

A Valuation-Based Interpretation of Stock Market Volatility and Return Predictability

Roland Clère *

Abstract

Building on Pástor, Sinha, and Swaminathan (2008), who establish the implied cost of capital as an informative proxy for latent expected returns, and Li, Ng, and Swaminathan (2013), who document its predictive content for subsequent market returns, this paper asks a broader question: whether the market-implied required return forms part of a valuation mechanism linking fundamentals, market volatility, and subsequent returns.

Using the S&P 500 since 2014, a Monte Carlo analysis first reconstructs annual return volatility from empirically calibrated capital accumulation, profitability, expected growth, and ex-ante required-return inputs. In an approach nearly the reverse of Pástor et al. (2008), volatility is not imposed on the data-generating process but emerges from their joint dynamics. Second, capital accumulation, dividend yield, and the return effect of recurrent volatility imply a required-return level remarkably close to the average market IRR estimated from expected cash flows; using realized volatility instead produces a very similar result, while longer-run evidence points to an additional conditioning role for inflation. Third, the annual returns through which this volatility is realized display an economically structured multidimensional Valuation–Cycle configuration in sample. This structure retains significant predictive content in an expanding-window one-year out-of-sample exercise, with the combined Valuation–Cycle specification outperforming the historical mean.

Finally, the results support a unified valuation-based interpretation linking market volatility, required-return levels, and subsequent market dynamics through a locally persistent and adaptively revised valuation convention.

JEL codes: G12, G17, G14, D84, D81.

* University of Lille, LUMEN Research Center.

1. Introduction

The magnitude of stock-market volatility calls for an explanation of its nature and origins. Shiller's (1981) excess-volatility argument suggests that revisions in expected future cash flows alone cannot account for observed fluctuations, while evidence from French (1980), French and Roll (1986), and Roll (1984) indicates that market activity itself may contribute to price variability. Explaining volatility therefore requires considering not only changes in expected cash flows but also changes in the way investors assess and price future risks.

This naturally leads to the distinction between measurable risk and fundamental uncertainty. Knight (1921) distinguished situations in which probabilities can meaningfully be assigned from those in which the relevant probabilities are themselves unknown, while Keynes (1921, 1936) stressed the limited basis on which expectations about an uncertain future can be formed. Yet securities must be valued before future cash flows are known. Required returns may therefore reflect both measurable variability and judgments about future conditions that cannot be reduced to stable probability distributions, so that risk and uncertainty interact within the valuation process rather than being priced separately.

The implied-cost-of-capital literature provides a natural starting point. Pástor, Sinha, and Swaminathan (2008) show that an implied cost of capital inferred from prices and expected cash flows can proxy the latent conditional expected return and use it to study the intertemporal relation between expected returns and market variance. Li, Ng, and Swaminathan (2013) subsequently establish its predictive content for future market returns beyond conventional valuation and business-cycle variables. We take these results as given but reverse and broaden the inquiry: rather than using the implied required return primarily to recover an expected-return–variance relation or to predict subsequent returns, we ask whether it belongs to a broader valuation mechanism from which market volatility can emerge, whose average required-return level can be economically rationalized, and whose configuration contains information about subsequent returns.

In the present study, we place this question within an intrinsic valuation framework applied to the S&P 500. Price changes are decomposed into book-equity accumulation and changes in the price-to-book multiple, the latter reflecting expected profitability, expected growth, and the market-implied ex-ante required return, proxied empirically by an expected internal rate of return (IRR). Throughout the paper, IRR is used as a generic term for an implied required return obtained by solving a valuation relation from observed prices and expected future

fundamentals, whether the valuation is expressed in discounted-cash-flow or residual-income form. This terminology therefore encompasses measures commonly referred to as implied costs of capital (ICC) in the literature*. The relevant fundamentals are therefore understood as the determinants of valuation: not only currently observed economic quantities, but also expectations about their future evolution and the rate at which those expected outcomes are capitalized. The required return is thus treated not merely as an empirical measure of expected returns, but as one dimension of a broader valuation configuration.

The first empirical exercise uses this framework to reconstruct annual S&P 500 returns through Monte Carlo simulation. Empirical distributions and dependence structures are assigned to book-equity growth, normalized profitability, expected growth, and the expected IRR, from which successive valuation multiples and returns are derived. The simulation therefore asks a question nearly opposite to that in Pástor et al. (2008). In their design, the expected return and annual market volatility are calibrated upstream in order to evaluate whether the IRR recovers the relation between them. Here, neither realized return volatility nor its target level is imposed. Volatility is instead an output of the joint evolution of fundamental and required-return inputs within a prevailing valuation regime. The resulting volatility is close to that realized over 2014–2022, even though realized market volatility is not used as a calibration target.

This result raises a second question that differs from the primary use of the IRR in both Pástor et al. (2008) and Li et al. (2013): what determines the economic level of the required return? Standard asset-pricing representations generally express the required return on equities as the sum of a risk-free rate and an equity risk premium, although the origin and magnitude of that premium remain difficult to explain, as illustrated by the equity premium puzzle (Mehra and Prescott, 1985). The framework developed here examines whether capital accumulation, distributed income, and recurrent market volatility imply an economically plausible period-average required return close to the market IRR inferred from expected cash flows.

Having examined the magnitude of market volatility and the economic level of the required return, the analysis turns to the annual return dynamics through which that volatility is realized. Rather than asking again whether the implied required return forecasts subsequent returns, as in Li et al. (2013), the third exercise asks whether it is sufficient to characterize the valuation state from which those returns emerge. The in-sample evidence points to a multidimensional structure in which valuation conditions, price dynamics, profitability, and capital accumulation

* Appendix B provides a glossary of the main concepts and notation used throughout the paper.

carry distinct information. This structure also retains significant predictive content in an expanding-window one-year OOS exercise, with the combined Valuation–Cycle specification outperforming the historical mean. The magnitude of this OOS performance is comparable to the one-month gains obtained in Clère (2026a), although that study did not use a directly estimated market IRR, and substantially exceeds the one-month aggregate-IRR gains reported by Li et al. (2013). These comparisons remain indicative rather than direct, given differences in samples and forecasting designs.

The results therefore support a broader interpretation in which valuation conventions provide temporary coordination under fundamental uncertainty while remaining subject to revision. Market volatility can arise both from movements within a prevailing valuation configuration and from shifts in that configuration, while subsequent returns remain partly organized by information already embedded in fundamentals, market prices, and required returns. The coexistence of structure, predictive persistence, and recurrent revision motivates the valuation-convention interpretation developed in Section 5.

Section 2 positions the research question in the relevant literature and develops the analytical framework. Section 3 presents the data and empirical design. Section 4 reports the three empirical exercises. Section 5 develops their implications for fundamental uncertainty, valuation conventions, and informational efficiency. Section 6 concludes.

2. Toward a Unified Valuation Framework for Market Volatility, Required Returns, and Subsequent Market Dynamics

The excess-volatility literature initially formalized the issue through variance-bound tests. LeRoy and Porter (1981) and Shiller (1981) showed that observed stock-price variability appeared difficult to reconcile with the variability of subsequently realized dividends under standard present-value restrictions. Campbell and Shiller (1988) subsequently linked this problem to return predictability by emphasizing that valuation ratios reflect both expected cash flows and expected returns. Once discount rates are allowed to vary, price fluctuations need not mirror fluctuations in realized cash flows; modern asset-pricing interpretations accordingly place time variation in discount rates at the center of movements in valuation ratios (Cochrane, 2011). The analysis below builds on this broader present-value perspective by treating expected cash flows and the required return jointly within the valuation state.

2.1. Risk, Uncertainty, and the Required Return

The distinction between measurable risk and fundamental uncertainty developed by Knight (1921) and Keynes (1921, 1936) has a direct implication for valuation: investors must form required returns even when future outcomes cannot be represented by a stable probability distribution. The required return may therefore reflect not only compensation for measurable variability, but also judgments formed under incomplete knowledge about future economic conditions.

When uncertainty can be reduced to measurable and hedgeable risk, financial theory provides powerful pricing tools. Black and Scholes (1973) show how dynamic replication permits contingent claims to be valued without introducing an asset-specific risk premium into the pricing equation, while Merton (1974) extends this logic to corporate securities by treating equity as a call option on the stochastic value of the firm's assets. This risk-neutral logic does not remove the need to value the risky underlying asset itself; it solves the pricing of claims written on that asset once its value and stochastic dynamics are specified.

