Oxford Bulletin of Economics and Statistics

ORIGINAL ARTICLE OPEN ACCESS

The Role of Consumer Sentiment in the Stock Market: A Multivariate Dynamic Mixture Model With Threshold Effects

Zacharias Psaradakis[1] | Francisco Rapetti[2] | Martin Sola[2] | Patricio Yunis[3]

[1]Birkbeck Business School, Birkbeck, University of London, London, UK | [2]Department of Economics, Universidad Torcuato di Tella, Buenos Aires, Argentina | [3]Department of Economics, University of Pennsylvania, Philadelphia, USA

Correspondence: Zacharias Psaradakis (z.psaradakis@bbk.ac.uk)

Received: 27 April 2024 | **Revised:** 9 April 2025 | **Accepted:** 9 April 2025

Keywords: consumer sentiment | mixture models | price–dividend ratio | threshold | time-varying mixing weights | volatility

ABSTRACT

We consider the relationship between stock prices, volatility and consumer sentiment. The analysis is based on a new multivariate model defined as a time-varying mixture of dynamic models in which contemporaneous relationships among variables are allowed and the mixing weights have a threshold-type structure. We discuss issues related to the stability of the model and the estimation of its parameters. Our empirical results show that consumer sentiment significantly affects the S&P 500 price–dividend ratio and market volatility in at least one of the model's two regimes, which are associated with endogenously determined low and high consumer sentiment.

JEL Classification: C32, C51, G12

1 | Introduction

Whether consumer and/or investor sentiment have predictive power for macroeconomic and financial variables, or their effects are already incorporated in such variables, are questions of long-standing interest in the economics and finance literature. Many studies have documented significant impacts of measures of sentiment on stock prices, consumer spending and economic growth (see, inter alia, Ludvigson [1], Baker and Wurgler [2], Schmeling [3], Stambaugh et al. [4], Shen et al. [5]). There is also evidence that the impact of sentiment is asymmetric, in the sense that sentiment affects economic and financial variables in different ways during optimistic 'high-sentiment' periods and pessimistic 'low-sentiment' periods (see, e.g., Desroches and Gosselin [6], Chen [7], Shen et al. [5]).

Empirical studies of the role of sentiment typically rely on econometric models in which measures of sentiment are included as covariates, the marginal effect of which on the response variable of interest (e.g., stock prices) may be assessed conditionally on other relevant covariates (e.g., market volatility). Such models are often static specifications, a rather unsatisfactory choice given that interactions between the response variable and the covariates are unlikely to be solely contemporaneous. Correct dynamic specification, therefore, in the absence of which the interpretation of the model and statistical inferences drawn from it are problematic, typically necessitates the use of autoregressive and/or distributed lag structures.

Beyond these considerations, the possibility that the parameters of the model may not be the same in high-sentiment and low-sentiment periods adds to the appeal of multiple-regime specifications in which different regimes represent states of nature associated with different sentiment levels. Within such a framework, one may think of the regime prevailing at each point in time as being chosen with a probability that is determined as

https://doi.org/10.1111/obes.12682

a function of some appropriate set of variables contained in the conditioning set relative to which the conditional distribution of the response variable is defined. In consequence, the conditional distribution in question takes the form of a mixture of distributions with time-varying mixing weights. Although there are various possible candidates for the probability-determining variables (e.g., those that optimise a suitable goodness-of-fit criterion), past values of a measure of sentiment are a natural choice when sentiment is viewed as the factor with which different regimes are associated.

Nevertheless, considerable care needs to be taken when estimating the parameters of models in which the regime-determining probabilities are functions of variables other than lagged responses. As Pouzo et al. [8] demonstrated, the common practice of relying on limited-information methods which ignore the generating mechanism of the variables that drive the evolution of the regime-determining probabilities is flawed and likely to result in biased and inconsistent parameter estimates. Valid statistical inference in such settings typically requires joint modelling of all relevant observable variables.

Our approach in this paper aims to address all three important modelling and inferential issues discussed in the previous paragraphs: correct dynamic specification, identification of economically interpretable regimes, and consistent and precise estimation of unknown parameters. Our starting point is the observation that relationships among variables such as the stock price–dividend ratio (or stock returns), market volatility and consumer/investor sentiment may reasonably be expected to be different in high-sentiment and low-sentiment periods. Therefore, we argue in favour of analysing such relationships using a nonlinear dynamic model in which regimes associated with low and high sentiment are determined endogenously, depending on the probability that some unobservable threshold value for a sentiment measure is exceeded.

Importantly, we adopt an explicitly multivariate approach in which all observable variables of interest are jointly modelled. More specifically, we use a multivariate mixture of autoregressive distributed lag models in which the mixing weights are time-varying and have a threshold structure similar to that of Dueker et al. [9]. This means that the regime (associated with a component of the mixture) that prevails at each date is determined by the (conditional) probability that the values of a regime-determining variable exceed or are smaller than an unspecified threshold value. The regime-determining variable in our setting, with which economically interpretable regimes can be associated, is a measure of consumer sentiment, thus distinguishing between periods of high and low sentiment.

The proposed multivariate model is a conditional model, in the sense that it allows for contemporaneous relationships among the variables of interest, in addition to dynamic interactions, but it is not a partial model since none of the variables is treated as unmodelled/exogenous. It may be viewed as either a recursively identified dynamic structural model or a by-product of conditioning operations in a multivariate mixture autoregressive model with parameters which specify a recursive Granger-causal ordering of the variables of interest. Importantly, the conditional specification has a threshold structure that is economically meaningful, in the sense that the regimes associated with the components of the mixture are directly related to high and low consumer sentiment. Moreover, the parameterisation of the model allows inferences about the short-run and long-run role of sentiment in the stock market to be made.

We discuss how local stability of the model may be analysed and consider estimation of unknown parameters by the method of maximum likelihood (ML). A simulation study demonstrates that likelihood-based inferential procedures are accurate in samples of sizes that are typical in many applications, provided a complete model for all observable stochastic variables is specified. Treating one of the variables as exogenous and specifying a partial model for the generating mechanism of the remaining variables conditional on the unmodelled variable (as is frequently done in the empirical setting of interest here) is costly. Such an approach results in severely biased parameter estimates and hypothesis tests which reject incorrectly with unacceptably high probability. What is more, these difficulties are likely to be more pronounced the larger is the sample size. Our results parallel those of Pouzo et al. [8], who investigated the consequences of model misspecification in a general multiple-regimes setting (which includes many types of mixture autoregressive models as special cases). As noted earlier, their findings also highlighted the importance of using full-information estimation methods when the variables that determine the probability of regime changes are not strictly exogenous.

Finally, the conditional mixture model is used to examine the low-frequency relationships between U.S. stock prices, market volatility and consumer sentiment. More specifically, we consider how the relationship between the S&P 500 price–dividend ratio and stock market volatility may be affected by different economic conditions characterised by high and low consumer sentiment, as measured by the well-known Index of Consumer Sentiment published by the University of Michigan. Our findings suggest that consumer sentiment has significant predictive power for both stock market volatility and the price–dividend ratio—while consumer sentiment affects market volatility when the economic outlook is pessimistic, it has a significant impact on the price–dividend ratio during high-sentiment periods. By contrast, consumer sentiment is found to have no significant effect on the price–dividend ratio in a model that treats market volatility as an unmodelled variable, highlighting the pitfalls of relying on partial models.

The rest of the paper is organised as follows. The basic idea behind multivariate models with mixture autoregressive dynamics is briefly discussed in Section 2. Section 3 introduces a multivariate dynamic conditional model with a threshold-type, time-varying mixture structure, examines its stability properties, and considers ML estimation of its parameters. Section 4 reports the results of simulation experiments that assess the finite-sample properties of the ML estimator and of related statistics in the complete model and in a partial model that treats one variable as unmodelled. Section 5 investigates the relationship between the price–dividend ratio in the U.S. stock market, a measure of the volatility of the market, and consumer sentiment. Section 6 summarises and concludes.

2 | Mixture Autoregressive Models

To convey the basic idea behind models with mixture autoregressive dynamics, consider a two-component, multivariate mixture autoregressive model for a d-dimensional time series $\{\boldsymbol{w}_t, t = 1, 2, \dots\}$, $d \geqslant 1$, formulated as

$$\boldsymbol{w}_t = \begin{cases} \boldsymbol{w}_{0,t}, & \text{with probability } \Lambda(\overline{\boldsymbol{w}}_{t-1}) =: \Lambda_t, \\ \boldsymbol{w}_{1,t}, & \text{with probability } 1 - \Lambda_t, \end{cases} \tag{1}$$

with

$$\boldsymbol{w}_{i,t} := \boldsymbol{c}_i + \sum_{j=1}^{p} \boldsymbol{A}_{i,j} \boldsymbol{w}_{t-j} + \boldsymbol{\Sigma}_i^{1/2} \boldsymbol{\varepsilon}_t, \; i = 0, 1, \tag{2}$$

and $\overline{\boldsymbol{w}}_{t-1} := (\boldsymbol{w}'_{t-1}, \dots, \boldsymbol{w}'_{t-p})'$, for some integer $p \geqslant 1$. Here, Λ is a real-valued function on the pd-dimensional real space $\mathbb{R}^{pd}$ whose range is contained in the interval $[0, 1]$, $\{\boldsymbol{\varepsilon}_t\}$ are independent, identically distributed d-dimensional random vectors with zero mean and identity covariance matrix, with $\boldsymbol{\varepsilon}_t$ independent of $\{\boldsymbol{w}_s, s < t\}$ for all t, and $\boldsymbol{c}_i$, $\boldsymbol{\Sigma}_i$ and $\boldsymbol{A}_{i,j}$ $(i = 0, 1; j = 1, \dots, p)$ are fixed parameters (with $\boldsymbol{\Sigma}_i$ symmetric and positive definite).[1] Hence, the conditional distribution function of $\boldsymbol{w}_t$ given $\{\boldsymbol{w}_s, s < t\}$ is the mixture $\Lambda_t F_\varepsilon(\boldsymbol{\Sigma}_0^{-1/2}\{\boldsymbol{w}_t - \boldsymbol{m}_{0,t}\}) + (1 - \Lambda_t) F_\varepsilon(\boldsymbol{\Sigma}_1^{-1/2}\{\boldsymbol{w}_t - \boldsymbol{m}_{1,t}\})$, where F_ε is the distribution function of $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{m}_{i,t} := \boldsymbol{c}_i + \sum_{j=1}^{p} \boldsymbol{A}_{i,j} \boldsymbol{w}_{t-j}$, $i = 0, 1$.

