Bootstrap Inference for Present-Value Models with Estimated Factors[†]

Xuanbin Yang, Tingting Cheng[*]

School of Finance, Nankai University, China

November 18, 2026

Abstract: Existing bootstrap inference for present-value models typically relies on directly observed predictors, while expected returns and dividend growth may depend on information dispersed across macroeconomic and financial variables. We extend the present-value framework by introducing common factors into its state-space representation to parsimoniously incorporate such high-dimensional information, and develop bootstrap likelihood-ratio inference with estimated factors. We establish consistency of the two-step estimator and show that replacing unobserved factors with principal-components estimates has no first-order effect on either the likelihood-ratio statistic or its innovation-bootstrap distribution. An empirical application to annual U.S. data from 1960 to 2023 illustrates the bootstrap procedure.

Keywords: Present-value models; Factor models; Bootstrap likelihood-ratio tests; Return predictability

JEL classification: C12; C32; C38; G12

[*]Corresponding author. E-mail address: Tingting.Cheng@nankai.edu.cn.

[†]We are grateful for the financial support from the National Natural Science Foundation of China (72673089, 72173068) and the Fundamental Research Funds for the Central Universities.

1 Introduction

Whether variation in aggregate valuation ratios reflects changes in expected returns or expected dividend growth is central to asset-pricing research. The Campbell–Shiller present-value relation links the price–dividend ratio to expectations of future returns and dividend growth (Campbell and Shiller 1988; Yu and Yan 2023). A prominent interpretation attributes most valuation-ratio variation to expected returns, because dividend yields forecast returns whereas aggregate dividend growth appears only weakly predictable (Fama and French 1988; Cochrane 2008; Yu and Chen 2024). Subsequent work, however, shows that this conclusion can be sensitive to the sample, forecasting specification, treatment of latent expectations, modeling of dividend dynamics, and temporal instability (Chen 2009; van Binsbergen and Koijen 2010; Chen et al. 2012; Pettenuzzo et al. 2020; Yu, Huang and Chen 2023; Yu, Chen and Li 2023). Reliable inference on both components is therefore important, especially in the short annual samples typically available for this analysis.

Latent-state present-value models provide a unified framework for jointly studying return and dividend-growth predictability, but conventional asymptotic inference can be unreliable in finite samples. Piatti and Trojani (2020) show that asymptotic likelihood-ratio (LR) tests can over-reject and propose an innovation-based bootstrap that delivers more reliable finite-sample critical values. Expected returns and expected dividend growth, however, are conditional expectations, so evidence about their predictability can depend on the information set used to model them (Lettau and Ludvigson 2005). Existing observed-predictor specifications based on a few individually selected variables may omit relevant external information and make conclusions sensitive to predictor choice. A natural way to address this information-set limitation without proliferating state-space parameters is to use a latent factor model: a small number of latent common factors can summarize the shared variation in a large panel and retain predictive signals dispersed across individual series (Kelly and Pruitt 2013; Cheng et al. 2024; Guo et al. 2026). This dimension reduction, however, creates a new inferential problem. Because the factors are unobserved and must be estimated, factor-estimation error enters both the state-space likelihood and the bootstrap procedure. The consistency of the resulting estimator and the validity of bootstrap inference therefore require separate theoretical justification.

This paper fills this gap by extending bootstrap likelihood-ratio inference to present-value models with estimated common factors. We first introduce latent common factors into the dynamics of expected returns and expected dividend growth and derive the resulting state-space representation. We then propose a two-step estimator that extracts the factors by principal components and estimates the state-space parameters by Kalman-filter maximum likelihood, and construct a feasible bootstrap procedure based on the estimated factors. The bootstrap accounts for varia-

tion from state-space innovations and factor estimation by resampling innovations, constructing a bootstrap panel, and re-estimating the bootstrap factors. We prove that the two-step estimator is consistent under standard approximate factor-model conditions and suitable regularity conditions for the infeasible state-space model. Under the additional rate condition $\sqrt{T}/N \to 0$, we show that replacing the rotated true factors with their estimates has no first-order effect on either the sample likelihood-ratio statistic or its bootstrap counterpart. Together, these results establish the first-order validity of asymptotic and bootstrap LR inference with estimated factors.

Monte Carlo experiments show that the proposed bootstrap substantially reduces the size distortions of asymptotic LR tests in short samples. Its rejection frequencies are close to the nominal level and broadly comparable to those of the observed-factor bootstrap in these short-sample designs.