This leaves the underlying valuation problem unresolved. In practice, professional equity valuation continues to rely extensively on discounted cash-flow methods, in which expected future cash flows are discounted at a rate intended to compensate investors for committing capital to an uncertain future (Bancel and Mittoo, 2014). Conventional models such as the CAPM allocate a market-wide equity risk premium across securities according to their exposure to systematic risk (Sharpe, 1964; Lintner, 1965), but do not determine the level of that aggregate premium itself.

This difficulty is illustrated by the equity premium puzzle. Mehra and Prescott (1985) showed that the historical excess return on equities is difficult to reconcile with standard consumption-based preferences without implausibly high risk aversion. The puzzle therefore bears directly on the issue considered here: explaining the level of the aggregate required return, rather than merely allocating a given market premium across individual assets.

The required return should also be distinguished from its empirical estimate. Implied-cost-of-capital approaches infer a discount rate by reversing the usual valuation problem: observed prices and expected future cash flows determine the internal rate of return consistent with them. Such an implied rate is therefore a model-based estimate rather than a direct observation of the underlying required return. Whether its economic level can be independently grounded is a separate question examined below.

More broadly, the market required return is a natural point of interaction between measurable risk and fundamental uncertainty. Changes in the required return can generate substantial movements in valuation multiples even when expected cash flows remain broadly unchanged, while a given change in the discount rate may reflect revisions in measurable risk, in confidence about future economic states, or in both. The empirical problem is therefore not only to measure variation in the required return, but also to understand the economic content of its level and its role in broader market dynamics. The next section reviews the evidence on implied required returns and uses it to formulate the research question pursued here.

2.2. Implied Required Returns: Existing Evidence and the Research Question

Pástor, Sinha, and Swaminathan (2008) provide the theoretical basis for interpreting an implied cost of capital as information about a latent expected return. In their framework, the true conditional expected return is unobservable, whereas the IRR is obtained by equating the observed price with expected future cash flows. Under their assumptions, the two move closely together, making the IRR potentially more informative about expected-return variation than subsequently realized returns, which contain large unexpected-return innovations. Their empirical use of the IRR consequently focuses on the intertemporal relation between expected returns and market variance.

The direction of their simulation exercise is important for the present paper. Pástor et al. specify the expected-return process and its relation to variance upstream, choosing parameters that imply an unconditional expected return of 15% and annual volatility of 18%, and then examine whether the IRR recovers the underlying relation. Volatility and expected return are therefore features of the simulated economic environment rather than quantities the IRR is required to explain. The analysis below turns this logic around. Fundamental and required-return inputs are introduced into an intrinsic valuation mechanism and market volatility is allowed to emerge from their joint dynamics rather than being calibrated as a target. The level of the required return is then considered separately as an economic object in its own right. The distinction is thus not primarily between alternative IRR formulas, but between two directions of inquiry: using the IRR to recover variation in an already specified expected-return–variance environment, or asking whether a valuation mechanism containing an implied required return can account for the observed level of volatility and provide an economic reference for the required return itself.

Li, Ng, and Swaminathan (2013) establish a second result that further changes the relevant question. Building on the IRR interpretation of Pástor et al. (2008), they show in sample that

aggregate implied required returns contain information about subsequent market returns over horizons from one month to several years and retain predictive content relative to conventional valuation ratios and business-cycle variables. They also obtain positive recursive out-of-sample performance at the one-month horizon. Hoang et al. (2021) extend this evidence by combining several implied-cost-of-capital measures through partial least squares and showing that the resulting common factor improves both in-sample and out-of-sample return predictability. These exercises therefore provide direct empirical support for the proposition that variation in aggregate implied required returns contains information about future realized returns. Importantly, however, the additional valuation and macro-financial variables are used principally as competing predictors, controls, or inputs into a common predictive factor: the central question remains whether implied required returns contain incremental predictive information beyond them.

Complementary evidence in Clère (2026a) points toward a broader interpretation. Using a valuation-based proxy for the market-implied required return rather than a directly estimated implied return, the analysis finds that aggregate U.S. stock returns are better characterized by a multidimensional Valuation–Cycle–Sentiment configuration than by a single predictive variable. Valuation levels are more prominent at longer horizons, while cyclical and dynamic information gains importance as the horizon shortens; the resulting configuration also retains substantial OOS predictive content at the one-month horizon. This evidence suggests that the information associated with required returns may depend on the broader configuration in which they are embedded and on the horizon over which market adjustment is observed.

The present analysis extends this broader evidence by asking whether a directly estimated market IRR is sufficient to characterize the valuation state associated with subsequent one-year returns. Valuation and Cycle variables are not treated as controls external to the required return, but as components of the same configuration. The empirical issue is therefore how predictive information is distributed within that configuration. Principal-component and orthogonalization tests examine whether future-return information collapses into the market IRR or a dominant common factor, or instead remains distributed across economically distinct dimensions. The expanding-window OOS exercise then asks whether this multidimensional structure retains predictive content at the substantially longer one-year horizon.

The preceding literature therefore leaves two complementary questions open. The first concerns levels: can the observed magnitude of market volatility and the average required return be economically rationalized within the same valuation structure rather than treated primarily as

imposed quantities or empirical proxies? The second concerns dynamics: is the predictive information associated with the implied required return a property of that scalar measure alone, or of the broader valuation configuration in which it is embedded? These questions motivate the three empirical exercises developed below: reconstructing the magnitude of annual market volatility, providing an economic reference for the average required-return level, and identifying the structure and predictive content of the successive annual returns through which that volatility is realized. The valuation decomposition in the next section provides a common framework for examining these three dimensions.

2.3. A Valuation-Based Decomposition of Market Returns

To examine these questions within a common framework, market value is decomposed into changes in the productive and accounting base of listed firms and changes in the valuation attached to that base. The starting point is the standard present-value principle according to which equity value reflects expected future cash flows discounted at the return required by shareholders (Williams, 1938; Gordon and Shapiro, 1956). The same valuation logic can be expressed in terms of book equity and the rents generated on that equity, an approach closely related to residual-income valuation (Ohlson, 1995) and equivalent to conventional discounted-cash-flow representations under consistent assumptions (Fernandez, 2003).

Under a steady-state representation in which normalized return on equity, expected growth, and the required return are summarized by ROE_t , g_t , and k_t , respectively, market value can be written as:

$$P_t = B_t + B_t \times \frac{ROE_t - k_t}{k_t - g_t} \quad (1)$$

where P_t denotes market value, B_t book equity, ROE_t normalized return on book equity, g_t the expected long-run growth rate of book equity, and k_t the required return on equity, with $k_t > g_t$. The first term represents the accounting value of the equity already accumulated, while the second capitalizes the economic rent generated when expected profitability differs from the required return. When $ROE_t > k_t$, the firm generates positive rents and market value exceeds book equity; when $ROE_t = k_t$, no such rent is capitalized and market value equals book equity.

Dividing **equation (1)** by book equity gives the corresponding price-to-book ratio:

$$P_t/B_t = 1 + \frac{ROE_t - k_t}{k_t - g_t} = \frac{ROE_t - g_t}{k_t - g_t} \quad (2)$$

This expression makes the valuation mechanism particularly transparent. For a given book-equity base, the market multiple increases with expected profitability and growth and decreases with the required return. Consequently, even relatively small revisions in k_t can generate substantial changes in market value when the spread between the required return and expected growth is narrow.

The same representation provides a useful decomposition of observed price changes. Since market value is identically equal to book equity multiplied by the price-to-book ratio,

$$P_t = B_t \times (P_t/B_t)$$

the logarithmic change in market price between $t - 1$ and t can be written as:

$$\ln\left(\frac{P_t}{P_{(t-1)}}\right) = \ln\left(\frac{B_t}{B_{(t-1)}}\right) + \ln\left(\frac{P_t/B_t}{P_{(t-1)}/B_{(t-1)}}\right). \quad (3)$$

Equation (3) separates two economically distinct sources of market-price performance. The first term measures book-equity accumulation. At the aggregate market level and over relatively short horizons, it can be interpreted as an accounting proxy for the expansion of the capital base supporting listed firms, provided that leverage and capital intensity do not change excessively. This interpretation becomes less restrictive over longer horizons, when financing structures, sector composition, and capital intensity may themselves change. The second term measures the change in valuation attached to that capital base. From **equation (2)**, it reflects revisions in normalized profitability, expected growth, and the required return. Unlike book-equity accumulation, this component is therefore intrinsically forward-looking: it changes when investors revise the parameters through which future economic outcomes are translated into current prices.