Various mixture autoregressive models can be obtained by different choices of the time-varying mixing weights Λ_t. A class of weights that is particularly relevant for our purposes is that considered in Dueker et al. [9] and Dueker et al. [10]. Such mixing weights imply that, at each date t, the probability that determines which of the two autoregressive components is chosen is given by the normalised conditional probability (given $\overline{\boldsymbol{w}}_{t-1}$) that the latent components $\boldsymbol{w}_{0,t}$ and $\boldsymbol{w}_{1,t}$ are below/above some threshold values. For example, assuming the noise $\boldsymbol{\varepsilon}_t$ is Gaussian, the mixing weights may be specified as

$$\Lambda_t = \frac{\Phi_d(\boldsymbol{\Sigma}_0^{-1/2}\{\boldsymbol{w}^* - \boldsymbol{m}_{0,t}\})}{\Phi_d(\boldsymbol{\Sigma}_0^{-1/2}\{\boldsymbol{w}^* - \boldsymbol{m}_{0,t}\}) + [1 - \Phi_d(\boldsymbol{\Sigma}_1^{-1/2}\{\boldsymbol{w}^* - \boldsymbol{m}_{1,t}\})]}, \tag{3}$$

where $\boldsymbol{w}^*$ is a d-dimensional location parameter and Φ_d is the distribution function of the d-variate normal distribution with zero mean and identity covariance matrix.[2] Another possibility is to define Λ_t in terms of the conditional probability (given $\overline{\boldsymbol{w}}_{t-1}$) of a linear function $\boldsymbol{a}'\boldsymbol{w}_{i,t}$ taking values smaller/greater than $\boldsymbol{a}'\boldsymbol{w}^*$, for some d-dimensional nonzero vector $\boldsymbol{a}$, so that

$$\Lambda_t = \frac{\Phi(\boldsymbol{a}'\{\boldsymbol{w}^* - \boldsymbol{m}_{0,t}\}/\{\boldsymbol{a}'\boldsymbol{\Sigma}_0\boldsymbol{a}\}^{1/2})}{\Phi(\boldsymbol{a}'\{\boldsymbol{w}^* - \boldsymbol{m}_{0,t}\}/\{\boldsymbol{a}'\boldsymbol{\Sigma}_0\boldsymbol{a}\}^{1/2}) + [1 - \Phi(\boldsymbol{a}'\{\boldsymbol{w}^* - \boldsymbol{m}_{1,t}\}/\{\boldsymbol{a}'\boldsymbol{\Sigma}_1\boldsymbol{a}\}^{1/2})]}, \tag{4}$$

with $\Phi := \Phi_1$. When $d \geqslant 2$, more general threshold structures that yield mixture models with d^2 components may also be considered, as in Dueker et al. [10]. Specifications of this type are well-suited to applications in which the probability of a state of nature (associated with a component of the mixture) prevailing depends on the values of the variables being modelled relative to unspecified threshold values. A threshold structure for time-varying mixing weights is also at the centre of the conditional mixture model that will be discussed in the next section.

Other time-varying mixing weights Λ_t that have been used in the literature include, among others: $\Lambda_t = [1 + \exp(-h(\overline{\boldsymbol{w}}_{t-1}))]^{-1}$ (for $d = 1$), where h is a real-valued affine function on $\mathbb{R}^p$ (Wong and Li [12]); $\Lambda_t = [1 + \exp(-\tilde{h}(|\overline{\boldsymbol{w}}_{t-1}|^{1/2}))]^{-1}$, where $\tilde{h}$ is an increasing, real-valued affine function on $\mathbb{R}$ and $|\overline{\boldsymbol{w}}_{t-1}|$ is the one-norm of $\overline{\boldsymbol{w}}_{t-1}$ (Bec et al. [13]); $\Lambda_t = \lambda\zeta_{pd,0}(\overline{\boldsymbol{w}}_{t-1})/[\lambda\zeta_{pd,0}(\overline{\boldsymbol{w}}_{t-1}) + (1 - \lambda)\zeta_{pd,1}(\overline{\boldsymbol{w}}_{t-1})]$, for some $0 < \lambda < 1$, with $\zeta_{pd,i}$, $i = 0, 1$, being the density function of a pd-dimensional random vector $(\boldsymbol{v}'_{i,t}, \dots, \boldsymbol{v}'_{i,t-p+1})'$ such that $\{\boldsymbol{v}^*_{i,t}\}$ obeys a d-variate, causal, pth-order autoregressive model with intercept $\boldsymbol{c}_i$, coefficients $\boldsymbol{A}_{i,1}, \dots, \boldsymbol{A}_{i,p}$, and Gaussian noise having mean zero and covariance matrix $\boldsymbol{\Sigma}_i$ (Kalliovirta et al. [14]). Unlike specifications such as (3) and (4), or variations thereof, these mixing weights are not consistent with the threshold interpretation that is natural in our context.

3 | Dynamic Mixture Model With Threshold-Type Weights

Motivated by the empirical setting highlighted in Section 1, involving the relationship between stock prices, volatility and consumer sentiment, we consider in the sequel a trivariate mixture model with threshold-type mixing weights analogous to those described in the previous section. Unlike conventional mixture autoregressive models, our model also allows for contemporaneous interactions among the three variables and takes the form of a mixture autoregressive distributed lag model. To the best of our knowledge, such models with threshold-type mixing weights have not been considered in the literature before.[3]

We begin by explaining how such a dynamic conditional system can be obtained from a trivariate mixture autoregressive model, the coefficients of which satisfy appropriate identification restrictions, and discuss a specification of the mixing weights which is natural and economically interpretable in the context of our empirical setting. The stability properties of the multivariate mixture model are then investigated and estimation of the parameters of the model is considered. For the sake of clarity and for ease of presentation, we focus on a model with a first-order dynamic structure, but higher-order dynamics can be accommodated in a straightforward manner.

3.1 | Model

Consider a trivariate time series $\{\boldsymbol{w}_t = (y_t, x_t, z_t)', t \geqslant 1\}$ satisfying a mixture autoregressive model of the form

$$\boldsymbol{w}_t = \begin{cases} \boldsymbol{c}_0 + \boldsymbol{A}_0\boldsymbol{w}_{t-1} + \boldsymbol{\Sigma}_0^{1/2}\boldsymbol{\varepsilon}_t, & \text{with probability } G_t, \\ \boldsymbol{c}_1 + \boldsymbol{A}_1\boldsymbol{w}_{t-1} + \boldsymbol{\Sigma}_1^{1/2}\boldsymbol{\varepsilon}_t, & \text{with probability } 1 - G_t, \end{cases} \tag{5}$$

where $\{\boldsymbol{\varepsilon}_t\}$ are independent Gaussian random vectors with zero mean and identity covariance matrix, $\boldsymbol{c}_0$ and $\boldsymbol{c}_1$ are vectors of intercepts, $\boldsymbol{A}_0$ and $\boldsymbol{A}_1$ are upper triangular coefficient matrices, $\boldsymbol{\Sigma}_0$ and $\boldsymbol{\Sigma}_1$ are symmetric, positive definite matrices, and G_t is a continuous function (to be specified later) of one or more of the components of $\boldsymbol{w}_{t-1}$ that takes values in $[0, 1]$. The restrictions implied by the triangularity of $\boldsymbol{A}_0$ and $\boldsymbol{A}_1$ may be viewed as identifying restrictions that specify a recursive Granger-causal ordering of the three variables. No triangularity or diagonality requirements are imposed on $\boldsymbol{\Sigma}_0$ and $\boldsymbol{\Sigma}_1$.

By a standard conditioning argument applied to each of the two components of (5), a recursive, multivariate conditional mixture model may be obtained. This is defined by

$$\boldsymbol{w}_t = \begin{cases} (y_{0,t}, x_{0,t}, z_{0,t})', & \text{with probability } G_t, \\ (y_{1,t}, x_{1,t}, z_{1,t})', & \text{with probability } 1 - G_t, \end{cases} \tag{6}$$

where, for $i = 0, 1$,

$$y_{i,t} := \mu_y^i + \beta_{yx}^i x_t + \beta_{yz}^i z_t + \phi_{yy}^i y_{t-1} + \phi_{yx}^i x_{t-1} + \phi_{yz}^i z_{t-1} + \sigma_\eta^i \eta_t, \tag{7}$$

$$x_{i,t} := \mu_x^i + \beta_{xz}^i z_t + \phi_{xx}^i x_{t-1} + \phi_{xz}^i z_{t-1} + \sigma_v^i v_t, \tag{8}$$

$$z_{i,t} := \mu_z^i + \phi_z^i z_{t-1} + \sigma_z^i \xi_t. \tag{9}$$

Here, $\{\boldsymbol{e}_t := (\eta_t, v_t, \xi_t)'\}$ are independent Gaussian random vectors, with zero mean and identity covariance matrix, such that $\boldsymbol{e}_t$ is independent of $\{\boldsymbol{w}_s, s < t\}$ for all t, and the parameter $\boldsymbol{\theta}_i := (\mu_y^i, \beta_{yx}^i, \beta_{yz}^i, \phi_{yy}^i, \phi_{yx}^i, \phi_{yz}^i, \sigma_\eta^i, \mu_x^i, \beta_{xz}^i, \phi_{xx}^i, \phi_{xz}^i, \sigma_v^i, \mu_z^i, \phi_z^i, \sigma_z^i)'$, with $\sigma_\eta^i, \sigma_v^i, \sigma_z^i > 0$, is a function of the parameters of (5).[4]