We apply the procedure to annual U.S. data from 1960 to 2023, comparing observed-predictor specifications with models based on common factors estimated from a large macroeconomic and financial panel. Consistent with Piatti and Trojani (2020), including observed predictors such as book-to-market or stock variance does not yield significant evidence of dividend-growth predictability under bootstrap inference in our sample. By contrast, the joint two-factor specification rejects the null of constant expected dividend growth at the 5% level while retaining evidence of return predictability. These findings suggest that incorporating high-dimensional information through common factors can reveal predictive relations that remain undetected in the observed-predictor specifications.

The remainder of the paper is organized as follows. Section 2 presents the present-value model with latent factors and the two-step estimator. Section 3 develops bootstrap likelihood-ratio inference with estimated factors. Section 4 presents Monte Carlo simulations. Section 5 reports the empirical analysis, and Section 6 concludes. Derivations, estimation details, assumptions, proofs, and additional empirical results are provided in the Online Supplementary Material.

2 Present-value model and estimation

2.1 Present-value model with latent factors

In the benchmark present-value model, expected dividend growth and expected returns follow first-order autoregressions without external predictors (van Binsbergen and Koijen 2010; Piatti and Trojani 2020). We augment the benchmark information set with latent common factors extracted from a large panel of macroeconomic and financial variables.

Let $r_{t+1} = \log\{(P_{t+1} + D_{t+1})/P_t\}$ be the cum-dividend log market return and $\Delta d_{t+1} = \log(D_{t+1}/D_t)$ aggregate log dividend growth. The log price–dividend ratio is $pd_t = \log(P_t/D_t)$.

Conditional on information available at time t, expected dividend growth and expected return are $g_t = E_t(\Delta d_{t+1})$ and $\mu_t = E_t(r_{t+1})$, respectively.

Let $X_t = (X_{1t}, \ldots, X_{Nt})'$ denote an $N \times 1$ large-dimensional vector of observed macroeconomic and financial variables at date t, and let $e_t = (e_{1t}, \ldots, e_{Nt})'$ denote the corresponding idiosyncratic component. The approximate factor model is

$$X_t = \Lambda F_t + e_t, \qquad t = 1, \ldots, T, \tag{2.1}$$

where F_t is an $r \times 1$ vector of common factors and $\Lambda = (\lambda_1, \ldots, \lambda_N)'$ is the $N \times r$ loading matrix. Stacking observations over t, let $X = (X_1, \ldots, X_T)'$, $F = (F_1, \ldots, F_T)'$, and $e = (e_1, \ldots, e_T)'$. Then $X = F\Lambda' + e$. The factors are normalized to have zero mean, $E[F_t] = 0$. We augment the benchmark dynamics with F_t as follows:

$$g_{t+1} = \gamma_0 + \gamma_1(g_t - \gamma_0) + \gamma_2' F_t + \varepsilon^g_{t+1}, \tag{2.2}$$

$$\mu_{t+1} = \delta_0 + \delta_1(\mu_t - \delta_0) + \delta_2' F_t + \varepsilon^\mu_{t+1}, \tag{2.3}$$

$$F_{t+1} = \Xi F_t + \varepsilon^F_{t+1}, \qquad \Delta d_{t+1} = g_t + \varepsilon^d_{t+1}. \tag{2.4}$$

Here γ_2 and δ_2 are $r \times 1$ coefficient vectors and Ξ is an $r \times r$ transition matrix. For $j, k \in \{g, \mu, d\}$, let σ_j denote the standard deviation of ε^j_t and ρ_{jk} its contemporaneous correlation with ε^k_t. Let $\rho_{gF}, \rho_{\mu F}$ and ρ_{dF} collect the correlations of the factor shocks with $\varepsilon^g_t, \varepsilon^\mu_t$ and ε^d_t, respectively. Writing $\widetilde{g}_t = g_t - \gamma_0$ and $\widetilde{\mu}_t = \mu_t - \delta_0$, the Campbell–Shiller log-linearization gives

$$pd_t = A - (B_1 \widetilde{\mu}_t + B_3' F_t) + (B_2 \widetilde{g}_t + B_4' F_t)\,, \tag{2.5}$$

where the present-value coefficients are defined in Supplementary Appendix A.

2.2 Estimation

We propose the following two-step estimation procedure.

Step 1: Factor estimation. The factors are estimated from the panel X by principal components under the normalization $\widetilde{F}'\widetilde{F}/T = I_r$ (Bai 2003). The number of factors r is treated as fixed and can be consistently estimated using methods such as those of Bai and Ng (2002), Onatski (2010), and Ahn and Horenstein (2013).

Step 2: State-space estimation. The estimated factors are treated as observed. The state-space parameters are estimated by maximizing the Gaussian Kalman-filter likelihood (see Supplementary Appendix B).