This distinction is central to the analysis that follows. If the valuation multiple were stationary on average, the central tendency of market-price growth should be anchored, over sufficiently long periods, by the accumulation of book equity. Departures from that path may arise from fluctuations in ROE_t , g_t , and k_t . In particular, a persistent decline in the required return raises the price-to-book ratio and produces realized price growth above that generated by capital accumulation alone; an increase has the opposite effect.

Equation (3) consequently provides a direct bridge between fundamentals and volatility. Market-price variability need not originate exclusively in fluctuations in the underlying capital base. It can also arise from revisions in the valuation assigned to that base. Some of these revisions may be associated with observable changes in profitability or expected growth, while others may operate through the required return. Because this component is latent, its variation is particularly relevant to the interaction between measurable risk and fundamental uncertainty developed above. The decomposition does not imply that every movement in the valuation multiple is fundamental; it provides an accounting and valuation framework for examining how changes in fundamentals and the required return are reflected in market prices.

The analysis that follows uses this decomposition to examine whether capital accumulation, distributed income, and recurrent return volatility can provide an economically grounded reference for the required return, and whether the resulting market dynamics exhibit an identifiable structure at the one-year horizon.

2.4. Capital Accumulation, Volatility, and the Required-Return Reference

Equation (3) decomposes aggregate market-price growth into book-equity accumulation and changes in the valuation multiple. Accordingly, the volatility of logarithmic market returns reflects the joint variability of these two components and their covariance:

$$\sigma_t^2 = Var_t(g_B) + Var_t[\Delta \ln(P/B)] + 2 Cov_t[g_B, \Delta \ln(P/B)]. \quad (4)$$

The σ_t introduced here therefore denotes the conditional volatility of the total logarithmic market-price return, rather than the volatility of changes in the valuation multiple alone. If the latter has no permanent trend under a sufficiently persistent valuation configuration, **equation (3)** suggests that the central tendency of market-price growth should ultimately be anchored by expected book-equity accumulation. To connect this economic decomposition with the return associated with the expected future market price, consider a conventional diffusion representation in which the logarithmic market-price change over a unit interval has conditional mean μ_t and variance σ_t^2 . If prices are conditionally lognormally distributed, the continuously compounded price return associated with the expected future market price can be written as:

$$\ln\left(\frac{E_t(P_{t+1})}{P_t}\right) = \mu_t + \frac{1}{2}\sigma_t^2. \quad (5)$$

The left-hand side differs from the expected logarithmic price return itself, $E_t[\ln(P_{t+1}/P_t)]$, which is equal to μ_t . The variance term therefore reflects the familiar lognormal relation between the expected future price and the mean logarithmic price change. Under the persistence condition discussed above, **equation (3)** suggests that the conditional drift μ_t may be related to expected book-equity accumulation.

If, from this point onward, the superscript $*$ denotes an ex-ante quantity associated with the prevailing valuation configuration rather than a trailing or subsequently realized quantity, one may write:

$$\mu_t \approx g_{B,t}^*, \quad (6)$$

where $g_{B,t}^*$ denotes expected book-equity accumulation. This relation is not an accounting identity. It reflects the proposition that market value cannot persistently grow faster than the underlying equity base unless profitability, growth expectations, or required returns themselves undergo sustained changes.

Equity investors also receive distributed income. Let DY_t^* denote the expected dividend yield associated with the prevailing valuation configuration, and define its logarithmic equivalent as

$$dy_t^{L*} \equiv \ln(1 + DY_t^*).$$

Let σ_t^* denote anticipated recurrent market volatility within the prevailing valuation configuration. The central conjecture considered here is that expected capital accumulation, distributed income, and recurrent volatility jointly provide an economically meaningful reference for the latent return required by equity investors. We refer to this quantity as the required-return reference (RRR). Let k^{L*} denote its continuously compounded form. Then:

$$k_t^{L*} \approx g_{B,t}^* + dy_t^{L*} + \frac{1}{2}(\sigma_t^*)^2. \quad (7)$$

In discrete annual terms:

$$k_t^* \approx \exp \left[g_{B,t}^* + dy_t^{L*} + \frac{1}{2}(\sigma_t^*)^2 \right] - 1. \quad (8)$$

RRR and the market-implied IRR represent two distinct empirical routes toward the same latent ex-ante required return. The market IRR is inferred by solving the valuation relation from observed prices and expected cash flows. RRR proceeds in the opposite direction: it constructs an economically motivated reference from expectations of reinvested and distributed income

and from recurrent market risk. The two quantities therefore need not coincide by construction; their empirical correspondence is itself an object of investigation.

The passage from **equation (6)** to **equation (7)** is economic rather than algebraic. **Equation (5)** relates the conditional logarithmic drift to the return associated with the expected future market price, while **equation (6)** proposes an economic anchor for that drift. **Equation (7)** adds expected distributed income and asks whether the resulting total shareholder return also provides information about the latent return investors require before committing capital. The resulting RRR, k_t^* , can therefore be compared with the estimated market IRR, k_t , as an alternative empirical representation of the same latent ex-ante required return.

The variance term in **equation (7)** should not be interpreted as an independently estimated price of risk. **Equation (4)** defines σ_t^2 as the conditional variance of the total logarithmic market-price return, while its appearance in **equation (7)** follows from the lognormal relation in **equation (5)** rather than from a separate risk-pricing assumption. The economic proposition is broader: expected accumulation anchors the central tendency of market-price growth, expected distributed income contributes directly to shareholder return, and recurrent market variability affects the relation between that central tendency and expected future market value. Together, these quantities may therefore help anchor the return investors require.

This interpretation differs from starting with the conventional decomposition $k_t = r_{f,t} + ERP_t$. It does not deny that investors compare equities with risk-free assets, nor that changes in the risk-free rate affect equity valuation. Rather, it asks whether the aggregate required return may possess an economically grounded reference level of its own, with the equity premium then emerging partly as the difference between that required return and the prevailing risk-free rate. In this representation, $ERP_t = k_t - r_{f,t}$ remains an accounting relation, but the level of k_t need not be mechanically determined by adding an independently given equity premium to the risk-free rate.

RRR is not an equilibrium condition but an economically motivated empirical reference for the same latent required return that the market IRR seeks to recover from prices and expected cash flows. Their correspondence need not be exact, but a close relation would indicate that the return revealed by market valuation is consistent with the income and risk characteristics of the prevailing valuation configuration. The resulting equity premium may then retain its conventional cross-sectional role in differentiating required returns across securities according

to systematic-risk exposure. Inflation may provide an additional constraint on the nominal level of the required return.

2.5. Empirical Implications

Taken together, Sections 2.1–2.4 suggest three empirical implications that can be confronted with the data. First, if the valuation mechanism described above captures an economically meaningful source of market variability, empirically calibrated movements in capital accumulation, profitability, expected growth, and the market IRR should be capable of generating a substantial share of observed annual market volatility. Any remaining difference between the resulting conditional distribution and realized outcomes could then be examined in relation to revisions in the valuation configuration, particularly revisions in the expected IRR. The secular movements in the market-implied expected IRR documented by Clère (2026b) suggest that such revisions may be economically significant.

Second, if RRR, k_t^* , provides an economically motivated reference for the same latent required return that the estimated market IRR, k_t , seeks to recover, the two should be quantitatively comparable when the underlying valuation configuration is sufficiently stable. A close correspondence would be consistent with the proposition that the required return is related to expected capital accumulation, distributed income, and recurrent market volatility. Longer-run evidence can then be used to examine whether inflation or long-term government yields also condition its nominal level.

Third, reproducing the magnitude of annual volatility in distribution would not by itself establish that the successive market movements through which that volatility is realized are economically structured. If the prevailing valuation configuration exhibits sufficient local persistence, however, predetermined information describing its Valuation and macro-financial Cycle dimensions should remain systematically associated with subsequent one-year market returns. Conversely, if annual market movements are largely undifferentiated short-horizon noise around fundamentals, little such structure should remain detectable.

These three implications organize the empirical analysis. Section 4.1 examines how much annual volatility can be generated within an empirically calibrated valuation configuration and whether departures from that distribution coincide with revisions in the expected IRR. Section 4.2 compares RRR, as defined by **equation (8)** with the separately estimated market IRR and considers whether inflation or long-term government yields provide additional nominal constraints. Section 4.3 turns from volatility in distribution to its temporal realization and tests

whether subsequent one-year returns exhibit systematic structure in relation to predetermined Valuation and macro-financial Cycle information. This final test also connects with Clère (2026a), in which realized returns reflect both the valuation configuration prevailing at t and subsequent revisions to it.