The model is completed by specifying the time-varying mixing weights G_t as

$$G_t = \frac{\Phi(\{z^* - \mu_z^0 - \phi_z^0 z_{t-1}\}/\sigma_z^0)}{\Phi(\{z^* - \mu_z^0 - \phi_z^0 z_{t-1}\}/\sigma_z^0) + [1 - \Phi(\{z^* - \mu_z^1 - \phi_z^1 z_{t-1}\}/\sigma_z^1)]}, \tag{10}$$

where z^* is an (unknown) location parameter. Hence, at any date t, there are two possible regimes corresponding to the components of the model governed by the parameters $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$. The mixing weights G_t and $1 - G_t$ represent the (conditional) probability of $\boldsymbol{w}_t$ being generated by the component governed by $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$, respectively, and are such that

$$G_t = \frac{\Pr(z_{0,t} \leqslant z^* | z_{t-1})}{\Pr(z_{0,t} \leqslant z^* | z_{t-1}) + \Pr(z_{1,t} > z^* | z_{t-1})},$$

and

$$1 - G_t = \frac{\Pr(z_{1,t} > z^* | z_{t-1})}{\Pr(z_{0,t} \leqslant z^* | z_{t-1}) + \Pr(z_{1,t} > z^* | z_{t-1})}.$$

The parameter z^* acts, therefore, as a threshold, in the sense that the regime governed by $\boldsymbol{\theta}_0$ ($\boldsymbol{\theta}_1$) is more (less) likely to prevail at each date t when, conditional on the value of z_{t-1}, the probability of the value of z_t associated with $(\mu_z^0, \phi_z^0, \sigma_z^0)$ not exceeding z^* is high, or the probability of the value of z_t associated with $(\mu_z^1, \phi_z^1, \sigma_z^1)$ exceeding z^* is low.

It is important to note that the specification of the mixing weights in (10) as a function of only z_{t-1} is motivated by the observation that in many empirical settings, including the one considered in Section 5, economically interpretable regimes can be characterised in terms of the conditional probability that the implied value of one of the variables under consideration is above/below some threshold value. This is a useful way of addressing the difficult problem of regime identification in multivariate mixture models with threshold-type mixing weights that may potentially depend on many variables. In such models, it is desirable to define mixture components (regimes) in terms of threshold structures which can be justified on economic grounds and which, given the data constraints often encountered in practice, imply a computationally manageable number of components. Failing to do so in large-dimensional mixture models can give rise to results that lack a coherent economic interpretation as well as robustness. In our empirical setting, it is natural to consider consumer sentiment as the sole driving variable for the time-varying mixing weights. The two components of the mixture are, correspondingly, associated with high-sentiment and low-sentiment regimes.[5] More generally, a sensible strategy will be to rely on economic reasoning to determine what an appropriate component/threshold structure may be in the context of the specific modelling problem under consideration.

Another point to note is that, in the recursive dynamic system defined in (6) to (10), the equations for y_t and x_t have a mixture autoregressive distributed lag structure, while the equation for z_t is a pure mixture autoregressive scheme. This structure is the result of assuming that the conditional distribution of $\boldsymbol{w}_t$ given $\boldsymbol{w}_{t-1}$ is that implied by a trivariate mixture autoregressive model with triangular coefficient matrices. Alternatively, however, the system (6) to (10) could be directly specified and viewed as a recursively identified dynamic structural model. In the context of the empirical setting discussed in Sections 1 and 5 (with y_t, x_t and z_t representing the stock price–dividend ratio, stock market volatility and a measure of consumer sentiment, respectively), the parametrization of the model is such that the two components are economically interpretable (as high and low sentiment regimes). The structural model could also be expressed in reduced form as a mixture autoregressive model like (5).

3.2 | Stability

In view of the recursive structure of the model and of the fact that the mixing weights depend solely on z_{t-1}, the stability of the model may be analysed by considering the stability of the dynamics of $\{z_t\}$ first. Then, under suitable conditions, it is reasonable to expect local stability of the trivariate model to be inherited from that of $\{z_t\}$ around some point of long-run equilibrium.

Local stability of the generating mechanism of $\{z_t\}$ may be examined by considering a noiseless version of it, often called the skeleton (see Tong [15]). The latter is defined here as

$$z_{t+1} = K(z_t), \quad t \geqslant 0,$$

where K is the function on $\mathbb{R}$ defined by

$$K(u) := G(u)(\mu_z^0 + \phi_z^0 u) + \{1 - G(u)\}(\mu_z^1 + \phi_z^1 u), \quad u \in \mathbb{R},$$

with

$$G(u) := \frac{\Phi(\{z^* - \mu_z^0 - \phi_z^0 u\}/\sigma_z^0)}{\Phi(\{z^* - \mu_z^0 - \phi_z^0 u\}/\sigma_z^0) + [1 - \Phi(\{z^* - \mu_z^1 - \phi_z^1 u\}/\sigma_z^1)]}.$$

A fixed point of K, said to be an equilibrium (or stationary) point for $\{z_t\}$, is any real number z_e such that $z_e = K(z_e)$. Hence, as discussed in Dueker et al. [9] and Dueker et al. [11], local stability at a fixed point z_e may be assessed by considering the first-order Taylor expansion of K about z_e, that is,

$$K(u) = z_e + \dot{K}(z_e)(u - z_e) + \mathcal{O}\left(\left(u - z_e\right)^2\right),$$

where

$$\dot{K}(z_e) := \frac{dK}{du}(z_e) = \phi_z^1 + (\phi_z^0 - \phi_z^1)G(z_e) + \{\mu_z^0 - \mu_z^1 + (\phi_z^0 - \phi_z^1)z_e\}\dot{G}(z_e),$$

$$\dot{G}(z_e) := \frac{dG}{du}(z_e) = -\frac{\{\phi_z^0\varphi(\tau_0)[1-\Phi(\tau_1)]/\sigma_z^0\} + \{\phi_z^1\varphi(\tau_1)\Phi(\tau_0)/\sigma_z^1\}}{\{\Phi(\tau_0) + 1 - \Phi(\tau_1)\}^2},$$

$$\tau_i := (z^* - \mu_z^i - \phi_z^i z_e)/\sigma_z^i,\ i = 0, 1,$$

φ is the standard normal density function, and $\mathcal{O}((u - z_e)^2)$ is a quantity that is asymptotically bounded by $(u - z_e)^2$ in magnitude as u tends to z_e. If $|\dot{K}(z_e)| < 1$, then z_e is an attracting fixed point for K and the equilibrium at z_e is locally stable.

Next, given an equilibrium point z_e, the implied equilibrium values x_e and y_e for $\{x_t\}$ and $\{y_t\}$, respectively, can be obtained by exploiting the recursive structure of the system. More specifically, in view of (6) to (10), we have

$$\begin{bmatrix} b_{11} & b_{12} & b_{13} \\ 0 & b_{22} & b_{23} \\ 0 & 0 & b_{33} \end{bmatrix} \begin{bmatrix} y_e \\ x_e \\ z_e \end{bmatrix} = \begin{bmatrix} \mu_y^0 G(z_e) + \mu_y^1\{1 - G(z_e)\} \\ \mu_x^0 G(z_e) + \mu_x^1\{1 - G(z_e)\} \\ \mu_z^0 G(z_e) + \mu_z^1\{1 - G(z_e)\} \end{bmatrix},$$

where

$$\begin{aligned}
b_{11} &:= 1 - \phi_{yy}^0 G(z_e) - \phi_{yy}^1\{1 - G(z_e)\}, \\
b_{12} &:= -(\beta_{yx}^0 + \phi_{yx}^0)G(z_e) - (\beta_{yx}^1 + \phi_{yx}^1)\{1 - G(z_e)\}, \\
b_{13} &:= -(\beta_{yz}^0 + \phi_{yz}^0)G(z_e) - (\beta_{yz}^1 + \phi_{yz}^1)\{1 - G(z_e)\}, \\
b_{22} &:= 1 - \phi_{xx}^0 G(z_e) - \phi_{xx}^1\{1 - G(z_e)\}, \\
b_{23} &:= -(\beta_{xz}^0 + \phi_{xz}^0)G(z_e) - (\beta_{xz}^1 + \phi_{xz}^1)\{1 - G(z_e)\}, \\
b_{33} &:= 1 - \phi_z^0 G(z_e) - \phi_z^1\{1 - G(z_e)\}.
\end{aligned}$$

Consequently,

$$x_e = \frac{\mu_x^0 G(z_e) + \mu_x^1\{1 - G(z_e)\} + [(\beta_z^0 + \phi_{xz}^0)G(z_e) + (\beta_z^1 + \phi_{xz}^1)\{1 - G(z_e)\}]z_e}{1 - \phi_{xx}^0 G(z_e) - \phi_{xx}^1\{1 - G(z_e)\}},$$

and

$$y_e = \frac{\mu_y^0 G(z_e) + \mu_y^1\{1 - G(z_e)\} + [(\beta_{yx}^0 + \phi_{yx}^0)G(z_e) + (\beta_{yx}^1 + \phi_{yx}^1)\{1 - G(z_e)\}]x_e}{1 - \phi_{yy}^0 G(z_e) - \phi_{yy}^1\{1 - G(z_e)\}} + \frac{[(\beta_{yz}^0 + \phi_{yz}^0)G(z_e) + (\beta_{yz}^1 + \phi_{yz}^1)\{1 - G(z_e)\}]z_e}{1 - \phi_{yy}^0 G(z_e) - \phi_{yy}^1\{1 - G(z_e)\}},$$

provided $|b_{11}b_{22}b_{33}| > 0$. The equilibrium point (y_e, x_e, z_e) is locally stable whenever $|\dot{K}(z_e)| < 1$, $|\phi_{yy}^0 G(z_e) + \phi_{yy}^1\{1 - G(z_e)\}| < 1$ and $|\phi_{xx}^0 G(z_e) + \phi_{xx}^1\{1 - G(z_e)\}| < 1$.