Let θ denote the parameter vector of the present-value model with latent factors. For any r-dimensional factor path $f = \{f_t\}_{t=1}^T$, define $Y_t(f) = (\Delta d_t, pd_t, f_t')'$, and let $\mathcal{L}_T(\theta; f)$ denote the

Gaussian Kalman-filter log-likelihood based on $\{Y_t(f)\}_{t=1}^T$. Given a null parameter space $\Theta_0 \subset \Theta$, compute

$$\widehat{\theta} = \underset{\theta\in\Theta}{\operatorname{argmax}}\, \mathcal{L}_T(\theta; \widetilde{F}), \qquad \widehat{\theta}_0 = \underset{\theta\in\Theta_0}{\operatorname{argmax}}\, \mathcal{L}_T(\theta; \widetilde{F}).$$

The principal-components estimator $\widetilde{F}_t$ estimates $H'F_t$ for some nonsingular $r \times r$ matrix H. Since $\gamma_2' F_t = (H^{-1}\gamma_2)'(H'F_t)$, applying the corresponding change of coordinates to all factor-related parameters gives the rotated true parameter vector θ^H. This reparameterization preserves the zero restrictions and LR statistics.

The following theorem establishes consistency and convergence rates for the two-step estimators.

Theorem 2.1. *Under Assumptions A, B and C(i)–(ii), as $N, T \to \infty$, $\widehat{\theta} - \theta^H = o_p(1)$. If C(iii) and D also hold and $\sqrt{T}/N \to 0$, then $\sqrt{T}(\widehat{\theta} - \theta^H) = O_p(1)$. The corresponding conclusions hold for the free parameters of the restricted estimator $\widehat{\theta}_0$ when $\theta^H \in \Theta_0$.*

Theorem 2.1 shows that the state-space estimator based on estimated factors is consistent for θ^H. Under the null, the restricted estimator $\widehat{\theta}_0$ is also consistent, which justifies its use in generating bootstrap samples. The root-T rate for both estimators allows us to use a local likelihood expansion to compare the LR statistics based on estimated and true factors in Theorem 3.1.

3 Bootstrap likelihood-ratio tests

3.1 Null hypotheses

Testing return and dividend-growth predictability amounts to testing whether expected returns and expected dividend growth vary over time. In the present-value model with latent factors, the null hypothesis of no dividend-growth predictability is

$$H_0^{\text{div}}: \quad \gamma_1 = 0, \quad \gamma_2 = 0_r, \quad \sigma_g = 0, \quad \rho_{g\mu} = 0, \quad \rho_{gF} = 0_r. \tag{3.1}$$

The null H_0^{div} imposes $m_d = 2r + 3$ restrictions and implies $g_t = \gamma_0$, so expected dividend growth is constant.

Similarly, the null hypothesis of no return predictability is

$$H_0^{\text{ret}}: \quad \delta_1 = 0, \quad \delta_2 = 0_r, \quad \sigma_\mu = 0, \quad \rho_{g\mu} = 0, \quad \rho_{\mu d} = 0, \quad \rho_{\mu F} = 0_r. \tag{3.2}$$

The null H_0^{ret} imposes $m_r = 2r + 4$ restrictions and implies $\mu_t = \delta_0$, so expected returns are constant.

Let m denote the number of restrictions under the null hypothesis. For a factor sequence f, the LR statistic is

$$LR_T(f) = 2\{\max_{\theta\in\Theta} \mathcal{L}_T(\theta; f) - \max_{\theta\in\Theta_0} \mathcal{L}_T(\theta; f)\}. \tag{3.3}$$

In the infeasible setting where the true factors F_t are observed, $LR_T(F) \Rightarrow \chi^2_m$ under a regular null and the likelihood conditions of Theorem 3.1. Piatti and Trojani (2020) document over-rejection with chi-square critical values and show that an innovation bootstrap improves finite-sample size control. We propose a bootstrap LR test for present-value models with estimated factors and give sufficient conditions for its first-order validity.

3.2 Bootstrap procedure

To obtain a bootstrap p-value for $LR_T(\widetilde{F})$, we write the present-value model in state-space form. The present-value identity (2.5) allows us to eliminate $\widetilde{\mu}_t$. The resulting state vector is $\mathcal{X}_t = \big(\widetilde{g}_{t-1}, \varepsilon^g_t, \varepsilon^\mu_t, \varepsilon^d_t, (\varepsilon^F_t)'\big)'$, where the last $r+3$ entries form the structural shock vector ε^X_t. In the original sample the model is fitted with $Y_t(\widetilde{F}) = (\Delta d_t, pd_t, \widetilde{F}_t')'$. The state transition and measurement equations of the fitted model are

$$\mathcal{X}_{t+1} = \mathcal{F}\mathcal{X}_t + \mathcal{B}\widetilde{F}_{t-1} + \Gamma\varepsilon^X_{t+1}, \tag{3.4}$$

$$Y_t(\widetilde{F}) = M_0 + M_1 Y_{t-1}(\widetilde{F}) + M_2\mathcal{X}_t. \tag{3.5}$$