3. Data and Empirical Design

The empirical analysis combines three complementary exercises that require different sample windows. The first calibrates a Monte Carlo reference distribution for annual market returns and volatility over the period for which a consistent estimate of the expected market IRR is available. The second examines whether capital accumulation and recurrent market volatility provide an economically meaningful reference for the latent ex-ante required return, with longer-run evidence used to assess the additional role of inflation and long-term government bond yields. The third uses a reduced-form specification to test whether variables describing the prevailing valuation configuration contain information about subsequent one-year S&P 500 returns and how predictive information is distributed across the economically distinct dimensions of the valuation configuration, including the expected IRR. The samples and sampling frequencies therefore differ by design rather than representing alternative estimations of the same specification.

3.1. Data and Timing Conventions

The analysis focuses on the U.S. equity market, represented by the S&P 500 index. The valuation framework requires market prices, earnings, book equity, normalized profitability, expected growth, and an expected market internal rate of return (IRR) serving as a proxy for the market-implied ex-ante required return. S&P 500 price and earnings data are drawn from the series compiled by Robert Shiller, aggregate book-equity data from Multiple.com, and expected nominal U.S. GDP growth from International Monetary Fund forecasts. The construction of the expected market IRR is described in Section 3.2.

Normalized profitability is measured by a ten-year moving average of annual aggregate S&P 500 return on equity, limiting the influence of transitory fluctuations in reported earnings. Thus, the normalized value used in 2014 incorporates observations from 2005 through 2014. Expected long-run growth is anchored in nominal U.S. GDP growth forecasts and enters the valuation multiple as an external growth assumption.

The Monte Carlo exercise is calibrated primarily over 2014–2022, the period covered by the consistent aggregate S&P 500 IRR series available for this study. The sample is used retrospectively to estimate the empirical distributions, dispersion, contemporaneous dependence, and persistence of valuation inputs that are themselves observable or estimable ex ante. The simulation therefore preserves not only the marginal variability of these inputs, but also their joint dependence and the serial persistence of the principal state variables, so that return volatility emerges from their combined dynamic structure within the valuation mechanism. Realized S&P 500 return volatility is not imposed as a calibration target and enters only as an ex-post reference against which the simulated distribution is assessed. The exercise should accordingly be interpreted as a conditional ex-ante reconstruction calibrated on the recent historical regime, rather than as a real-time forecast of the realized return path. The expected IRR, being inferred partly from contemporaneous market prices, is treated as a valuation input rather than as a price-independent fundamental. Earlier observations are used where longer histories are required, notably for normalized profitability and alternative assumptions about book-equity accumulation.

Sampling frequency differs across the empirical exercises by design. The Monte Carlo simulation is annual, consistent with the natural frequency of several fundamental inputs and with evidence that aggregate stock returns exhibit different forms of serial dependence across horizons. Positive dependence at short horizons and mean reversion at longer horizons make the one-year interval a useful intermediate horizon (Poterba and Summers, 1988; Fama and French, 1988), while annual sampling also limits the influence of transitory short-run price movements. The return specification addresses a different question: whether the valuation configuration observed at a given date contains information about market performance over the following year. Its explanatory variables are therefore observed quarterly while the return horizon remains one year.

The one-year specification uses quarter-end observation dates and yields 37 usable observations over the period covered by the IRR series. Explanatory variables precede the corresponding return period, with accounting quantities aligned to the appropriate information lag. Some fundamental and macroeconomic inputs nevertheless rely on subsequently revised historical series, so the exercise should not be interpreted as a strict real-time forecast; Section 3.4 discusses this issue in detail. Quarterly sampling implies that successive one-year returns overlap over three quarters. Inference therefore relies on Newey–West heteroskedasticity- and autocorrelation-consistent standard errors with three lags.

The longer-run analysis of the required return uses a separate historical extension. Published estimates of the U.S. market required return are combined with the more recent series and aggregated by decade from the 1930s through the 2010s. This evidence is used only to examine whether the recent relation is compatible with broader historical patterns in inflation and long-term government yields; it enters neither the Monte Carlo calibration nor the one-year return specification.

These timing conventions reflect the distinct purposes of the three exercises: the Monte Carlo analysis examines the level of annual market volatility that can be generated by an ex-ante valuation configuration; the required-return analysis examines the economic determinants of the level of the market required return; and the quarterly specification turns to dynamics, testing whether predetermined valuation information is related to the subsequent one-year market returns through which realized volatility unfolds.

3.2. Market-Implied Required Return

The required return on equity is not directly observable. Realized returns provide an ex-post outcome, but they do not directly measure the return investors required before the investment period. Between the valuation date and the realization of returns, prices may change because the ex-ante required return itself is revised as perceptions of risk and uncertainty evolve, but also because expectations of future cash flows, growth, or other elements of the valuation state are updated. Realized returns therefore combine the return initially embedded in market valuation with subsequent revisions in that configuration. For the present purpose, a forward-looking measure constructed at the valuation date is consequently more closely aligned with the latent ex-ante required return than historical measures or subsequently realized returns.

Implied-cost-of-capital approaches follow this logic by reversing the valuation problem: observed market prices and expected future cash flows are used to infer the discount rate consistent with them (Claus and Thomas, 2001; Gebhardt et al., 2001; Easton, 2004). The resulting internal rate of return (IRR) equates the observed price with the present value of a specified expected cash-flow stream and can therefore be interpreted as an ex-ante estimate of the return expected by investors and, under valuation consistency, as a proxy for the latent required return embedded in market prices. Pástor et al. (2008) provide a direct theoretical basis for this interpretation by showing that an implied required-return measure can closely track the latent conditional expected return. The IRR nevertheless remains a model-based estimate whose quality depends critically on the expected cash flows used as inputs. Both the academic implied-

cost-of-capital literature and professional valuation practice commonly use sell-side analyst forecasts as observable proxies for prevailing expectations about future corporate performance.

The present analysis uses the expected U.S. market IRR produced by the Fairness Finance model, available consistently from December 2014 and introduced in Clère (2026b). Fairness Finance is a bottom-up, stock-flow-consistent valuation framework applied to the constituents of the S&P 500. Companies are valued individually using contemporaneous sell-side forecasts and internally consistent projections of operating performance, investment, financing, and cash generation. Beyond the explicit forecast horizon, growth assumptions are progressively converged toward expected nominal U.S. GDP growth from IMF forecasts, while operating margins revert toward sectoral long-run averages. Company-level IRRs are then aggregated using market-capitalization weights. This construction allows differences across firms and sectors to enter the market-level estimate rather than relying on a representative firm.

Because sell-side earnings and cash-flow forecasts exhibit persistent optimism, the Fairness Finance methodology adjusts the raw implied IRR downward using an expanding-window estimate of the average historical forecast bias, measured from differences between forecasts and subsequent realizations. Over the period considered, this adjustment amounts to approximately two percentage points. The series used here is therefore the optimism-adjusted expected market IRR. Throughout the empirical analysis, k_t denotes this adjusted estimate, corresponding to $E_t(k | S)$, the expected required return conditional on survival, unless otherwise specified.

The expected IRR should therefore be interpreted as a forward-looking, model-based proxy for the latent required return, subject to measurement error rather than as an independently observed or exogenous causal variable. It enters the empirical analysis through its distribution and persistence in the Monte Carlo exercise, its average level in the comparison with the required-return reference (RRR) developed in Section 2.4, and its role, alongside other predetermined variables, in the structure of subsequent one-year S&P 500 returns.

3.3. Monte Carlo Design

The Monte Carlo exercise estimates the distribution of annual S&P 500 price changes generated by the valuation framework developed in Section 2.3 when its main inputs evolve according to empirically calibrated distributions. As discussed in Section 3.1, it is a conditional ex-ante reconstruction calibrated retrospectively on the recent historical regime rather than a real-time forecast of the realized return path.

For each simulated annual transition, two successive valuation states, $t - 1$ and t , are generated. At each date, the price-to-book multiple is reconstructed from **equation (2)** using normalized profitability, expected long-run growth, and the expected market IRR k_t . Book equity evolves between the two dates according to the selected distribution of book-equity accumulation. Expected long-run growth, denoted g' , is anchored to 75% of the IMF five-year forecast of nominal U.S. GDP growth. This empirical convention, also used in the Fairness Finance valuation framework, was identified by comparing the growth rate implied by observed S&P 500 price-to-book ratios with available long-run growth proxies and therefore predates the Monte Carlo calibration. The annual logarithmic price return is then obtained from **equation (3)** as the sum of book-equity accumulation and the change between the two simulated valuation multiples. Dividend income does not enter the simulation itself and is introduced separately in Section 4.2.1, where expected capital accumulation, distributed income, and the recurrent-volatility adjustment are combined to construct the required-return reference (RRR).

