# Do Dividend Strips Force a Bubble? Long-Horizon Discounting and the No-Bubble Restriction

Yukun Cao*

May 2026

*Department of Economics and Related Studies, University of York, United Kingdom.

**Abstract**

We test whether a flexible no-bubble present-value model can accommodate equity-index valuations once traded dividend strips pin down the near-end of cash-flow claims. The test exploits a decomposition in which observed strips anchor the near end, an affine no-arbitrage model prices the untraded fundamental tail, and an intrinsic bubble component absorbs any remaining long-end valuation. We estimate the model on monthly S&P 500 data using rank-one time-varying prices of risk and compare three information sets: a spot-only likelihood, a joint-strips likelihood, and a maturity hold-out design that estimates on one-to-five-year strips and validates on unseen six-to-ten-year strips. In spot data alone, the unrestricted and no-bubble specifications are observationally equivalent (LR = 0.003, $p$ = 0.956). Once strips enter the likelihood, the no-bubble restriction is rejected (LR = 113.55, $p < 0.001$), and strip-pricing RMSE falls from 0.214 to 0.120. The rejection also appears in the maturity hold-out design (LR = 153.68, $p < 0.001$). The unrestricted model's gain is concentrated in the residual long-end wedge, the transversality margin, and the term structure of expected strip excess returns, not in short-run return fit. The evidence suggests that traded strips provide identifying information about long-horizon discounting that spot data alone do not reveal.

**Keywords:** Dividend strips; intrinsic bubbles; long-horizon discounting; present-value models; expected returns; equity term structure
**JEL classification:** G12; G13; G14; C58

# 1 Introduction

A spot equity index price is a claim on an infinite stream of future dividends. That aggregation is useful for valuation, but it creates a sharp identification problem: the spot price alone does not reveal how much of valuation is generated by cash flows near the front of the horizon and how much is generated by discounting at the long end. This distinction matters for any present-value restriction that depends on the behaviour of distant cash flows. In particular, it matters for the no-bubble restriction. This paper asks whether a flexible no-bubble present-value model remains empirically viable once traded dividend strips are used to discipline the near end of the valuation decomposition.

The question is motivated by the modern interpretation of the excess-volatility puzzle. Shiller (1981) shows that stock prices appear too volatile relative to subsequent realised dividends. The response developed by Campbell & Shiller (1988), Fama & French (1988), and Cochrane (2011) is that valuation ratios can move substantially because expected returns vary over time. A high price-dividend ratio is therefore not, by itself, evidence of a bubble; it may reflect low discount rates, especially at long horizons. This logic imposes an important discipline on any empirical bubble test. A rejection of the no-bubble restriction is informative only if the benchmark has already been given enough discount-rate flexibility that the rejection cannot be dismissed as a mechanical consequence of a rigid expected-return specification.

Dividend strips make this discipline operational. A strip is a traded claim to dividends over a specific future horizon. A panel of strips therefore provides market prices for individual parts of the cash-flow tree rather than only for their aggregate. Once these strips are combined with the spot index price, the first decade of the present-value decomposition is observed from market prices rather than inferred from the model (van Binsbergen et al., 2012, 2013; van Binsbergen & Koijen, 2017). The remaining valuation object is then a long-end residual. In the model developed below,

$$\frac{P_t}{D_t} = N_t + \mathrm{Tail}(X_t) + \mathrm{Bub}(X_t),$$

where $N_t$ is the observed near-end component constructed from traded one-to-ten-year dividend strips, $\mathrm{Tail}(X_t)$ is the model-implied present value of cash flows beyond the traded strip horizon, and $\mathrm{Bub}(X_t)$ is an intrinsic bubble component in the spirit of Froot & Obstfeld (1991). Under the no-bubble restriction, the bubble scaling coefficient is zero, and the fundamental tail alone must explain $P_t/D_t - N_t$. The test is therefore not whether a nonlinear term can improve a spot-price regression. It is whether the fundamental tail can absorb the residual long-end wedge after the near end has been anchored by traded claims.

We estimate this restriction in a dynamic affine no-arbitrage model with time-varying prices of risk. The state vector is four-dimensional and contains the short rate and principal components extracted from a predictor panel. The stochastic discount factor has rank-one time-varying prices of risk, so expected returns are allowed to vary through a persistent state-dependent channel.

The restricted model is not a constant-discount-rate straw man. It is a no-bubble benchmark with a flexible discount-rate channel; traded strips are used not to compensate for a weak benchmark, but to discipline a benchmark that already has substantial freedom to move expected returns.

The empirical design holds the structural model fixed and varies only the information set used in estimation. The first design is spot-only: the likelihood uses market excess returns, dividend growth, and the spot valuation ratio, but excludes the strip panel. This is the environment most favourable to the no-bubble restriction because all long-end discipline must come indirectly through spot-market observables. The second design is joint-strips: the strip panel enters the likelihood directly, so the near end of valuation is anchored by market prices during estimation. The third design is a maturity hold-out exercise: the model is estimated on one-to-five-year strips and evaluated on unseen six-to-ten-year strips. The hold-out design asks whether the strip-based conclusion survives in maturity space rather than merely fitting the maturities used in estimation.

The central result is that the no-bubble conclusion changes once strips enter the information set. In the spot-only design, the unrestricted and restricted models have virtually identical likelihoods: LR = 0.003 with $p = 0.956$. The no-bubble restriction therefore survives when the model is estimated only on spot-market information. In the joint-strips design, the same restriction is rejected: LR = 113.55 with $p < 0.001$, and the strip-pricing RMSE falls from 0.214 under the restricted model to 0.120 under the unrestricted model. The result is not an artefact of fitting the observed maturities. In the maturity hold-out design, the model is estimated without the six-to-ten-year strips, yet those unseen longer maturities continue to favour the unrestricted specification: LR = 153.68, $p < 0.001$, and the validation-set strip RMSE falls from 0.224 to 0.139.

The mechanism is concentrated at the long end. The unrestricted model does not primarily improve short-run return fit, instead, it changes the configuration of long-horizon discounting. First, once strips anchor the near end, the restricted model is pushed close to the transversality boundary: in the joint-strips and hold-out designs, the restricted model's transversality margin is approximately $10^{-6}$, whereas the unrestricted model has margins of about $1.5 \times 10^{-3}$ and $2.0 \times 10^{-3}$, respectively. Second, the unrestricted specification implies higher short- and medium-horizon expected strip excess returns than the restricted specification. In the joint-strips design, for example, the twelve-month expected strip excess return is 0.0048 under the unrestricted model and 0.0006 under the no-bubble restriction. Third, the improvement in valuation fit is concentrated in the residual wedge $P_t/D_t - N_t$, the part of valuation that remains after the observed near end has been removed. These patterns suggest that strips alter the no-bubble conclusion through long-horizon discounting and transversality pressure rather than through a generic improvement in short-run moments.

The evidence does not establish that a literal rational bubble of a fixed size exists in every month. It suggests that traded strips reveal a large and persistent long-end component that a flexible no-bubble model struggles to absorb without relaxing the no-bubble restriction. In this sense, the estimated intrinsic bubble is best viewed as a disciplined way to parameterise residual long-end valuation after the first decade of cash-flow claims has been priced in the market.

A natural concern is that traded dividend-strip data are available over a short and potentially unrepresentative sample. The paper addresses this concern in several ways. The baseline comparison holds the sample fixed and varies only whether strips enter the likelihood, so the central object is the marginal contribution of strip discipline rather than the unconditional level or slope of the equity term structure. A maturity hold-out design shows that the result extends to unseen longer maturities. We investigate a long-sample strip-free external validation exercise to estimate the model on January 2000 to December 2024 spot data and evaluate it against the 2017–2024 strip panel; extending the spot-only sample does not substitute for direct strip discipline. Additional diagnostics vary the treatment of strip measurement errors, shut down time variation in prices of risk, sweep the state-vector dimension, and lengthen the model-implied tail. These exercises change the sharpness of the test, but they do not restore a clean no-bubble conclusion once strips discipline the near end directly.

The paper contributes to three strands of literature. First, relative to the excess-volatility and return-predictability literature, it shows that a flexible discount-rate channel that is sufficient to protect the no-bubble restriction in spot data alone is not sufficient once the near end is observed from traded strips. The point is not to reject discount-rate variation; it is to show where discount-rate flexibility becomes constrained by maturity-specific market prices. Second, relative to the rational-bubble literature, the paper embeds the Froot & Obstfeld (1991) logic in a dynamic affine system with observed strips and a model-implied fundamental tail. The identifying variation comes from the residual long-end wedge rather than from a spot-only nonlinear valuation relation. Third, relative to the dividend-strip literature, the paper uses strips not primarily to describe the equity term structure, but to test the survival of a structural restriction. The closest comparison is Giglio et al. (2024), who show that rich affine models can recover realistic equity term structures without using strip data in estimation. The present paper asks the complementary question: once traded strips themselves enter the information set, does a no-bubble restriction that survives without strip discipline remain viable? We find that the restriction that survives without strip discipline no longer survives with it.

The rest of the paper proceeds as follows. Section 2 positions the paper relative to the main literatures. Section 3 develops the affine pricing model, the strip-pricing recursion, and the tail-plus-bubble decomposition. Section 4 describes the data, state-vector construction, and information designs. Section 5 presents the core identification results. Section 6 reports the validation hierarchy and design diagnostics. Section 7 studies how strips alter long-horizon discounting. Section 8 interprets the price decomposition and the estimated long-end component. Section 9 concludes.

# 2 Related literature and contribution

This paper contributes to three strands of literature: the excess-volatility and discount-rate-variation literature, the literature on rational bubbles and structural restrictions, and the literature on dividend strips and equity term structures. In each case, the contribution is not to add one more empirical finding to an existing body of evidence. It is to show that the conclusion one draws about a structural restriction depends materially on whether traded strips are allowed to discipline long-horizon discounting.

## 2.1 Excess volatility, return predictability, and long-horizon discounting

The starting point for this paper is the modern resolution of the excess-volatility puzzle. Shiller (1981) documents that stock prices appear too volatile relative to subsequent realised dividends. Campbell & Shiller (1988) and Fama & French (1988) recast that observation in present-value form by linking the dividend-price ratio to expected returns and expected dividend growth. The synthesis in Cochrane (2011) argues that the bulk of variation in valuation ratios reflects time-varying discount rates rather than expected cash-flow growth. Within this framework, a high valuation ratio is not prima facie evidence for a bubble; it may instead reflect a low long-horizon discount rate.

This literature implies a demanding standard for any paper that claims to identify a non-fundamental price component. A serious test of the no-bubble restriction must first grant the benchmark model enough discount-rate flexibility to absorb valuation movements through expected returns. If the benchmark is too restrictive, rejection tells the reader about model misspecification rather than about bubbles. The present paper takes this lesson as its starting point. The no-bubble benchmark is given a dynamic affine SDF with rank-one time-varying prices of risk, and the state vector is chosen to be large enough that the spot-only design cannot reject the restriction. The contribution to this literature is therefore not to revisit whether discount-rate variation matters. It is to ask whether a benchmark that already incorporates such variation can still survive once the near end of valuation is directly observed from traded dividend strips.

That question connects to the long-run risk literature. Bansal & Yaron (2004) and subsequent work (Bansal et al., 2012, 2014, 2016) emphasise that persistent expected-growth and volatility components can have first-order effects on asset prices. The present paper does not evaluate long-run risk models per se. Rather, it adopts the discipline this literature has established: any structural test that is sensitive to the far end of discounting must embed discount-rate dynamics of sufficient richness. The paper then asks whether that richness remains sufficient once strips constrain the near end, converting a broad present-value argument into a narrower identification exercise in which the residual long-end wedge is the object of interest.

## 2.2 Rational bubbles and the identification problem

A second strand concerns the theoretical possibility and empirical identification of rational bubbles. Tirole (1985) establishes conditions under which speculative price components can exist in equilibrium. The econometric difficulty is that standard time-series tests have limited power to distinguish explosive bubble components from highly persistent fundamentals, a point made forcefully by Diba & Grossman (1988) and Evans (1991). Within this literature, Froot & Obstfeld (1991) introduce a class of bubbles that are nonlinear functions of dividends rather than extraneous sunspots, allowing persistent deviations from a linear present-value benchmark without severing the link to cash-flow processes.

Our paper builds on Froot & Obstfeld (1991) but differs along three dimensions that are relevant for identification. First, the empirical environment is substantially richer: rather than working with the aggregate price-dividend relation alone, the model is estimated as a joint no-arbitrage system for market returns, dividend growth, the valuation ratio, and a panel of ten traded strip prices. Second, the intrinsic-bubble component is embedded in an affine framework with time-varying prices of risk, so the no-bubble benchmark is not a constant-discount-rate straw man. Third, the near end of the valuation decomposition is anchored by observed strip prices rather than inferred from spot data. These differences shift the identification problem. The question is no longer whether a nonlinear bubble term can improve the fit of a spot-only regression. It is whether the no-bubble restriction survives once the first decade of cash-flow claims is removed from the model's discretion and the residual long-end wedge must be absorbed by the fundamental tail alone.

This identification approach also separates the paper from a growing literature that uses dividend derivatives to measure risk premia, expected growth, or bubble-like residuals in a reduced-form manner. Avino et al. (2021), Casta (2022), and Branger et al. (2024) show that derivative-implied quantities can be informative about valuation components. The present paper uses strips differently. Rather than computing a residual wedge and interpreting it after the fact, it embeds traded strips in a structural likelihood and tests whether a specific restriction, the no-bubble condition, remains viable in the presence of strip discipline. That distinction matters because a reduced-form wedge can always be computed; what is more demanding is to show that a flexible structural model can no longer accommodate the restriction once the near end is observed.

## 2.3 Dividend strips, equity term structures, and the short-sample problem

The third strand is the literature on dividend strips and the equity term structure. van Binsbergen et al. (2012, 2013); van Binsbergen & Koijen (2017) show that traded claims to dividends at specific horizons provide a direct empirical window into maturity-specific discounting. That development moved the term structure of risky discount rates from an indirect inference problem to an object that can, over part of the maturity spectrum, be measured from market prices.

At the same time, the strip literature has generated a substantive empirical debate rather than a settled fact pattern. Bansal et al. (2021) argue that inference from traded strip data must be interpreted with care because the available futures sample is short, begins only in the mid-2000s, and may oversample recessionary conditions. Golez & Jackworth (2024) extend the options-based sample and show that measurement error and holding-period choice materially affect inferences about strip returns and Sharpe ratios. These concerns are legitimate and directly relevant to the present paper. Any strip-based conclusion that ignores the short-sample problem risks overstating the information content of a limited and potentially unrepresentative data window.

The present paper is designed to remain informative in the presence of these concerns rather than to dismiss them. The empirical strategy does not ask strips to recover the unconditional shape of the entire equity term structure. Nor does it treat the strip panel as a stand-alone description of long-horizon discounting. Instead, it uses strips in a narrower role: to anchor the near end of the present-value decomposition and to test whether that incremental information alters the viability of a structural restriction. The information-design framework is constructed so that the comparison between strip-free and strip-using estimation isolates the marginal contribution of strip discipline. This approach limits the paper's exposure to the short-sample critique, because the claim is not about the unconditional level or slope of the equity term structure. It is about whether the no-bubble restriction survives a shift in the information set that disciplines the near end.

## 2.4 Relation to Giglio et al. (2024)

The study more closer to the present paper is Giglio et al. (2024). They estimate a rich affine model on a large cross section of equity portfolios without using strip data in estimation, and then validate the model-implied term structures against observed traded strips. Their approach addresses several limitations of the direct strip sample simultaneously: the short time span, the sparse cross section, and concerns about liquidity. The result is an equity term-structure object that extends back to the 1970s and covers a much broader maturity range than is directly traded.

The two papers are complementary rather than competing. Giglio et al. (2024) ask whether one can recover realistic term structures without strip data and then validate them externally. The present paper asks a different question: once traded strips themselves enter the information set, does a structural no-bubble restriction remain viable? The designs therefore operate in opposite directions. Giglio et al. (2024) demonstrate that strip-free estimation can reproduce strip prices. The present paper demonstrates that strip-using estimation changes a structural conclusion that would otherwise survive in spot data alone.

This contrast also clarifies the role of the long-sample strip-free external validation exercise in Section 6.4. That exercise mirrors the logic of Giglio et al. (2024) by estimating on a 25-year spot-only sample and evaluating against observed strips. Its purpose is not to serve as the main empirical pillar of the paper. It is to show that extending the spot-only sample does not substitute for direct strip discipline along both the level and the dynamics of the strip term structure. The paper's identification advantage lies elsewhere: in the information-design comparison that holds the sample fixed and varies only whether strips enter the likelihood.

### 2.5 Decomposition and measurement

Finally, the paper is related to studies that use dividend derivatives or other market-based observables to decompose price movements across horizons. Gormsen & Koijen (2020) use stock and dividend futures to separate growth expectations across horizons during the COVID-19 crisis. Knox & Vissing-Jorgensen (2022) combine real yields, option-implied premia, and dividend futures to decompose realised returns period by period. Gormsen (2021) shows that the slope of the equity term structure is countercyclical and difficult to reconcile with one-factor environments, while Gonçalves (2021) proposes reinvestment risk as a mechanism for rationalising low long-maturity premia.

These papers share the idea that maturity-specific claims help identify which part of the valuation object is moving and why. The present paper contributes to this literature by shifting the use of strips from decomposition to structural testing. Rather than using strip prices to infer expected growth or to apportion realised returns ex post, the paper uses them to determine whether a structural no-bubble model remains empirically viable once horizon-specific prices are observed. The bubble application is the focus here because it is especially sensitive to transversality and to the treatment of the far end. But the broader methodological contribution is that traded strips can be used not only to describe term structures, but also to test the survival of structural restrictions that would otherwise appear less demanding in spot data alone.

## 3 Model

The model must accomplish three tasks. First, it must price dividend claims at all horizons within a no-arbitrage affine framework that permits time-varying expected returns. Second, it must decompose the spot price into a component observed from traded strips and a component inferred by the model. Third, it must nest both the no-bubble restriction and an intrinsic bubble alternative within the same pricing kernel, so that the restriction can be tested without changing the underlying discount-rate specification. The subsections below develop these elements in sequence.

### 3.1 State dynamics, dividends, and the stochastic discount factor

Time is discrete at a monthly frequency, $t = 0, 1, 2, \ldots$. The economy is driven by a $K$-dimensional predictive state vector $X_t \in \mathbb{R}^K$, which evolves under the physical measure $\mathbb{P}$ as a

Gaussian VAR(1):

$$X_{t+1} = \Phi_x X_t + \varepsilon_{x,t+1}, \qquad \varepsilon_{x,t+1} \sim N(0, \Sigma), \qquad \rho(\Phi_x) < 1. \tag{1}$$

The one-month nominal short rate is affine in the state,

$$y_t = \mu_y + e_1' X_t, \tag{2}$$

where $e_1$ selects and rescales the first element of $X_t$ so that $y_t$ remains in original units.

The nominal stochastic discount factor takes the standard log-normal affine form:

$$M_{t+1} = \exp\left(-y_t - \frac{1}{2}\lambda_t' \Sigma \lambda_t - \lambda_t' \varepsilon_{x,t+1}\right), \tag{3}$$

with the market price of state risk affine in the state,

$$\lambda_t = \lambda_0 + \Lambda_1 X_t. \tag{4}$$

To retain parsimony while allowing economically meaningful time variation in expected returns, we impose

$$\lambda_0 = \begin{bmatrix} \lambda_{0y} & 0 & \cdots & 0 \end{bmatrix}', \qquad \Lambda_1 = \ell\ell', \qquad \|\ell\| = 1. \tag{5}$$

Time variation in risk premia is thus driven by a single index $\ell' X_t$. This rank-one restriction serves a specific purpose in the identification design: it ensures that the no-bubble benchmark already possesses a non-trivial channel for time-varying expected returns before any bubble component is introduced, so that a rejection of the restriction cannot be attributed to an overly rigid discount-rate specification.

Define the risk-neutral transition matrix

$$\Phi_x^Q \equiv \Phi_x - \Sigma\Lambda_1. \tag{6}$$

Log dividends follow a random walk with drift:

$$d_{t+1} = d_t + \mu_d + \varepsilon_{d,t+1}, \qquad \varepsilon_{d,t+1} \sim N(0, \sigma_d^2), \tag{7}$$

where $d_t \equiv \log D_t$ and $D_t$ denotes the trailing twelve-month dividend level. The covariance between state and dividend innovations is

$$\gamma_d \equiv \mathrm{Cov}(\varepsilon_{x,t+1}, \varepsilon_{d,t+1}) \in \mathbb{R}^K. \tag{8}$$

The law of motion for $D_t$ is a reduced form for the annual dividend level; it does not model within-year payout timing. The within-year allocation of dividends across calendar months is handled through deterministic seasonality weights introduced in the next subsection, so that all economic uncertainty is carried by $D_t$ while the month-of-year mapping serves only as a timing device.

## 3.2 Pricing monthly dividend claims

Observed dividend futures refer to annual constant-maturity horizons, but the state vector and the stochastic discount factor evolve monthly. The present value of the annual dividend level at a monthly horizon $u$:

$$\Pi_{u,t} \equiv E_t\left[\left(\prod_{s=1}^{u} M_{t+s}\right) D_{t+u}\right], \qquad u = 1, 2, \ldots \tag{9}$$

Conjecture the exponential-affine form

$$\log \Pi_{u,t} = d_t + a_u + b_u' X_t. \tag{10}$$

Monthly cash dividends are then obtained from the annual level through a deterministic calendar-month mapping,

$$D_{t+u}^{(1m)} = \omega_{m(t+u)} D_{t+u}, \qquad \omega_m > 0, \qquad \sum_{m=1}^{12} \omega_m = 1, \tag{11}$$

where $m(t+u) \in \{1, \ldots, 12\}$ denotes the calendar month of date $t+u$ and $\omega_m$ is the typical share of annual dividends paid in month $m$, estimated from within-year cash-dividend data. The model price of the monthly dividend claim is then

$$P_t^{(um)} = \omega_{m(t+u)} \Pi_{u,t} = D_t\, \omega_{m(t+u)} \exp(a_u + b_u' X_t). \tag{12}$$

Substituting the stochastic discount factor, the state dynamics, and the dividend process into the one-step Euler recursion $\Pi_{u+1,t} = E_t[M_{t+1}\Pi_{u,t+1}]$ yields exact exponential-affine recursions for the pricing coefficients:

$$a_{u+1} = a_u + \mu_d - \mu_y - \gamma_d'\lambda_0 + b_u'(\gamma_d - \Sigma\lambda_0) + \frac{1}{2}\sigma_d^2 + \frac{1}{2} b_u' \Sigma b_u, \tag{13}$$

$$b_{u+1}' = b_u' \Phi_x^Q - e_1' - \gamma_d' \Lambda_1. \tag{14}$$

The full derivation and the one-month base coefficients are reported in Appendix A. The model thus delivers closed-form prices for monthly dividend claims at all horizons, and these claims serve as the building blocks for both the observed strip panel and the unobserved long-end tail.

## 3.3 From monthly claims to observed strips

Listed dividend futures on equity indices settle at fixed calendar dates—typically the third Friday of December each year—so the time to maturity of any given contract shrinks day by day. The raw futures prices therefore do not correspond to constant-maturity claims. To obtain strip observables that are comparable across time and that match the model's horizon-indexed pricing structure, the raw contracts must be interpolated into constant-maturity annual strips. Following Kragt et al. (2020), this interpolation accounts for the uneven within-year pattern of aggregate dividend payments; the detailed construction is deferred to Section 4 and Appendix D. The result is a panel of constant-maturity strip present values $P^{\text{strip,obs}}_{t,k}$ for $k = 1, \ldots, 10$ future dividend years.