The coefficient matrices are given in Supplementary Appendix B. Let $\mathcal{X}_{t,t-1}$ denote the initialized Kalman predictor of $\mathcal{X}_t$ given $\{Y_s(\widetilde{F})\}_{s=1}^{t-1}$, and let $\eta_t = \eta_t(\theta; \widetilde{F})$ denote the corresponding prediction error for $Y_t(\widetilde{F})$. The Kalman filter yields the following innovations representation of the present-value model.

$$\mathcal{X}_{t+1,t} = \mathcal{F}\mathcal{X}_{t,t-1} + \mathcal{B}\widetilde{F}_{t-1} + \mathcal{F}\mathcal{K}_t\eta_t, \tag{3.6}$$

$$Y_t(\widetilde{F}) = M_0 + M_1 Y_{t-1}(\widetilde{F}) + M_2\mathcal{X}_{t,t-1} + \eta_t, \tag{3.7}$$

where $\mathcal{K}_t$ is the Kalman gain and $S_t(\theta)$ denotes the covariance of η_t.

Our bootstrap likelihood-ratio test for the state-space model with estimated factors in (3.4)–(3.5) is based on the following six-step algorithm.

1. Using the parameter estimate $\widehat{\theta}_0$ obtained by fitting the model under the null with $\widetilde{F}$, construct the series of standardized innovations and center it by setting
$$\widehat{e}_{0t} = S_t(\widehat{\theta}_0)^{-1/2}\eta_t(\widehat{\theta}_0; \widetilde{F}), \qquad \widehat{e}^c_{0t} = \widehat{e}_{0t} - T^{-1}\sum_{s=1}^{T}\widehat{e}_{0s}.$$
Let $\widehat{\Sigma}_{e,0} = T^{-1}\sum_{s=1}^{T}\widehat{e}^c_{0s}\,\widehat{e}^{c\prime}_{0s}$ be their sample covariance.
2. Draw $\{\widehat{e}^*_{0t}\}_{t=1}^T$ independently and uniformly with replacement from the vectors $\{\widehat{\Sigma}^{-1/2}_{e,0}\widehat{e}^c_{0s}\}_{s=1}^T$. Set $\eta^*_t = S_t(\widehat{\theta}_0)^{1/2}\widehat{e}^*_{0t}$.
3. Using the innovations representation (3.6)–(3.7) and the draws from Step 2, construct the bootstrap sample $\{Y^*_t\}_{t=1}^T$ as follows.
$$\mathcal{X}^*_{t+1,t} = \widehat{\mathcal{F}}\mathcal{X}^*_{t,t-1} + \widehat{\mathcal{B}}F^*_{t-1} + \widehat{\mathcal{F}}\widehat{\mathcal{K}}_t\eta^*_t, \tag{3.8}$$
$$Y^*_t = \widehat{M}_0 + \widehat{M}_1 Y^*_{t-1} + \widehat{M}_2\mathcal{X}^*_{t,t-1} + \eta^*_t, \tag{3.9}$$

where $Y_t^* = (\Delta d_t^*, pd_t^*, F_t^{*\prime})'$, $\widehat{\mathcal{F}}$, $\widehat{\mathcal{B}}$, $\widehat{\mathcal{K}}_t$, $\widehat{M}_0$, $\widehat{M}_1$, and $\widehat{M}_2$ are evaluated at the restricted estimate $\widehat{\theta}_0$, and the initial conditions are $Y_0^* = Y_0$ and $\mathcal{X}_{1,0}^* = 0$.

4. Draw external multipliers z_{it} iid across (i,t) from a distribution with mean zero and unit variance (e.g., standard normal). Apply the wild bootstrap to the residuals $\widetilde{e}_{it} = X_{it} - \widetilde{\lambda}_i'\widetilde{F}_t$ by setting

$$e_{it}^* = \widetilde{e}_{it} z_{it}, \qquad X_{it}^* = \widetilde{\lambda}_i' F_t^* + e_{it}^*. \tag{3.10}$$

   Re-estimate the factors from the bootstrap panel X^* to obtain $\widetilde{F}^*$.

5. Using the bootstrap sample $\{(\Delta d_t^*, pd_t^*, \widetilde{F}_t^{*\prime})'\}_{t=1}^T$, obtain the restricted and unrestricted maximum likelihood estimates $\widehat{\theta}_0^*$ and $\widehat{\theta}^*$, respectively. The bootstrap LR statistic is

$$LR_T^*(\widetilde{F}^*) = 2\left\{\mathcal{L}_T^*(\widehat{\theta}^*; \widetilde{F}^*) - \mathcal{L}_T^*(\widehat{\theta}_0^*; \widetilde{F}^*)\right\},$$

   where $\mathcal{L}_T^*$ denotes the bootstrap log-likelihood.