The use of the market-implied IRR as a simulation input requires an explicit qualification. Following Pástor, Sinha, and Swaminathan (2008), variation in an implied cost of capital is treated as informative about variation in the latent conditional expected return. This provides a basis for using the observed distribution and persistence of the IRR as features of the valuation state, but it does not make the exercise independent of market prices, since the IRR is itself inferred from observed prices and expected cash flows. More importantly, neither the result in Pástor et al. nor the present calibration implies that introducing this IRR process into the valuation mechanism must generate the observed level of annual market volatility. The volatility produced by the simulation is therefore an output of the joint valuation dynamics, not a calibration target or an accounting identity. Its proximity to realized volatility is assessed only ex post in Section 4.1.

The simulation preserves both contemporaneous dependence and intertemporal persistence among the valuation inputs. The contemporaneous correlation structure among book-equity accumulation, normalized profitability, expected growth, and the expected IRR is calibrated from the available sample. Persistence is introduced separately through the observed first-order correlations of normalized profitability, expected growth, and the expected IRR between $t - 1$ and t . It is therefore not imposed directly on the price-to-book ratio: successive multiples inherit persistence from their underlying valuation inputs. This structure allows movements in profitability, growth expectations, accumulation, and the required return to offset or reinforce one another rather than treating them as independent draws. The full correlation matrix and

persistence parameters are reported in **Appendix A**. The extent to which this input-level dependence reproduces the observed persistence of valuation multiples is examined ex post in Section 4.1.

Because the calibration period is short, particular attention is given to the distribution of book-equity accumulation. Three configurations are considered. The “optimistic-prudent” configuration combines the 5.0% average annual accumulation rate observed over 2014–2019 with the broader 7.3% standard deviation observed over 2004–2022. The “pessimistic-myopic” configuration instead extrapolates the recent 2014–2022 regime, combining a 3.7% mean with a 2.9% standard deviation. Finally, the “pessimistic-prudent” configuration retains the lower 3.7% mean while restoring the broader 7.3% dispersion. The three cases therefore separate assumptions about the central rate of capital accumulation from assumptions about the variability surrounding it.

For each configuration, 20,000 joint draws are generated. The principal outputs are the annual logarithmic price return, its volatility, and the resulting price-to-book multiple; simple annual returns are obtained from logarithmic returns when required. The scenarios are not intended as competing structural models, but as alternative reference configurations for assessing how much return and volatility can arise within a given valuation configuration. In particular, the broader-dispersion cases provide a measure of recurrent volatility that does not require an unanticipated shift in the underlying valuation parameters.

Realized outcomes can therefore be compared with the simulated distributions to assess whether departures are consistent with revisions in the valuation configuration, particularly in the expected required return, rather than with autonomous market noise. Section 4.1 examines this possibility by comparing simulated and realized S&P 500 outcomes with the observed evolution of the expected IRR. Section 4.2.1 then combines recurrent volatility with observed book-equity accumulation and distributed income to assess the economic level of the expected required return. Because the simulated volatility partly reflects variation in the market-implied IRR itself, that comparison is interpreted as a consistency exercise rather than as independent identification; recomputing the required-return reference with realized rather than simulated volatility provides an additional check by removing this specific simulation channel. The following subsection turns to the quarterly specification relating the prevailing valuation configuration to subsequent one-year market performance.

3.4. Ex-Ante Valuation Information and Subsequent One-Year Returns

The third empirical exercise examines directly whether the valuation configuration observed at the investment date contains information about subsequent one-year S&P 500 performance. It complements the preceding exercises in a different direction. Given the predictive evidence already established by Li et al. (2013), the purpose is not to test again whether the implied required return forecasts market returns in isolation, but whether predictive information is distributed across the broader valuation configuration. Whereas the Monte Carlo analysis asks how much market variability can be generated within a given configuration of fundamentals and the required return, the present specification asks whether these predetermined conditions are systematically related to the return subsequently realized by the market.

The dependent variable is the logarithmic change in the S&P 500 price over the following four quarters. With t now indexing quarter-end observations, the one-year forward return is:

$$R_{t,t+4} \equiv \ln \left(\frac{P_{t+4}}{P_t} \right)$$

The economic horizon remains one year; only the sampling frequency changes. The valuation configuration is observed quarterly, from end-December 2014 through end-December 2023, yielding 37 observations over the period for which a consistent aggregate S&P 500 expected-IRR series is available. This quarterly design increases the number of usable observations and the information available for estimation without shortening the return horizon.

The predictor set is deliberately limited and follows the economic structure developed in the preceding sections rather than an unrestricted search over candidate variables. The Valuation block contains the expected IRR, the central latent parameter examined throughout the paper, together with P/E as a conventional market-valuation measure. The Cycle block captures three economically motivated dimensions of adjustment around that valuation structure: changes in recent price growth, the divergence between price growth and underlying profitability, and changes in book-equity accumulation relative to the broader economy. This emphasis on changes and deviations rather than levels alone is consistent with evidence that recent trends and departures from past levels can contain predictive information not captured by contemporaneous predictor levels (Fama and French, 1988; Poterba and Summers, 1988; Han et al., 2026). The purpose is not to identify a structural return-generating equation or an exhaustive forecasting model, but to test whether these economically distinct dimensions jointly contain information about subsequent market performance.

The first variable is k_t , the optimism-adjusted expected market IRR defined in Section 3.2. It is used as a proxy for the ex-ante required return embedded in the prevailing market valuation, subject to the measurement qualifications discussed above. It should not be interpreted as a direct forecast of the subsequent realized return. For a given valuation at t , a higher k_t implies a higher ex-ante expected return, but the return ultimately realized over the following year also reflects subsequent revisions in profitability, growth, capital accumulation, and the required return itself.

The second variable is the trailing market price-to-earnings ratio, PE_t , based on earnings reported over the preceding twelve months and aligned with the observation date. Unlike a cyclically adjusted multiple, it captures the current relation between market price and recently reported earnings. Its coefficient is not constrained a priori, since a high multiple may reflect either a demanding entry valuation or favorable expectations regarding future profitability and growth.

Two variables describe recent market-price dynamics. Let $g_{P,t}$ denote logarithmic price growth over the twelve months ending at quarter t . The first measures the change in this trailing twelve-month growth rate relative to its value four quarters earlier:

$$\Delta g_{P,t} \equiv g_{P,t} - g_{P,t-4}.$$

Although observed quarterly, $\Delta g_{P,t}$ therefore remains an annual acceleration measure, updated every three months. A positive value indicates that annual price growth has accelerated relative to the rate prevailing one year earlier, while a negative value indicates deceleration.

The second price-related variable compares trailing annual market-price growth with aggregate profitability:

$$GAP_t \equiv g_{P,t} - ROE_t$$

where ROE_t denotes aggregate return on equity, measured using trailing annual earnings relative to the previously observed book-equity base. GAP_t measures the extent to which recent market-price appreciation has moved ahead of, or fallen behind, the profitability of the underlying equity base. A large positive gap therefore characterizes a configuration in which market prices have advanced rapidly relative to profitability. The construction is consistent with the residual-income valuation logic in which market value is assessed relative to the profitability generated by the underlying book-equity base (Ohlson, 1995).

The final variable captures changes in book-equity accumulation relative to the broader economy, consistent with evidence that real and financial cycles are closely linked (Jawadi et al., 2022). Let $g_{B,t}$ denote the annual growth rate of aggregate book equity measured at quarter t , computed from quarterly observations using exponentially declining weights with a decay factor of 0.90. This recency weighting provides a smooth trailing annual measure while giving moderately greater weight to the most recent information. The decay factor is a fixed smoothing convention and is not calibrated on subsequent return performance. Let $g_{Y,t}$ denote the corresponding unweighted annual growth rate of U.S. real GDP per capita. Their difference, $d_t \equiv g_{B,t} - g_{Y,t}$, measures the rate at which the listed corporate equity base has accumulated relative to the broader economy over the preceding year. The variable entering the regression is the change in this differential relative to its value twelve months earlier:

$$\Delta g_{B,t} \equiv d_t - d_{t-4}.$$

Updated quarterly, $\Delta g_{B,t}$ remains a twelve-month acceleration or deceleration measure of book-equity accumulation relative to per-capita output, rather than the level of accumulation used in the Monte Carlo exercise and the required-return analysis of Section 2.4.