As a simple numerical illustration of the preceding analysis, consider the model defined by equations (6) to (10) with the parameter values contained in Table 1. (Here and in all subsequent tables, the headings 'Below Threshold' and 'Above Threshold' refer to the regimes governed by $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$, respectively.) The values of the parameters are chosen so as to obtain a clear separation of the two regimes and a unique attracting fixed point

TABLE 1 | Parameter values used in the stability and Monte Carlo exercises: ϕ_{ab}^i (β_{ab}^i) is the regime-i coefficient on the lagged (contemporaneous) variable a in the equation for variable b; ϕ_b^i, μ_b^i and σ_b^i are the regime-i autoregressive coefficient, intercept and noise standard deviation, respectively, in the equation for variable b.

Below threshold		**Above threshold**	
Equation for y_t			
μ_y^0	−0.70	μ_y^1	0.80
β_{yx}^0	−0.70	β_{yx}^1	0.50
β_{yz}^0	−0.78	β_{yz}^1	0.40
ϕ_{yy}^0	0.39	ϕ_{yy}^1	0.27
ϕ_{yx}^0	0.47	ϕ_{yx}^1	0.29
ϕ_{yz}^0	0.73	ϕ_{yz}^1	−0.19
σ_η^0	1.08	σ_η^1	0.55
Equation for x_t			
μ_x^0	0.50	μ_x^1	0.10
β_{xz}^0	−0.35	β_{xz}^1	0.51
ϕ_{xx}^0	0.08	ϕ_{xx}^1	0.09
ϕ_{xz}^0	0.14	ϕ_{xz}^1	0.03
σ_υ^0	0.32	σ_υ^1	0.07
Equation for z_t			
μ_z^0	−0.10	μ_z^1	0.40
ϕ_z^0	0.77	ϕ_z^1	0.67
σ_z^0	0.27	μ_z^1	0.22
Threshold			
z^*	−0.22		

for the skeleton of $\{z_t\}$. In this case, there is a unique equilibrium point $(y_e, x_e, z_e) = (0.749, 0.153, 0.029)$, which is a point of local stability since $\dot{K}(z_e) = 0.93$, $\phi_{yy}^0 G(z_e) + \phi_{yy}^1\{1 - G(z_e)\} = 0.30$ and $\phi_{xx}^0 G(z_e) + \phi_{xx}^1\{1 - G(z_e)\} = 0.09$. We shall consider this parameter configuration further in the simulation experiments in Section 4.

Lastly, a few remarks on the global stability/stationarity of the model are worth making. Noting that the noise ξ_t in (9) is Gaussian and the mixing weights G_t in (10) are a continuous, positive function of z_{t-1}, $\{z_t\}$ can be shown to be a V-uniformly ergodic (and hence geometrically ergodic) Markov chain on $\mathbb{R}$, for the function V defined by $V(u) := 1 + |u|^r$ for $u \in \mathbb{R}$ and any $r > 1$, provided $\max\{|\phi_z^0|, |\phi_z^1|\} < 1$; furthermore, $\{z_t\}$ has a unique stationary distribution with finite moments of all orders (cf. Carvalho and Skoulakis [16]).[6] Therefore, if the distribution of the initial value z_0 is the stationary distribution, then $\{z_t\}$ is a (globally) stationary process. Interestingly, however, results in Bec et al. [13] suggest that geometric ergodicity of $\{z_t\}$ can hold under the weaker requirement that only one of the parameters ϕ_z^0 and ϕ_z^1 has absolute value less than one, with no restriction placed on the other.[7] Nevertheless, it should be pointed out that analogous ergodicity results do not necessarily hold for the trivariate process $\{\boldsymbol{w}_t\}$ satisfying (5) and (10)—even if stability restrictions are imposed on both components of the mixture model. The reason is that the mixing weights in (10) are a function of z_{t-1} and not of all three variables in $\boldsymbol{w}_{t-1}$. In this case, geometric ergodicity of $\{\boldsymbol{w}_t\}$ typically requires the spectral radius of at least one of the coefficient matrices $\boldsymbol{A}_0$ and $\boldsymbol{A}_1$ to be less than one and each element of $\boldsymbol{A}_0$ that is associated with (y_{t-1}, x_{t-1}) to equal the corresponding element of $\boldsymbol{A}_1$ (cf. Bec et al. [13]).[8] The latter requirement, which permits only the coefficients on z_{t-1} to differ across the two mixture components, is restrictive and unlikely to hold in general. Bearing in mind that conditions of this type are generally sufficient, but not necessary, for geometric ergodicity and global stationarity, an analysis of the skeleton of the model can provide useful insights into its local stability characteristics about potential equilibrium points.

3.3 | Estimation

The parameters of the model defined by equations (6) to (10) can be estimated by the ML method. With $\boldsymbol{\theta} := (\boldsymbol{\theta}_0', \boldsymbol{\theta}_1', z^*)'$, the log-likelihood function corresponding to observations $\boldsymbol{w}_0, \boldsymbol{w}_1, \ldots, \boldsymbol{w}_T$ from the model (conditional on $\boldsymbol{w}_0$) takes the form

$$\mathcal{L}(\boldsymbol{\theta}) := \sum_{t=1}^{T} \ln\left(G_t f_{0,t} + (1 - G_t) f_{1,t}\right), \tag{11}$$

where G_t is given in (10) and

$$f_{i,t} := (2\pi)^{-3/2} (\det \boldsymbol{S}_i)^{-1/2} \times \exp\left(-\frac{1}{2}\{\boldsymbol{B}_i \boldsymbol{w}_t - \boldsymbol{\mu}_i - \boldsymbol{D}_i \boldsymbol{w}_{t-1}\}' \boldsymbol{S}_i^{-1} \{\boldsymbol{B}_i \boldsymbol{w}_t - \boldsymbol{\mu}_i - \boldsymbol{D}_i \boldsymbol{w}_{t-1}\}\right),$$

for $i = 0, 1$. Here,

$$\boldsymbol{B}_i := \begin{bmatrix} 1 & -\beta_{yx}^i & -\beta_{yz}^i \\ 0 & 1 & -\beta_{xz}^i \\ 0 & 0 & 1 \end{bmatrix}, \qquad \boldsymbol{D}_i := \begin{bmatrix} \phi_{yy}^i & \phi_{yx}^i & \phi_{yz}^i \\ 0 & \phi_{xx}^i & \phi_{xz}^i \\ 0 & 0 & \phi_z^i \end{bmatrix},$$

$$\boldsymbol{S}_i := \mathrm{diag}((\sigma_\eta^i)^2, (\sigma_\upsilon^i)^2, (\sigma_z^i)^2), \qquad \boldsymbol{\mu}_i := (\mu_y^i, \mu_x^i, \mu_z^i)'.$$

Maximisation of the function (11) with respect to $\boldsymbol{\theta}$ yields the ML estimator $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$. Under conventional regularity conditions, $\hat{\boldsymbol{\theta}}$ is a consistent estimator for the true value of $\boldsymbol{\theta}$, say $\boldsymbol{\theta}^*$, and $T^{1/2}\boldsymbol{J}(\boldsymbol{\theta}^*)^{1/2}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)$ is asymptotically normal, as T tends to infinity, with zero mean and identity covariance matrix, where $\boldsymbol{J}(\boldsymbol{\theta}^*) := -\mathrm{plim}_{T\to\infty} T^{-1}(\partial^2\mathcal{L}/\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}')(\boldsymbol{\theta}^*)$.[9] In practice, the average observed information $-T^{-1}\ddot{\mathcal{L}}(\hat{\boldsymbol{\theta}})$ may be used in place of $\boldsymbol{J}(\boldsymbol{\theta}^*)$, where $\ddot{\mathcal{L}}(\hat{\boldsymbol{\theta}}) := (\partial^2\mathcal{L}/\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}')(\hat{\boldsymbol{\theta}})$, and inference on $\boldsymbol{\theta}$ be based on the approximate normality of $\hat{\boldsymbol{\theta}}$, with mean $\boldsymbol{\theta}^*$ and covariance matrix $\{-\ddot{\mathcal{L}}(\hat{\boldsymbol{\theta}})\}^{-1}$. In the next section of the paper, we use Monte Carlo methods to assess the quality of such large-sample asymptotic approximations.

4 | Monte Carlo Simulations

In this section, simulation methods are used to explore the finite-sample properties of the ML estimator and related test statistics in the class of models under consideration. In particular, we are interested in the properties of the ML estimator in a fully specified three-equation model and in a partial model which excludes one of the equations.

4.1 | Experimental Design and Simulation

The model defined by Equations (6) to (10) is used as the data-generating mechanism in the Monte Carlo experiments, with the same parameter values as those used in the numerical analysis in Section 3.2 (see Table 1). The sample sizes selected, that is, $T \in \{100, 200, 400, 800, 1600, 3200\}$, are representative of data sets that are typically used in empirical work (samples of 3200 or more observations are not uncommon in studies using weekly or daily data). In all experiments, $50 + T$ data points for $\boldsymbol{w}_t$ are generated, starting with $\boldsymbol{w}_0 = (-0.9, 0.7, -0.22)'$, but only the last T of these points are used for inference in each Monte Carlo replication in order to attenuate the effect of initial values.

The ML estimate $\hat{\boldsymbol{\theta}} := (\hat{\boldsymbol{\theta}}_0', \hat{\boldsymbol{\theta}}_1', \hat{z}^*)'$ of the 31-dimensional parameter $\boldsymbol{\theta}$ associated with the three-equation model in (6) to (10) is obtained via direct maximisation of the log-likelihood function given in (11) by means of the Broyden–Fletcher–Goldfarb–Shanno quasi-Newton method (an analogous procedure is used in the case of a two-equation partial model discussed later). Starting points for the iterations of the numerical optimization algorithm are obtained from a grid of seven values for each parameter (including the true value), with those points that result in the highest value of the log-likelihood function being ultimately selected.[10] Since ML estimation is computationally costly, especially in the context of simulations, 1000 Monte Carlo replications per experiment are carried out.