The model counterpart of the $k$-year strip is obtained by summing the twelve monthly claims spanning that dividend year. Substituting the closed-form monthly claim price from Section 3.2, gives:

$$P^{(ky)}_{t,k} = \sum_{u=12(k-1)+1}^{12k} P^{(um)}_t = D_t \sum_{u=12(k-1)+1}^{12k} \omega_{m(t+u)} \exp(a_u + b_u' X_t), \tag{15}$$

where $\omega_{m(t+u)}$ denotes the seasonality weight for the calendar month of date $t + u$.

The observable entering the likelihood is the log strip ratio:

$$\widetilde{p}^{\,obs}_{k,t} \equiv \log\left(\frac{P^{\text{strip,obs}}_{t,k}}{D_t}\right), \qquad \widetilde{p}^{\,model}_{k,t} \equiv \log\left(\frac{P^{(ky)}_{t,k}}{D_t}\right). \tag{16}$$

This year–month bridge is the channel through which traded strips enter the identification. Once a panel of constant-maturity strips is combined with the spot price, the first ten years of the present-value decomposition are no longer inferred by the model—they are pinned down by market prices. The algebraic aggregation from monthly claims to annual strip objects appears in Appendix B.

## 3.4 From the near end to the spot price: tail and intrinsic bubble

Let $P_t$ denote the observed spot price of the equity index. Define the observed near-end component as

$$N_t \equiv \frac{1}{D_t} \sum_{k=1}^{10} P^{\text{strip,obs}}_{t,k}, \tag{17}$$

the dividend-scaled present value of the first ten years of cash flows, directly observed from traded strips.

The model-implied fundamental tail beyond ten years is

$$Tail(X_t) \equiv \frac{1}{D_t} \sum_{u=121}^{\infty} P^{(um)}_t = \sum_{u=121}^{\infty} \omega_{m(t+u)} \exp(a_u + b_u' X_t). \tag{18}$$

Under the no-bubble restriction, the fundamental price satisfies

$$\frac{P_t^f}{D_t} = N_t + Tail(X_t). \tag{19}$$

The unrestricted specification augments this fundamental decomposition with an intrinsic bubble component in the spirit of Froot & Obstfeld (1991):

$$B_t = cD_t \exp(\kappa' X_t), \qquad c \geq 0. \tag{20}$$

When $c = 0$, the model collapses to the no-bubble benchmark. When $c > 0$, the bubble is not an arbitrary residual: it remains tied to the same state vector and pricing kernel that governs the strips and the tail.

The full price decomposition is therefore

$$\frac{P_t}{D_t} = N_t + Tail(X_t) + Bub(X_t), \qquad Bub(X_t) \equiv c \exp(\kappa' X_t). \tag{21}$$

No-arbitrage requires the bubble to satisfy its own Euler equation, $B_t = E_t[M_{t+1} B_{t+1}]$, which pins down the bubble loading:

$$\kappa' = -b_1'(\Phi_x^Q - I)^{-1}, \tag{22}$$

together with a scalar bubble-closure condition:

$$a_1 + \frac{1}{2} b_1'(\Phi_x^Q - I)^{-1} \Sigma (\Phi_x^Q - I)^{-1\prime} b_1 - b_1'(\Phi_x^Q - I)^{-1}(\gamma_d - \Sigma \lambda_0) = 0, \tag{23}$$

further detail is reported in Appendix B. The bubble must therefore be consistent with the same pricing kernel that prices every strip maturity and the entire fundamental tail.

To ensure that the fundamental tail converges, the model imposes a sufficient transversality restriction. The corresponding condition and the linearisation of the adjusted dividend–price ratio are reported in Appendix C.

Equation (21) is the core identification device. Because $N_t$ is observed from traded strips, the residual wedge between the spot price and the near end must be absorbed entirely by $Tail(X_t)$ and, in the unrestricted specification, by $Bub(X_t)$. Strips thus convert the bubble question from an unconstrained decomposition problem into a constrained long-horizon discounting problem.

## 3.5 Measurement equations, information designs, and the null hypothesis

The model is linked to the data through four measurement blocks. The same structural model underlies all specifications; what varies across information designs is which subset of these blocks enters the likelihood.

The first measurement block is the one-month market excess return $r_{t+1}^x$. The affine SDF implies

$$r^x_{t+1} = -\frac{1}{2}\sigma_r^2 + \gamma_r'(\lambda_0 + \Lambda_1 X_t) + \varepsilon_{r,t+1}, \qquad \varepsilon_{r,t+1} \sim N(0, \sigma_r^2), \tag{24}$$

where

$$\gamma_r \equiv \mathrm{Cov}(\varepsilon_{x,t+1}, \varepsilon_{r,t+1}). \tag{25}$$

The second measurement block is annual log dividend growth,

$$\Delta d_{t+1} = \mu_d + \varepsilon_{d,t+1}, \qquad \varepsilon_{d,t+1} \sim N(0, \sigma_d^2). \tag{26}$$

The third block is an adjusted dividend–price ratio. Because the near-end $N_t$ is observed, its direct contribution can be netted out before confronting the model with the valuation ratio. Define

$$\widetilde{dp}_t \equiv dp_t + \frac{N_t - \bar{N}}{\overline{PD}(1 + \overline{PD})}, \tag{27}$$

where $dp_t \equiv \log(1 + D_t/P_t)$, $\bar{N} \equiv E[N_t]$, and $\overline{PD}$ is the linearisation anchor. A first-order approximation yields

$$\widetilde{dp}_t = h_0^{dp} + (h_1^{dp})' X_t + \varepsilon_{dp,t}, \qquad \varepsilon_{dp,t} \sim N(0, \sigma_{dp}^2). \tag{28}$$

The closed-form loadings $h_0^{dp}$ and $(h_1^{dp})'$ are reported in Appendix C. The adjusted ratio loads on $X_t$ only through the tail and bubble gradients, with the near-end contribution removed by construction.

The fourth measurement block is the strip panel. For maturities $k = 1, \ldots, 10$,

$$\tilde{p}_{k,t}^{obs} = \tilde{p}_{k,t}^{model}(X_t; \theta) + \varepsilon_{k,t}. \tag{29}$$

In the preferred implementation, the strip-residual covariance matrix is profiled out of the likelihood rather than parameterised directly.

These four blocks are combined differently across the three information designs. In the spot-only design, the strip block is excluded from the likelihood, though observed strips remain available for external validation. In the joint-strips design, all four blocks enter the likelihood, so that traded strips discipline the near end directly inside estimation. In the maturity hold-out design, only the one-to-five-year strips enter the likelihood, while the six-to-ten-year strips are reserved for out-of-sample validation in maturity space.

The tested structural restriction is

$$H_0: \; c = 0. \tag{30}$$

The unrestricted model allows $c \geq 0$; the restricted model imposes $c = 0$. The central empirical question is whether this restriction survives once the model is confronted with strip-based discipline on the near end and, by implication, on the configuration of long-horizon discounting.

# 4 Data, information design, and empirical objects

This section describes the data, the state-vector construction, the three information designs, and the key empirical objects used throughout the paper. The organising principle is that traded strips are not an auxiliary robustness check but rather market prices that alter the model's degrees of freedom for long-horizon discounting by anchoring the near end of valuation.

## 4.1 Data

The baseline estimation sample is monthly and spans January 2017 to December 2024, the common window over which the strip panel, the spot-market valuation block, and the predictor panel can be aligned at month-end. The strip panel itself begins in December 2016, yielding 95 monthly transitions. A longer spot-only sample from January 2000 to December 2024 is reserved for the strip-free external-validation exercise reported in Appendix I and is not used to redefine the baseline.

The data consist of four blocks, summarised in Table 1. The first is the observed strip panel. The second is the spot-market block used to construct prices, returns, and dividends for the S&P 500. The third is the risk-free rate, which enters both as the first element of the state vector and as the short rate in the stochastic discount factor. The fourth is the predictor panel from which the state vector is extracted. The dividend futures, zero-coupon swap yields, equity index, and predictor panel data are sourced from Refinitiv Datastream.

S&P 500 annual dividend futures are listed for up to eleven future dividend years. Each contract settles on the third Friday of December in the relevant year, so its time to maturity decreases day by day. To align these contracts with the model's horizon-indexed pricing structure, we construct constant-maturity annual strip present values following Kragt et al. (2020). The construction involves a front-contract look-back adjustment for the first maturity, a seasonally weighted interpolation across adjacent contracts for maturities two through ten, and discounting at market-specific zero-coupon swap yields. The front-contract adjustment consumes one listed maturity, so eleven annual contracts yield ten constant-maturity strip present values spanning the one-to-ten-year segment. Full details appear in Appendix D.

The spot-market block is built from the daily S&P 500 price index $P_t^{idx}$ and the daily total-return index $TR_t^{idx}$. Daily cash dividends are inferred from the identity

$$Div_t^{(1d)} = P_{t-1}^{idx} \left( \frac{TR_t^{idx}}{TR_{t-1}^{idx}} - \frac{P_t^{idx}}{P_{t-1}^{idx}} \right), \tag{31}$$

with negative values truncated at zero and implausibly large observations (exceeding 5% of the previous day's index level) filtered before aggregation. Summing daily dividends within a calendar month yields monthly cash dividends $Div_t^{(1m)}$. The annual dividend level entering the model is the trailing twelve-month sum,

$$D_t = \sum_{j=0}^{11} Div_{t-j}^{(1m)}. \tag{32}$$

The calendar-month seasonality weights $\omega_m$ that allocate $D_t$ across months are estimated separately from within-year cash-dividend data over January 2000 to December 2023. The estimation procedure is reported in Appendix A.4. Monthly market excess returns are computed as the difference between the log total return on the S&P 500 index and the log risk-free return within the same month. The risk-free rate $y_t$ used as the first state variable is the continuously compounded transformation of the monthly risk-free return, computed as $y_t = \log(1 + R_t^f)$, where $R_t^f$ is the monthly realisation of the three-month Treasury bill rate.

The predictor panel is referred from the Goyal et al. (2024), they provide a comprehensive reassessment of equity premium predictors. We select twelve predictors that jointly span the main economic channels through which expected returns may vary, including interest rates, credit conditions, risk compensation, equity supply, and real activity. The monthly predictors are the long-term government bond yield (lty), the term spread (tms), the default yield spread (dfy), realised stock-market variance (svar), the equity variance-risk premium (vrp), net equity issuance (ntis), the Treasury bill surplus ratio (tchi), and the short-interest ratio (shtint). The quarterly predictors are real personal-consumption growth (pce), the change in credit standards (crdstd), and the investment-to-capital ratio (i/k). The sole annual predictor is growth of the PCE price index (gpce). All series are aligned to month-end. Whenever a predictor is observed at quarterly or annual frequency, the last available observation is carried forward so that the predictor panel is balanced at the monthly frequency before state-vector extraction.

## 4.2 State-vector construction

The state vector $X_t$ is observed and constructed directly from the data in three steps. The first step standardises the risk-free rate to form the first element of the state vector, $X_{1,t} = (y_t - \bar{y})/\sigma_y$, where $\bar{y}$ and $\sigma_y$ are the in-sample mean and standard deviation of $y_t$. This element enters the short-rate equation of the model through the selector $e_1$ described in Section 3.1.

The second step orthogonalises the remaining predictors against $X_{1,t}$. Each standardised predictor $z_{j,t}$ is regressed on a constant and $X_{1,t}$, and the residual is retained and re-standardised to unit variance. This orthogonalisation ensures that the principal components extracted in the next step capture variation in the predictor panel that is not already spanned by the short rate.

The third step extracts principal components from the orthogonalised predictor panel. The baseline specification retains the first $r = 3$ principal components, so that the state vector is $X_t = (X_{1,t}, \mathrm{PC}_{1,t}, \mathrm{PC}_{2,t}, \mathrm{PC}_{3,t})'$ with $K = 4$. All principal component scores are z-scored before entering the state vector, so that every element of $X_t$ has zero mean and unit variance in the

Table 1: Data sources and variable definitions

| Block | Name | Description | Freq. | Reference |
|---|---|---|---|---|
| Strip panel | Dividend futures | S&P 500 annual dividend index futures | D | |
| | Discount curve | Zero-coupon interest rate swap yields | D | |
| Spot-market block | $P_t^{idx}$ | S&P 500 price index | D | |
| | $TR_t^{idx}$ | S&P 500 total-return index | D | |
| | $D_t$ | Trailing 12-month cash dividend level | M | |
| Risk-free rate | $y_t$ | $\log(1 + R_t^f)$; risk-free return from 3-month T-bill | M | |
| Predictor panel | lty | Long-term (10-year) government bond yield | M | Fama & French (1989) |
| | tms | Term spread | M | Campbell & Thompson (2008) |
| | dfy | Default yield spread (BAA − AAA) | M | Keim & Stambaugh (1986) |
| | svar | Realised stock-market variance (rolling 1-month) | M | Bekaert & Hoerova (2014) |
| | vrp | Equity variance-risk premium ($\text{VIX}^2$ − svar) | M | Bollerslev, Tauchen & Zhou (2009) |
| | ntis | Net equity issuance | M | Baker & Wurgler (2000) |
| | tchi | Treasury bill surplus (T-bill supply / GDP) | M | Greenwood, Hanson & Stein (2015) |
| | shtint | Short-interest ratio (shares short / shares outstanding) | M | Rapach et al. (2016) |
| | pce | Real personal-consumption growth (annualised) | Q | Cochrane (1991) |
| | crdstd | Change in credit-standards index | Q | Lopez-Salido, Stein & Zakrajšek (2017) |
| | i/k | Investment-to-capital ratio | Q | Chen, Novy-Marx & Zhang (2011) |
| | gpce | Growth of PCE price index (inflation) | A | Goyal & Welch (2008) |

estimation sample.

The choice of $r = 3$ is not governed by a variance-explained threshold. A 95% explained-variance criterion applied to the orthogonalised panel would suggest a larger state dimension, but the design-diagnostic evidence in Section 6 shows that the relevant question is not how much predictor variance one captures but rather which state dimension is sufficient to make the three information designs deliver a coherent story. The four-state specification ($K = 4$) is the lowest-dimensional design in which spot-only survival of the no-bubble restriction coexists with strip-based rejection in an economically coherent valuation region. This claim is substantiated by the state-vector dimension sweep reported in Appendix J, which re-estimates the full model for $r \in \{0, 1, 3, 4, 5, 6\}$ under both the rank-one and constant price-of-risk specifications.

## 4.3 Information designs

Section 3.5 introduced the four measurement blocks and the three information designs at the level of the model. This subsection restates the designs in terms of the economic questions they are intended to answer.

The spot-only design excludes the strip block from the likelihood entirely. It asks the counterfactual question of whether the no-bubble restriction would survive if traded strips were unavailable and the model already permits time-varying expected returns. The strip panel remains available for external validation but plays no role in estimation.

The joint-strips design includes all four measurement blocks, with the strip-residual covariance matrix profiled out of the likelihood rather than parameterised directly. This is the preferred baseline because it allows traded strips to discipline the near end inside estimation without introducing additional measurement parameters. By anchoring $N_t$, the joint design forces the model to confront the residual long-end wedge and thereby sharpens inference on $c$.

The maturity hold-out design estimates the model on the one-to-five-year strip block and evaluates it on unseen six-to-ten-year strips. It asks whether the strip-based conclusion extends to longer maturities that the likelihood has not seen, thereby providing a maturity-space out-of-sample test.

Taken together, the three designs separate survival from identification. Spot-only survival is informative only if it persists once the model is confronted with traded strips. The joint and hold-out designs address exactly that question.

## 4.4 Key empirical objects

The empirical results are reported in terms of five objects that recur throughout the paper. This subsection defines each object and explains its role in the identification argument.

The price decomposition in Equation (21) partitions the observed valuation ratio $P_t/D_t$ into three components. The first is the traded near end $N_t$, which is the dividend-scaled present value of the first ten years of cash flows, constructed directly from observed strip prices. The second is the model-implied tail $Tail(X_t)$, which collects the discounted cash flows beyond the ten-year traded horizon and is the sole model-dependent component of the fundamental valuation. The third is the bubble-like component $Bub(X_t)$, which appears only in the unrestricted specification and absorbs any residual valuation that neither the near end nor the tail can explain. Once $N_t$ is anchored by market prices, the difference $P_t/D_t - N_t$ defines what we refer to as the residual long-end wedge. The central empirical question is whether the no-bubble model can explain this wedge through the tail alone, or whether a non-zero $Bub(X_t)$ is required.

Beyond these three valuation components, two diagnostic objects play a central role in interpreting the results. The first is the transversality margin, defined as the gap between the unconditional mean of the short rate and the risk-adjusted unconditional growth rate of discounted dividends implied by the estimated parameters (Appendix C). When this margin is close to zero, the restricted model sustains a finite fundamental tail only by driving long-run discounting to the boundary of the admissible region, which signals that the no-bubble restriction is binding in an economically meaningful sense. The second diagnostic object is the term structure of expected strip excess returns, which plots the model-implied risk premium as a function of dividend-claim maturity. Comparing this term structure across information designs and across restricted and unrestricted specifications reveals how strips reshape the configuration of long-horizon discount rates.

## 4.5 Validation hierarchy

The empirical evidence is organised in three layers, each addressing a progressively broader concern about the robustness of the main finding.

The first layer is the information-set comparison across the spot-only, joint-strips, and maturity hold-out designs. This comparison addresses the paper's central claim directly. If the no-bubble restriction survives in the spot-only design but fails once traded strips enter the likelihood, the conclusion that strips alter the viability of the restriction follows from a within-sample change in the information set rather than from a change in the model itself.

The second layer is a long-sample external-validation exercise in the spirit of Giglio et al. (2024). A strip-free model is estimated on the January 2000 to December 2024 spot-only sample and then confronted with observed strips that played no role in estimation (Appendix I). This exercise asks whether a model trained on a longer history without strip discipline can nonetheless reproduce the traded strip panel ex post. It is informative about whether the short-sample conclusion reflects genuine information in strips or merely a sample-specific artefact, but it remains secondary because the paper's main contribution is identification rather than strip-free recovery.

The third layer is a set of design diagnostics reported in Section 6. These diagnostics ask whether the main result is fragile with respect to specific modelling choices. They include an alternative implementation of the joint-strips likelihood in which strip measurement variances are estimated directly rather than profiled out, a specification in which time variation in the price of risk is shut down entirely, a state-vector dimension sweep across $r \in \{0, 1, 3, 4, 5, 6\}$ principal components, and a tail-horizon sensitivity analysis that varies the truncation point of the model-implied fundamental tail. The relevant standard across all of these exercises is not numerical invariance but rather whether any economically plausible simplification restores a clean no-bubble conclusion once traded strips constrain long-horizon discounting.

# 5 Core identification results

This section reports the central empirical finding of the paper. The three information designs introduced in Section 4.3 are estimated on the baseline sample and compared along the dimensions defined in Section 4.4. The comparison is organised around a central question: does the no-bubble restriction $H_0 : c = 0$ survive once traded strips discipline the near end of the price decomposition in Equation (21)?

Table 2 reports the results. In the spot-only design, the unrestricted and restricted models are nearly observationally equivalent, and a likelihood-ratio test cannot reject $c = 0$. Once the strip panel enters the likelihood in the joint-strips design, the same restriction is rejected ($p < 0.001$). The rejection persists in the maturity hold-out design, in which the model is estimated on the one-to-five-year strip block alone and evaluated on unseen six-to-ten-year strips. The subsections below develop each of these findings in turn, beginning with the data fact that motivates the information-design comparison.

Table 2: Likelihood-ratio tests and fit diagnostics across information designs

| Design | Model | $\log L$ | LR$(c=0)$ | $p$-value | Corr. $P_t/D_t$ | Corr. long end | Strip RMSE | Mean $B_t/P_t$ |
|---|---|---|---|---|---|---|---|---|
| Spot-only | Unrestricted | 1010.16 | 0.003 | 0.956 | 0.929 | 0.926 | 0.416 | 0.150 |
| | Restricted | 1010.16 | | | 0.929 | 0.926 | 0.446 | – |
| Joint strips | Unrestricted | 3838.22 | 113.55 | 0.000 | 0.323 | 0.420 | 0.120 | 0.412 |
| | Restricted | 3781.45 | | | 0.320 | 0.424 | 0.214 | – |
| Maturity hold-out | Unrestricted | 2306.95 | 153.68 | 0.000 | 0.231 | 0.266 | 0.139 | 0.468 |
| | Restricted | 2230.11 | | | 0.323 | 0.404 | 0.224 | – |

The table reports maximum-likelihood estimates, likelihood-ratio test statistics, and fit diagnostics for the three information designs described in Section 4.3. The estimation sample is monthly from January 2017 to December 2024. Each design is estimated under both the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications of the model in Section 3. $\log L$ is the maximised log-likelihood. LR$(c=0)$ is twice the difference in log-likelihood between the unrestricted and restricted models, distributed as $\chi^2(1)$ under the null hypothesis that the bubble scaling coefficient is zero. Corr. $P_t/D_t$ is the time-series correlation between the observed spot valuation ratio and its model-implied counterpart. Corr. long end is the corresponding correlation for the residual long-end component $P_t/D_t - N_t$, where $N_t$ is the traded near end defined in Equation (21). Strip RMSE is the root mean squared error of the log strip ratios $\widetilde{p}_{k,t}$; for the spot-only design this is an external-validation measure because strips do not enter the likelihood, for the joint-strips design it is the in-sample strip RMSE across all ten maturities, and for the maturity hold-out design it is the out-of-sample RMSE on the unseen six-to-ten-year strip block. Mean $B_t/P_t$ is the sample average of the estimated bubble component as a fraction of the spot price, reported only for the unrestricted specification. An alternative implementation of the joint-strips design in which strip measurement variances are estimated directly rather than profiled out is reported in Appendix E.

## 5.1 The empirical wedge between spot and strip valuations

Before turning to the structural estimates, it is useful to isolate the raw data fact that motivates the entire information-design comparison. Figure 1 plots two valuation signals side by side over the baseline sample: the spot-market dividend-price ratio and a strip-implied valuation signal constructed from the traded one-to-ten-year strip panel, and these two series comove negatively.

This divergence is not generated by a model. It is a property of the data that precedes any structural estimation. If the spot market and the strip market conveyed the same information about long-horizon discounting, a flexible spot-only present-value model would already be sufficient, because strips would add nothing that the spot price had not already revealed. The negative comovement suggests that traded strips may carry incremental information about the term structure of discount rates that the spot price alone does not capture. The remainder of this section asks whether this incremental information is sufficient to overturn the no-bubble conclusion that emerges from the spot-only design.

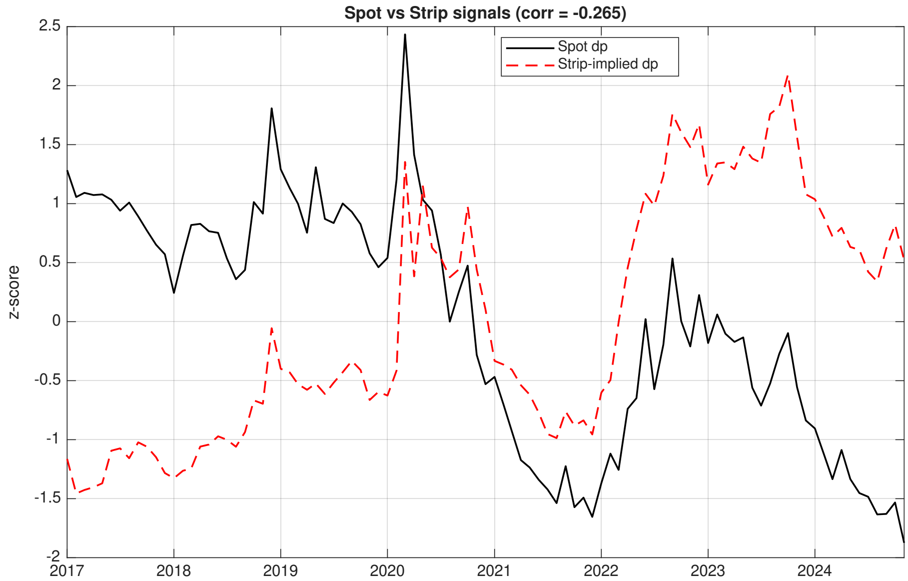


Figure 1: Spot-market and strip-implied valuation signals

The figure plots two $z$-scored valuation signals over the baseline sample (January 2017 to December 2024). The first signal is the spot-market log dividend-price ratio $\log(1 + D_t/P_t)$. The second is a strip-implied valuation signal constructed from the observed constant-maturity one-to-ten-year dividend strip panel described in Section 4.1. Both series are standardised to zero mean and unit variance for visual comparability. The negative comovement between the two signals is a feature of the data rather than a model implication and provides the empirical wedge that motivates the information-design comparison reported in the subsequent subsections.