Variables used in the regression and their definitions are summarized in **Table 1**.

The empirical one-year return specification is:

$$R_{t,t+4} = \alpha + \beta_1 k_t + \beta_2 PE_t + \beta_3 \Delta g_{P,t} + \beta_4 GAP_t + \beta_5 \Delta g_{B,t} + \varepsilon_{t,t+4}. \quad (9)$$

No structural interpretation is imposed on the individual coefficients in isolation. Several variables can transmit more than one economic effect, and their partial coefficients also depend on the allocation of common information across correlated predictors. **Equation (9)** instead tests whether a parsimonious set of predetermined variables describing required return, valuation, profitability, market dynamics, and capital accumulation jointly provides explanatory content for subsequent one-year market performance.

Some accounting and macroeconomic inputs are measured using subsequently revised historical series. At each quarter-end observation date, however, three of the four quarters entering the trailing annual measures were already observed, while the final quarter would have had to be estimated in real time. The retrospectively constructed variables may therefore differ from the corresponding real-time estimates. This limitation therefore affects only the final component of some trailing annual measures but prevents the specification from being interpreted as a strict real-time forecasting exercise. All five explanatory variables are

standardized before estimation, while the dependent variable remains expressed as a one-year logarithmic return.

Table 1. Variables Used in the One-Year Return Specification

Symbol	Calculation / definition	Economic meaning
$R_{t,t+4}$	$R_{t,t+4} \equiv \ln(P_{t+4}/P_t)$	Subsequent one-year logarithmic S&P 500 price return measured from quarter-end t to $t+4$. Dependent variable.
k_t	Optimism-adjusted aggregate expected market IRR produced by Fairness Finance; bottom-up company IRRs aggregated using market-capitalization weights. Corresponds to $E_t(k S)$.	Proxy for the market-implied ex-ante required return embedded in the prevailing valuation configuration.
PE_t	Contemporaneous S&P 500 trailing price-to-earnings ratio at quarter-end t .	Current market valuation relative to reported earnings.
$\Delta g_{p,t}$	$\Delta g_{p,t} \equiv g_{p,t} - g_{p,t-4}$, where $g_{p,t}$ is the logarithmic S&P 500 price growth over the preceding twelve months.	Annual price-growth acceleration. Positive values indicate that trailing annual price growth has accelerated relative to one year earlier.
GAP_t	$GAP_t \equiv g_{p,t} - ROE_t$, where ROE_t is trailing annual aggregate profitability relative to the previously observed book-equity base.	Extent to which recent market-price growth has moved ahead of, or fallen behind, underlying profitability.
$\Delta g_{B,t}$	$\Delta g_{B,t} \equiv d_t - d_{t-4}$ with $d_t \equiv g_{B,t} - g_{Y,t}$. $g_{B,t}$ is annual aggregate book-equity growth using quarterly observations with exponentially declining weights; $g_{Y,t}$ is annual real U.S. GDP-per-capita growth.	Annual acceleration or deceleration in book-equity accumulation relative to per-capita output.

Notes: t denotes a quarter-end observation date. All explanatory variables are standardized before estimation; the dependent variable remains expressed as a one-year logarithmic return. Δg_p and Δg_B are annual measures updated quarterly. The construction of the expected market IRR is described in Section 3.2; detailed timing and measurement conventions are reported in Section 3.4.

Quarterly sampling of a one-year forward return necessarily generates overlapping observations: adjacent dependent-variable observations share three of their four quarterly return intervals. Conventional OLS standard errors would therefore be inappropriate even in the absence of other specification problems. Coefficients are estimated by OLS, but inference is based throughout on Newey–West heteroskedasticity- and autocorrelation-consistent standard errors with three lags. Section 4.3 subsequently evaluates the relation through residual and specification diagnostics, leave-one-out estimation, alternative price measurement, variance-decomposition and orthogonalization exercises, and an expanding-window one-year out-of-sample exercise.

4. Empirical Results

The empirical results are organized around the three complementary exercises introduced in Section 3. The first examines the return and volatility generated by an ex-ante valuation configuration, the second assesses whether the expected required return itself can be economically grounded in observable economic quantities, and the third explores whether the prevailing valuation configuration helps explain subsequent one-year market dynamics.

4.1. Monte Carlo Distribution and Valuation Revisions

The first exercise proceeds in two steps. It first assesses the return and volatility generated by the valuation framework, and then examines whether departures from realized outcomes are consistent with revisions in the expected IRR and the resulting change in the valuation multiple.

4.1.1. Simulated Market Volatility

Table 2 reports the main outcomes of the three Monte Carlo configurations described in Section 3.3 and compares them with realized S&P 500 outcomes over 2014–2022.

Table 2. Simulated and Realized S&P 500 Outcomes, 2014–2022

Configuration	Mean book -equity accumulation	Std. dev. of book-equity accumulation	Simulated log price return	Simulated volatility	Volatility gap vs realized	Main interpretation
Optimistic-prudent	5.0%	7.3%	6.2%	13.7%	-0.5 pp	Higher accumulation, broad instability
Pessimistic-myopic	3.7%	2.9%	4.5%	11.1%	-3.1 pp	Recent regime extrapolated
Pessimistic-prudent	3.7%	7.3%	4.7%	13.7%	-0.5 pp	Lower accumulation, broad instability
Realized S&P 500	—	—	8.1%	14.2%	0.0 pp	Observed outcome

Notes: Simulated log price return is computed from the expected simulated future price as $R^L \equiv \ln(E(P_t)/P_{t-1})$; it therefore differs from the mean of the simulated logarithmic returns. Simulated volatility is the standard deviation of annual logarithmic price returns. The realized S&P 500 return is the annualized logarithmic price return over 2014–2022. Dividend income is excluded throughout. The volatility gap is the difference between simulated volatility and the realized S&P 500 volatility of 14.2%, expressed in percentage points (pp). Mean book-equity accumulation and its standard deviation correspond to the assumptions used in each Monte Carlo configuration.

The simulations provide a reference for market returns and volatility within the calibrated regime, using retrospectively calibrated distributions of valuation inputs that are themselves observable or estimable ex ante. They are therefore intended to characterize the return and volatility associated with a given valuation configuration rather than to reproduce the realized sequence of market outcomes.

The “optimistic-prudent” configuration combines a 5.0% average rate of book-equity accumulation with a 7.3% standard deviation. It produces a simulated log price return of

approximately 6.2% and market volatility of 13.7%. The “pessimistic-myopic” configuration, which extrapolates the lower and more stable 2014–2022 accumulation regime, produces a markedly lower simulated log price return of 4.5% and volatility of 11.1%. Finally, combining the lower 3.7% accumulation rate with the broader 7.3% dispersion raises simulated volatility back to 13.7%, while leaving the simulated log price return comparatively low at 4.7%.

These results show that the assumed dispersion of book-equity accumulation has a substantial effect on simulated market volatility but a much smaller effect on its central return tendency. Moving from the pessimistic-myopic to the pessimistic-prudent configuration raises simulated volatility by 2.6 percentage points, from 11.1% to 13.7%, while the corresponding simulated return increases by only about 0.2 percentage points. Conversely, raising the mean accumulation rate from 3.7% to 5.0%, while retaining the same broader dispersion, primarily affects the central return outcome rather than simulated volatility.

The comparison with realized outcomes is particularly informative. Over 2014–2022, the S&P 500 generated an annualized logarithmic price return of approximately 8.1% and volatility of 14.2%. The recent-regime, “pessimistic-myopic” calibration therefore understates the realized return by about 3.6 percentage points and volatility by 3.1 percentage points. By contrast, once the broader historical dispersion of book-equity accumulation is retained, simulated volatility reaches 13.7%, only 0.5 percentage points below the realized value. The corresponding simulated returns nevertheless remain below the realized outcome, by approximately 1.9 percentage points in the “optimistic-prudent” configuration and 3.4 percentage points in the “pessimistic-prudent” configuration.

This close correspondence should be interpreted as conditional quantitative consistency within the calibrated regime rather than as an out-of-sample reconstruction from price-independent fundamentals. The result remains informative, however, because market-return volatility is not itself imposed on the simulation, and the expected IRR series was developed independently for valuation purposes rather than calibrated to reproduce market volatility.