6. Repeat Steps 2–5 independently B times to obtain the bootstrap LR statistics $\{LR_{T,b}^*(\widetilde{F}_b^*)\}_{b=1}^B$. Their empirical distribution approximates the null distribution of $LR_T(\widetilde{F})$.

The bootstrap accounts for variation induced by the state-space innovations, which are resampled to generate the financial variables and latent factors jointly under the null in Steps 2–3 (Piatti and Trojani 2020). Following Gonçalves and Perron (2014), it also accounts for factor estimation uncertainty by generating X^* with a residual wild bootstrap and re-estimating the factors in Step 4. We next examine the validity of the bootstrap test.

3.3 Bootstrap validity

For a given bootstrap draw, let $LR_T^*(F^*)$ denote the LR statistic obtained by fitting the models in Step 5 with the generated factors F^* in place of $\widetilde{F}^*$. Write P^* for probability conditional on the original sample.

Theorem 3.1. *Under the null hypothesis, suppose Assumptions A, B, C(i)–(iii) and D hold. Let $N, T \to \infty$ with $\sqrt{T}/N \to 0$.*

(i) *The sample LR statistics satisfy*

$$LR_T(\widetilde{F}) - LR_T(F) = o_p(1). \tag{3.11}$$

(ii) *If the multipliers z_{it} satisfy $E^*|z_{it}|^4 \le M < \infty$ uniformly in i, t, N, T, then*

$$LR_T^*(\widetilde{F}^*) - LR_T^*(F^*) = o_{p^*}(1) \quad \textit{in probability.} \tag{3.12}$$

(iii) *If C(iv) also holds, then $LR_T(\widetilde{F}) \Rightarrow \chi_m^2$ and, under the multiplier condition in (ii),*

$$\sup_x \left| P^*\{LR_T^*(\widetilde{F}^*) \le x\} - P\{LR_T(\widetilde{F}) \le x\} \right| \to_p 0. \tag{3.13}$$

Under the stated assumptions, parts (i) and (ii) show that factor estimation has no first-order effect on the LR statistic, either in the original sample or conditionally in the bootstrap.

Under the additional conditions in part (iii), Theorem 3.1 justifies using the conditional bootstrap distribution to approximate the null distribution of $LR_T(\widetilde{F})$. With B independent replications, the corresponding bootstrap p-value is

$$\widehat{p}^*_{T,B} = \frac{1}{B}\sum_{b=1}^{B} \mathbf{1}\{LR^*_{T,b}(\widetilde{F}^*_b) > LR_T(\widetilde{F})\}.$$

At the nominal level α, the null hypothesis is rejected if $\widehat{p}^*_{T,B} < \alpha$.

4 Simulation

We examine the finite-sample size of the tests of no return and no dividend-growth predictability. This is the central inference problem identified by Piatti and Trojani (2020, Section 4.4), who document substantial over-rejection by asymptotic LR tests in short samples. In each Monte Carlo replication, we generate data from (2.2)–(2.5) and the factor model below, with one factor following $F_{t+1} = \xi_1 F_t + \varepsilon^F_{t+1}$. We calibrate the two DGPs using the restricted parameter estimates reported by Piatti and Trojani (2020) for the model with the book-to-market (BM) ratio as predictor: Panel A of their Table E.3 for H_0^{ret} and Panel A of their Table E.2 for H_0^{div}. The predictor mean ξ_0 is set to zero in both DGPs. Structural innovations are iid zero-mean Gaussian with the calibrated covariance matrix.[1]

The panel follows

$$X_{it} = \lambda_i F_t + e_{it}, \qquad \lambda_i \overset{\text{iid}}{\sim} U[-1,1], \qquad e_{it} \overset{\text{iid}}{\sim} N(0, 0.05^2), \tag{4.1}$$

with λ_i, e_{it}, and structural innovations mutually independent. We consider $N, T \in \{50, 100, 200\}$.

We compare the asymptotic LR test and the estimated- and observed-factor bootstraps in Table 4.1, with $S = 200$ Monte Carlo replications, $B = 100$ bootstrap draws and $\alpha = 0.05$. The proposed test uses standard normal wild multipliers. Let $LR_{NT,s} = LR_T(\widetilde{F}_s)$ and $LR^*_{NT,b,s}$ be the sample and bth bootstrap LR statistics in replication s. We compute

$$p^*(LR_{NT,s}) = \frac{1}{B}\sum_{b=1}^{B} \mathbf{1}\{LR^*_{NT,b,s} > LR_{NT,s}\}, \qquad \widehat{\alpha}^*_{NT} = \frac{1}{S}\sum_{s=1}^{S} \mathbf{1}\{p^*(LR_{NT,s}) < \alpha\}. \tag{4.2}$$

For the other two tests, empirical size is computed using their respective p-values.