## 5.2 Spot-only survival of the no-bubble restriction

The analysis begins with the environment that is least likely to reject the no-bubble model. In the spot-only design, the likelihood uses market excess returns, dividend growth, and the dividend-price ratio, which excludes the strip panel entirely. The exercise asks a sharp counterfactual question: does the no-bubble restriction survive even when expected returns are already allowed to vary over time through the rank-one price-of-risk specification described in Section 3.1?

In Table 2, the unrestricted and restricted log-likelihoods are virtually identical (LR = 0.003, $p$ = 0.956), and the no-bubble restriction cannot be rejected. The two specifications are also nearly indistinguishable in the spot valuation block. The correlation between observed and model-implied $P_t/D_t$ exceeds 0.92 for both, and the long-end valuation correlations are similarly close.

This survival result should nevertheless be interpreted as strip-free survival only. The spot-only model is not exposed to the empirical wedge documented in Figure 1, and Appendix F shows that the same spot-only estimates fail to reproduce the observed strip term structure when confronted with strips ex post. In particular, the model-implied average strip equity yield is negative at every maturity, whereas the observed average yield is uniformly positive. Spot-only survival is therefore a necessary benchmark rather than a sufficient conclusion; it establishes that the restriction can survive in the absence of strip discipline, but it says nothing about whether the restriction can withstand the incremental information that traded strips provide.

## 5.3 Strip-based rejection of the no-bubble restriction

The preferred joint-strips design adds the observed strip panel to the likelihood and profiles out the strip-residual covariance matrix rather than parameterising it directly. An alternative implementation in which the strip measurement variances are estimated as free parameters is reported in Appendix E and delivers a qualitatively identical rejection, though with less interpretable valuation diagnostics.

Once traded strips discipline the near end of the price decomposition, the no-bubble restriction is rejected. As Table 2 shows, the unrestricted model improves the log-likelihood by 56.8 points (LR = 113.55, $p < 0.001$) and reduces the in-sample strip RMSE from 0.214 to 0.120.

The economic content of this rejection is visible in Figure 2. Panel (a) compares the observed spot valuation ratio to the model-implied total valuation under both specifications. Panel (b) isolates the residual long-end component, defined as $P_t/D_t - N_t$, and compares it to the restricted model's fundamental tail and to the unrestricted model's tail plus bubble. The unrestricted specification shifts the model-implied long-end component toward the data-implied residual wedge that remains after the near end is anchored by market prices. The unrestricted model's gain is therefore concentrated in the residual value that lies beyond the directly observed near end, which is the dimension along which the information-design comparison is intended to discriminate.

At the same time, the unrestricted model does not dominate every measurement block uniformly. In Table 2, the total-valuation and long-end correlations remain positive under both specifications, but the unrestricted model's advantage is concentrated in likelihood and strip pricing rather than in a broad improvement across all fit dimensions. This pattern is consistent with the intended identification margin. The joint-strips design is constructed so that the strip block disciplines the near end, and the residual long end is the margin along which the bubble component either contributes or does not. The fact that the unrestricted model's improvement is concentrated along this margin is consistent with the interpretation that strips alter the viability of the no-bubble restriction through their effect on long-horizon discounting, rather than through an incidental improvement in the fit of the spot valuation block. Appendix G reports supplementary measurement-block figures for all three designs, showing that the unrestricted and restricted specifications are not visibly separated in the short-run equity-risk-premium or adjusted dividend-price blocks.

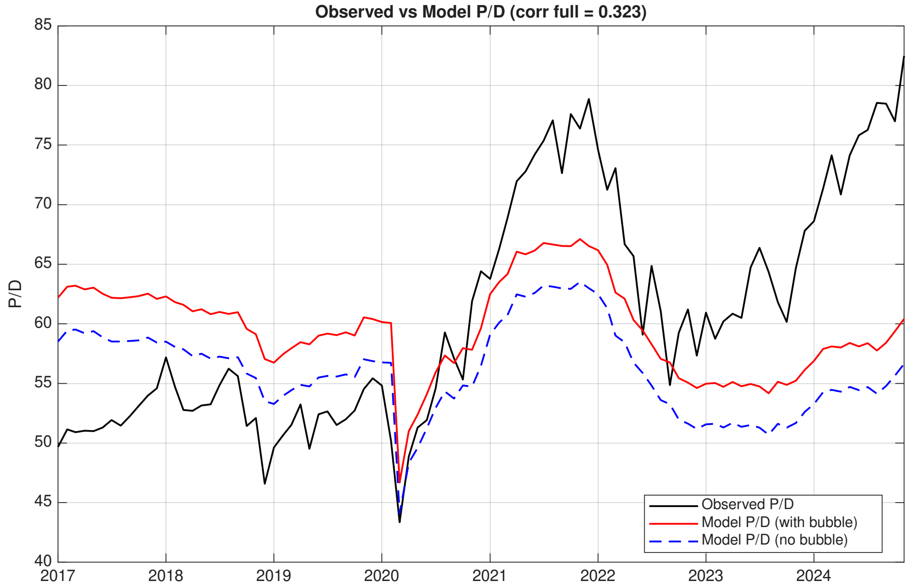


(a) Observed and model-implied spot valuation ratio $P_t/D_t$.

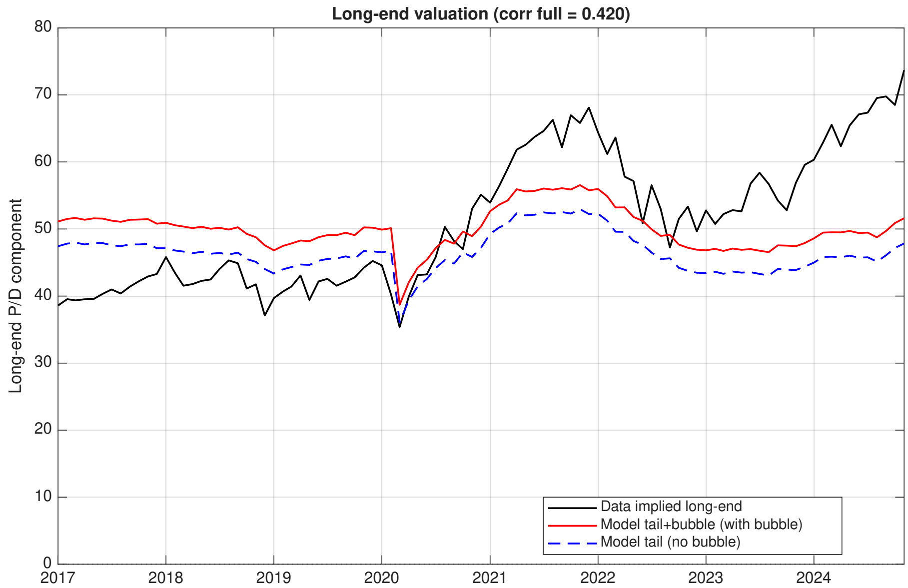


(b) Data-implied residual long end and model-implied counterparts.

Figure 2: Joint-strips design diagnostics

The figure reports diagnostics for the preferred joint-strips design over the baseline sample (January 2017 to December 2024). Panel (a) compares the observed spot valuation ratio $P_t/D_t$ to the model-implied total valuation under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. Panel (b) compares the data-implied residual long-end component, defined as $P_t/D_t - N_t$ where $N_t$ is the traded near end from Equation (21), to the restricted model's fundamental tail $Tail(X_t)$ and to the unrestricted model's tail plus bubble $Tail(X_t) + Bub(X_t)$. The unrestricted model's gain relative to the restricted specification is concentrated in the residual long-end wedge that remains after the near end has been anchored by observed strips, consistent with the interpretation that the bubble component absorbs long-horizon valuation that the fundamental tail alone does not account for.

## 5.4 Maturity hold-out confirmation

The joint-strips result could in principle reflect overfitting to the specific maturities included in the strip block. The maturity hold-out design addresses this concern. The model is estimated using only the one-to-five-year strips, and the six-to-ten-year strips are reserved for out-of-sample evaluation. The exercise is therefore a maturity-space validation rather than a purely time-series one.

The result is qualitatively consistent with the joint-strips finding. The unrestricted model retains a likelihood advantage of comparable magnitude (LR = 153.68, $p < 0.001$) and reduces the validation-set strip RMSE from 0.224 to 0.139. The unseen longer-maturity strips continue to favour the unrestricted specification even though they played no role in estimation.

This evidence narrows the set of admissible interpretations. The unrestricted model is not merely fitting the strip maturities it was shown during estimation; it is also pricing unseen longer maturities more accurately, which suggests that the bubble component captures a feature of the long-end discount-rate configuration rather than a maturity-specific fitting artefact. At the same time, the hold-out design does not dominate every block. The restricted model retains modestly tighter spot-valuation correlations, so the hold-out evidence suggests a confirmation of the joint-strips rejection rather than as an independently dominant result. Its contribution is to show that longer maturities carry incremental information about the viability of the no-bubble restriction, and that this information points in the same direction as the joint-strips result.

# 6 Validation hierarchy and design diagnostics

The finding in Section 5 is deliberately narrow: the no-bubble restriction survives in the strip-free spot-only design but is rejected once traded strips discipline the near end of valuation. This section asks whether that contrast is robust to alternative modelling choices. The evidence is organised as a validation hierarchy. The first layer examines whether the rejection depends on the specific treatment of risk-price dynamics or the strip measurement block. The second layer reports a long-sample strip-free external validation exercise in the spirit of Giglio et al. (2024). The third layer varies the state-vector dimension and the tail-horizon truncation. Throughout, the relevant criterion is not numerical invariance across specifications. It is whether any economically plausible simplification restores a no-bubble conclusion once traded strips constrain long-horizon discounting.

## 6.1 Constant price-of-risk diagnostic

All alternative specifications in this subsection are estimated on the same baseline sample with the same $r = 3$ state vector. The data construction, sample window, and state-vector extraction are held fixed; only the specific feature under examination changes across specifications.

The first diagnostic shuts down time variation in prices of risk by setting $\Lambda_1 = 0$. This specification removes the rank-one dynamic channel described in Section 3.1 and retains only the constant component $\lambda_0$ of the market price of risk. The question is not whether such a model can reject the no-bubble restriction in a likelihood comparison. It is whether it can do so while preserving a valuation decomposition in which the total-valuation and long-end correlations are positive.

Table 3, PanelA shows, in the joint-strips design, the unrestricted constant-price model still rejects $c = 0$ (LR = 102.89, $p < 0.001$), but the total-valuation and long-end correlations are both negative. The weaker SDF can therefore produce a statistical rejection, but in a parameter region where the model no longer generates a positive relationship between the observed and model-implied valuation ratio.

The spot-only constant-price specification is informative from a different angle. The unrestricted and restricted log-likelihoods are virtually identical (LR = −0.65), the implied bubble component is approximately zero, and the external strip validation RMSE exceeds 17 for the unrestricted model and 37 for the restricted model. Once time-varying prices of risk are removed, the spot-only model neither rejects nor recovers the strip curve.

Table 3, Panel B reports the mechanism-level diagnostics. Because $\Lambda_1 = 0$ by construction, the latent risk-price factor is absent and the model loses the dynamic discount-rate channel that drives the baseline design. The expected-return term structure becomes the residual object through which the model attempts to absorb the strip evidence. In the spot-only constant-price specification, this absorption generates implausibly large expected strip excess returns at medium and long horizons (Table 3, Panel B), whereas in the joint-strips constant-price specification the transversality margin separates unrestricted from restricted models but the valuation decomposition remains in a region with negative correlations.

## 6.2 Diagonal strip-covariance diagnostic

The second diagnostic retains the rank-one time-varying price-of-risk channel but simplifies the strip measurement block by replacing the profiled full covariance with a diagonal strip-covariance matrix. The question is whether the profiled covariance treatment is responsible for the rejection in Section 5.3 or whether the rejection persists under a simpler strip-weighting scheme.

Table 3, Panel A shows that the diagonal-covariance specification also rejects the no-bubble restriction (LR = 115.21, $p < 0.001$). The strip-based rejection is therefore not specific to the profiled covariance treatment. However, the same table shows that the diagonal specification achieves a strip RMSE of 0.081 while leaving the total-valuation and long-end correlations negative. Relative to the preferred baseline, the diagonal treatment fits the strip panel with lower error but does not deliver positive cross-block valuation correlations.

Panel B of the same table shows that, under diagonal covariance, the unrestricted and restricted expected-return term structures are closer to one another than in the preferred baseline, and the unrestricted transversality margin is smaller (0.000666 versus 0.001505 in the preferred design). The diagonal specification therefore does not overturn the strip-based rejection, but it compresses the economic separation between the unrestricted and restricted models along the mechanism dimensions.

## 6.3 Implications for the baseline design

The diagnostic battery in the preceding subsections informs the rationale for the preferred baseline specification along two dimensions.

A model that lacks dynamic discount-rate variation ($\Lambda_1 = 0$) can produce a statistical rejection in the strip-using design, but in a parameter region where the valuation decomposition exhibits negative correlations. The same model can survive in the spot-only design, but that survival is accompanied by strip validation errors that exceed those of the preferred specification by more than two orders of magnitude. The time-varying price-of-risk channel is therefore not an auxiliary feature of the design; it is the channel through which strip-based rejection occurs in a region where the valuation decomposition retains a positive relationship with the data.

A model that simplifies the strip measurement block (diagonal covariance) also produces a statistical rejection, but it compresses the economic distinction between unrestricted and restricted specifications along the mechanism dimensions. The preferred profiled covariance preserves both the rejection and the economic separation.

The preferred baseline is therefore best understood as the specification in which three conditions hold simultaneously: the strip-based rejection of $c = 0$ is present, the long-end valuation correlations are positive, and the mechanism-level diagnostics separate the unrestricted from the restricted model along economically interpretable margins. The full diagnostic-battery tables are reported in Appendix H.

## 6.4 Long-sample strip-free external validation

The diagnostics above vary the model specification while holding the sample and information set fixed. This subsection varies both. The question it addresses is whether a spot-only model estimated on a longer sample, one that extends backward to January 2000 and therefore encompasses substantially more business-cycle variation, can reproduce the traded strip term structure even though strips play no role in estimation. The exercise follows the external-validation logic of Giglio et al. (2024): the model is estimated without strip data, and the resulting parameter estimates are then applied, without re-estimation, to the 2017–2024 overlap period in which traded strips are observed. The comparison between model-implied and observed strip equity yields is therefore entirely out-of-sample from the perspective of the strip panel.

The exercise proceeds in two stages. In the first stage, the state vector is reconstructed by re-running the principal-component extraction on the predictor panel over the extended January 2000 to December 2024 window, after which the spot-only likelihood (market excess returns, dividend growth, and the raw log dividend-price ratio, with no strip anchor and no adjusted dividend-price ratio) is maximised under both the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. Table 4, Panel A reports the training-sample log-likelihoods. The restricted specification is favoured (LR $= -19.04$, $p = 1.000$), indicating that a bubble component does not emerge from the longer spot-only sample when strips are excluded.

Table 3: Design diagnostics for alternative specifications

*Panel A. Fit and valuation diagnostics*

| Specification | Model | $\log L$ | LR | Strip RMSE | Corr. total $P/D$ | Corr. long end | Mean $B_t/P_t$ |
|---|---|---|---|---|---|---|---|
| Joint strips, constant prices of risk | Unrestr. | 3749.19 | 102.89 | 0.132 | −0.135 | −0.114 | 0.410 |
| | Restr. | 3697.75 | | 0.196 | −0.141 | −0.114 | – |
| Spot-only, constant prices of risk | Unrestr. | 920.03 | −0.65 | 17.361 | −0.189 | −0.114 | 0.000 |
| | Restr. | 920.35 | | 37.040 | −0.214 | −0.116 | – |
| Joint strips, diagonal strip covariance | Unrestr. | 2098.50 | 115.21 | 0.081 | −0.112 | −0.177 | 0.350 |
| | Restr. | 2040.89 | | 0.081 | −0.126 | −0.183 | – |

*Panel B. Mechanism diagnostics*

| Specification | Model | TransMargin | AR(1) of risk-price factor | ER 12m | ER 60m | ER 120m | $c$ |
|---|---|---|---|---|---|---|---|
| Joint strips, constant prices of risk | Unrestr. | 0.001397 | 0.000 | 0.0046 | 0.0043 | 0.0047 | 24.123 |
| | Restr. | 0.000001 | 0.000 | 0.0018 | 0.0031 | 0.0034 | 0.000 |
| Spot-only, constant prices of risk | Unrestr. | 0.000002 | 0.000 | 0.0804 | 0.3608 | 0.1133 | 0.0005 |
| | Restr. | 0.000005 | 0.000 | 0.3603 | 0.6715 | 0.2098 | 0.000 |
| Joint strips, diagonal strip covariance | Unrestr. | 0.000666 | 0.477 | 0.0045 | 0.0047 | 0.0045 | 20.491 |
| | Restr. | 0.000001 | 0.475 | 0.0045 | 0.0045 | 0.0040 | 0.000 |

The table reports alternative-specification diagnostics estimated on the same baseline sample (January 2017 to December 2024) with the same $r = 3$ state vector as the preferred design in Table 2. Each specification is estimated under both the unrestricted ($c \geq 0$) and restricted ($c = 0$) models. Panel A reports fit and valuation diagnostics. $\log L$ is the maximised log-likelihood. LR is twice the difference in log-likelihood between unrestricted and restricted models. Strip RMSE is the root mean squared error of the log strip ratios; for the joint-strips specifications this is in-sample, and for the spot-only specification it is an external-validation measure because strips do not enter the likelihood. Corr. total $P/D$ is the time-series correlation between observed and model-implied $P_t/D_t$. Corr. long end is the corresponding correlation for $P_t/D_t - N_t$. Mean $B_t/P_t$ is the average estimated bubble share, reported only for the unrestricted specification. Panel B reports mechanism diagnostics. TransMargin is the gap between the unconditional short rate and the risk-adjusted growth rate that ensures tail convergence (Appendix C). AR(1) of risk-price factor is the first-order autocorrelation of the latent index $\ell' X_t$; in the constant-price specifications this is zero by construction because $\Lambda_1 = 0$. ER 12m, ER 60m, and ER 120m are model-implied expected strip excess log returns at the indicated horizons. The constant-price specifications remove the dynamic risk-price channel, while the diagonal-covariance specification retains it but simplifies the strip-weighting scheme. Full diagnostic results appear in Appendix H.

In the second stage, the estimated parameters from the first stage are applied to the 2017–2024 overlap data to compute model-implied strip equity yields, defined as $\mathrm{EY}_{k,t} = -\log(\widetilde{p}_{k,t})/k$ for each maturity $k$. These model-implied yields are compared to the observed constant-maturity strip yields along three dimensions: average level across the maturity curve, time-series root mean squared error at each maturity, and time-series correlation at each maturity.

Table 4, Panel B and Figure 3 report the external-validation results. The observed average equity-yield curve is mildly upward-sloping and close to zero at all maturities (ranging from 0.005 at the one-year horizon to 0.011 at the ten-year horizon). The unrestricted strip-free model generates a term structure that is also upward-sloping, but at a level that exceeds the observed values by more than an order of magnitude (0.177 at the one-year horizon, 0.329 at the ten-year horizon). The restricted strip-free model keeps the average level closer to the data at short maturities (0.010 at the one-year horizon), but turns negative at the long end (−0.007 at the ten-year horizon).

The time-series diagnostics reveal a trade-off. The unrestricted model achieves positive correlations between model-implied and observed equity yields, reaching 0.65 at the ten-year horizon, but with RMSE values that exceed 0.30. The restricted model keeps RMSE below 0.09 at all maturities, but its correlations are negative at most horizons (−0.31 at one year, −0.25 at ten years). Neither specification jointly recovers both the level and the dynamics of the traded strip panel. The unrestricted model captures some of the time-series comovement at the cost of level accuracy, while the restricted model achieves level proximity at the cost of time-series dynamics. The mean term-structure slope (10y minus 1y) further illustrates this trade-off: the observed slope is 0.006, the unrestricted model produces 0.152, and the restricted model produces −0.017.

This outcome suggests that extending the spot-only sample to 25 years does not discipline long-horizon discounting with sufficient precision to reproduce the strip curve along both margins simultaneously. The exercise does not overturn the strip-based rejection in Section 5.3; rather, it illustrates from a different angle why direct strip discipline may carry information about long-horizon discount rates that is not recoverable from spot data alone. The full maturity-by-maturity and slope-dynamics results are reported in Appendix I.

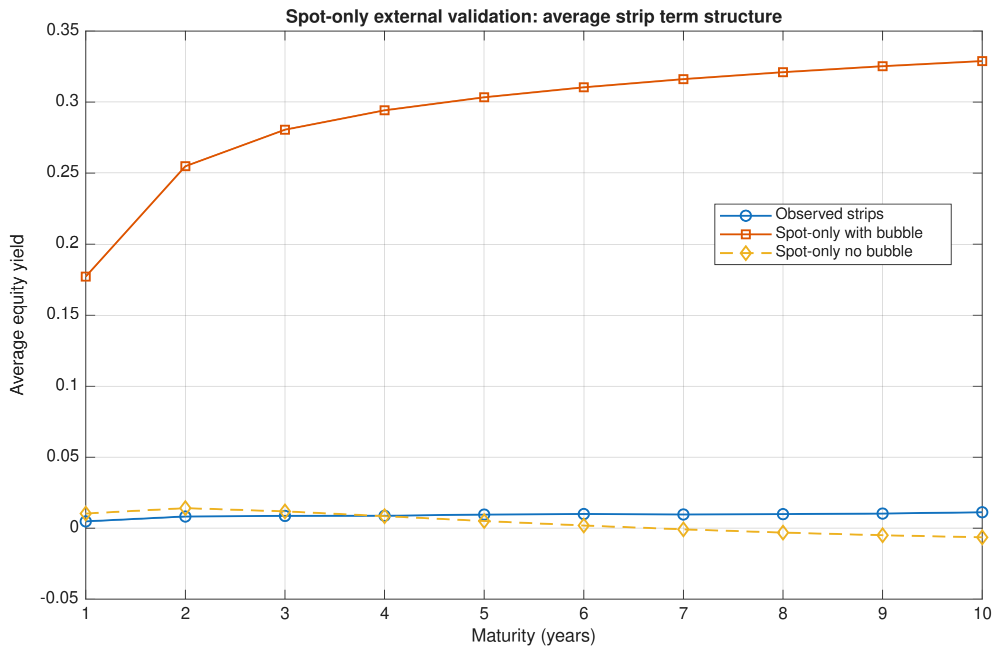


Figure 3: Long-sample strip-free external validation of the average strip term structure
The figure plots the average strip equity yield as a function of maturity (one to ten years) over the 2017–2024 overlap sample. The line with markers is the average observed strip equity yield from the constant-maturity panel described in Section 4.1. The remaining lines are the average model-implied strip equity yields under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications estimated on the long January 2000 to December 2024 spot-only sample without strips in the likelihood. The unrestricted specification generates an upward-sloping term structure at a level far above the observed values. The restricted specification keeps the average level closer to the data at short maturities but turns negative at the long end. Neither specification jointly reproduces the level and shape of the observed strip term structure.

