. and Watson, M. W. (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33.

Theil, H. and Goldberger, A. S. (1961). On pure and mixed statistical estimation in economics. *International Economic Review*, 2(1):65–78.

Yule, G. and Kendall, M. (1965). *An introduction to the theory of statistics*. C. Griffin & Co., Belmont, California; 14th ed.