Table 4.1 shows that asymptotic LR tests substantially over-reject in short samples, consistent with Piatti and Trojani (2020). At $T = 50$ (closest to our empirical sample size), rejection frequencies are 23.5–27.0% for no return predictability and 11.0% for no dividend-growth predictability at the nominal 5% level. The proposed bootstrap yields 4.0–5.5% and 4.0–4.5%, respectively, both close to the nominal 5% level. The observed-factor bootstrap yields broadly similar rejection frequencies. As T increases, the size distortions of the asymptotic LR tests decline. For no return

[1]The Kalman filter is initialized with the unconditional mean and covariance of the centered state. Population log-linearization constants remain fixed throughout estimation and resampling: $\rho = e^{-h}$, $A = -\log(e^h - 1)$, and $\kappa = \log(1 + e^A) - \rho A$, where $h = \delta_0 - \gamma_0$.

Table 4.1: Finite-sample sizes of asymptotic and bootstrap tests

		No return predictability			No dividend-growth predictability		
T	N	Asymptotic	Est. factor	Obs. factor	Asymptotic	Est. factor	Obs. factor
50	50	27.0	5.5	4.5	11.0	4.5	4.5
50	100	25.0	4.0	4.5	11.0	4.0	4.5
50	200	23.5	4.0	4.5	11.0	4.5	4.5
100	50	17.0	5.0	5.5	8.5	4.5	5.5
100	100	21.0	4.5	5.5	9.0	4.5	5.5
100	200	18.0	6.5	5.5	9.0	4.0	5.5
200	50	15.5	4.5	6.0	4.5	3.5	4.0
200	100	13.5	4.5	6.0	4.5	4.0	4.0
200	200	15.5	5.5	6.0	4.0	2.5	4.0

Notes: Entries report empirical size, measured as rejection frequencies in percent under the null at the 5% nominal level. We use $S = 200$ Monte Carlo replications and $B = 100$ resamples for each bootstrap test. Asymptotic denotes the LR test using estimated factors and chi-square critical values. Est. factor denotes the proposed bootstrap with estimated factors. Obs. factor denotes the innovation bootstrap of Piatti and Trojani (2020) with the true factor observed.

predictability, however, rejection frequencies remain at 13.5–15.5% when $T = 200$. For tests based on estimated factors, rejection frequencies vary modestly with N. This suggests that increasing the cross-sectional dimension alone offers limited gains in size control, highlighting the importance of bootstrap inference even at larger N. The proposed bootstrap therefore improves finite-sample size control for tests based on estimated factors, particularly in short samples, reducing the probability of finding predictive relations when none exist.

5 Empirical analysis

5.1 Empirical design

We obtain monthly returns with and without dividends on the value-weighted portfolio of all New York Stock Exchange, Amex, and Nasdaq stocks from January 1959 through December 2023 from the Center for Research in Security Prices (CRSP). We construct annual time series of aggregate dividends and prices, assuming that monthly dividends are cash-reinvested at the 30-day T-bill rate, and retain the period 1960–2023. Data on 30-day T-bill rates are also obtained from CRSP. We use annual observations to reduce the influence of seasonal dividend payments on inference and facilitate comparison with earlier present-value studies (van Binsbergen and Koijen 2010; Piatti and Trojani 2020).

We extract factors from 226 quarterly FRED-QD series (McCracken and Ng 2021), after annual averaging, stationarity transformations, and standardization. Three observed predictors from Welch and Goyal (2008) are book-to-market (BM), stock variance (SVAR), and the consumption–

wealth–income ratio (CAY). We report the same observed predictors as Piatti and Trojani (2020) for direct comparison.

The IC_{p1}–IC_{p3} criteria of Bai and Ng (2002) all select eight factors, while the ER and GR estimators of Ahn and Horenstein (2013) select one and the ED estimator of Onatski (2010) selects two. Such differences are unsurprising in finite samples, as these methods rely on penalized fit or different features of the eigenvalue spectrum. Moreover, factor-model theory does not guarantee that factors explaining panel variation are equally informative about dividend growth or returns. We therefore consider a parsimonious sequence of specifications, cumulatively adding the first four principal components in descending order of their contribution to panel variation. Given the short annual sample, this sequence balances information coverage against model complexity and allows us to assess the sensitivity of predictability inference to factor dimension.