Table 4: Long-sample strip-free external validation

| Metric | Observed strips | Unrestricted ($c \geq 0$) | Restricted ($c = 0$) |
|---|---|---|---|
| *Panel A. Training sample (2000–2024, spot-only estimation)* | | | |
| Log-likelihood | – | 2849.727 | 2859.249 |
| LR statistic / $p$-value | – | −19.044 / 1.000 | |
| *Panel B. External validation (2017–2024 overlap with traded strips)* | | | |
| Average EY, 1-year | 0.0047 | 0.1771 | 0.0102 |
| Average EY, 5-year | 0.0096 | 0.3034 | 0.0050 |
| Average EY, 10-year | 0.0112 | 0.3289 | −0.0065 |
| RMSE of EY, 1-year | – | 0.1857 | 0.0830 |
| RMSE of EY, 10-year | – | 0.3182 | 0.0292 |
| Correlation of EY, 1-year | – | 0.2544 | −0.3095 |
| Correlation of EY, 10-year | – | 0.6488 | −0.2469 |
| Mean slope (10y–1y) | 0.0064 | 0.1518 | −0.0167 |

The table reports selected statistics for the long-sample strip-free external validation exercise. Panel A shows the training-sample log-likelihood for the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications estimated on January 2000 to December 2024 spot-only data without strips in the likelihood. The negative LR statistic indicates that the restricted model is favoured in the training sample. Panel B evaluates the estimated models against the observed traded strip panel over the 2017–2024 overlap period. EY denotes the strip equity yield at the indicated maturity, defined as minus the log strip price ratio divided by maturity. RMSE and Correlation compare model-implied and observed equity yields in the time series. Mean slope (10y–1y) is the average difference between the ten-year and one-year equity yields. The unrestricted specification achieves positive time-series correlations at longer maturities but with level errors exceeding 0.30. The restricted specification keeps RMSE below 0.09 but with correlations that are negative at most horizons. Full maturity-by-maturity and slope results appear in Appendix I.

## 6.5 State-vector dimension sweep

The preceding subsections vary the risk-price dynamics and the strip-measurement treatment. This subsection asks a different question: how does the state-vector dimension affect the internal coherence of the three information designs?

For each value of $r \in \{0, 1, 3, 4, 5, 6\}$, the state vector is reconstructed by rerunning the principal-component extraction on the same baseline sample, after which the spot-only, joint-strips, and maturity hold-out designs are re-estimated. The goal is not to maximise a single fit metric but to identify the smallest state dimension at which three conditions hold simultaneously: the spot-only design does not reject $c = 0$, the strip-using designs reject $c = 0$, and the strip-using valuation correlations are positive.

Table 5 reports the sweep under the preferred rank-one price-of-risk specification. At $r = 0$, the spot-only design already rejects the no-bubble restriction ($p < 0.001$), indicating that the state vector is too weak to sustain the restriction even without strip discipline. At $r = 1$, the spot-only design no longer rejects, but the strip-using designs remain in a region with negative valuation correlations. At $r = 3$, the three designs align: the spot-only design does not reject ($p = 0.705$), the joint-strips design rejects (LR = 115.8) with positive valuation correlations (0.278 and 0.350), and the maturity hold-out design also rejects (LR = 147.2) with positive correlations (0.237 and 0.277).

Increasing the state dimension beyond $r = 3$ does not produce a monotone improvement. At $r = 4$, the joint-strips correlations are high (0.843, 0.852), but the hold-out correlations turn negative. At $r = 5$, the hold-out correlations become positive (0.913, 0.919), but the joint-strips LR statistic weakens relative to $r = 3$. At $r = 6$, the hold-out correlations are again negative. These patterns suggest that additional factors redistribute fit across blocks rather than producing a uniform incremental gain.

The baseline choice of $r = 3$ is therefore motivated by a minimum-sufficiency criterion: it is the smallest dimension at which the information-design comparison is internally coherent under the rank-one specification. Appendix J reports the full $r \times \lambda_1$-mode sweep, including the interaction with the constant-price-of-risk specification. Those results show that increasing the state dimension does not substitute for time-varying prices of risk: when $\Lambda_1 = 0$, the strip-using designs reject the restriction statistically but the valuation correlations remain negative across the entire dimension grid.

## 6.6 Tail-horizon sensitivity

A remaining concern is whether the strip-based rejection depends on a particular truncation of the model-implied fundamental tail. To address this concern, the baseline designs are re-estimated with the additional tail horizon $N_{\text{TAIL\_ADD}}$ set to 240, 480, and 960 months (20, 40, and 80 years beyond the observed strip segment). The baseline specification uses 480 months.

Table 6 reports the results. The spot-only design does not reject $c = 0$ at any of the three tail

Table 5: State-vector dimension sweep under rank-one time-varying prices of risk

| $r$ | Spot-only | | Joint strips | | | Maturity hold-out | | | |
|---|---|---|---|---|---|---|---|---|---|
| | $p(c=0)$ | Strip RMSE (restricted) | LR | Corr. total $P/D$ | Corr. long end | LR | Val. strip RMSE (unrestricted) | Corr. total $P/D$ | Corr. long end |
| 0 | < 0.001 | 0.230 | 95.1 | −0.269 | −0.296 | 119.1 | 0.161 | −0.268 | −0.296 |
| 1 | 0.272 | 0.271 | 107.0 | −0.205 | −0.259 | 134.1 | 0.120 | −0.186 | −0.248 |
| **3** | **0.705** | **0.261** | **115.8** | **0.278** | **0.350** | **147.2** | **0.141** | **0.237** | **0.277** |
| 4 | 0.937 | 0.276 | 38.2 | 0.843 | 0.852 | 99.8 | 0.099 | −0.377 | −0.490 |
| 5 | 0.930 | 0.305 | 36.0 | 0.855 | 0.854 | 136.3 | 0.166 | 0.913 | 0.919 |
| 6 | 1.000 | 0.297 | 74.4 | 0.855 | 0.846 | 74.9 | 0.085 | −0.383 | −0.494 |

The table reports the state-vector dimension sweep for the preferred rank-one time-varying price-of-risk specification over the baseline sample (January 2017 to December 2024). For each $r \in \{0, 1, 3, 4, 5, 6\}$, the state vector is reconstructed from the same predictor panel described in Section 4.2, and the three information designs are re-estimated. The spot-only columns report the likelihood-ratio $p$-value for testing $c = 0$ and the restricted model's external strip-validation RMSE. The joint-strips columns report the LR statistic for testing $c = 0$ and the unrestricted model's total-valuation and long-end correlations. The hold-out columns report the LR statistic, the unrestricted model's validation-set RMSE on the unseen six-to-ten-year strips, and the unrestricted valuation correlations. The highlighted row ($r = 3$) is the smallest dimension at which the three designs align: the spot-only design does not reject $c = 0$, both strip-using designs reject, and the strip-using valuation correlations are positive. The full $r \times \lambda_1$-mode sweep appears in Appendix J.

horizons ($p = 0.77, 0.86, 0.97$). This stability indicates that the spot-only survival result is not driven by a particular tail truncation.

In the strip-using designs, the LR statistics decline monotonically as the tail horizon lengthens: from 341 to 119 to 19 in the joint-strips design, and from 362 to 154 to 16 in the hold-out design. All remain statistically significant at the 5% level. The pattern is consistent with the interpretation that a longer model-implied tail gives the restricted model more flexibility to absorb the residual long-end wedge without a bubble component, which attenuates the test but does not reverse it.

The source of this attenuation can be isolated. In the joint-strips design, the unrestricted log-likelihood is nearly flat across the three tail horizons (moving from 3839.3 to 3841.0 to 3840.9), whereas the restricted log-likelihood improves from 3668.7 to 3781.4 to 3831.5. The same asymmetry appears in the hold-out design. Tail-horizon sensitivity operates primarily through the restricted model, whose fit improves as the tail lengthens, rather than through a redefinition of what the unrestricted model does.

The unrestricted model's valuation diagnostics are stable throughout the sweep. In the joint-strips design, total-valuation correlations remain positive (0.278, 0.313, 0.311) and the hold-out validation-set RMSE is approximately 0.14 at all three horizons. This stability suggests that the attenuation of the LR statistics at longer tail horizons does not reflect a deterioration of the strip-based design itself. Appendix K reports the full likelihood values.

Table 6: Tail-horizon sensitivity of the strip-based rejection

| Additional tail horizon (years / months) | Spot-only | Joint strips | | | Maturity hold-out | |
|---|---|---|---|---|---|---|
| | $p(c=0)$ | LR $(c=0)$ | Corr. total $P/D$ | Corr. long end | LR $(c=0)$ | Val. strip RMSE |
| 20 (240m) | 0.7653 | 341.252 | 0.278 | 0.350 | 362.018 | 0.1394 |
| **40 (480m, baseline)** | **0.8603** | **119.075** | **0.313** | **0.404** | **153.825** | **0.1403** |
| 80 (960m) | 0.9670 | 18.861 | 0.311 | 0.399 | 16.106 | 0.1394 |

The table reports the tail-horizon sensitivity of the strip-based rejection under the preferred baseline specification ($r$ = 3, rank-one time-varying prices of risk, January 2017 to December 2024). Additional tail horizon refers to $N_{\mathrm{TAIL\_ADD}}$, the number of extra monthly claims used to construct the model-implied fundamental tail beyond the ten-year observed strip segment. The 480-month row is the baseline. The spot-only column reports the likelihood-ratio $p$-value for testing $c = 0$. The joint-strips columns report the LR statistic and the unrestricted model's total-valuation and long-end correlations. The hold-out columns report the LR statistic and the unrestricted model's validation-set strip RMSE on the unseen six-to-ten-year strips. LR statistics decline monotonically with tail length as the restricted model gains flexibility to absorb the residual long-end wedge through the model-implied tail. The unrestricted model's valuation diagnostics are approximately stable across the sweep. Full likelihood decompositions appear in Appendix K.

# 7 The mechanism through which strips alter long-horizon discounting

The preceding two sections show that the information-design comparison is not sensitive to a single baseline choice. This section turns to mechanism. The question is not whether the unrestricted specification simply adds a positive bubble term. It is whether strip-informed estimation changes the configuration of long-horizon discounting in an economically interpretable way. The analysis is organised around three objects: the latent price-of-risk factor, the transversality margin, and the term structure of expected excess returns.

## 7.1 Price-of-risk reconfiguration

Table 7 reports moments of the latent price-of-risk factor across the three information designs. The strip-using designs produce larger unrestricted-versus-restricted differences than the spot-only design. In the spot-only specification, the unrestricted and restricted models have similar latent-factor moments and both operate with transversality margins close to zero. Once strips enter the likelihood, the unrestricted model moves away from the transversality boundary while the restricted model remains near it.

This pattern is visible in the preferred joint and hold-out designs. The unrestricted model yields positive transversality margins of $1.5 \times 10^{-3}$ and $2.0 \times 10^{-3}$, respectively, whereas the restricted model in both cases sits at approximately $10^{-6}$. This asymmetry suggests that the no-bubble model is pushed toward the boundary of admissible long-run discounting once strips anchor the near end and the residual wedge must be absorbed by the fundamental tail alone.

Table 7: Price-of-risk and long-horizon discounting diagnostics

| Design | Model | $E[\ell' X_t]$ | $\text{Std}(\ell' X_t)$ | AR(1) | $\rho(\Phi_x^Q)$ | TransMargin | $c$ |
|---|---|---|---|---|---|---|---|
| Spot-only | Unrestricted | −0.0139 | 1.0008 | 0.9117 | 0.9676 | 0.000008 | 8.952 |
| | Restricted | −0.0139 | 1.0009 | 0.9082 | 0.9666 | 0.000073 | 0.000 |
| Joint strips | Unrestricted | −0.0172 | 1.0031 | 0.7484 | 0.9728 | 0.001505 | 24.503 |
| | Restricted | 0.0170 | 1.0033 | 0.7386 | 0.9732 | 0.000001 | 0.000 |
| Maturity hold-out | Unrestricted | −0.0164 | 1.0040 | 0.6832 | 0.9775 | 0.002039 | 27.797 |
| | Restricted | −0.0186 | 1.0020 | 0.9454 | 0.9737 | 0.000001 | 0.000 |

The table reports moments of the latent price-of-risk factor $\ell' X_t$ and related long-horizon discounting diagnostics for the three information designs described in Section 4.3, estimated on the baseline sample (January 2017 to December 2024). Each design is estimated under both the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. $E[\ell' X_t]$ and $\text{Std}(\ell' X_t)$ are the sample mean and standard deviation of the latent risk-price index from the rank-one specification in Section 3.1. AR(1) is its first-order autocorrelation. $\rho(\Phi_x^Q)$ is the spectral radius of the risk-neutral state transition matrix. TransMargin is the gap between the unconditional short rate and the risk-adjusted growth rate that ensures tail convergence (Appendix C); a positive value indicates room for a finite model-implied tail, whereas a value near zero indicates that the model operates close to the boundary of admissible long-run discounting. $c$ is the estimated bubble scaling coefficient.

## 7.2 Transversality pressure

The transversality margin is tied directly to the finiteness of the model-implied tail. A positive margin indicates that the model can sustain a finite long-run fundamental component. A margin that collapses toward zero indicates that the model operates near the boundary of admissible long-run discounting.

The strip-using designs exhibit a separation along this margin. In the preferred joint design, the unrestricted model has a transversality margin of approximately $1.5 \times 10^{-3}$, while the restricted model sits at approximately $10^{-6}$. The hold-out design shows the same pattern with an unrestricted margin of $2.0 \times 10^{-3}$. The spot-only design does not generate a comparable separation: both specifications operate with margins close to zero. This asymmetry suggests that the strip-using restricted model is pushed toward the admissibility boundary by the need to absorb the residual long-end wedge through the fundamental tail alone, a pressure that the unrestricted model alleviates through the bubble component.

A near-zero transversality margin is not by itself evidence for a bubble. The relevant observation is that, in the strip-using designs, the restricted model is systematically closer to the boundary than the unrestricted model, whereas in the spot-only design this separation does not arise. The margin therefore serves as a diagnostic for the mechanism through which strip discipline alters the viability of the no-bubble restriction.

Table 8: Expected strip excess returns across information designs

| Design | Model | 12m | 60m | 120m |
|---|---|---|---|---|
| Spot-only | Unrestricted | −0.0178 | 0.0102 | 0.0070 |
| | Restricted | −0.0141 | 0.0067 | 0.0057 |
| Joint strips | Unrestricted | 0.0048 | 0.0045 | 0.0049 |
| | Restricted | 0.0006 | 0.0029 | 0.0035 |
| Maturity hold-out | Unrestricted | 0.0046 | 0.0048 | 0.0054 |
| | Restricted | 0.0014 | 0.0031 | 0.0035 |

The table reports model-implied expected excess log returns for dividend strips at selected horizons (12, 60, and 120 months) across the three information designs described in Section 4.3. Each design is estimated on the baseline sample (January 2017 to December 2024) under both the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. The entries are monthly expected excess log returns. In the preferred joint-strips and hold-out designs, the unrestricted model implies higher short- and medium-horizon strip premia than the restricted model. The spot-only design does not produce a comparable separation, which is consistent with the finding that spot data alone do not generate the same pressure on the no-bubble restriction.

### 7.3 Expected-return term structure

If the strip-based rejection is driven by long-horizon discounting rather than by short-run return fit, then the unrestricted and restricted models should differ in the shape of horizon-specific expected excess returns. Table 8 reports expected strip excess returns at three horizons.

In the preferred joint design, the unrestricted model implies higher short- and medium-horizon strip premia than the restricted model (0.0048 versus 0.0006 at the 12-month horizon). In the hold-out design, the unrestricted model produces a steeper long-end term structure (0.0054 versus 0.0035 at the 120-month horizon). The spot-only design does not display a comparable separation across horizons, which is consistent with the finding that spot data alone do not generate the same pressure on the no-bubble restriction.

The mechanism evidence across all three objects points in a consistent direction. The unrestricted model does not improve fit primarily through short-run realised returns. Instead, it alters the configuration of long-horizon expected returns and reduces the transversality pressure that the no-bubble model faces once traded strips discipline the near end.

## 8 Price decomposition and interpretation

This section returns to the paper's central object,

$$P_t = P_t^f + B_t.$$

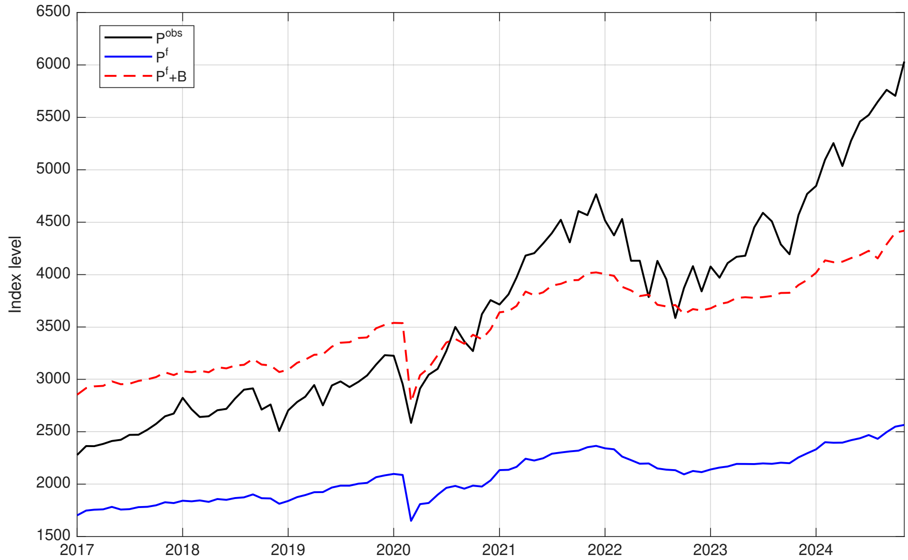


Figure 4: Price decomposition under the preferred joint-strips design

The figure plots the observed equity index level $P_t^{obs}$, the model-implied fundamental component $P_t^f$, and the unrestricted total $P_t^f + B_t$ under the preferred joint-strips design over the baseline sample (January 2017 to December 2024). The model is estimated with the near end anchored by the observed one-to-ten-year constant-maturity strip panel described in Section 4.1. The fundamental component captures the broad direction of prices, including the 2020 decline and the subsequent recovery, while systematically leaving a residual gap. Adding the estimated bubble component $B_t$ closes a substantial part of that gap, particularly around the 2021–2022 valuation peak and the 2023–2024 period.

The decomposition is most informative once the near end has been anchored by observed strips, because any residual gap between price and the near-plus-tail fundamental component is a long-end object.

Figure 4 plots the observed equity index level alongside the model-implied fundamental component and the unrestricted total. The fundamental component $P_t^f$ tracks the broad direction of prices, including the 2020 decline and the subsequent recovery. The unrestricted component $P_t^f + B_t$ closes a substantial part of the residual gap, particularly around the 2021–2022 valuation peak and the 2023–2024 recovery period. The fundamental component is not uninformative; it captures a substantial share of medium-run price movement. And it shows is that once the near end is observed from traded strips, a residual wedge remains in the long end.

Figure 5 isolates the residual long-end component. The unrestricted model's gain relative to the restricted specification is concentrated in the wedge that lies beyond the directly observed near end. This pattern suggests that the strip-based result is a long-end identification finding rather than a fit improvement across the full valuation decomposition.

## 8.1 Long-end valuation diagnostics across information designs

Table 9 reports the long-end valuation diagnostics for the unrestricted specification across the three information designs. These objects, including the traded near end, the model-implied tail, the residual long-end wedge, the transversality margin, and the term structure of expected strip


Figure 5: Residual long-end component after anchoring the near end

The figure plots the data-implied residual long-end component, computed as observed $P_t/D_t$ minus the near-end component $N_t$ implied by traded strips (Equation (21)), alongside the unrestricted model's tail-plus-bubble component $Tail(X_t) + Bub(X_t)$ under the preferred joint-strips design over the baseline sample (January 2017 to December 2024). The figure isolates the part of valuation that remains after the first decade of cash-flow claims has been anchored by market prices. The unrestricted model's gain relative to the restricted specification is concentrated in this residual long-end wedge.

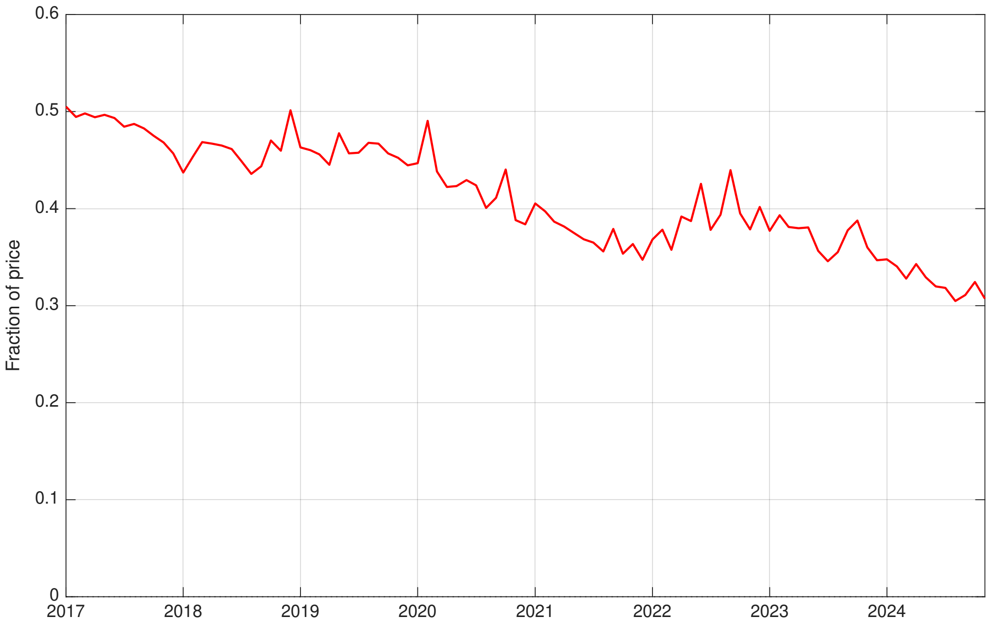


Figure 6: Estimated bubble share under the preferred joint-strips design

The figure plots the implied bubble share $B_t/P_t$ under the unrestricted preferred joint-strips specification over the baseline sample (January 2017 to December 2024). The share is economically large throughout the sample and varies gradually over time. The bubble share should be interpreted as the fraction of total price attributed to the non-fundamental long-end component under the model specification, rather than as a literal measure of a rational bubble of fixed magnitude in every month.

excess returns, provide a compact summary of the price decomposition. The mean bubble share rises from approximately 0.15 in the spot-only design to approximately 0.41 in the preferred

Table 9: Long-end valuation diagnostics across information designs

| Design | Mean valuation components | | | | Correlations | | ER slope (120m–12m) | Mean $B_t/P_t$ |
|---|---|---|---|---|---|---|---|---|
| | Near end used | Tail | Bubble | Total $P/D$ | Total $P/D$ | Long end | | |
| Spot-only | 13.48 | 37.89 | 9.04 | 60.41 | 0.929 | 0.926 | 0.02482 | 0.150 |
| Joint strips | 9.56 | 25.44 | 24.54 | 59.54 | 0.323 | 0.420 | 0.00014 | 0.412 |
| Maturity hold-out | 9.56 | 22.36 | 27.87 | 59.79 | 0.231 | 0.266 | 0.00080 | 0.468 |

The table reports long-end valuation diagnostics for the unrestricted ($c \geq 0$) specification under each of the three information designs described in Section 4.3, estimated on the baseline sample (January 2017 to December 2024). Near end used is the sample mean of the near-end component $N_t$; in the strip-using designs this is anchored by traded strips, while in the spot-only design it is inferred internally from the model. Tail is the sample mean of the model-implied fundamental tail $Tail(X_t)$. Bubble is the sample mean of $Bub(X_t)$. Total $P/D$ is the sum of all components. Corr. total $P/D$ and Corr. long end are defined as in Table 2. ER slope (120m–12m) is the difference between model-implied expected strip excess log returns at 120 months and 12 months. Mean $B_t/P_t$ is the average estimated bubble share. The mean bubble share increases monotonically from the spot-only to the hold-out design, which is consistent with the view that additional strip discipline shifts residual long-end valuation into the non-fundamental component.

joint design and to approximately 0.47 in the hold-out design. This monotone pattern across designs suggests that, as the model is exposed to more long-horizon information through traded strips, the no-bubble restriction becomes harder to sustain and the residual wedge is pushed into a larger non-fundamental long-end component.