A second consistency result concerns the persistence of the valuation multiple. The correlation between successive simulated price-to-book ratios is approximately 0.90, compared with about 0.86 in the observed data. Unlike the serial dependence of the underlying valuation inputs, this persistence is not imposed directly on the multiple but emerges from their joint dynamics through the valuation equation. The simulation therefore reproduces not only a level of annual return volatility close to that observed, but also a key dynamic property of the valuation state.

The volatility result is therefore more supportive of the valuation framework than the recent low-dispersion regime alone would suggest. Most of the observed annual market volatility can be generated by the joint variation of the fundamental and required-return inputs without calibrating an additional residual volatility component. The remaining difference is small under the broader-dispersion configurations and cannot by itself be interpreted as evidence of irrational or autonomous excess volatility.

The more pronounced discrepancy concerns the simulated price-return level. Even the “optimistic-prudent” configuration produces an annualized return below the 8.1% subsequently realized, while the recent-regime calibration falls substantially further short. This suggests that the valuation configuration did not remain stationary over the period. According to **equation (3)**, sustained price growth above book-equity accumulation must be associated with an increase in the valuation multiple. The next subsection examines whether the observed evolution of the expected IRR can account for a material part of this re-rating.

This distinction also matters for the interpretation of volatility. A persistent directional shift in the required return, rather than ordinary fluctuations around a stationary distribution, changes the valuation multiple over several years and can affect both the average realized return and the dispersion of annual price changes relative to the simulated distribution. The relevant question is therefore not whether the Monte Carlo model reproduces every observed fluctuation, but whether the remaining divergence can be related to an economically identifiable revision in the valuation state.


4.1.2. Changes in the Required Return and Market Valuation

The Monte Carlo results identify a remaining discrepancy, particularly in the average realized price return. The next step is therefore to examine whether this divergence is associated with a revision in the valuation configuration itself. A natural candidate is the expected market IRR. Between 2014 and 2022, IRR_t declined from approximately 6.8% to 5.9%, while normalized ROE increased and expected long-run growth remained broadly stable. According to **equation (2)**, this combination should raise the price-to-book multiple even without any acceleration in book-equity accumulation. The observed evolution of S&P 500 valuation is consistent with this mechanism.

Figure 1 places the change in the estimated IRR alongside annual S&P 500 price returns and book-equity accumulation. Although annual market returns fluctuate considerably, the estimated IRR displays a downward tendency over the period. Book-equity accumulation, by

contrast, remains much closer to the lower-growth regime used in the pessimistic Monte Carlo configurations. Beyond its role in documenting the decline in the estimated IRR, **Figure 1** also suggests that realized annual market fluctuations may reflect a more structured interaction between capital accumulation, valuation, and the required return. This possibility is examined more formally in Section 4.3.

Figure 1. S&P 500 Annual Price Returns, Book-Equity Accumulation, and Estimated Market IRR, 2014–2022



Notes: S&P 500 returns and book-equity accumulation are annual logarithmic changes and are reported on the left-hand scale. IRR_t , reported on the right-hand scale, is the optimism-adjusted estimate of the market-implied IRR described in Section 3.2. The series are shown to illustrate the evolution of realized market performance, capital accumulation, and the expected valuation rate over the period.

The cumulative effect is more clearly seen through the valuation decomposition reported in **Table 3**.

Table 3. S&P 500 Valuation Parameters and Counterfactual Required-Return Scenario

	2014	2022	2022 counterfactual
S&P 500 price	2,054	3,912	2,969
Annualized log price return, 2014–2022	—	8.1%	4.6%
Book equity	726	974	974
Price-to-book ratio	2.83	4.02	3.05
Normalized ROE	13.7%	14.7%	14.7%
Estimated IRR	6.8%	5.9%	6.8%
Expected long-term growth – g'	3.10%	2.98%	2.98%

Notes: The counterfactual column retains the 2022 normalized ROE, expected long-term growth g' , and book-equity level while holding the estimated market IRR at its 2014 level of 6.8%. The corresponding price-to-book ratio is obtained from **equation (2)**, and the counterfactual S&P 500 level is obtained by applying this multiple to 2022 book equity. Annualized log price returns are computed relative to the 2014 S&P 500 level and exclude dividends. Figures reproduce the original calibration. The 2022 endpoint was estimated using information available at the end of March 2022; consistently with the original exercise, the end-March S&P 500 level is treated as the 2022 endpoint.

Between year-end 2014 and the 2022 endpoint used in the original calibration, aggregate book equity increases from 726 to 974, corresponding to annualized logarithmic accumulation of

approximately 3.7%. Over the same period, the S&P 500 price rises from 2,054 to 3,912. The resulting annualized logarithmic price return of approximately 8.1% can therefore be decomposed, using **equation (3)**, as:

$$\frac{1}{8} \ln \left(\frac{P_{2022}}{P_{2014}} \right) = \frac{1}{8} \ln \left(\frac{B_{2022}}{B_{2014}} \right) + \frac{1}{8} \ln \left(\frac{(P/B)_{2022}}{(P/B)_{2014}} \right) \approx 3.7\% + 4.4\% = 8.1\% \quad (10)$$

Almost half of the realized annual price appreciation therefore comes from book-equity accumulation, while slightly more than half reflects an increase in the valuation multiple.

The counterfactual in the final column isolates the valuation effect associated with the decline in the expected IRR. Holding IRR_t at its 2014 level of 6.8%, while retaining the 2022 values of normalized profitability and expected growth, produces a price-to-book ratio of approximately 3.05 rather than the observed 4.02. Applied to the 2022 book-equity base, this implies a counterfactual S&P 500 level of approximately 2,969. The corresponding annualized logarithmic price return is approximately 4.6%, remarkably close to the 4.5% simulated log price return generated by the pessimistic-myopic Monte Carlo configuration. The comparison is suggestive: the return implied by a broadly unchanged required-return configuration is close to that generated from the recent capital-accumulation regime, whereas the substantially higher realized return coincides with a decline in the market IRR and the associated expansion of the valuation multiple.

This counterfactual should not be interpreted as identifying the decline in the required return as the unique cause of the observed market appreciation. IRR_t is inferred from prices and expected cash flows, and the valuation parameters jointly characterize the prevailing market configuration. The exercise instead quantifies the sensitivity of market value to the discount-rate component of **equation (2)**. Conditional on the other valuation inputs, the difference between an expected IRR of 6.8% and 5.9% is sufficient to account for a large part of the observed increase in the S&P 500 multiple.

The same mechanism can also contribute to the volatility gap identified in Section 4.1. A persistent revision in the required return can shift the valuation multiple relative to the distribution generated under a stationary calibration. The 11.1% volatility generated by the Monte Carlo simulation for the recent low-dispersion regime therefore need not capture price variation associated with a persistent revision in IRR. Conversely, the 13.7% volatility generated under the broader-dispersion configurations comes within 0.5 percentage point of the 14.2% realized volatility.

The evidence does not allow the remaining volatility to be allocated precisely between revisions in the required return and market amplification mechanisms. Nor does the decline in IRR_t by itself identify its underlying cause. It is, however, consistent with a change in the valuation state that may have included a reduction in perceived uncertainty as post-financial-crisis systemic concerns receded. More generally, the result illustrates why uncertainty is difficult to isolate empirically: when investors change the compensation they require for bearing an uncertain future, that change becomes observable indirectly through valuation and realized returns rather than as a separately measured source of risk.

The next section reverses the perspective and examines the level of the required return itself.

4.2. Direct Foundations of the Required Return

The analysis first constructs the required-return reference (RRR) from capital accumulation, distributed income, and recurrent market volatility and compares it with the market-implied IRR. It then examines whether inflation provides an additional long-run constraint on the nominal level of the required return.

4.2.1. Capital Accumulation, Volatility, and the Required Return

Section 4.1 showed that a valuation configuration calibrated from observable or estimable valuation inputs can generate annual market volatility close to that subsequently realized. We now use the recurrent-volatility estimate from that exercise to evaluate the required-return reference (RRR) developed in Section 2.4. Under a relatively persistent valuation configuration, **equation (7)** combines expected capital accumulation, distributed income, and recurrent market volatility. For the period-average exercise conducted here, these ex-ante target quantities are approximated by their corresponding average or recurrent values over the recent regime. We therefore combine the period-average proxy for expected logarithmic book-equity accumulation, $\bar{g}_B^*$, the average logarithmic dividend yield, $\bar{d}_y^{L*} \equiv \ln(1 + \overline{DY})$, and the recurrent-volatility adjustment, $(\sigma^*)^2/2$.