4.2 | Complete Model

For each of the parameters associated with the equation for y_t, Table 2 reports: (i) the Monte Carlo estimate of the finite-sample bias of the ML estimator of the parameter; (ii) the ratio of the finite-sample standard deviation of the estimator to its estimated standard error obtained from $\{-\ddot{\mathcal{L}}(\hat{\boldsymbol{\theta}})\}^{-1}$ (averaged across Monte Carlo replications); (iii) the rejection frequency of a t-type two-sided test of the null hypothesis that the parameter equals its true value, using the 0.975 standard-normal quantile as critical value. Corresponding results for the parameters associated with the equations for x_t and z_t are reported in Tables 3 and 4, respectively.

The ML estimator exhibits modest bias only in the case of a small number of parameters and for the smallest of the sample sizes considered. Even for these parameters, however, bias is a decreasing function of the sample size and becomes negligible for $T \geqslant 200$. In addition, estimated standard errors obtained from the observed information matrix provide accurate approximations to the sampling standard deviation of the ML estimators, especially in samples of 200 or more observations. Encouraging results are also obtained in the context of hypothesis testing: the deviation of the estimated rejection probabilities of tests from the nominal 0.05 level rarely is substantial enough to make the tests unattractive for application when $T \geqslant 200$. In summary, the overall pattern of the simulation results is in accord with conventional ML asymptotic theory, although it does suggest that at least 200 observations are typically needed before large-sample approximations provide an accurate guide for inference.

TABLE 2 | Equation for y_t in complete model.

	Below threshold							Above threshold						
T	μ_y^0	β_{yx}^0	β_{yz}^0	ϕ_{yy}^0	ϕ_{yx}^0	ϕ_{yz}^0	σ_η^0	μ_y^1	β_{yx}^1	β_{yz}^1	ϕ_{yy}^1	ϕ_{yx}^1	ϕ_{yz}^1	σ_η^1
Bias														
100	−0.077	0.037	0.018	−0.052	−0.112	−0.195	−0.060	−0.088	0.091	−0.008	−0.161	−0.193	0.155	−0.114
200	−0.007	−0.038	0.007	−0.028	−0.025	−0.130	−0.031	−0.013	0.039	0.002	−0.058	−0.022	0.076	−0.054
400	−0.034	0.039	0.002	−0.009	−0.001	0.021	−0.011	0.001	0.007	−0.015	−0.042	−0.014	0.060	−0.026
800	0.036	−0.028	−0.001	−0.006	0.002	−0.011	−0.007	−0.004	0.010	0.002	−0.021	−0.024	0.012	−0.012
1600	−0.012	0.016	0.001	−0.005	−0.004	−0.006	−0.003	−0.005	−0.006	−0.002	−0.012	−0.015	0.002	−0.004
3200	−0.003	0.000	0.002	−0.003	−0.016	−0.004	−0.002	−0.002	0.004	−0.003	−0.010	−0.011	0.011	−0.003
Ratio of sampling standard deviation to estimated standard error														
100	1.033	1.062	1.097	1.044	1.051	1.065	1.141	1.253	1.180	1.255	1.240	1.226	1.220	1.221
200	1.043	1.052	1.009	0.997	1.024	1.033	1.063	0.983	0.973	1.068	1.074	1.072	1.034	1.085
400	1.020	1.017	0.993	0.998	0.998	1.017	1.054	0.991	1.042	0.997	1.017	1.016	1.010	1.006
800	0.958	0.990	1.029	1.036	0.989	1.009	0.990	0.993	0.991	1.013	0.982	1.001	0.981	1.026
1600	0.956	0.943	0.965	0.991	0.941	0.991	0.988	1.021	0.993	1.030	0.991	0.989	1.024	1.020
3200	1.000	0.994	0.983	0.981	0.992	0.974	1.010	0.993	1.011	1.031	1.025	0.977	1.001	0.998
Test rejection frequency (nominal level 0.05)														
100	0.063	0.067	0.082	0.073	0.064	0.066	0.163	0.107	0.091	0.119	0.110	0.107	0.109	0.253
200	0.071	0.077	0.060	0.060	0.074	0.068	0.116	0.081	0.076	0.094	0.068	0.078	0.090	0.143
400	0.065	0.067	0.048	0.057	0.059	0.061	0.075	0.057	0.058	0.055	0.058	0.067	0.049	0.084
800	0.069	0.071	0.045	0.056	0.045	0.082	0.054	0.054	0.045	0.059	0.042	0.034	0.047	0.071
1600	0.047	0.038	0.037	0.053	0.041	0.057	0.057	0.047	0.051	0.057	0.046	0.055	0.049	0.063
3200	0.053	0.050	0.047	0.050	0.059	0.044	0.054	0.044	0.052	0.052	0.059	0.047	0.047	0.058

Note: Results from Monte Carlo experiments discussed in Section 4. The Monte Carlo design can be found in Table 1.

4.3 | Partial Model

In a second set of experiments, we consider ML estimation of the parameters of the equations for y_t and z_t in a partial model which does not include the equation for x_t. The data-generating mechanism used is the same as that described in Section 4.1.

Results for the ML estimators of parameters associated with the equations for y_t and z_t are summarised in Tables 5 and 6, respectively. It is immediately obvious that the consequences of treating x_t as an unmodelled variable are rather severe. In sharp contrast to the results obtained from the complete three-equation model (cf. Table 2), many parameters of the equation for y_t are now estimated with substantial bias. What is more, bias increases with the sample size for the estimators of some parameters. In cases where the average estimated parameter values deviate considerably from the corresponding true value, the simulation-estimated level of a test for the hypothesis that the parameter equals its true value is, unsurprisingly, considerably larger than the nominal 0.05 level. Moreover, these level distortions increase with the sample size, rendering tests manifestly unreliable. Such severe over-rejection is evidently associated more with the bias of coefficient estimators than with that of estimated standard errors; with a few exceptions, the inaccuracy of standard errors obtained from the observed information matrix is fairly modest.

The cost of not utilizing the information contained in the equation for x_t is not quite so high when estimating the parameters of the equation for z_t. Results for the ML estimator of these parameters are generally good for samples of more than 400 observations. For smaller sample sizes, however, inference based on the two-equation partial model is less accurate than inference based on the complete three-equation model, the differences being particularly striking in the case of the threshold parameter z^*.

In summary, the simulation evidence demonstrates that, despite the recursive nature of the system, focusing on the equations for y_t and z_t alone is perilous. Even if only the parameters of the equations for y_t and z_t are of interest, joint estimation of these parameters together with those of the equation for x_t is required in order to ensure that ML estimates have desirable statistical properties and inference is accurate. The equation for x_t contains useful information about the variation of parameters between the two regimes and the mixture structure of the dynamics of the three variables. Ignoring this information by focusing on a model

TABLE 3 | Equation for x_t in complete model.

	Below threshold					Above threshold				
T	μ_x^0	β_{xz}^0	ϕ_{xx}^0	ϕ_{xz}^0	σ_ψ^0	μ_x^1	β_{xz}^1	ϕ_{xx}^1	ϕ_{xz}^1	σ_ψ^1
Bias										
100	0.009	−0.005	0.000	0.043	−0.063	0.024	−0.070	−0.339	−0.234	−0.083
200	0.040	0.000	−0.003	0.012	−0.029	0.005	−0.006	−0.171	−0.064	−0.041
400	0.026	−0.001	−0.007	−0.032	−0.011	0.006	0.002	−0.067	0.003	−0.018
800	−0.012	0.000	0.000	0.013	−0.004	0.002	0.005	−0.035	0.015	−0.011
1600	0.005	0.000	−0.003	−0.014	−0.002	0.003	0.001	0.001	0.019	−0.006
3200	0.008	0.000	−0.001	−0.006	0.000	−0.001	−0.001	0.002	−0.003	−0.003
Ratio of sampling standard deviation to estimated standard error										
100	1.072	1.045	1.099	1.047	2.507	1.154	1.187	1.140	1.144	1.149
200	1.004	0.994	1.005	1.009	2.517	1.065	1.063	1.034	1.074	1.046
400	1.012	1.023	1.027	1.021	1.035	1.012	1.024	0.948	0.999	1.033
800	1.011	0.975	1.044	1.013	1.050	0.974	0.999	0.993	0.968	1.005
1600	1.036	0.971	1.044	0.984	0.987	0.999	1.016	0.948	0.993	0.969
3200	0.992	1.003	0.978	1.013	0.972	1.003	1.002	0.989	1.001	0.974
Test rejection frequency (nominal level 0.05)										
100	0.063	0.067	0.075	0.059	0.118	0.104	0.089	0.103	0.082	0.180
200	0.050	0.057	0.052	0.049	0.098	0.070	0.073	0.067	0.074	0.116
400	0.053	0.060	0.054	0.052	0.086	0.055	0.067	0.049	0.062	0.079
800	0.055	0.044	0.062	0.052	0.068	0.049	0.056	0.069	0.069	0.068
1600	0.059	0.039	0.058	0.055	0.051	0.058	0.046	0.042	0.046	0.052
3200	0.049	0.049	0.047	0.049	0.044	0.052	0.044	0.049	0.050	0.061

Note: Results from Monte Carlo experiments discussed in Section 4. The Monte Carlo design can be found in Table 1.

TABLE 4 | Equation for z_t in complete model.