## 8.2 Bubble share and interpretation

Figure 6 plots the estimated bubble share $B_t/P_t$ under the preferred joint-strips specification. The share is economically large throughout the sample and varies gradually over time rather than appearing only in isolated episodes.

The contribution of the paper is not to establish that a literal rational bubble of a specific magnitude exists in every month. The more defensible interpretation is that, once traded strips pin down the near end and discipline long-horizon discounting, a simple no-bubble present-value model becomes substantially harder to sustain. The residual wedge is pushed into a large and persistent non-fundamental long-end component.

# 9 Conclusion

This paper asks whether a no-bubble present-value model remains empirically viable once long-horizon discounting is disciplined by traded dividend strips. The organising comparison is across information sets rather than across technical likelihood variants. When strips are excluded from the likelihood, the no-bubble restriction survives. Once traded strips anchor the near end of valuation, the restriction becomes materially harder to sustain, both in the preferred joint-strips design and in the maturity hold-out design.

The auxiliary evidence sharpens rather than dilutes that message. Weaker discount-rate dynamics, a simplified strip-measurement block, alternative state-vector dimensions, a longer strip-free training sample, and more generous long-end tails all change the sharpness of the test.

None, however, restores a clean no-bubble conclusion once traded strips are allowed to discipline long-horizon discounting directly.

The interpretation of the paper is correspondingly restrained. The results do not establish that a literal rational bubble of fixed size is present in every period. A more cautious reading is that traded strips reveal a large and persistent residual long-end component that a flexible no-bubble model struggles to absorb without relaxing the no-bubble restriction. Read in that way, the paper contributes evidence on bubbles, but its broader contribution is methodological: traded strips are useful not only as descriptive term-structure objects, but also as a disciplined empirical device for studying long-horizon discounting and structural restrictions in equity markets.

# 10 Appendix

# A Exact pricing of monthly dividend claims

This appendix derives the exponential-affine pricing recursion for monthly dividend claims used in Section 3.2. The model is solved on a monthly grid, so the natural primitive is the present value of the annual dividend level discounted to a monthly horizon. The derivation yields exact recursions for the pricing coefficients $(a_u, b_u)$ and reports the one-month base case in closed form.

## A.1 Recursion for the present value of the annual dividend level

Let

$$\Pi_{u,t} \equiv E_t\left[\left(\prod_{s=1}^{u} M_{t+s}\right) D_{t+u}\right], \qquad u = 1, 2, \ldots,$$

and conjecture the exponential-affine form

$$\log \Pi_{u,t} = d_t + a_u + b_u' X_t.$$

The one-step Euler recursion is

$$\Pi_{u+1,t} = E_t[M_{t+1} \Pi_{u,t+1}].$$

Substituting the log-normal SDF, the Gaussian state dynamics, and the random-walk dividend process, and evaluating the resulting conditional expectation via the Gaussian moment-generating function, gives

$$a_{u+1} = a_u + \mu_d - \mu_y - \gamma_d' \lambda_0 + b_u'(\gamma_d - \Sigma\lambda_0) + \frac{1}{2}\sigma_d^2 + \frac{1}{2} b_u' \Sigma b_u, \tag{33}$$

$$b_{u+1}' = b_u' \Phi_x^Q - e_1' - \gamma_d' \Lambda_1, \tag{34}$$

where $\Phi_x^Q = \Phi_x - \Sigma\Lambda_1$ is the risk-neutral transition matrix defined in the main text.

## A.2 One-month base coefficients

Setting $u = 0$ in (33)–(34) and noting that $a_0 = 0$ and $b_0 = 0$ (since $\Pi_{0,t} = D_t$) yields the base case:

$$a_1 = \mu_d - \mu_y - \gamma_d' \lambda_0 + \frac{1}{2}\sigma_d^2, \qquad b_1' = -e_1' - \gamma_d' \Lambda_1. \tag{35}$$

The coefficient $a_1$ collects the one-month risk-adjusted expected dividend growth net of the short rate, while each element of $b_1$ captures the one-month exposure of the discounted dividend level to the corresponding state variable.

## A.3 From the annual level to monthly claim prices

Using the deterministic seasonality mapping $D_{t+u}^{(1m)} = \omega_{m(t+u)} D_{t+u}$, where $\omega_{m(t+u)}$ is the share of annual dividends paid in calendar month $m(t+u)$, the model price of the claim to the single monthly cash flow at horizon $u$ is

$$P_t^{(um)} = \omega_{m(t+u)}\,\Pi_{u,t} = D_t\,\omega_{m(t+u)}\exp(a_u + b_u' X_t).$$

These monthly claims are the building blocks aggregated into annual constant-maturity strips in Section 3.3 and summed over the long-end tail in Section 3.4.

## A.4 Estimation of calendar-month seasonality weights

The seasonality weights $\omega_m$ allocate the annual dividend level across calendar months. They are estimated from monthly cash-dividend data and should not be confused with the contract-roll weights $\varpi_{i,t}$ used in the constant-maturity interpolation of Appendix D.

The raw inputs are the daily S&P 500 price index and total-return index. Daily cash dividends are inferred from the identity

$$Div_t^{(1d)} = P_{t-1}^{idx}\left(\frac{TR_t^{idx}}{TR_{t-1}^{idx}} - \frac{P_t^{idx}}{P_{t-1}^{idx}}\right).$$

Negative values are set to zero, and implausibly large observations are filtered. The resulting daily cash dividends are summed within each calendar month to obtain monthly cash dividends $Div_t^{(1m)}$.

Let $m(t) \in \{1, \ldots, 12\}$ denote the calendar month of date $t$. Monthly cash dividends are related to the annual level by

$$D_t^{(1m)} = \omega_{m(t)} D_t, \qquad \omega_m > 0, \qquad \sum_{m=1}^{12} \omega_m = 1, \tag{36}$$

where

$$D_t = \sum_{j=0}^{11} Div_{t-j}^{(1m)}$$

is the trailing twelve-month annual dividend level.

Define the realised monthly share

$$s_t \equiv \frac{Div_t^{(1m)}}{D_t}.$$

The seasonality parameter is the conditional mean of this share:

$$\omega_m = E\left[\frac{Div_t^{(1m)}}{D_t}\middle| m(t) = m\right]. \tag{37}$$

Using monthly cash-dividend data from 2000 to 2023, we estimate

$$\widehat{\omega}_m = \frac{\frac{1}{T_m}\sum_{t:m(t)=m}(Div_t^{(1m)}/D_t)}{\sum_{\widetilde{m}=1}^{12}\frac{1}{T_{\widetilde{m}}}\sum_{t:m(t)=\widetilde{m}}(Div_t^{(1m)}/D_t)},$$

and renormalise so that $\sum_{m=1}^{12}\widehat{\omega}_m = 1$.

Conditioning on the calendar month is essential. The unconditional mean of $Div_t^{(1m)}/D_t$ would collapse the within-year profile towards $1/12$, destroying the very seasonality the mapping is designed to capture. Under (36), all economic uncertainty resides in the annual level $D_t$; the within-year allocation is deterministic conditional on the calendar month.

# B From monthly claims to annual strips, the long-end tail, and bubble closure

This appendix supports Sections 3.3 and 3.4. It aggregates monthly claims into the annual strip objects that correspond to observed constant-maturity dividend strips, derives the model-implied tail beyond the traded horizon, and reports the bubble-closure condition implied by no-arbitrage.

## B.1 Year–month bridge

The $k$-year annual strip is the sum of the twelve monthly claims spanning the $k$-th future dividend year:

$$P_{t,k}^{(ky)} = \sum_{u=12(k-1)+1}^{12k} P_t^{(um)} = D_t \sum_{u=12(k-1)+1}^{12k} \omega_{m(t+u)} \exp(a_u + b_u' X_t),$$

where $\omega_{m(t+u)}$ and $(a_u, b_u)$ are the seasonality weights and pricing coefficients from Appendix A. The model-implied log strip ratio is therefore

$$\widetilde{p}_{k,t}^{\text{model}} = \log\left(\sum_{u=12(k-1)+1}^{12k} \omega_{m(t+u)} \exp(a_u + b_u' X_t)\right).$$

Because each annual strip aggregates twelve monthly claims with horizon-specific loadings, the log strip ratio is not itself affine in $X_t$. It becomes approximately affine only when the within-year variation in $(a_u, b_u)$ is small relative to the cross-year variation.

## B.2 Tail decomposition

The model-implied fundamental tail beyond the ten-year traded horizon is

$$Tail(X_t) = \sum_{u=121}^{\infty} \omega_{m(t+u)} \exp(a_u + b_u' X_t).$$

Under the no-bubble restriction, the fundamental price satisfies

$$\frac{P_t^f}{D_t} = N_t + Tail(X_t), \qquad N_t \equiv \frac{1}{D_t} \sum_{k=1}^{10} P_{t,k}^{\text{strip,obs}}.$$

The near-end component $N_t$ is directly observed from traded strips. The tail is the sole model-dependent component of the fundamental valuation.

## B.3 Bubble Euler equation and closure

The intrinsic bubble specification is

$$B_t = c D_t e^{\kappa' X_t}, \qquad c \geq 0.$$

No-arbitrage requires the bubble to satisfy its own Euler equation:

$$B_t = E_t[M_{t+1} B_{t+1}].$$

Substituting the SDF, the state dynamics, and the dividend process, and matching terms that must hold for all realisations of $X_t$, yields two restrictions. Matching the $X_t$-dependent terms gives the bubble loading:

$$\kappa' = -b_1'(\Phi_x^Q - I)^{-1}.$$

Matching the constant terms yields the scalar bubble-closure condition:

$$a_1 + \frac{1}{2} b_1'(\Phi_x^Q - I)^{-1} \Sigma (\Phi_x^Q - I)^{-1'} b_1 - b_1'(\Phi_x^Q - I)^{-1}(\gamma_d - \Sigma\lambda_0) = 0. \tag{38}$$

# C Transversality and the adjusted dividend–price equation

This appendix supports Sections 3.4 and 3.5. It reports a sufficient transversality condition for the convergence of the model-implied tail and derives the first-order approximation for the adjusted dividend–price ratio $\widetilde{dp}_t$ used in the likelihood.

## C.1 Sufficient transversality restriction

The fundamental tail $Tail(X_t) = \sum_{u=121}^{\infty} \omega_{m(t+u)} \exp(a_u + b_u' X_t)$ converges if the risk-adjusted expected dividend growth rate is strictly below the long-run short rate. A sufficient condition is

$$\mu_y > \mu_d + \frac{1}{2}\sigma_d^2 - \gamma_d'\lambda_0 + b_1'(I - \Phi_x^Q)^{-1}\left[(\gamma_d - \Sigma\lambda_0) + \frac{1}{2}\Sigma(I - \Phi_x^{Q\prime})^{-1} b_1\right]. \tag{39}$$

The left-hand side is the unconditional mean of the short rate; the right-hand side collects the risk-adjusted unconditional growth rate of discounted dividends, inclusive of Jensen and covariance corrections that accumulate through the state dynamics. The transversality margin reported in the empirical section is the difference between the two sides of (39). A margin close to zero signals that the restricted model sustains finite valuation only by pushing long-run discounting to the boundary of the admissible region.

## C.2 Adjusted dividend–price ratio

Define the price–dividend ratio and its log transform:

$$PD_t \equiv \frac{P_t}{D_t} = N_t + Tail(X_t) + Bub(X_t), \qquad dp_t \equiv \log\left(1 + \frac{1}{PD_t}\right).$$

A first-order Taylor expansion of $dp_t$ around the unconditional anchor

$$\overline{PD} = \bar{N} + Tail(0) + c$$

gives

$$dp_t \approx \log\left(1 + \frac{1}{\overline{PD}}\right) - \frac{1}{\overline{PD}(1 + \overline{PD})}\big[(N_t - \bar{N}) + Tail(X_t) - Tail(0) + Bub(X_t) - c\big].$$

Because $N_t$ is observed from traded strips, its contribution can be moved to the left-hand side to yield the adjusted dividend–price ratio:

$$\widetilde{dp}_t \equiv dp_t + \frac{N_t - \bar{N}}{\overline{PD}(1 + \overline{PD})} \approx h_0^{dp} + (h_1^{dp})' X_t + \varepsilon_{dp,t}, \qquad \varepsilon_{dp,t} \sim N(0, \sigma_{dp}^2).$$

The intercept is

$$h_0^{dp} = \log\left(1 + \frac{1}{\overline{PD}}\right).$$

The state loading requires the derivatives of the tail and bubble components evaluated at $X_t = 0$:

$$\left.\frac{\partial Tail(X)}{\partial X'}\right|_{X=0} = \sum_{u=121}^{\infty} \bar{\omega}_u e^{a_u} b_u', \qquad \left.\frac{\partial Bub(X)}{\partial X'}\right|_{X=0} = c\kappa',$$

where $\bar{\omega}_u \equiv \omega_{m(t+u)}$ evaluated at the reference calendar position. Combining these expressions yields

$$(h_1^{dp})' = -\frac{1}{\overline{PD}(1+\overline{PD})}\left(\sum_{u\geq 121} \bar{\omega}_u e^{a_u} b_u' + c\kappa'\right).$$

The adjusted dividend–price ratio thus loads on $X_t$ only through the tail and bubble gradients, with the near-end contribution removed by construction.

# D Construction of constant-maturity dividend futures and strip present values

This appendix documents the data-construction step that maps listed annual dividend futures contracts into the constant-maturity strip present values used in the empirical analysis. The construction follows Kragt et al. (2020), because annual dividend futures expire on fixed calendar dates, a constant-maturity series requires interpolation across adjacent contracts, with the interpolation weight reflecting the uneven intra-year pattern of aggregate dividend payments. This observed-data construction is conceptually distinct from the model-implied year–month bridge in Appendix B.

## D.1 Annual contracts and the front-contract adjustment

Let $F^{ann}_{i,t,n}$ denote the time-$t$ futures price on the $n$-th future dividend year for market $i$, and let $DI_{i,t}$ denote realised dividends-to-date in the current dividend year. The front contract partly overlaps with dividends already paid, so a constant-maturity one-year-ahead claim requires a look-back adjustment:

$$F^{CM}_{i,t,1} = F^{ann}_{i,t,1} - DI_{i,t} + \frac{DI_{i,t}}{F^{ann}_{i,t,1}} F^{ann}_{i,t,2}. \tag{40}$$

The first two terms remove the already-realised portion from the front contract; the third term replaces it with the corresponding fraction of the next contract, thereby constructing a claim that always spans one full future dividend year.

## D.2 Interpolation for maturities beyond one year

For $n \geq 2$, constant-maturity futures are formed by linear interpolation between adjacent annual contracts:

$$F^{CM}_{i,t,n} = (1-\varpi_{i,t})F^{ann}_{i,t,n} + \varpi_{i,t} F^{ann}_{i,t,n+1}, \qquad n \geq 2. \tag{41}$$

The roll weight $\varpi_{i,t} \in [0,1]$ measures where date $t$ lies within the current dividend-year cycle. Following Kragt et al. (2020), $\varpi_{i,t}$ is not linear in calendar time but instead reflects the cumulative within-year pattern of dividend payments, so that the interpolation correctly accounts for the uneven intra-year payout profile.

### D.3 Construction of the roll weight $\varpi_{i,t}$

The roll weight is constructed from dividend point data rather than from monthly cash dividends. Let $d^{pt}_{i,Y,k}$ denote the dividend points realised on trading day $k$ of dividend year $Y$, and define the cumulative payout fraction:

$$c_{i,Y,k} \equiv \frac{\sum_{j=1}^{k} d^{pt}_{i,Y,j}}{\sum_{j=1}^{N_{i,Y}} d^{pt}_{i,Y,j}}.$$

Because dividend-year lengths differ across years, the cumulative pattern is first normalised to $[0, 1]$, then averaged across years for a given market, and finally evaluated at the fractional position of day $t$ within its current dividend year. By construction,

$$\varpi_{i,t} = 0 \text{ at the start of the dividend year,} \qquad \varpi_{i,t} = 1 \text{ at expiry.}$$

The roll weight $\varpi_{i,t}$ is a contract-interpolation object derived from dividend point data. It should not be confused with the calendar-month seasonality weights $\omega_m$ used inside the model to allocate the annual dividend level across months.

### D.4 From constant-maturity futures to strip present values

Let $y_{i,t,n}$ denote the market-specific zero-coupon yield for maturity $n$. The present value of the $n$-year constant-maturity strip is

$$P^{strip,obs}_{i,t,n} = F^{CM}_{i,t,n} \exp(-n y_{i,t,n}), \qquad n = 1, \dots, 10. \tag{42}$$

These present values are the observed strip objects that enter the near-end component $N_t$ and the strip-panel measurement block in the likelihood.

## E Alternative implementation of the joint-strips likelihood

The preferred joint-strips design reported in Section 5.3 profiles out the strip-residual covariance matrix rather than estimating strip measurement variances as free parameters. This appendix reports an alternative implementation of the same information set in which the strip measurement variances are estimated directly. Because the underlying structural model and the information set are identical, the exercise serves only as a technical diagnostic and is not presented as a separate empirical design.

Table 10 summarises the results. The unrestricted model rejects the no-bubble restriction (LR = 115.01, $p < 0.001$) and fits the strip panel with an overall RMSE below 0.08. However, the total-valuation and long-end correlations are negative under both specifications. This sign reversal suggests that the directly parameterised strip variances absorb information in a way that distorts the broader valuation decomposition, even though the strip block itself is fitted with low RMSE. The profiled implementation avoids this pathology by allowing the data to determine the economic weighting of the strip panel within the likelihood, which is why it serves as the preferred baseline in the main text.

Figures 7 and 8 make this contrast transparent. The price decomposition in Figure 7 shows that the unrestricted model improves likelihood without delivering a coherent spot-valuation fit. The supplementary diagnostics in Figure 8 confirm that the adjusted dividend-price ratio and the equity-premium equation remain individually sensible, but the overall valuation interpretation is compromised by the negative long-end correlation. These figures suggest that the choice between profiled and directly parameterised strip variances affects the interpretability of the valuation decomposition, which motivates the use of the profiled implementation in the main text.

Table 10: Fit summary for the alternative joint-strips implementation

| Metric | Unrestricted ($c \geq 0$) | Restricted ($c = 0$) |
|---|---|---|
| Log-likelihood | 2015.83 | 1958.32 |
| LR statistic | 115.01 | – |
| $p$-value | $< 0.001$ | – |
| Correlation of total $P_t/D_t$ fit | $-0.212$ | $-0.225$ |
| Correlation of long-end fit | $-0.264$ | $-0.272$ |
| Overall strip RMSE | 0.0783 | 0.0784 |
| Mean $B_t/P_t$ | 0.364 | – |

The table reports maximum-likelihood estimates and fit diagnostics for an alternative implementation of the joint-strips design in which strip measurement variances are estimated as free parameters rather than profiled out of the likelihood. The estimation sample, structural model, and information set are identical to the preferred joint-strips design in Table 2. Log-likelihood is the maximised value of the sample log-likelihood function. LR statistic is twice the difference in log-likelihood between the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications, distributed as $\chi^2(1)$ under the null. Correlation of total $P_t/D_t$ fit is the time-series correlation between the observed and model-implied spot valuation ratio. Correlation of long-end fit is the corresponding correlation for the residual component $P_t/D_t - N_t$, where $N_t$ is the traded near end defined in Equation (21). Overall strip RMSE is the root mean squared error of the log strip ratios $\widetilde{p}_{k,t}$ across all ten maturities. Mean $B_t/P_t$ is the sample average of the estimated bubble share, reported only for the unrestricted specification. The negative valuation correlations indicate that the directly parameterised strip variances distort the broader valuation decomposition, motivating the profiled implementation adopted in the main text.

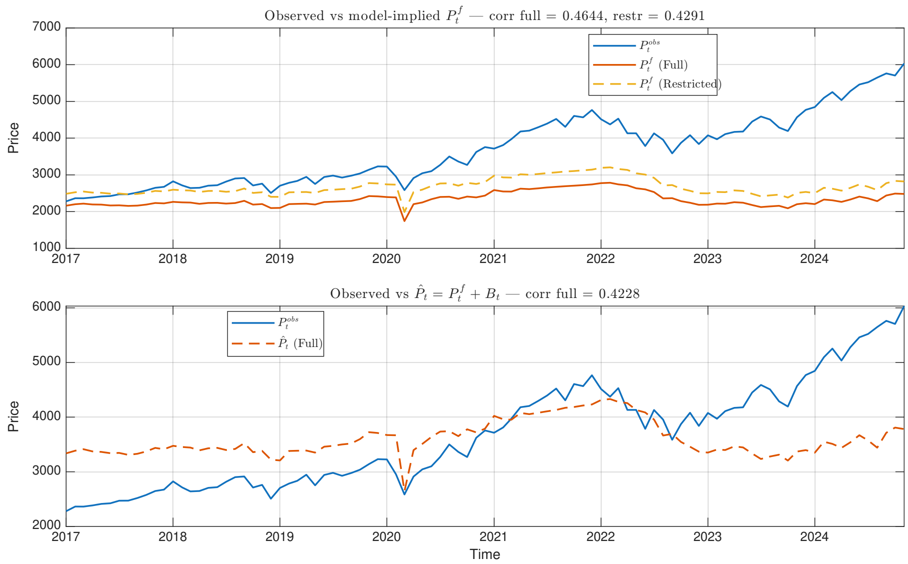


Figure 7: Price decomposition under the alternative joint-strips implementation

The figure reports the price decomposition under the alternative joint-strips implementation in which strip measurement variances are estimated directly rather than profiled out. The top panel compares the observed spot valuation ratio $P_t/D_t$ to the model-implied fundamental component $N_t + Tail(X_t)$ under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. The bottom panel adds the estimated bubble component $Bub(X_t)$ for the unrestricted specification, so that the model-implied total is $N_t + Tail(X_t) + Bub(X_t)$. Although the strip block is fitted tightly (overall RMSE of 0.078, Table 10), the implied total valuation exhibits a negative correlation with the observed series, indicating that the directly parameterised strip variances distort the broader valuation decomposition.

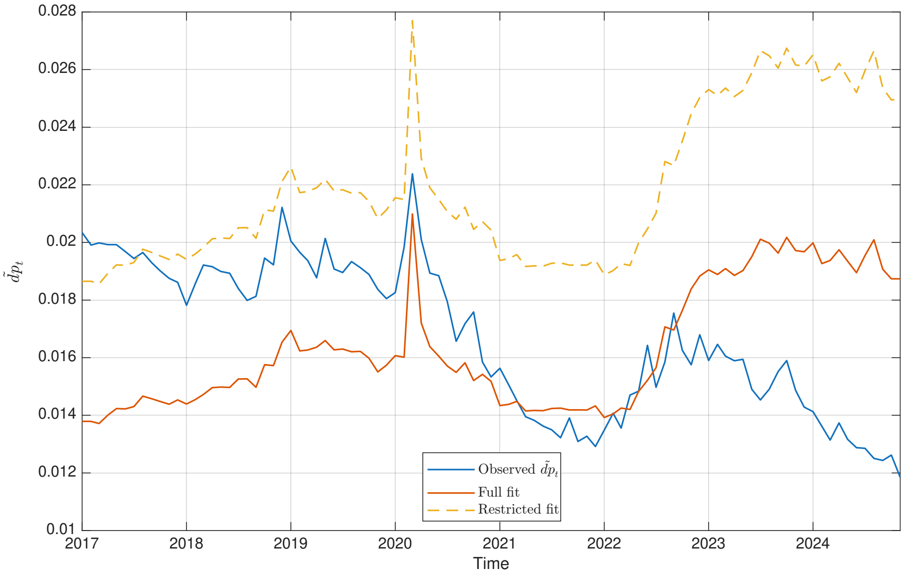


(a) Adjusted dividend-price ratio fit.