We test the null hypotheses of no dividend-growth predictability and no return predictability, reporting LR statistics and their Chi-square and bootstrap p-values. We use the innovation bootstrap of Piatti and Trojani (2020) for the benchmark and models with observed predictors, and our proposed bootstrap for models with estimated factors. Both use $B = 1{,}000$ replications.

5.2 Empirical results

Supplementary Tables E.1–E.2 report the benchmark and unrestricted extended-model estimates; Table E.3 reports the full constrained estimates. The benchmark has fitted R^2 values of 7.15% for returns and 15.37% for dividend growth. Its bootstrap p-values are 5.70% and 13.80%, respectively, so neither null is rejected at 5%, although both are rejected by the asymptotic LR tests. This difference is consistent with Piatti and Trojani (2020), who document that asymptotic LR tests over-reject the null of no predictability in finite samples.

We report results for the same observed predictors as Piatti and Trojani (2020), who select BM and SVAR for the dividend-growth predictability test and BM and CAY for the return predictability test based on the largest model-implied R^2 gains in their sample. In our sample, bootstrap inference does not reject constant expected dividend growth for BM or SVAR at 5% (Table 5.1), with p-values of 15.30% and 22.30%, in line with their findings. The bootstrap test rejects the null of no dividend-growth predictability at the 10% level for F_1 ($p = 9.50\%$) and at the 5% level for (F_1, F_2) ($p = 4.40\%$), although the three- and four-factor specifications do not reject at 5%. This finding extends the observed-predictor evidence of Piatti and Trojani (2020) by showing that common macroeconomic and financial factors can reveal dividend-growth predictability that observed-predictor tests do not detect.

Table 5.2 gives bootstrap p-values of 3.70% for BM, 1.40% for CAY, 0.90% for F_1, and 0.60% for (F_1, F_2). The BM and CAY tests reject constant expected returns at 5%, in line with Piatti and Trojani (2020), and the one- and two-factor models reject at 1%. Across the four factor

Table 5.1: Test of no dividend-growth predictability

$H_0^{\text{div}}: \gamma_1 = \gamma_{2,j} = \sigma_g = \rho_{g\mu} = \rho_{gz_j} = 0 \ (j = 1, \ldots, r)$

	BM	SVAR	F_1	(F_1, F_2)
Panel A: Selected constrained estimates				
γ_0	0.058	0.059	0.057	0.057
δ_0	0.084	0.084	0.082	0.082
γ_1	0	0	0	0
δ_1	0.793	0.936	0.899	0.910
$\gamma_{2,1}$	0	0	0	0
$\gamma_{2,2}$	–	–	–	0
$\delta_{2,1}$	0.005	−0.035	0.032	0.040
$\delta_{2,2}$	–	–	–	0.007
σ_g	0	0	0	0
σ_μ	0.027	0.047	0.035	0.035
σ_d	0.067	0.067	0.067	0.067
$\rho_{g\mu}$	0	0	0	0
$\rho_{\mu d}$	0.367	−0.143	−0.321	−0.274
Panel B: Asymptotic test				
LR	13.70	14.58	12.27	18.97
p-value (%)	1.76	1.23	3.12	0.83
Panel C: Bootstrap test				
p-value (%)	15.30	22.30	9.50	4.40

Notes: BM is book-to-market and SVAR stock variance. z_j denotes the jth predictor, either an observed variable or an estimated factor. Bootstrap p-values are based on 1000 replications.

Table 5.2: Test of no return predictability

$H_0^{\text{ret}}: \delta_1 = \delta_{2,j} = \sigma_\mu = \rho_{g\mu} = \rho_{\mu d} = \rho_{\mu z_j} = 0 \ (j = 1, \ldots, r)$

	BM	CAY	F_1	(F_1, F_2)
Panel A: Selected constrained estimates				
γ_0	0.059	0.062	0.061	0.062
δ_0	0.085	0.085	0.084	0.084
γ_1	0.888	0.995	0.993	0.991
δ_1	0	0	0	0
$\gamma_{2,1}$	−0.003	−0.001	−0.003	−0.009
$\gamma_{2,2}$	–	–	–	0.003
$\delta_{2,1}$	0	0	0	0
$\delta_{2,2}$	–	–	–	0
σ_g	0.015	0.006	0.006	0.012
σ_μ	0	0	0	0
σ_d	0.066	0.069	0.067	0.067
$\rho_{g\mu}$	0	0	0	0
$\rho_{\mu d}$	0	0	0	0
Panel B: Asymptotic test				
LR	24.77	26.11	23.72	28.71
p-value (%)	0.04	0.02	0.06	0.04
Panel C: Bootstrap test				
p-value (%)	3.70	1.40	0.90	0.60

Notes: BM is book-to-market and CAY the consumption–wealth–income ratio. z_j denotes the jth predictor, either an observed variable or an estimated factor. Bootstrap p-values are based on 1000 replications.

specifications, the bootstrap p-values are 0.90%, 0.60%, 0.70%, and 5.00%, respectively. All four therefore reject constant expected returns at 10%, with the first three also rejecting at 1%. The two-factor specification thus preserves the rejection of no return predictability while adding new evidence of dividend-growth predictability.