	Below the threshold				Above the threshold		
T	z^*	μ_z^0	ϕ_z^0	σ_z^0	μ_z^1	ϕ_z^1	σ_z^1
Bias							
100	−0.016	−0.031	−0.037	−0.019	0.007	−0.058	−0.046
200	0.010	−0.032	−0.012	−0.013	0.023	−0.035	−0.020
400	0.000	−0.004	−0.009	−0.007	0.003	−0.013	−0.010
800	0.001	−0.016	−0.004	−0.004	0.008	−0.007	−0.004
1600	0.004	−0.002	−0.002	−0.002	0.007	−0.004	−0.002
3200	−0.002	0.000	−0.002	−0.001	−0.002	−0.001	−0.002
Ratio of sampling standard deviation to estimated standard error							
100	1.003	1.104	1.047	0.964	1.114	1.042	1.088
200	1.034	1.019	1.054	1.029	1.093	1.016	1.044
400	1.060	1.036	1.025	0.986	1.054	1.021	1.048
800	0.969	1.000	0.998	1.011	1.039	1.026	1.050
1600	0.997	0.971	1.007	0.989	1.008	1.035	1.036
3200	0.976	0.991	1.037	1.047	0.972	1.011	0.990
Test rejection frequency (nominal level 0.05)							
100	0.055	0.057	0.078	0.062	0.078	0.060	0.120
200	0.065	0.041	0.065	0.065	0.070	0.060	0.075
400	0.065	0.057	0.064	0.055	0.062	0.064	0.073
800	0.042	0.044	0.055	0.056	0.060	0.058	0.066
1600	0.051	0.043	0.047	0.051	0.048	0.063	0.059
3200	0.046	0.052	0.060	0.062	0.049	0.051	0.049

Note: Results from Monte Carlo experiments discussed in Section 4. The Monte Carlo design can be found in Table 1.

that treats x_t as an unmodelled or exogenous variable, conditionally on which the generation process of (y_t, z_t) may be analysed, has a clear deleterious effect on the accuracy of inferential procedures. Such findings are analogous to those reported in Pouzo et al. [8] for a large class of Markov regime-switching models with time-varying transition probabilities (which includes mixture autoregressive models as a special case).

5 | Consumer Sentiment and the Stock Market

The possibility of a causal role for consumer sentiment in driving observed changes in stock prices has a long history in economics and finance (see, inter alia, Keynes [20], Fisher and Statman [21], Case et al. [22], Shiller [23]). Our aim in this section is to examine whether movements in the price–dividend ratio, which may be thought of as a proxy for movements in expected future dividend payments, are affected by a measure of consumer sentiment once market volatility is taken into account.[11] More specifically, using an index of consumer sentiment as a signal for the state the economy is in, we estimate the parameters of a two-component mixture model for the price–dividend ratio, market volatility and consumer sentiment with a threshold-type structure which reflects high and low sentiment. Within this framework, we also examine the implications of treating volatility as an unmodelled variable (as is often the case in the empirical literature). In light of the simulation results reported in Section 4.3, the omission of the volatility equation from the system is likely to affect estimates for the price–dividend equation significantly since volatility is used as a covariate in the latter equation and volatility dynamics may be reasonably expected to be different depending on whether a high-sentiment or low-sentiment regime is more likely to prevail.

5.1 | Data

Our empirical analysis is based on monthly data for: (i) the price–dividend ratio for the S&P 500 stock index (y_t), computed as the ratio of the index to the associated cumulative nominal dividends over the past 12 months (obtained from Welch and Goyal [25]);[12] (ii) the 3-month moving average of the market variance risk premium for the returns on the S&P 500 index (x_t), using the definition and data of Zhou [26];[13] (iii) the 12-month change in the Index of Consumer Sentiment (z_t) published by the Survey Research Center of the University of Michigan.[14] These variables are part of the set of predictor variables used by Lansing et al. [27] in their analysis of the predictability of excess returns. The data cover the period from March 1990 to December 2009 (238 observations in total for each time series).[15] The time series of the market variance risk premium and the change in consumer sentiment are rescaled, with the maximum value normalised to unity; this does not affect results beyond the rescaling of estimated coefficients.[16]

5.2 | Empirical Results

Table 7 contains ML estimates of the parameters of the complete model defined in (6) to (10), together with corresponding estimated standard errors. In the price–dividend equation,

TABLE 5 | Equation for y_t in partial model.

T	Below threshold							Above threshold						
	μ_y^0	β_{yx}^0	β_{yz}^0	ϕ_{yy}^0	ϕ_{yx}^0	ϕ_{yz}^0	σ_η^0	μ_y^1	β_{yx}^1	β_{yz}^1	ϕ_{yy}^1	ϕ_{yx}^1	ϕ_{yz}^1	σ_η^1
Bias														
100	−1.908	1.192	0.100	−0.074	0.014	−0.210	−0.172	1.379	0.057	−0.627	−0.205	−0.013	−0.004	−0.239
200	−2.499	1.513	0.068	−0.046	0.346	−0.255	−0.075	1.314	0.201	−0.623	−0.093	0.098	0.067	−0.115
400	−3.285	2.097	0.048	−0.009	0.515	−0.094	−0.021	1.239	0.292	−0.648	−0.062	0.130	0.194	−0.043
800	−3.832	2.413	0.037	−0.006	0.599	−0.280	0.009	1.196	0.332	−0.617	−0.032	0.115	0.155	−0.005
1600	−4.110	2.608	0.033	−0.002	0.648	−0.277	0.026	1.161	0.350	−0.612	−0.015	0.116	0.171	0.018
3200	−4.253	2.707	0.030	0.001	0.659	−0.288	0.033	1.137	0.381	−0.609	−0.007	0.121	0.208	0.027
Ratio of sampling standard deviation to estimated standard error														
100	1.042	0.559	0.620	0.572	0.215	0.239	1.138	1.227	1.176	1.197	1.262	0.959	1.032	1.487
200	2.184	1.550	0.994	1.047	1.069	1.073	1.293	1.250	1.109	1.245	1.111	1.088	1.021	1.224
400	1.947	1.483	0.950	0.940	0.956	0.976	1.155	1.126	1.045	1.035	1.007	0.967	0.939	1.072
800	1.517	1.281	0.946	0.956	0.950	0.874	1.053	1.112	0.976	0.937	0.962	0.899	0.910	1.078
1600	1.175	1.000	0.917	0.830	0.882	0.662	0.925	1.083	0.864	0.787	0.916	0.882	0.968	0.989
3200	1.096	1.051	0.938	0.881	0.930	0.780	0.970	1.034	0.888	0.920	0.999	0.937	0.926	0.976
Test Rejection Frequency (Nominal Level 0.05)														
100	0.599	0.345	0.131	0.145	0.157	0.135	0.319	0.609	0.230	0.437	0.196	0.152	0.169	0.490
200	0.674	0.475	0.083	0.089	0.127	0.099	0.210	0.686	0.161	0.407	0.091	0.092	0.094	0.288
400	0.785	0.722	0.062	0.045	0.138	0.085	0.162	0.763	0.127	0.410	0.085	0.082	0.074	0.176
800	0.909	0.874	0.081	0.064	0.260	0.145	0.101	0.899	0.142	0.572	0.060	0.065	0.065	0.107
1600	0.947	0.932	0.096	0.051	0.447	0.120	0.135	0.965	0.211	0.846	0.044	0.057	0.100	0.119
3200	0.986	0.983	0.135	0.048	0.716	0.124	0.241	0.989	0.424	0.983	0.052	0.061	0.131	0.200

Note: Results from Monte Carlo experiments discussed in Section 4. The Monte Carlo design can be found in Table 1.

the estimate of the impact multiplier β_{yx}^0 is negative and significant, implying a negative effect of volatility on stock prices in the regime associated with pessimistic times, that is, when the ex ante probability of the index of consumer sentiment being below the threshold value is relatively high. The long-run multiplier of volatility is positive in both regimes, reflecting the well-known mean–variance trade-off. Moreover, the estimate of β_{yz}^1 is positive and significant, suggesting that consumer sentiment directly affects stock prices in optimistic times, that is, when the sentiment indicator is likely to exceed the threshold value. This result is in line with the results obtained in Otoo [28]. The results also show the importance of allowing for dynamics in the price–dividend equation, with lagged values of at least two of the variables having statistically significant coefficients in both regimes. Consumer sentiment is found to have a significant contemporaneous effect on volatility when the former is likely to be below the threshold, but the long-run effect is small. Finally, the mixture autoregressive equation for consumer sentiment implies a clear separation of the two regimes.

The time-varying mixing weights G_t implied by the estimated parameters are shown in Figure 1, along with time series of the three variables. As seen in the top panel, the growth of the price–dividend ratio is inversely related to the sentiment indicator, suggesting a separation of regimes consistent with the fitted model and supporting the hypothesis that consumer sentiment affects stock prices and returns. The bottom left panel shows that volatility is higher when the economic outlook is not particularly good (i.e., when G_t is large), while the bottom right panel shows that the mixing weights G_t move, as expected, inversely with the index of consumer sentiment z_t since higher values of the latter are associated with a better economic outlook.

Finally, Table 8 reports ML estimation results for a partial model which excludes the volatility equation. There are substantial differences between these estimates and those presented in Table 7 for the complete model. The first notable difference is that the estimated threshold value changes sign and increases from −0.3401 to 0.1875. As a consequence, the time-varying mixing weights and estimated dynamics in each regime are very different. In the price–dividend equation, estimates of some parameters change substantially in magnitude compared to the complete model. For example, the estimated value of the impact multiplier of the price–dividend ratio with respect to volatility in the high-sentiment regime (β_{yx}^1) increases from −10.4918 to −4.0135. What is more, neither contemporaneous nor lagged sentiment appear to have a statistically significant effect on the price–dividend ratio in either of the two regimes. The differences in the estimated parameters of the consumer sentiment equation in the complete and partial models are smaller by comparison.