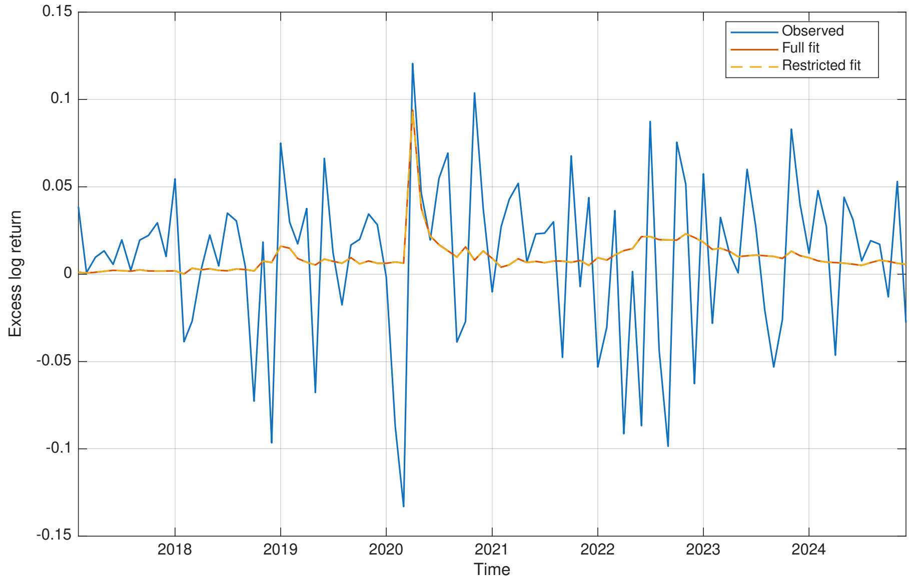


(b) Equity risk premium equation fit.

Figure 8: Supplementary fit diagnostics for the alternative joint-strips implementation
The figure reports supplementary fit diagnostics for the alternative joint-strips implementation. Panel (a) compares the observed adjusted dividend-price ratio $\widetilde{dp}_t$ to the fitted series under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. Panel (b) compares observed monthly excess log returns to the model-implied equity risk premium. Both measurement blocks are fitted comparably across the two specifications, indicating that the sign reversal in the total-valuation correlation documented in Table 10 arises from the interaction between the strip block and the valuation decomposition rather than from a failure of the return or dividend-price blocks individually.

# F Spot-only external strip validation

This appendix reports an external strip validation exercise for the short-sample spot-only model estimated in Section 5.2. The model is estimated without the strip panel in the likelihood, and the resulting parameter estimates are then used to compute model-implied strip equity yields maturity by maturity. These model-implied yields are compared to the observed traded strip panel, which played no role in estimation.

The exercise is informative for two reasons. First, it mirrors the external-validation logic of Giglio et al. (2024) by confronting a strip-free model with strip data ex post. Second, because the exercise uses the same short sample and the same structural model as the baseline designs in Section 5, the comparison is not confounded by differences in sample length, model specification, or state-vector construction. Any discrepancy between model-implied and observed strip yields can therefore be attributed directly to the absence of strip discipline in the estimation.

The evidence indicates that spot-only survival of the no-bubble restriction does not translate into strip-pricing accuracy. Figure 9 plots the average observed strip equity yield alongside the model-implied average yields under both specifications. The observed average term structure is positive and mildly upward-sloping across all ten maturities, whereas the model-implied term structures are negative at every maturity. Table 11 quantifies the mismatch maturity by maturity, and the results indicate that the level, sign, and cross-maturity shape of the observed strip term structure are not well reproduced by the strip-free model.

The time-series dimension of the mismatch is documented in Figure 10, which plots observed and model-implied strip equity yields at selected maturities over the sample. The strip-free model not only misses the average level but also fails to track the time variation of individual strip yields, particularly after 2022 when the divergence between spot and strip valuations widens. Correlations between observed and model-implied yields decline steadily with maturity and turn negative beyond the five-year horizon (Table 11), suggesting that the spot-only model has difficulty capturing the dynamics of longer-maturity dividend claims.

Figure 11 and Table 12 extend the comparison to strip-slope dynamics, defined as differences in equity yields across maturity pairs. Although some short-run slope variation is partially captured, the average level and medium-run evolution of several cross-maturity slopes remain materially misspecified. This evidence is informative because slope dynamics summarise the relative pricing of dividend claims at different horizons. A model that misses both the level and the slope of the strip term structure is unlikely to deliver a reliable decomposition of the spot price into near-end and long-end components.

Taken together, these results suggest that the spot-only model's survival of the no-bubble restriction in Section 5.2 reflects the absence of strip discipline rather than consistency with the traded strip panel. The model that passes the no-bubble test in the spot-only design is the same model that misprices the strip term structure along multiple dimensions when confronted with strips ex post. This pattern is consistent with the interpretation of the joint-strips rejection in

Section 5.3: traded strips appear to contain information about long-horizon discounting that the spot price alone does not reveal, and once this information enters the estimation, the no-bubble restriction is rejected.

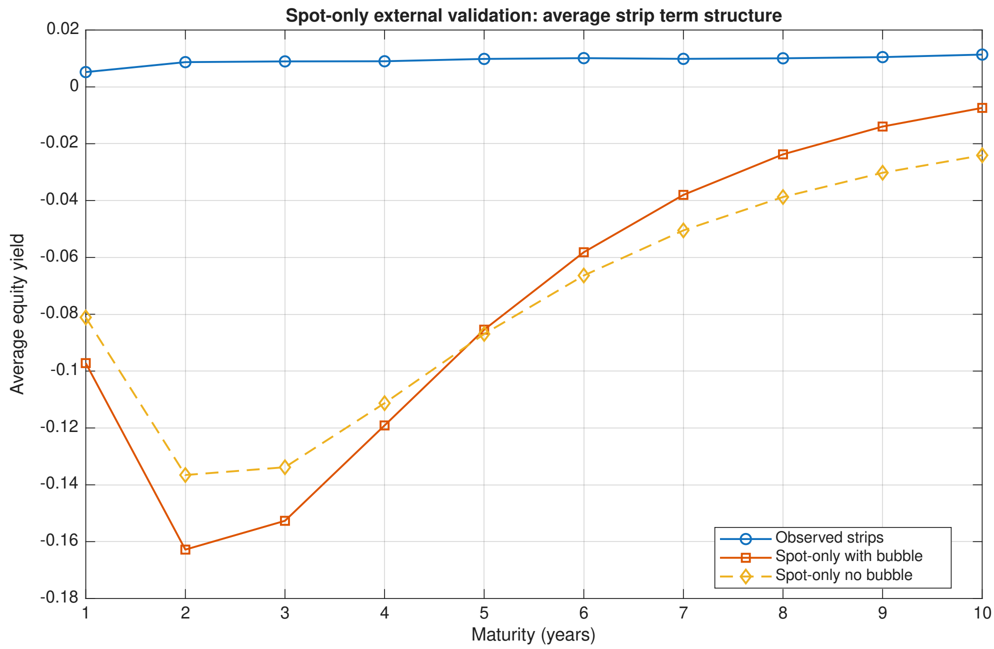


Figure 9: Spot-only external validation of the average strip term structure

The figure plots the average strip equity yield as a function of maturity (one to ten years) over the baseline sample (January 2017 to December 2024). The line with markers is the average observed strip equity yield, computed from the constant-maturity strip panel described in Section 4.1. The remaining lines are the average model-implied strip equity yields under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications of the short-sample spot-only model, which excludes strips from the likelihood. The observed average term structure is positive across all maturities, whereas both model-implied term structures are uniformly negative, suggesting a level and sign mismatch between the strip-free model and the traded strip panel.

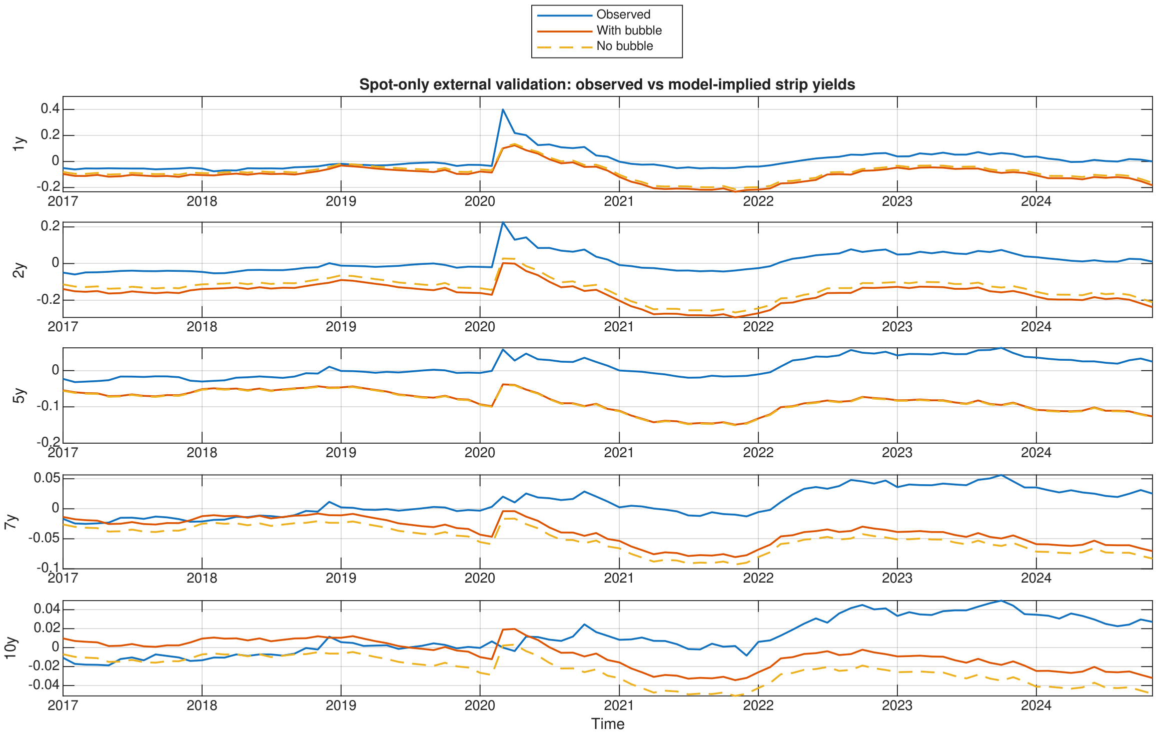


Figure 10: Spot-only external validation at selected strip maturities

The figure plots observed and model-implied strip equity yields over time for selected maturities (one, two, five, seven, and ten years). Each panel corresponds to a single constant-maturity strip. The model-implied series are computed from the short-sample spot-only parameter estimates, which exclude strips from the likelihood. The spot-only model misses not only the average level documented in Figure 9 but also the time-series dynamics of individual strip yields. The discrepancy is most pronounced after 2022, when the divergence between spot and strip valuations widens.

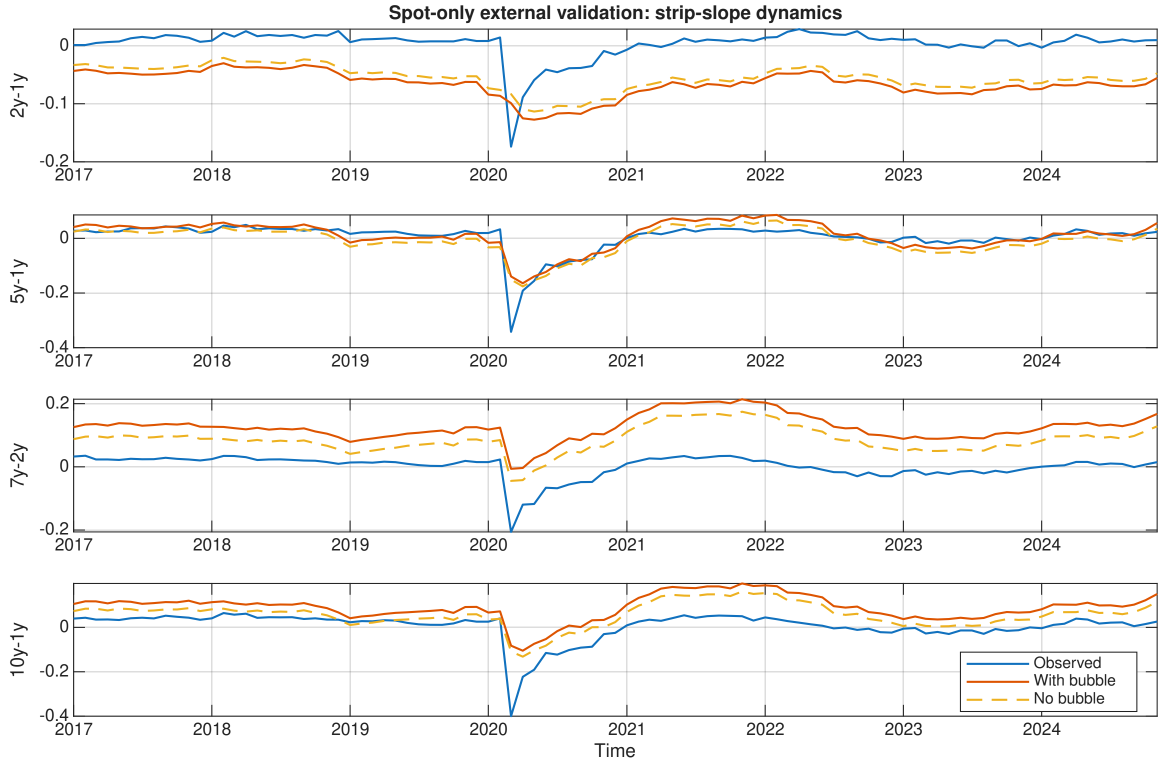


Figure 11: Spot-only external validation of strip-slope dynamics

The figure plots observed and model-implied strip-slope dynamics for selected maturity pairs over the baseline sample. Each slope is defined as the difference in strip equity yields between a longer and a shorter maturity (for example, the 10y−1y slope is the ten-year yield minus the one-year yield). The model-implied slopes are computed from the short-sample spot-only estimates. Although some short-run slope variation is partially captured, the average level and medium-run dynamics of several slopes remain systematically misspecified, indicating that the strip-free model fails to reproduce the relative pricing of dividend claims across maturities as well as their absolute levels.

Table 11: Spot-only external validation by strip maturity

| Maturity | Mean EY | | | RMSE | | Correlation | |
|---|---|---|---|---|---|---|---|
| | Obs. | Unrestr. | Restr. | Unrestr. | Restr. | Unrestr. | Restr. |
| 1y | 0.0052 | −0.0972 | −0.0811 | 0.1157 | 0.1015 | 0.6915 | 0.6918 |
| 2y | 0.0087 | −0.1628 | −0.1366 | 0.1805 | 0.1557 | 0.4724 | 0.4764 |
| 3y | 0.0090 | −0.1527 | −0.1338 | 0.1700 | 0.1522 | 0.2078 | 0.2115 |
| 4y | 0.0090 | −0.1191 | −0.1113 | 0.1366 | 0.1293 | 0.0267 | 0.0286 |
| 5y | 0.0098 | −0.0854 | −0.0869 | 0.1041 | 0.1055 | −0.0942 | −0.0936 |
| 6y | 0.0101 | −0.0582 | −0.0664 | 0.0779 | 0.0852 | −0.1834 | −0.1833 |
| 7y | 0.0099 | −0.0380 | −0.0505 | 0.0590 | 0.0695 | −0.2643 | −0.2643 |
| 8y | 0.0101 | −0.0237 | −0.0387 | 0.0463 | 0.0582 | −0.3304 | −0.3304 |
| 9y | 0.0105 | −0.0140 | −0.0302 | 0.0384 | 0.0503 | −0.3925 | −0.3924 |
| 10y | 0.0114 | −0.0074 | −0.0241 | 0.0340 | 0.0454 | −0.4503 | −0.4502 |

The table reports maturity-by-maturity external validation statistics for the short-sample spot-only model over the baseline sample (January 2017 to December 2024). The unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications are estimated without strips in the likelihood and then evaluated against the observed traded strip panel ex post. Mean EY is the time-series average of the strip equity yield at each maturity, defined as the log ratio of the strip present value to the current dividend level. RMSE is the root mean squared error between observed and model-implied strip equity yields over the sample. Correlation is the time-series correlation between observed and model-implied equity yields at each maturity. The observed average strip term structure is positive across all maturities, whereas both model-implied term structures are uniformly negative. Correlations are positive at short maturities, decline steadily with horizon, and turn negative beyond five years, suggesting that the ability of the spot-only model to capture time variation in strip yields diminishes as the horizon extends.

Table 12: Spot-only external validation of strip-slope dynamics

| Slope | Mean | | | RMSE | | Correlation | |
|---|---|---|---|---|---|---|---|
| | Obs. | Unrestr. | Restr. | Unrestr. | Restr. | Unrestr. | Restr. |
| 2y–1y | 0.0035 | −0.0656 | −0.0555 | 0.0717 | 0.0621 | 0.6832 | 0.6615 |
| 5y–1y | 0.0046 | 0.0118 | −0.0058 | 0.0316 | 0.0324 | 0.8236 | 0.8227 |
| 7y–2y | 0.0012 | 0.1248 | 0.0861 | 0.1272 | 0.0899 | 0.7130 | 0.7167 |
| 10y–1y | 0.0062 | 0.0898 | 0.0570 | 0.0929 | 0.0649 | 0.7838 | 0.7835 |

The table reports strip-slope external validation statistics for the short-sample spot-only model. Each slope is defined as the difference in strip equity yields between the longer-maturity and shorter-maturity endpoints indicated in the first column (for example, 2y−1y is the two-year equity yield minus the one-year equity yield). Mean is the time-series average of the slope over the baseline sample. RMSE and Correlation compare the observed and model-implied slope series. The table complements the level evidence in Table 11 by documenting that the strip-free model also struggles with the relative pricing of dividend claims across maturities, even when some short-run slope variation is partially captured.

# G Supplementary diagnostic figures for the information designs

This appendix reports design-specific fit figures for the three baseline information designs. These figures complement the main-text diagnostics in Section 5 by reporting the price-decomposition, adjusted dividend-price, and equity-risk-premium fit blocks for each design. Across all three designs, the unrestricted and restricted specifications are not visibly separated in the short-run equity-risk-premium block. The separation between specifications is concentrated in the long-end decomposition and strip-pricing blocks, which are the intended identification margins of the paper.

## G.1 Spot-only supplementary diagnostics

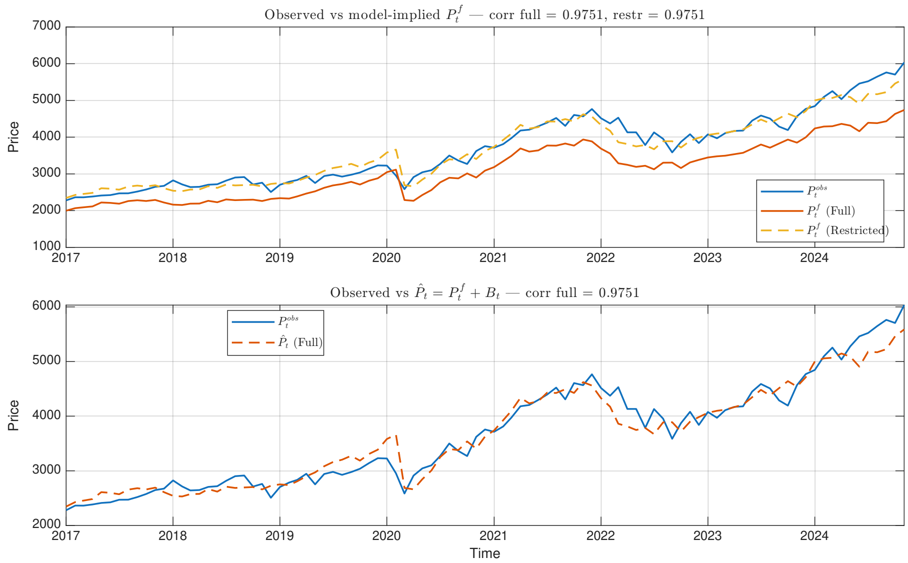


Figure 12: Spot-only price decomposition

The figure reports the price decomposition under the spot-only design over the baseline sample (January 2017 to December 2024). The top panel compares the observed spot valuation ratio $P_t/D_t$ to the model-implied fundamental component $N_t + Tail(X_t)$ under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. The bottom panel adds the estimated bubble component $Bub(X_t)$ for the unrestricted specification. The two specifications track the observed series with similar accuracy, which is consistent with the finding in Table 2 that the no-bubble restriction cannot be rejected in the spot-only design.

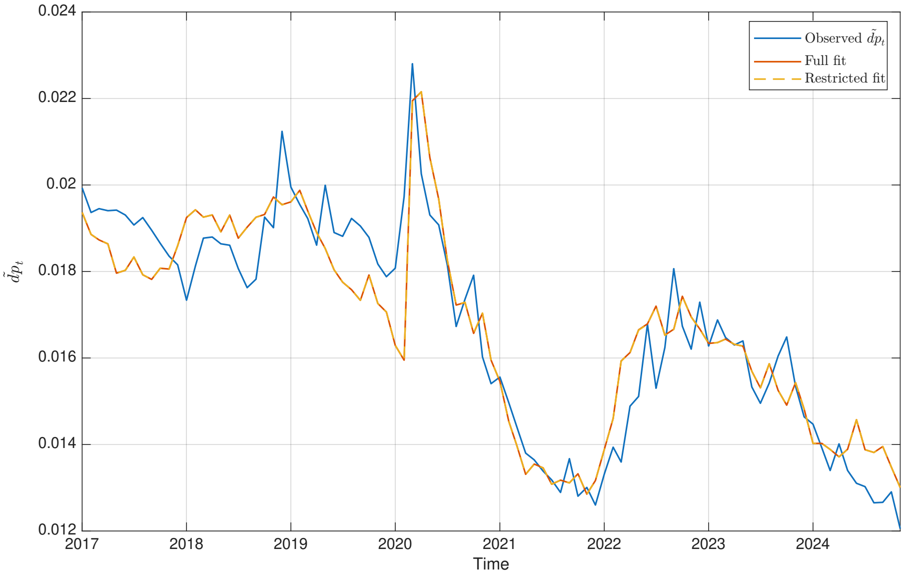


(a) Adjusted dividend-price ratio fit.

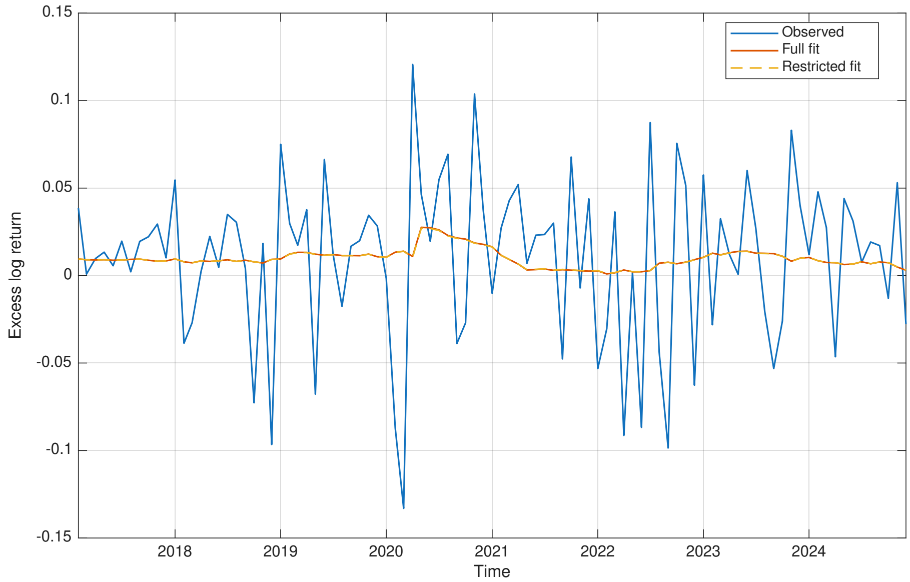


(b) Equity risk premium equation fit.