6 Conclusions

This paper develops bootstrap LR inference for present-value models with common factors estimated from a large macroeconomic and financial panel. Under suitable regularity and rate conditions, the two-step estimator is consistent, and factor estimation has no first-order effect on the LR statistic or its bootstrap distribution. Simulations show improved finite-sample size control relative to asymptotic LR tests, especially in short samples. Using annual U.S. data from 1960 to 2023, the bootstrap rejects constant expected dividend growth at 5% in the joint two-factor specification, but not in the BM or SVAR specifications. The two-factor specification also retains evidence of return predictability. Overall, the paper provides a theoretical foundation for bootstrap inference with estimated factors within the present-value framework and shows empirically that incorporating high-dimensional information through a factor model can reveal dividend-growth predictability undetected in individual observed-predictor specifications.

References

Ahn, S. C. and Horenstein, A. R. (2013), 'Eigenvalue ratio test for the number of factors', *Econometrica* **81**(3), 1203–1227.

Bai, J. (2003), 'Inferential theory for factor models of large dimensions', *Econometrica* **71**(1), 135–171.

Bai, J. and Ng, S. (2002), 'Determining the number of factors in approximate factor models', *Econometrica* **70**(1), 191–221.

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.

Chen, L. (2009), 'On the reversal of return and dividend growth predictability: A tale of two periods', *Journal of Financial Economics* **92**(1), 128–151.

Chen, L., Da, Z. and Priestley, R. (2012), 'Dividend smoothing and predictability', *Management Science* **58**(10), 1834–1853.

Cheng, T., Yang, X. and Zhao, A. B. (2024), 'Predicting the predictable: Decomposing and forecasting stock returns in a data-rich environment', *SSRN Electronic Journal* .

Cochrane, J. H. (2008), 'The dog that did not bark: A defense of return predictability', *Review of Financial Studies* **21**(4), 1533–1575.

Fama, E. F. and French, K. R. (1988), 'Dividend yields and expected stock returns', *Journal of Financial Economics* **22**(1), 3–25.

Gonçalves, S. and Perron, B. (2014), 'Bootstrapping factor-augmented regression models', *Journal of Econometrics* **182**(1), 156–173.

Guo, M., Xu, Q. and Yang, X. (2026), 'Market expectations in the cross-section of leading present values', *SSRN Electronic Journal* .

Kelly, B. and Pruitt, S. (2013), 'Market expectations in the cross-section of present values', *Journal of Finance* **68**(5), 1721–1756.

Lettau, M. and Ludvigson, S. C. (2005), 'Expected returns and expected dividend growth', *Journal of Financial Economics* **76**(3), 583–626.

McCracken, M. W. and Ng, S. (2021), 'FRED-QD: A quarterly database for macroeconomic research', *Federal Reserve Bank of St. Louis Review* **103**(1), 1–44.

Onatski, A. (2010), 'Determining the number of factors from empirical distribution of eigenvalues', *The Review of Economics and Statistics* **92**(4), 1004–1016.

Pettenuzzo, D., Sabbatucci, R. and Timmermann, A. (2020), 'Cash flow news and stock price dynamics', *Journal of Finance* **75**(4), 2221–2270.

Piatti, I. and Trojani, F. (2020), 'Dividend growth predictability and the price–dividend ratio', *Management Science* **66**(1), 130–158.

van Binsbergen, J. H. and Koijen, R. S. J. (2010), 'Predictive regressions: A present-value approach', *Journal of Finance* **65**(4), 1439–1471.

Welch, I. and Goyal, A. (2008), 'A comprehensive look at the empirical performance of equity premium prediction', *Review of Financial Studies* **21**(4), 1455–1508.

Yu, D. and Chen, L. (2024), 'Local predictability of stock returns and cash flows', *Journal of Empirical Finance* **77**, 101485.

Yu, D., Chen, L. and Li, L. (2023), 'Time-varying predictability of the long horizon equity premium based on semiparametric regressions', *Economics Letters* **224**, 111033.

Yu, D., Huang, D. and Chen, L. (2023), 'Stock return predictability and cyclical movements in valuation ratios', *Journal of Empirical Finance* **72**, 36–53.

Yu, D. and Yan, Y. (2023), 'Joint dynamics of stock returns and cash flows: A time-varying present-value framework', *Financial Management* **52**(3), 513–541.