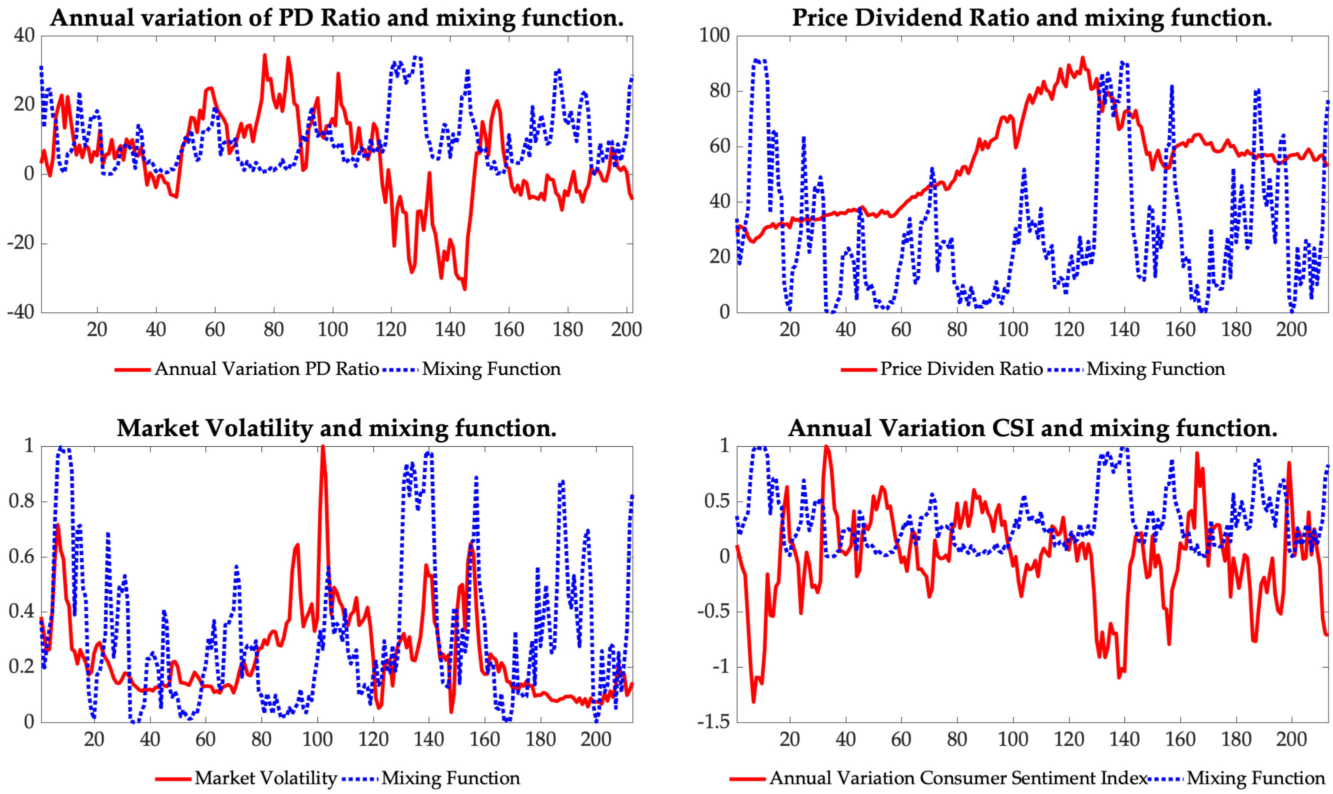

FIGURE 1 | Mixing function and data.

TABLE 6 | Equation for z_t in partial model.

	Below threshold				Above threshold		
T	z^*	μ_z^0	ϕ_z^0	σ_z^0	μ_z^1	ϕ_z^1	σ_z^1
Bias							
100	−0.420	0.326	−0.052	−0.041	−0.126	−0.056	−0.080
200	−0.265	0.241	−0.026	−0.025	−0.054	−0.033	−0.042
400	−0.144	0.152	−0.019	−0.012	−0.019	−0.014	−0.023
800	−0.052	0.036	−0.005	−0.006	0.006	−0.009	−0.009
1600	0.003	0.007	−0.002	−0.002	0.008	−0.004	−0.003
3200	0.018	−0.005	−0.001	−0.002	0.014	−0.003	−0.001
Ratio of sampling standard deviation to estimated standard error							
100	1.289	0.681	0.734	1.160	0.671	0.726	1.063
200	1.362	1.104	1.084	1.225	1.073	1.033	1.043
400	1.336	1.059	1.003	0.962	1.010	0.960	1.078
800	1.148	0.960	0.916	0.921	0.958	0.960	1.049
1600	0.952	0.943	0.992	0.954	0.896	0.939	0.991
3200	0.987	0.993	1.014	0.996	0.931	0.973	0.975
Test rejection frequency (nominal level 0.05)							
100	0.312	0.136	0.120	0.138	0.151	0.121	0.201
200	0.271	0.112	0.097	0.103	0.083	0.080	0.144
400	0.208	0.086	0.079	0.070	0.070	0.052	0.110
800	0.133	0.049	0.048	0.050	0.064	0.059	0.091
1600	0.084	0.052	0.055	0.055	0.052	0.045	0.063
3200	0.075	0.052	0.058	0.055	0.048	0.046	0.051

Note: Results from Monte Carlo experiments discussed in Section 4. The Monte Carlo design can be found in Table 1.

Such results are not very surprising and are entirely consistent with the findings from the Monte Carlo experiments discussed in Section 4.3. As pointed out there, results from a partial model should be interpreted very cautiously as they are likely to be unreliable.[17]

6 | Summary

In this article, we have considered the role of consumer sentiment in stock markets. More specifically, we have focused on some questions that have attracted attention in the literature, namely whether the effect of consumer sentiment is already incorporated in stock prices, whether consumer sentiment may be regarded as a variable with which different economic regimes are associated, whether financial variables such as stock prices/returns and volatility are simultaneously affected by consumer sentiment, and how to account for such joint relationships.

To answer these questions, we have considered a mixture of multivariate dynamic models with time-varying mixing weights. The latter have a threshold-type structure so that the (normalised conditional) probability that the implied value of a specific variable (consumer sentiment, in our setting) is below/above a threshold determines which component of the mixture (or regime) is chosen at each date. We have discussed the stability properties of such a model and ML estimation of its parameters. Monte Carlo experiments have demonstrated that likelihood-based inference is accurate in sample sizes typical in applications, provided the

TABLE 7 | ML parameter estimates for complete model (standard errors in parentheses): ϕ^i_{ab} (β^i_{ab}) is the regime-i coefficient on the lagged (contemporaneous) variable a in the equation for variable b; ϕ^i_b, μ^i_b and σ^i_b are the regime-i autoregressive coefficient, intercept and noise standard deviation, respectively, in the equation for variable b.

Below threshold			Above threshold		
Price–dividend ratio (y)					
μ^0_y	−1.2314	(2.0668)	μ^1_y	0.5060	(0.5673)
β^0_{yx}	−4.1047	(2.3223)	β^1_{yx}	−10.4918	(4.4356)
β^0_{yz}	0.2996	(1.5834)	β^1_{yz}	1.2586	(0.6094)
ϕ^0_{yy}	0.9957	(0.0260)	ϕ^1_{yy}	0.9801	(0.0110)
ϕ^0_{yx}	7.4692	(2.4047)	ϕ^1_{yx}	13.9885	(4.1716)
ϕ^0_{yz}	0.2226	(1.7368)	ϕ^1_{yz}	−0.8931	(0.6064)
σ^0_η	3.2559	(0.3035)	σ^1_η	1.6411	(0.1103)
Market volatility (x)					
μ^0_x	0.1202	(0.0483)	μ^1_x	0.0130	(0.0054)
β^0_{xz}	−0.1778	(0.0879)	β^1_{xz}	0.0015	(0.0124)
ϕ^0_{xx}	0.7173	(0.0947)	ϕ^1_{xx}	0.9240	(0.0249)
ϕ^0_{xz}	0.1745	(0.0950)	ϕ^1_{xz}	0.0011	(0.0127)
σ^0_ψ	0.1755	(0.0162)	σ^1_ψ	0.0322	(0.0023)
Consumer sentiment (z)					
μ^0_z	−0.0916	(0.0498)	μ^1_z	0.0254	(0.0196)
ϕ^0_z	0.8025	(0.0715)	ϕ^1_z	0.7407	(0.0563)
σ^0_z	0.3016	(0.0318)	σ^1_z	0.2181	(0.0141)
z^*	−0.3401	(0.0718)			

TABLE 8 | ML parameter estimates for partial model (standard errors in parentheses): ϕ^i_{ab} (β^i_{ab}) is the regime-i coefficient on the lagged (contemporaneous) variable a in the equation for variable b; ϕ^i_b, μ^i_b and σ^i_b are the regime-i autoregressive coefficient, intercept and noise standard deviation, respectively, in the equation for variable b.

Below threshold			Above threshold		
Price–dividend ratio (y)					
μ^0_y	1.5134	(0.8355)	μ^1_y	−2.7734	(0.7429)
β^0_{yx}	−5.3104	(1.6200)	β^1_{yx}	−4.0135	(3.0259)
β^0_{yz}	0.5729	(1.1979)	β^1_{yz}	0.4441	(0.6497)
ϕ^0_{yy}	0.9578	(0.0131)	ϕ^1_{yy}	1.0535	(0.0141)
ϕ^0_{yx}	6.0238	(1.8019)	ϕ^1_{yx}	9.2113	(2.9583)
ϕ^0_{yz}	−0.3800	(1.2426)	ϕ^1_{yz}	−0.3720	(0.5402)
σ^0_η	2.2703	(0.1661)	σ^1_η	1.1996	(0.1446)
Consumer sentiment (z)					
μ^0_z	−0.0582	(0.0258)	μ^1_z	0.1570	(0.0445)
ϕ^0_z	0.8629	(0.0442)	ϕ^z_1	0.4840	(0.0853)
σ^0_z	0.2109	(0.0164)	σ^1_z	0.2424	(0.0243)
z^*	0.1875	(0.1086)			

complete system is used. Treating one variable (volatility, in our setting) as an unmodelled conditioning variable has severe adverse effects on the accuracy of inferential procedures.

An analysis of the relationship between the S&P 500 price–dividend ratio, volatility and consumer sentiment using the proposed model has illustrated its practical usefulness. The model identifies two regimes associated with above-threshold and below-threshold values of consumer sentiment, with the latterhaving predictive power and a significant effect on both the price–dividend ratio and volatility, at least in the short run.

Acknowledgements

The authors are grateful to Jonathan Temple (Associate Editor) and two anonymous referees for their constructive comments and suggestions, which led to significant improvements.