Figure 13: Spot-only adjusted dividend-price and equity-premium fit
The figure reports supplementary measurement-block diagnostics for the spot-only design. Panel (a) compares the observed adjusted dividend-price ratio $\widetilde{dp}_t$ to the fitted series under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. Panel (b) compares observed monthly excess log returns to the model-implied equity risk premium. The two specifications are not visibly separated in either block, consistent with the near-zero LR statistic reported in Table 2.

## G.2 Preferred joint-strips supplementary diagnostics

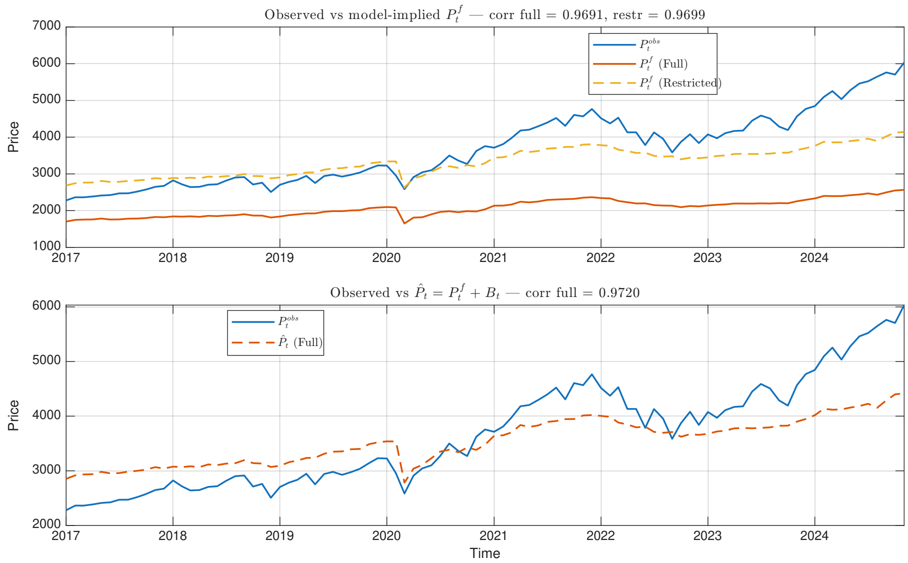


Figure 14: Joint-strips price decomposition

The figure reports the price decomposition under the preferred joint-strips design over the baseline sample (January 2017 to December 2024). The top panel compares the observed spot valuation ratio $P_t/D_t$ to the model-implied fundamental component under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. The bottom panel adds the estimated bubble component for the unrestricted specification. Relative to the spot-only design in Figure 12, the unrestricted model lifts the model-implied total valuation toward the observed series through the bubble component once the near end is anchored by traded strips.


(a) Adjusted dividend-price ratio fit.

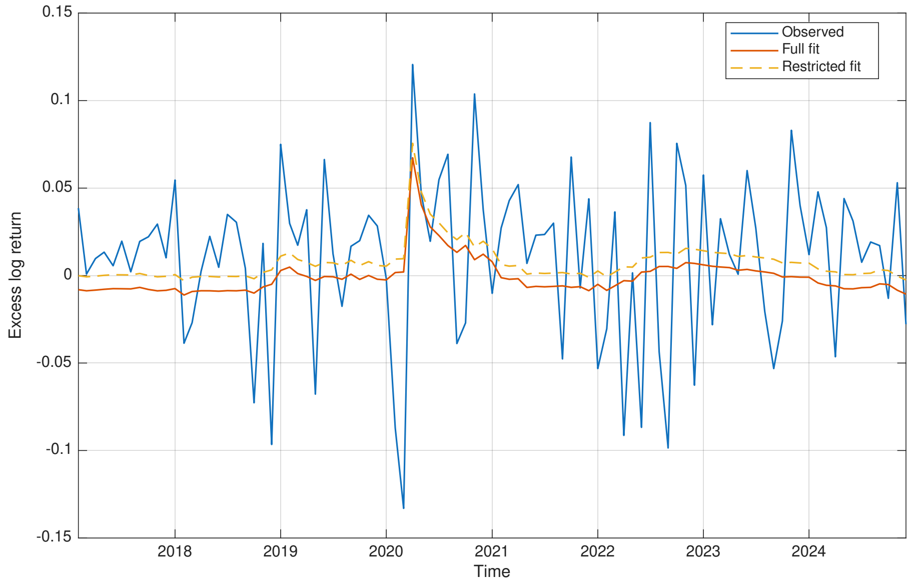


(b) Equity risk premium equation fit.

Figure 15: Joint-strips adjusted dividend-price and equity-premium fit

The figure reports supplementary measurement-block diagnostics for the preferred joint-strips design. Panel (a) compares the observed adjusted dividend-price ratio $\widetilde{dp}_t$ to the fitted series under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. Panel (b) compares observed monthly excess log returns to the model-implied equity risk premium. The separation between specifications in this design is concentrated in the strip and long-end blocks (Figure 2) rather than in the short-run return or dividend-price blocks shown here.

## G.3 Maturity hold-out supplementary diagnostics

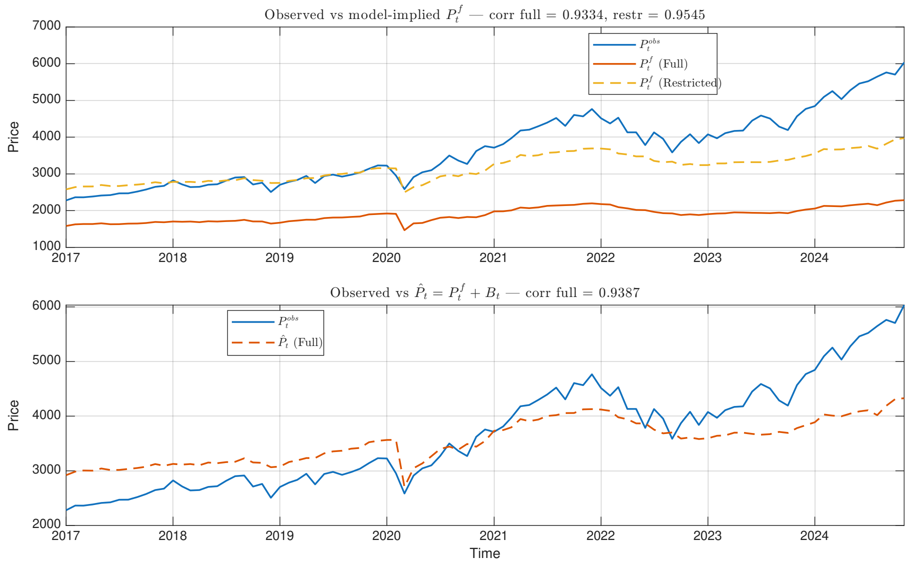


Figure 16: Maturity hold-out price decomposition

The figure reports the price decomposition under the maturity hold-out design over the baseline sample (January 2017 to December 2024). The model is estimated on the one-to-five-year strip block, and the six-to-ten-year strips are reserved for validation. The top panel compares the observed $P_t/D_t$ to the model-implied fundamental component under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. The bottom panel adds the estimated bubble component for the unrestricted specification. The restricted model retains modestly tighter spot-valuation correlations, while the unrestricted model is supported by its performance on the unseen longer maturities (Table 2).

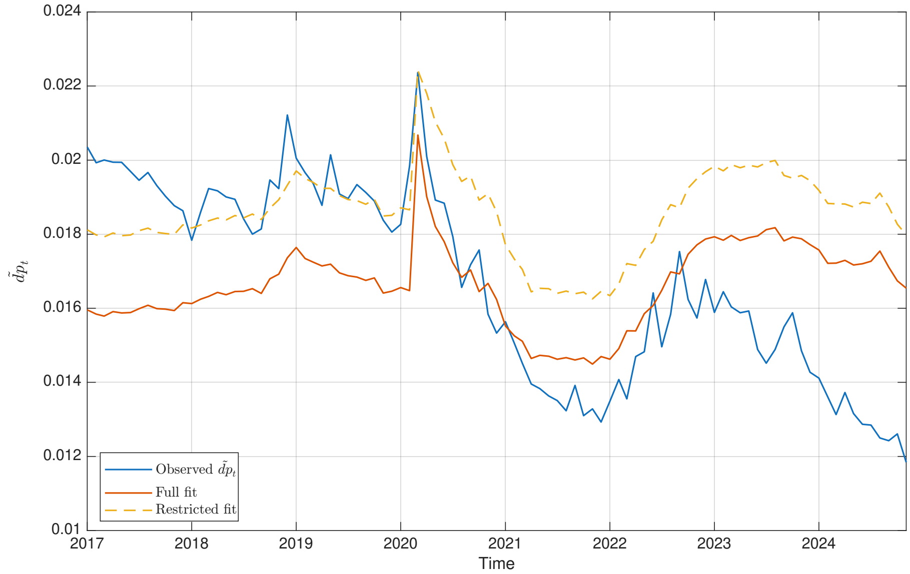


(a) Adjusted dividend-price ratio fit.

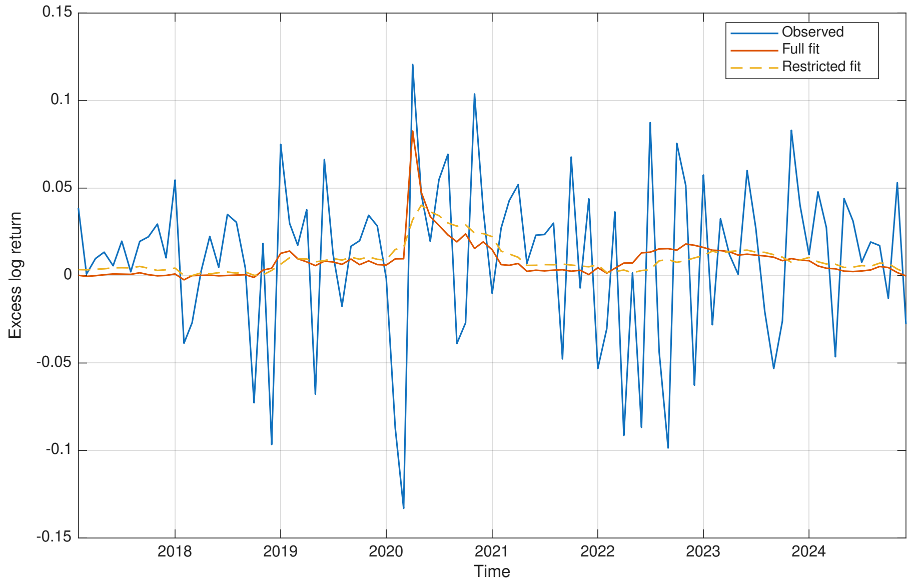


(b) Equity risk premium equation fit.

Figure 17: Maturity hold-out adjusted dividend-price and equity-premium fit
The figure reports supplementary measurement-block diagnostics for the maturity hold-out design. Panel (a) compares the observed adjusted dividend-price ratio $\widetilde{dp}_t$ to the fitted series under the unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications. Panel (b) compares observed monthly excess log returns to the model-implied equity risk premium. The hold-out result is driven by the unrestricted model's performance on unseen strip maturities (Table 2) rather than by a separation in the short-run return or dividend-price blocks.

# H Alternative design diagnostics

This appendix reports the full table-based results for the alternative-specification diagnostics summarised in Table 3. The battery contains two constant-price specifications (joint-strips and spot-only with $\Lambda_1 = 0$) and one diagonal strip-covariance specification, all estimated on the same baseline sample with the same $r = 3$ state vector.

## H.1 Constant price-of-risk specifications

Table 13 reports the full fit summary. The constant-price joint-strips specification rejects $c = 0$ (LR = 102.89, $p < 0.001$) but with negative total-valuation and long-end correlations. The constant-price spot-only specification does not reject and produces external strip-validation RMSE exceeding 17 (unrestricted) and 37 (restricted).

Table 14 reports the price-of-risk metrics. In the constant-price specifications, the latent risk-price moments are identically zero because $\Lambda_1 = 0$ by construction. Table 15 reports the corresponding parameter loadings.

## H.2 Diagonal strip-covariance specification

The diagonal-covariance specification rejects $c = 0$ (LR = 115.21, $p < 0.001$) with overall strip RMSE of 0.081. Table 13 shows that this specification achieves low strip RMSE while leaving total-valuation and long-end correlations negative. Table 16 reports the expected-return term structure: under diagonal covariance, the unrestricted and restricted expected-return profiles are closer to one another than in the preferred baseline, which is one manifestation of the compressed economic separation noted in Section 6.2. Table 17 reports the long-end valuation toolkit.

Table 13: Full fit summary for the diagnostic battery

| Specification | Model | $\log L$ | LR | $p$-value | Strip RMSE | | Correlations | | | $B_t/P_t$ (unrestr. only) | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | Est. | Val. | DP | Total $P/D$ | Long end | Mean | Max |
| Joint strips, constant prices of risk | Unrestr. | 3749.191 | 102.888 | $< 0.001$ | 0.132 | – | −0.176 | −0.135 | −0.114 | 0.410 | 0.671 |
| | Restr. | 3697.747 | | | 0.196 | – | −0.182 | −0.141 | −0.114 | – | – |
| Spot-only, constant prices of risk | Unrestr. | 920.027 | −0.649 | 1.000 | – | 17.361 | −0.213 | −0.189 | −0.114 | $7.9 \times 10^{-6}$ | $1.3 \times 10^{-5}$ |
| | Restr. | 920.352 | | | – | 37.040 | −0.236 | −0.214 | −0.116 | – | – |
| Joint strips, diagonal strip covariance | Unrestr. | 2098.496 | 115.209 | $< 0.001$ | 0.081 | – | −0.141 | −0.112 | −0.177 | 0.350 | 0.486 |
| | Restr. | 2040.892 | | | 0.081 | – | −0.156 | −0.126 | −0.183 | – | – |

The table reports the full fit summary for the diagnostic battery estimated on the baseline sample (January 2017 to December 2024) with the $r = 3$ state vector. Each specification is estimated under both the unrestricted ($c \geq 0$) and restricted ($c = 0$) models. Est. and Val. refer to the estimation-set and validation-set strip RMSE, respectively. For the spot-only specification, strips are excluded from the likelihood, so performance is evaluated ex post on the full one-to-ten-year panel. DP is the correlation between the observed adjusted dividend-price ratio and the model-implied fit. Total $P/D$ and Long end are as defined in Table 2. $B_t/P_t$ is the model-implied bubble share, reported only for the unrestricted specification. The constant-price specifications remove the $\Lambda_1$ channel. The diagonal-covariance specification retains it but uses a diagonal rather than profiled strip-residual covariance.

Table 14: Price-of-risk metrics for the diagnostic battery

| Specification | Model | Raw $\lambda_{0y}$ | $\rho(\Phi_x^Q)$ | TransMargin | Latent risk-price moments | | | $c$ | Mean $B_t/P_t$ |
|---|---|---|---|---|---|---|---|---|---|
| | | | | | $E[\ell' X_t]$ | $\mathrm{Std}(\ell' X_t)$ | AR(1) | | |
| Joint strips, constant prices of risk | Unrestr. | −47050.006 | 0.981244 | 0.001397 | 0.000000 | 0.000000 | 0.000000 | 24.123 | 0.410 |
| | Restr. | −1.400 | 0.981244 | 0.000001 | 0.000000 | 0.000000 | 0.000000 | 0.000 | 0.000 |
| Spot-only, constant prices of risk | Unrestr. | −8.550 | 0.981244 | 0.000002 | 0.000000 | 0.000000 | 0.000000 | 0.0005 | $7.9 \times 10^{-6}$ |
| | Restr. | −6.188 | 0.981244 | 0.000005 | 0.000000 | 0.000000 | 0.000000 | 0.0000 | 0.000 |
| Joint strips, diagonal strip covariance | Unrestr. | −91.863 | 0.982989 | 0.000666 | −0.012787 | 1.006359 | 0.476761 | 20.491 | 0.350 |
| | Restr. | 539.662 | 0.983074 | 0.000001 | −0.012733 | 1.006381 | 0.474526 | 0.000 | 0.000 |

The table reports price-of-risk metrics for the diagnostic battery. Raw $\lambda_{0y}$ is the estimated level parameter of the market price of risk (not interpretable in isolation). $\rho(\Phi_x^Q)$ is the spectral radius of the risk-neutral transition matrix. TransMargin is the transversality margin defined in Appendix C. The latent risk-price moments refer to the index $\ell' X_t$ from the rank-one specification; in the constant-price rows these are identically zero because $\Lambda_1 = 0$ by construction. The diagonal-covariance specification retains the rank-one channel, and the corresponding AR(1) of approximately 0.48 indicates moderate persistence in the risk-price factor.

Table 15: Price-of-risk loadings for the diagnostic battery

| Specification | Model | Risk-price direction ($\ell$) | | | | Return loadings ($\gamma_r$) | | | | Dividend loadings ($\gamma_d$) | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | $\ell_1$ | $\ell_2$ | $\ell_3$ | $\ell_4$ | $\gamma_{r,1}$ | $\gamma_{r,2}$ | $\gamma_{r,3}$ | $\gamma_{r,4}$ | $\gamma_{d,1}$ | $\gamma_{d,2}$ | $\gamma_{d,3}$ | $\gamma_{d,4}$ |
| Joint strips, constant prices of risk | Unrestr. | 0.0000 | 0.0000 | 0.0000 | 0.0000 | −0.0001 | 0.0001 | −0.0001 | −0.0002 | 0.0000 | −8.9059 | −12.9765 | 39.9997 |
| | Restr. | 0.0000 | 0.0000 | 0.0000 | 0.0000 | −0.8405 | 0.0001 | −0.0002 | 0.0001 | 6.5276 | −27.4515 | −44.9259 | 130.6100 |
| Spot-only, constant prices of risk | Unrestr. | 0.0000 | 0.0000 | 0.0000 | 0.0000 | −0.7056 | −0.0001 | −0.0000 | 0.0000 | 63.3236 | 151.2744 | 283.2193 | −211.8696 |
| | Restr. | 0.0000 | 0.0000 | 0.0000 | 0.0000 | −0.8862 | −0.0005 | −0.0010 | −0.0007 | 118.1376 | 165.1857 | 341.9446 | 27.0208 |
| Joint strips, diagonal strip covariance | Unrestr. | 0.0006 | −0.5347 | 0.5158 | −0.6694 | −0.0724 | 0.0058 | −0.0056 | 0.0072 | −0.0163 | 0.0348 | −0.1211 | −0.0247 |
| | Restr. | −0.0004 | −0.5316 | 0.5151 | −0.6724 | 0.0123 | 0.0053 | −0.0057 | 0.0075 | 0.0037 | 0.0511 | −0.1310 | −0.0443 |

The table reports the rank-one direction $\ell$, the return covariance loadings $\gamma_r$, and the dividend covariance loadings $\gamma_d$ for the diagnostic battery. In the constant-price specifications, $\ell$ is identically zero because $\Lambda_1 = 0$ is imposed; the $\gamma_d$ loadings in those rows absorb the model's attempt to match the strip evidence through the dividend-innovation channel rather than through dynamic risk prices. The diagonal-covariance specification retains the rank-one channel and produces $\ell$ loadings of economically meaningful magnitude.

Table 16: Expected strip excess returns for the diagnostic battery

| Specification | Model | 12m | 24m | 60m | 120m |
|---|---|---|---|---|---|
| Joint strips, constant prices of risk | Unrestr. | 0.0046 | 0.0043 | 0.0043 | 0.0047 |
| | Restr. | 0.0018 | 0.0021 | 0.0031 | 0.0034 |
| Spot-only, constant prices of risk | Unrestr. | 0.0804 | 0.4546 | 0.3608 | 0.1133 |
| | Restr. | 0.3603 | 0.9173 | 0.6715 | 0.2098 |
| Joint strips, diagonal strip covariance | Unrestr. | 0.0045 | 0.0045 | 0.0047 | 0.0045 |
| | Restr. | 0.0045 | 0.0047 | 0.0045 | 0.0040 |

The table reports model-implied expected excess log returns for dividend strips at selected horizons (monthly). The constant-price spot-only specification generates expected excess returns exceeding 0.36 per month at the 24-month and 60-month horizons, which is a consequence of the model's attempt to absorb strip evidence through distorted horizon-specific expected returns once the dynamic risk-price channel is absent. The constant-price joint-strips specification produces more moderate values because the strip block constrains the model within estimation. The diagonal-covariance specification produces expected-return profiles that are similar across unrestricted and restricted models, reflecting the compressed economic separation noted in Section 6.2.

Table 17: Long-end valuation components for the diagnostic battery

| Specification | Model | Mean valuation components | | | | | Correlations | | ER slope (120m–12m) | Mean $B_t/P_t$ |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Near end (observed) | Near end (model) | Tail | Bubble | Total $P/D$ | Total $P/D$ | Long end | | |
| Joint strips, constant prices of risk | Unrestr. | 9.560 | 9.530 | 26.391 | 24.287 | 60.237 | −0.135 | −0.114 | 0.00003 | 0.410 |
| | Restr. | 9.560 | 10.853 | 44.476 | 0.000 | 54.035 | −0.141 | −0.114 | 0.00153 | 0.000 |
| Spot-only, constant prices of risk | Unrestr. | 59.341 | 59.341 | 0.000 | 0.0005 | 59.342 | −0.189 | −0.114 | 0.03293 | $7.9 \times 10^{-6}$ |
| | Restr. | 59.330 | 59.330 | 0.000 | 0.0000 | 59.330 | −0.214 | −0.116 | −0.15042 | 0.000 |
| Joint strips, diagonal strip covariance | Unrestr. | 9.560 | 9.529 | 30.862 | 20.703 | 61.124 | −0.112 | −0.177 | 0.00005 | 0.350 |
| | Restr. | 9.560 | 9.565 | 36.477 | 0.000 | 46.037 | −0.126 | −0.183 | −0.00047 | 0.000 |

The table reports the long-end valuation decomposition for the diagnostic battery. Near end (observed) is the sample mean of $N_t$ from traded strips. Near end (model) is the model-implied counterpart. Tail is the sample mean of $Tail(X_t)$. Bubble is the sample mean of $Bub(X_t)$, applicable only to the unrestricted specification. Total $P/D$ is the sum of all components. Correlations are defined as in Table 2. ER slope (120m–12m) is the difference between the 120-month and 12-month expected strip excess returns. In the constant-price spot-only specification, the model-implied near end absorbs all of $P_t/D_t$ and the tail collapses to zero, which is consistent with the finding that this specification does not separate near-end from long-end valuation.

# I Long-sample strip-free external validation

This appendix reports the full results for the long-sample strip-free external validation exercise summarised in Section 6.4. The model is estimated on January 2000 to December 2024 spot-only data without strips in the likelihood, and the resulting parameter estimates are evaluated against the observed strip panel over the 2017–2024 overlap period.

## I.1 Maturity-by-maturity strip levels

Table 18 reports the maturity-by-maturity comparison. The unrestricted specification achieves positive time-series correlations at all maturities, ranging from 0.25 at the one-year horizon to 0.65 at the ten-year horizon. However, the level errors are large: average model-implied equity yields exceed 0.17 at the one-year maturity and 0.33 at the ten-year maturity, compared to observed values below 0.012. The restricted specification reduces level errors (RMSE below 0.09 at all maturities) but produces correlations that are negative beyond the two-year horizon.

Figure 18 plots selected maturities over time. The unrestricted series is persistently above the observed strip yields. The restricted series is closer in level but often moves in the opposite direction relative to observed yields, particularly around the 2020 episode and its aftermath.

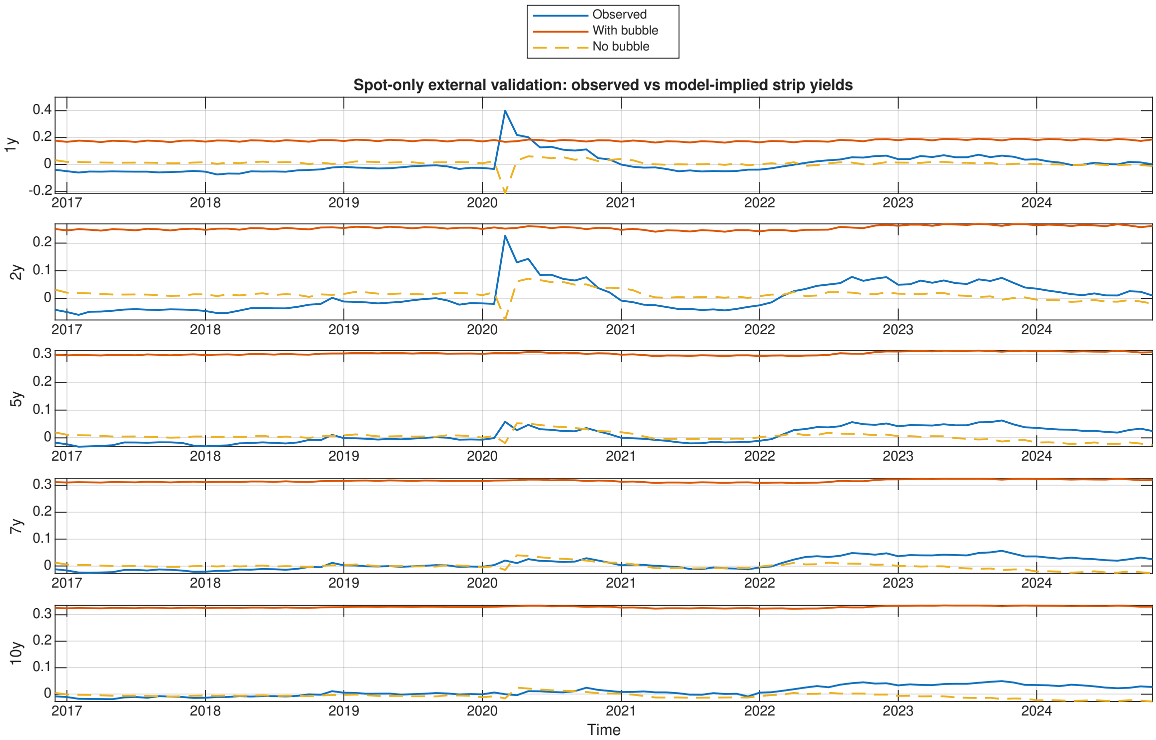


Figure 18: Long-sample strip-free external validation at selected maturities
The figure plots observed and model-implied strip equity yields over time for selected maturities (one, two, five, seven, and ten years) over the 2017–2024 overlap sample. Both specifications are estimated on January 2000 to December 2024 spot-only data without strips in the likelihood. The unrestricted specification ($c \geq 0$) generates equity yields that track some of the time-series variation in the observed series but at a level far above the data. The restricted specification ($c = 0$) is closer in average level but often fails to reproduce the direction of observed movements, particularly around 2020.

## I.2 Slope dynamics

Table 19 extends the comparison to strip-slope dynamics. The unrestricted specification produces mean slopes that are systematically too large (for example, 0.15 for the 10y–1y slope versus 0.006 observed). The restricted specification produces slopes that are generally too compressed or negative (−0.017 for the same pair). Slope correlations are weak or negative in both cases.

Figure 19 shows that the observed strip term structure inverts around 2020 and subsequently normalises. Neither strip-free specification reproduces this episode. The unrestricted model generates slopes that are too positive and too smooth. The restricted model compresses the slope and often moves in the opposite direction.

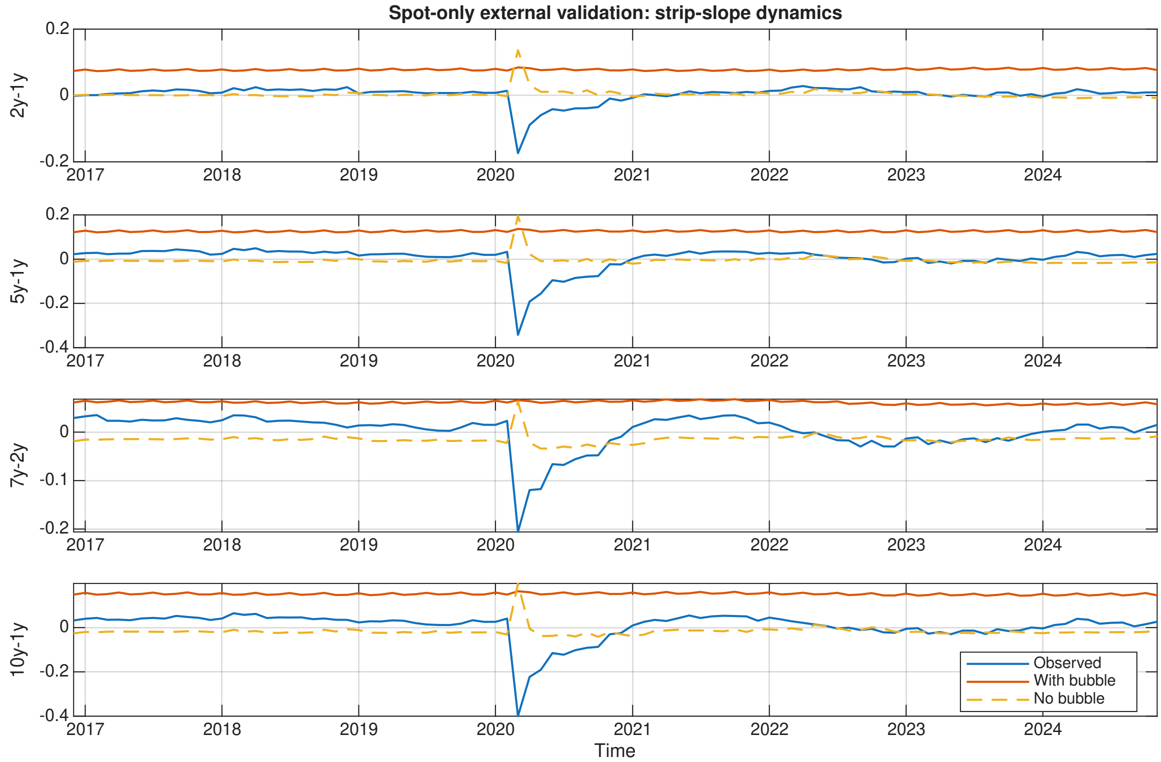


Figure 19: Long-sample strip-free external validation of strip-slope dynamics
The figure plots observed and model-implied strip-slope dynamics for four maturity pairs (2y–1y, 5y–1y, 7y–2y, and 10y–1y) over the 2017–2024 overlap sample. The long-sample strip-free model is estimated on January 2000 to December 2024 spot data without strips in the likelihood. The observed slope shows a sharp inversion around 2020 followed by gradual normalisation. The unrestricted specification remains too positive and too smooth, while the restricted specification is too compressed and often negative. Neither specification reproduces the crisis episode or the subsequent recovery.

## I.3 Interpretation

The appendix results show that the long-sample strip-free model does not jointly recover the level and the dynamics of the traded strip term structure. The unrestricted specification recovers some of the time-series comovement at longer maturities but at level errors that exceed the observed values by more than an order of magnitude. The restricted specification keeps levels closer to the data but loses time-series dynamics and slope behaviour.

This outcome is consistent with the interpretation that long-sample spot-only estimation does not discipline long-horizon discounting with sufficient precision to reproduce the strip curve along both margins simultaneously. The exercise supports the view that direct strip discipline is informative for the evaluation of the no-bubble restriction, even when the spot-only sample is extended to 25 years.

Table 18: Long-sample strip-free external validation by maturity

| Maturity | Mean EY | | | RMSE | | Correlation | |
|---|---|---|---|---|---|---|---|
| | Obs. | Unrestr. | Restr. | Unrestr. | Restr. | Unrestr. | Restr. |
| 1y | 0.0047 | 0.1771 | 0.0102 | 0.1857 | 0.0830 | 0.2544 | −0.3095 |
| 2y | 0.0082 | 0.2548 | 0.0141 | 0.2512 | 0.0533 | 0.5370 | 0.0769 |
| 3y | 0.0086 | 0.2806 | 0.0118 | 0.2742 | 0.0408 | 0.6566 | 0.1034 |
| 4y | 0.0087 | 0.2942 | 0.0084 | 0.2869 | 0.0349 | 0.7008 | 0.0559 |
| 5y | 0.0096 | 0.3034 | 0.0050 | 0.2948 | 0.0312 | 0.7232 | 0.0074 |
| 6y | 0.0099 | 0.3104 | 0.0018 | 0.3012 | 0.0293 | 0.7226 | −0.0371 |
| 7y | 0.0096 | 0.3162 | −0.0009 | 0.3071 | 0.0286 | 0.7111 | −0.0916 |
| 8y | 0.0098 | 0.3211 | −0.0032 | 0.3117 | 0.0282 | 0.6996 | −0.1469 |
| 9y | 0.0103 | 0.3253 | −0.0050 | 0.3155 | 0.0284 | 0.6783 | −0.2013 |
| 10y | 0.0112 | 0.3289 | −0.0065 | 0.3182 | 0.0292 | 0.6488 | −0.2469 |

The table reports maturity-by-maturity external validation statistics for the long-sample strip-free model. The unrestricted ($c \geq 0$) and restricted ($c = 0$) specifications are estimated on January 2000 to December 2024 spot-only data without strips in the likelihood and then evaluated against the observed traded strip panel over the 2017–2024 overlap period. EY denotes strip equity yield, defined as minus the log strip price ratio divided by maturity. RMSE is the root mean squared error between observed and model-implied equity yields. Correlation is the time-series correlation between observed and model-implied yields at each maturity. The unrestricted specification achieves positive correlations at all maturities but with average equity yields far above the data. The restricted specification reduces level errors but produces correlations that are negative beyond the two-year horizon.

Table 19: Long-sample strip-free external validation of strip slopes

| Slope | Mean | | | RMSE | | Correlation | |
|---|---|---|---|---|---|---|---|
| | Obs. | Unrestr. | Restr. | Unrestr. | Restr. | Unrestr. | Restr. |
| 2y–1y | 0.0034 | 0.0777 | 0.0038 | 0.0788 | 0.0382 | −0.3034 | −0.7599 |
| 5y–1y | 0.0048 | 0.1263 | −0.0053 | 0.1330 | 0.0704 | −0.3301 | −0.6683 |
| 7y–2y | 0.0014 | 0.0614 | −0.0150 | 0.0697 | 0.0426 | 0.1046 | −0.2325 |
| 10y–1y | 0.0064 | 0.1518 | −0.0167 | 0.1592 | 0.0823 | −0.1895 | −0.5232 |

The table reports strip-slope external validation statistics for the long-sample strip-free model. Each slope is defined as the difference in strip equity yields between the longer-maturity and shorter-maturity endpoints. The unrestricted specification produces slopes that are systematically too large, while the restricted specification produces slopes that are too compressed or negative. Slope correlations are weak or negative in both cases, indicating that neither strip-free specification reproduces the relative pricing of dividend claims across maturities in the overlap period.

# J State-vector dimension sweep across risk-price specifications

This appendix reports the full state-vector dimension sweep across both dimension ($r \in \{0, 1, 3, 4, 5, 6\}$) and risk-price specification ($\lambda_1$-mode $\in$ {rank1, none}). The sweep is estimated on the same baseline sample with the same predictor panel as in Section 4.2. For each combination ($r, \lambda_1$-mode), the state vector is reconstructed by re-running the principal-component extraction, after which the spot-only, joint-strips, and maturity hold-out designs are re-estimated.

Table 20 reports the results. Among the rank-one rows, the pattern documented in Table 5 of the main text is visible: $r = 3$ is the smallest dimension at which the spot-only design does not reject, both strip-using designs reject, and the strip-using valuation correlations are positive. At $r \geq 4$, individual fit moments can be higher in one block, but the hold-out valuation correlations alternate in sign across dimensions, suggesting that additional factors redistribute fit rather than producing a monotone gain.

Among the constant-price rows ($\lambda_1$-mode = none), the spot-only design never rejects the no-bubble restriction (all $p$-values equal 1.000), but the associated strip validation errors can be large. The strip-using designs reject the restriction at all dimensions, but the total-valuation and long-end correlations remain negative throughout the grid. This pattern indicates that increasing the state dimension does not substitute for time-varying prices of risk in producing an economically interpretable identification result.

Table 20: Full state-vector dimension sweep across $r$ and risk-price specification

| $\lambda_1$ mode | $r$ | Spot-only | | | Joint strips | | | | Maturity hold-out | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | $p(c=0)$ | $\log L_F$ | Strip RMSE (restricted) | LR | Corr. total $P/D$ | Corr. long end | $\log L_F$ | LR | Val. strip RMSE (unrestricted) | Corr. total $P/D$ | Corr. long end |
| none | 0 | 1.000 | 883.903 | 0.242 | 105.5 | −0.216 | −0.259 | 3741.102 | 112.5 | 0.126 | −0.216 | −0.259 |
| | 1 | 1.000 | 919.824 | 37.851 | 123.3 | −0.259 | −0.229 | 3748.927 | 124.0 | 0.154 | −0.259 | −0.229 |
| | 3 | 1.000 | 920.091 | 33.331 | 102.2 | −0.135 | −0.114 | 3749.194 | 116.8 | 0.152 | −0.135 | −0.114 |
| | 4 | 1.000 | 920.982 | 1023.822 | 108.8 | −0.581 | −0.715 | 3766.660 | 109.6 | 0.125 | −0.581 | −0.715 |
| | 5 | 1.000 | 921.045 | 111.876 | 132.8 | −0.582 | −0.716 | 3782.189 | 85.0 | 0.123 | −0.582 | −0.716 |
| | 6 | 1.000 | 920.933 | 55.780 | 128.5 | −0.513 | −0.685 | 3762.565 | 31.0 | 0.125 | −0.513 | −0.685 |
| rank1 | 0 | $< 0.001$ | 924.828 | 0.230 | 95.1 | −0.269 | −0.296 | 3750.495 | 119.1 | 0.161 | −0.268 | −0.296 |
| | 1 | 0.272 | 925.964 | 0.271 | 107.0 | −0.205 | −0.259 | 3767.502 | 134.1 | 0.120 | −0.186 | −0.248 |
| | **3** | **0.705** | **1010.217** | **0.261** | **115.8** | **0.278** | **0.350** | **3839.356** | **147.2** | **0.141** | **0.237** | **0.277** |
| | 4 | 0.937 | 1016.101 | 0.276 | 38.2 | 0.843 | 0.852 | 3868.560 | 99.8 | 0.099 | −0.377 | −0.490 |
| | 5 | 0.930 | 1019.375 | 0.305 | 36.0 | 0.855 | 0.854 | 3812.998 | 136.3 | 0.166 | 0.913 | 0.919 |
| | 6 | 1.000 | 1017.816 | 0.297 | 74.4 | 0.855 | 0.846 | 3833.491 | 74.9 | 0.085 | −0.383 | −0.494 |

The table reports the full state-vector dimension sweep across $r \in \{0, 1, 3, 4, 5, 6\}$ and $\lambda_1$-mode $\in$ {rank1, none} over the baseline sample (January 2017 to December 2024). For each combination, the state vector is reconstructed from the same predictor panel (Section 4.2) and the three information designs are re-estimated. The spot-only columns report the LR $p$-value for $c = 0$, the unrestricted spot-only log-likelihood, and the restricted model's external strip-validation RMSE. The joint-strips columns report the LR statistic, the unrestricted valuation correlations, and the unrestricted log-likelihood. The hold-out columns report the LR statistic, the unrestricted validation-set RMSE on unseen six-to-ten-year strips, and the unrestricted valuation correlations. The highlighted row ($r = 3$, rank1) is the baseline specification. Among the rank-one rows, this is the smallest dimension at which the three designs are internally coherent. Among the constant-price rows, all dimensions produce negative valuation correlations in the strip-using designs, indicating that a richer state vector does not substitute for time-varying prices of risk.

# K Tail-horizon sensitivity

This appendix reports the full likelihood decomposition for the tail-horizon sensitivity exercise summarised in Table 6. The preferred baseline specification ($r = 3$, rank-one time-varying prices of risk) is held fixed, and only the additional tail horizon $N_{\text{TAIL\_ADD}}$ varies across rows.

Table 21 reports the unrestricted and restricted log-likelihoods for each design at each tail horizon. The spot-only result is stable: the unrestricted and restricted log-likelihoods are virtually identical at all three horizons, and the LR statistics are close to zero. In the strip-using designs, the unrestricted log-likelihoods are approximately flat across the sweep (for example, in the joint-strips design, 3839.3, 3841.0, 3840.9), whereas the restricted log-likelihoods improve monotonically as the tail horizon lengthens (3668.7, 3781.4, 3831.5). The monotone attenuation of the LR statistics therefore reflects increased flexibility of the restricted model at longer tail horizons rather than a deterioration of the unrestricted model's fit.

The unrestricted model's valuation diagnostics are approximately stable: the joint-strips correlations remain positive (0.278, 0.313, 0.311 for total $P/D$) and the hold-out validation-set strip RMSE is approximately 0.14 at all three horizons. These patterns suggest that the tail-horizon choice affects the severity of the test on the restricted model but does not alter the qualitative finding that the strip-using designs favour the unrestricted specification.

Table 21: Full tail-horizon sensitivity results

| Additional tail horizon (years / months) | Spot-only | | | | Joint strips | | | | | Maturity hold-out | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | $\log L_F$ | $\log L_R$ | LR | $p$-value | $\log L_F$ | $\log L_R$ | LR | Corr. total $P/D$ | Corr. long end | $\log L_F$ | $\log L_R$ | LR | Val. strip RMSE |
| 20 years (240m) | 1010.180 | 1010.135 | 0.089 | 0.7653 | 3839.343 | 3668.717 | 341.252 | 0.278 | 0.350 | 2306.954 | 2125.945 | 362.018 | 0.1394 |
| **40 years (480m, baseline)** | **1010.172** | **1010.156** | **0.031** | **0.8603** | **3840.964** | **3781.426** | **119.075** | **0.313** | **0.404** | **2307.054** | **2230.141** | **153.825** | **0.1403** |
| 80 years (960m) | 1010.149 | 1010.148 | 0.002 | 0.9670 | 3840.912 | 3831.482 | 18.861 | 0.311 | 0.399 | 2306.948 | 2298.895 | 16.106 | 0.1394 |

The table reports the full tail-horizon sensitivity results for the preferred specification ($r = 3$, rank-one time-varying prices of risk, January 2017 to December 2024). Additional tail horizon refers to $N_{\mathrm{TAIL\_ADD}}$, the number of extra monthly claims in the model-implied fundamental tail beyond the ten-year observed strip segment. $\log L_F$ and $\log L_R$ are the maximised log-likelihoods of the unrestricted and restricted models, respectively. LR is twice their difference. The spot-only columns show that the no-bubble restriction is not rejected at any tail horizon. The joint-strips and hold-out columns show that the unrestricted log-likelihood is approximately stable while the restricted log-likelihood improves monotonically as the tail lengthens, producing the observed attenuation of the LR statistics. The unrestricted valuation correlations in the joint-strips design remain positive throughout, and the hold-out validation-set RMSE is approximately constant, indicating that the tail choice affects the restricted model's absorptive capacity rather than the unrestricted model's fit quality.

# References


Avino, D. E., Stancu, A., & Simen, C. W. (2021). The Predictive Power of the Dividend Risk Premium. *Journal of Financial and Quantitative Analysis*, *56*(8), 2843–2869.

Bansal, R., Kiku, D., Shaliastovich, I., & Yaron, A. (2014). Volatility, the Macroeconomy, and Asset Prices. *The Journal of Finance*, *69*(6), 2471–2511.

Bansal, R., Kiku, D., & Yaron, A. (2012). An Empirical Evaluation of the Long-Run Risks Model for Asset Prices. *Critical Finance Review*, *1*(1), 183–221.

Bansal, R., Kiku, D., & Yaron, A. (2016). Risks for the Long Run: Estimation with Time Aggregation. *Journal of Monetary Economics*, *82*, 52–69.

Bansal, R., Miller, S., Song, D., & Yaron, A. (2021). The Term Structure of Equity Risk Premia. *Journal of Financial Economics*, *142*(3), 1209–1228.

Bansal, R. & Yaron, A. (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance*, *59*(4), 1481–1509.

Branger, N., Trede, M., & Wilfling, B. (2024). Extracting Stock-Market Bubbles from Dividend Futures.

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

Casta, M. (2022). Deriving Equity Risk Premium Using Dividend Futures. *The North American Journal of Economics and Finance*, *60*, 101667.

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

Diba, B. T. & Grossman, H. I. (1988). Explosive Rational Bubbles in Stock Prices? *The American Economic Review*, *78*(3), 521–530.

Evans, G. W. (1991). Pitfalls in Testing for Explosive Bubbles in Asset Prices. *The American Economic Review*, *81*(4), 922–930.

Fama, E. F. & French, K. R. (1988). Dividend Yields and Expected Stock Returns. *Journal of Financial Economics*, *22*(1), 3–25.

Froot, K. A. & Obstfeld, M. (1991). Intrinsic Bubbles: The Case of Stock Prices. *American Economic Review*, *81*(5), 1189–1214.

Giglio, S., Kelly, B., & Kozak, S. (2024). Equity Term Structures Without Dividend Strips Data. *The Journal of Finance*, *79*(6), 4143–4196.

Golez, B. & Jackwerth, J. (2024). Holding Period Effects in Dividend Strip Returns. *The Review of Financial Studies*, *00*(0).

Gonçalves, A. S. (2021). Reinvestment Risk and the Equity Term Structure. *The Journal of Finance*, *76*(5), 2153–2197.

Gormsen, N. J. (2021). Time variation of the equity term structure. *The Journal of Finance*, *76*(4), 1959–1999.

Gormsen, N. J. & Koijen, R. S. J. (2020). Coronavirus: impact on stock prices and growth expectations. *The Review of Asset Pricing Studies*, *10*(4), 574–597.

Goyal, A., Welch, I., & Zafirov, A. (2024). A Comprehensive 2022 Look at the Empirical Performance of Equity Premium Prediction. *The Review of Financial Studies*, *37*(11), 3490–3557.

Knox, B. & Vissing-Jorgensen, A. (2022). A Stock Return Decomposition Using Observables. *Finance and Economics Discussion Series*, *2022*(010), 1–58.

Kragt, J., Jong, F. d., & Driessen, J. (2020). The Dividend Term Structure. *Journal of Financial and Quantitative Analysis*, *55*(3), 829–867.

Shiller, R. (1981). Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends? *The American Economic Review*, (71(3):), 421–436.

Tirole, J. (1985). Asset Bubbles and Overlapping Generations. *Econometrica*, *53*(6), 1499–1528.

van Binsbergen, J., Brandt, M., & Koijen, R. (2012). On the Timing and Pricing of Dividends. *American Economic Review*, *102*(4), 1596–1618.

van Binsbergen, J., Hueskes, W., Koijen, R., & Vrugt, E. (2013). Equity Yields. *Journal of Financial Economics*, *110*(3), 503–519.

van Binsbergen, J. H. & Koijen, R. S. (2017). The Term Structure of Returns: Facts and Theory. *Journal of Financial Economics*, *124*(1), 1–21.