Endnotes

[1] Here and elsewhere, $\boldsymbol{Q}^{1/2}$ denotes the symmetric, positive definite square root of a symmetric, positive definite matrix $\boldsymbol{Q}$. Unless otherwise indicated, all vectors are presented as column vectors.

[2] Note that, for $d = 1$, (3) reduces to the specification used in Dueker et al. [9]. Also note that, although conditional Gaussianity is used as a convenient assumption in much of what follows, non-Gaussian distributions with a continuous density function may be considered instead. Dueker et al. [11], for instance, considered distributions belonging to the Student-t family and, in the univariate case, replaced Φ_d in (3) with the distribution function of a (rescaled) Student-t distribution.

[3] In a univariate context, Wong and Li [12] considered mixture autoregressions (with logistic-type mixing weights) that may also include a distributed lag of strictly exogenous variables.

[4] It is perhaps worth pointing out that uncorrelatedness and Gaussianity of the components of $\boldsymbol{e}_t$ are not additional assumptions but a consequence of the conditioning operations in the context of (5).

[5] By comparison, a more general threshold structure of the type considered in Dueker et al. [10] would result in a trivariate conditional mixture model with 9 components and 138 parameters.

[6] For a definition and discussion of these concepts of ergodicity of Markov chains, the reader is referred to Douc et al. [17, Ch. 15].

[7] This parallels results for stochastic-unit-root and Markov-switching autoregressive processes, for which stability within all regimes is not necessary for geometric ergodicity and global stationarity (e.g., Gourieroux and Robert [18], Stelzer [19]).

[8] It should be noted that such results are obtained in Bec et al. [13] for a special case of mixture autoregressive models like (5) in which intercepts are zero and the noise covariance matrices are identical across regimes.

[9] Typical (high-level) conditions include identifiability of the model, compactness of the parameter space to whose interior $\boldsymbol{\theta}^*$ belongs, nonsingularity of $\lim_{T\to\infty} T^{-1}\mathrm{Var}[(\partial\mathcal{L}/\partial\boldsymbol{\theta})(\boldsymbol{\theta}^*)]$, asymptotic normality of $T^{-1/2}(\partial\mathcal{L}/\partial\boldsymbol{\theta})(\boldsymbol{\theta}^*)$, and uniform convergence of $T^{-1}(\partial^2\mathcal{L}/\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}')$ in a neighbourhood of $\boldsymbol{\theta}^*$ (cf. Dueker et al. [10]). The framework of Pouzo et al. [8] may also be used to establish consistency and asymptotic normality of $\widehat{\boldsymbol{\theta}}$, exploiting the fact that their general setting includes mixture autoregressive models as special cases.

[10] We note that estimation results appear to be robust with respect to the choice of starting values.

[11] This is in line with specifications that allow for mean–variance trade-off (e.g., Yu and Yuan [24]).

[12] Updated data are available at https://sites.google.com/view/agoyal145.

[13] The data are available at https://sites.google.com/site/haozhouspersonalhomepage.

[14] The data are available at http://www.sca.isr.umich.edu/tables.html.

[15] The end-date of the sample period is chosen so as to exclude post-2009 developments (e.g., proximity of interest rates to the zero lower bound, Covid-19) whose potential effects on stock prices are not directly related to those discussed in our empirical application. Such distinct periods could be accommodated by including additional components/regimes in the mixture model, along the lines of Dueker et al. [10]. This, however, would complicate the structure of the model unnecessarily and divert attention from the modelling and inferential issues highlighted in Section 1, on which we wish to focus.

[16] Before considering multivariate conditional mixture models for (y_t, x_t, z_t), we tested for neglected nonlinearity (of unknown form) in a single-regime version of (6) to (9) using a general portmanteau-type test. Specifically, we carried out a likelihood-ratio test of the hypothesis that $(y_{t-1}^2, x_t^2, z_t^2)$, (x_{t-1}^2, z_t^2) and z_{t-1}^2 have zero coefficients when added to the right-hand side of the equations for y_t, x_t and z_t, respectively (these test variables were selected so as to avoid collinearity problems). The P-value of the test is 0.0197, suggesting that there is strong evidence against the linear specification.

[17] A logistic analogue of the model (6) to (10) can be obtained by replacing the mixing weights in (10) with $G_t = [1 + \exp(-\gamma\{z_{t-1} - z^*\})]^{-1}, \gamma > 0$. The local stability properties of such a model can be analysed using the approach outlined in Section 3.2, with $G(u) = [1 + \exp(-\gamma\{u - z^*\})]^{-1}$ and $\dot{G}(z_e) = \gamma \exp(-\gamma\{z_e - z^*\})/[1 + \exp(-\gamma\{z_e - z^*\})]^2$. As is the case with our model, it is important that the parameters of the complete three-equation system be estimated jointly when utilizing these mixing weights; otherwise, results will be unreliable. ML estimates of the parameters of the logistic threshold specification based on our data are qualitatively similar to those reported in Table 7, although the time-varying mixing weights implied by these estimates are somewhat different from those shown in Figure 1. The full set of results is available upon request.

References

1. S. C. Ludvigson, "Consumer Confidence and Consumer Spending," *Journal of Economic Perspectives* 18, no. 2 (2004): 29–50.

2. M. Baker and J. Wurgler, "Investor Sentiment in the Stock Market," *Journal of Economic Perspectives* 21, no. 2 (2007): 129–151.

3. M. Schmeling, "Investor Sentiment and Stock Returns: Some International Evidence," *Journal of Empirical Finance* 16, no. 3 (2009): 394–408.

4. R. F. Stambaugh, J. Yu, and Y. Yuan, "The Short of It: Investor Sentiment and Anomalies," *Journal of Financial Economics* 104, no. 2 (2012): 288–302.

5. J. Shen, J. Yu, and S. Zhao, "Investor Sentiment and Economic Forces," *Journal of Monetary Economics* 86 (2017): 1–21.

6. B. Desroches and M.-A. Gosselin, "Evaluating Threshold Effects in Consumer Sentiment," *Southern Economic Journal* 70, no. 4 (2004): 942–952.

7. S.-S. Chen, "Lack of Consumer Confidence and Stock Returns," *Journal of Empirical Finance* 18, no. 2 (2011): 225–236.

8. D. Pouzo, Z. Psaradakis, and M. Sola, "Maximum Likelihood Estimation in Markov Regime-Switching Models With Covariate-Dependent Transition Probabilities," *Econometrica* 90, no. 4 (2022): 1681–1710.

9. M. J. Dueker, M. Sola, and F. Spagnolo, "Contemporaneous Threshold Autoregressive Models: Estimation, Testing and Forecasting," *Journal of Econometrics* 141, no. 2 (2007): 517–547.

10. M. J. Dueker, Z. Psaradakis, M. Sola, and F. Spagnolo, "Multivariate Contemporaneous Threshold Autoregressive Models," *Journal of Econometrics* 160, no. 2 (2011): 311–325.

11. M. J. Dueker, Z. Psaradakis, M. Sola, and F. Spagnolo, "State-Dependent Threshold Smooth Transition Autoregressive Models," *Oxford Bulletin of Economics and Statistics* 75, no. 6 (2013): 835–854.

12. C. S. Wong and W. K. Li, "On a Logistic Mixture Autoregressive Model," *Biometrika* 88, no. 3 (2001): 833–846.

13. F. Bec, A. Rahbek, and N. Shephard, "The ACR Model: A Multivariate Dynamic Mixture Autoregression," *Oxford Bulletin of Economics and Statistics* 70, no. 5 (2008): 583–618.

14. L. Kalliovirta, M. Meitz, and P. Saikkonen, "Gaussian Mixture Vector Autoregression," *Journal of Econometrics* 192, no. 2 (2016): 485–498.

15. H. Tong, *Non-Linear Time Series: A Dynamical System Approach* (Oxford University Press, 1990).

16. A. Carvalho and G. Skoulakis, "Ergodicity and Existence of Moments for Local Mixtures of Linear Autoregressions," *Statistics & Probability Letters* 71, no. 4 (2005): 313–322.

17. R. Douc, E. Moulines, P. Priouret, and P. Soulier, *Markov Chains* (Springer, 2018).

18. C. Gourieroux and C. Y. Robert, "Stochastic Unit Root Models," *Econometric Theory* 22, no. 6 (2006): 1052–1090.

19. R. Stelzer, "On Markov-Switching ARMA Processes–Stationarity, Existence of Moments, and Geometric Ergodicity," *Econometric Theory* 25, no. 1 (2009): 43–62.

20. J. M. Keynes, *The General Theory of Employment, Interest, and Money* (Macmillan, 1936).

21. K. L. Fisher and M. Statman, "Consumer Confidence and Stock Returns," *Journal of Portfolio Management* 30, no. 1 (2003): 115–127.

22. K. E. Case, J. M. Quigley, and R. J. Shiller, "Comparing Wealth Effects: The Stock Market Versus the Housing Market," *B.E. Journal of Macroeconomics* 5, no. 1 (2005): 1–34.

23. R. J. Shiller, "Narrative Economics," *American Economic Review* 107, no. 4 (2017): 967–1004.

24. J. Yu and Y. Yuan, "Investor Sentiment and the Mean–Variance Relation," *Journal of Financial Economics* 100, no. 2 (2011): 367–381.

25. I. Welch and A. Goyal, "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," *Review of Financial Studies* 21, no. 4 (2008): 1455–1508.

26. H. Zhou, "Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty," *Annual Review of Financial Economics* 10, no. 1 (2018): 481–497.

27. K. J. Lansing, S. F. LeRoy, and J. Ma, "Examining the Sources of Excess Return Predictability: Stochastic Volatility or Market Inefficiency?," *Journal of Economic Behavior & Organization* 197 (2022): 50–72.

28. M. W. Otoo, "Consumer Sentiment and the Stock Market. Finance and Economics Discussion Series 1999-60," Board of Governors of the Federal Reserve System (U.S.) (1999).