Over the recent sample, average logarithmic book-equity accumulation is approximately 3.66%. The average S&P 500 dividend yield, calculated from the Shiller series as annual dividends relative to average monthly prices and then averaged over the period, is approximately 1.86%, corresponding to a logarithmic income yield of about 1.84%. The Monte Carlo results in Section 4.1 imply an ex-ante recurrent annual volatility of approximately 13.7% under the broader-dispersion configurations, close to but slightly below the 14.2% volatility

subsequently realized. Using this 13.7% recurrent-volatility estimate gives a lognormal adjustment of about 0.94%.

Applying the RRR relation to these period-average inputs gives:

$$\bar{k}^{L*} \approx 0.0366 + 0.0184 + \frac{0.137^2}{2} \approx 0.0644. \quad (11)$$

Converted back to a discrete annual return, RRR is 6.65%.

$$\bar{k}^* \approx e^{0.0644} - 1 \approx 6.65\%. \quad (12)$$

The resulting RRR of 6.65% is close to the average adjusted market-implied IRR of approximately 6.40% estimated over the same period, a difference of about 25 basis points. The comparison nevertheless contains a limited degree of endogenous information because the market-implied IRR, which is inferred partly from contemporaneous prices, contributes to the simulated return volatility subsequently used in equation (7). This dependence does not, however, create an accounting identity between IRR and RRR: book-equity accumulation and distributed income are measured separately, while volatility enters RRR only through the term $\frac{1}{2}\sigma^2$. Quantitatively, replacing the simulated volatility of 13.7% by the realized volatility of 14.2% raises the volatility adjustment from about 0.94% to 1.01% and RRR by less than eight basis points. The close correspondence with the market-implied IRR therefore does not hinge materially on feeding the IRR through the Monte Carlo volatility estimate.

As anticipated in Section 3.3, the result should therefore be interpreted as a conditional consistency finding rather than as independent identification of the required return. Pástor et al. (2008) provide a basis for treating variation in the implied required return as informative about variation in the latent expected return, but their result does not impose either the level of volatility generated here or equality between the market-implied and economically motivated required returns. The present exercise adds a different quantitative observation: a recurrent-volatility measure generated partly from IRR dynamics yields, together with separately measured accumulation and distributed income, an RRR close to the market-implied IRR, and essentially the same conclusion is obtained when simulated volatility is replaced by realized market volatility. Given the short sample and the possibility that these quantities respond jointly to broader economic conditions, this correspondence remains evidence of economic coherence rather than an invariant structural relation.

This interpretation differs from treating the aggregate equity premium as an independent primitive. The conventional decomposition remains valid as an accounting representation, but the direction of interpretation changes: if RRR provides an economically meaningful reference for the market required return, the equity premium can be viewed partly as the difference between that required return and the prevailing risk-free rate, rather than as an exogenously given quantity from which the required return is constructed.

The close correspondence observed over the recent period also raises the question of what conditions the required return over longer horizons and across different nominal environments. Inflation, in particular, may affect the nominal return investors require to preserve purchasing power and therefore constrain the level around which the required return is formed. The next subsection examines whether the historical evolution of market required returns is related not only to long-term government bond yields but also to inflation.

4.2.2. Inflation as an Additional Constraint on the Required Return

Capital accumulation, distributed income, and recurrent volatility may not exhaust the conditions associated with the required return. Because both expected corporate cash flows and required returns are expressed in nominal terms, inflation may affect the nominal scale of the valuation process. The relevant question is not simply whether equity required returns exceed inflation, but whether their spread over inflation displays greater historical stability than the conventional spread over the risk-free rate. Inflation may therefore provide an additional long-run constraint on the nominal level of the required return, without entering the relation developed in Section 2.4 as a separate additive component.

To examine this possibility, the recent Fairness Finance estimates are placed in a longer historical perspective using published estimates of the U.S. equity required return from Jagannathan et al. (2000), Ke and Liu (2014), and Mishra and O'Brien (2019). These historical estimates, and the adjustments used to place them on a broadly comparable basis, are presented in greater detail in Clère (2026b), from which the decadal series used here is drawn. Because the underlying estimates rely on different methodologies and are not continuously available over the full period, they are aggregated at the decadal level and should be interpreted as broad historical comparisons rather than as a homogeneous annual time series. Estimates based on analyst forecasts are adjusted, where applicable, for the forecast-optimism correction described in Section 3.2. The historical comparison reveals a marked change in the relation between the estimated required return and the nominal risk-free rate. Until the 1970s, the spread between

the two was generally large. It then compressed sharply during the high-interest-rate environment of the 1980s, when the expected IRR exceeded long-term Treasury yields by only about one percentage point. The subsequent decline in government bond yields did not produce an equivalent decline in the market IRR, and the conventional equity premium therefore widened again. **Table 4** compares the resulting IRR estimates with long-term U.S. Treasury yields and consumer-price inflation from the 1930s through the 2010s. The comparison with inflation produces a different pattern. Since the 1980s, the difference between the expected IRR and consumer-price inflation has remained within a comparatively narrow range of approximately 5–6 percentage points.

Table 4. Expected U.S. Market IRR, Long-Term Treasury Yields, and Inflation by Decade

Decade	Expected IRR	Long-term Treasury yield	IRR minus Treasury yield	CPI inflation	IRR minus inflation
1930s	11.9%	3.0%	8.9%	-2.0%	13.8%
1940s	13.6%	2.2%	11.4%	5.6%	8.0%
1950s	12.0%	3.1%	8.9%	2.1%	10.0%
1960s	10.0%	4.8%	5.2%	2.3%	7.7%
1970s	11.5%	7.5%	4.0%	7.1%	4.4%
1980s	11.6%	10.6%	1.0%	5.6%	6.0%
1990s	8.9%	6.7%	2.3%	3.0%	5.9%
2000s	7.8%	4.5%	3.4%	2.6%	5.2%
2010s	7.2%	2.4%	4.8%	1.8%	5.5%


Notes: Expected IRRs combine the historical estimates reported by Jagannathan et al. (2000), Ke and Liu (2014), and Mishra and O’Brien (2019), supplemented by Fairness Finance estimates for the most recent period. Estimates based on analyst forecasts are adjusted where applicable for forecast optimism. Long-term Treasury yields correspond to 20-year U.S. government bonds before 1970 and 10-year Treasury yields thereafter. All figures are decadal averages.

Figure 2 makes the contrast between the two spreads more visible. What is comparatively stable over the recent decades is therefore not the conventional equity premium, but the spread between the estimated market required return and inflation.

The relative stability of the IRR–inflation spread since the 1980s is suggestive, but should not be interpreted as evidence of a structural inflation-pricing rule.

The number of decadal observations is small, the underlying IRR estimates come from several sources, and monetary regimes changed substantially over the period. The historical evidence nevertheless indicates that the required return did not move one-for-one with long-term government bond yields. In particular, the sharp increase in long-term Treasury yields around the Volcker period was associated with a compression of the conventional equity premium rather than with a comparable rise in the market IRR.

Figure 2. Expected U.S. Market IRR Relative to Long-Term Treasury Yields and Inflation



Notes: The solid series reports the difference between the expected market IRR and long-term Treasury yields. The second series reports the difference between the expected market IRR and consumer-price inflation. Values are decadal averages corresponding to **Table 4**.

This pattern is consistent with the possibility that inflation conditions the required-return relation developed in Section 2.4 and examined empirically in Section 4.2.1. If investors require a return that remains economically meaningful in real terms, the nominal required return may depend not only on capital accumulation, distributed income, and recurrent market variability, but also on the broader inflation environment. This does not imply that inflation should simply be added to the relation, since nominal accumulation and expected cash flows may themselves incorporate part of its effect. Rather, inflation may constrain the level around which the required return is formed. In that representation, the observed equity premium remains the difference between k_t and $r_{f,t}$, but its magnitude need not be independently fixed.

The evidence remains exploratory. It does not establish whether inflation affects the required return directly, through nominal cash-flow growth, through perceptions of risk and uncertainty, or through some combination of these channels. Its role here is more limited: the long-run comparison suggests that the relation between required returns and risk-free rates is historically unstable, whereas inflation may provide an additional conditioning variable alongside the valuation-based determinants examined over the recent period.

Having examined the economic anchoring of the required return, we return to the volatility question raised in Section 4.1. The Monte Carlo exercise showed that the joint dynamics of fundamental and valuation inputs can generate annual market volatility close to that observed. A more demanding question is whether the successive one-year price movements through which that volatility is realized are themselves economically structured, rather than merely constituting short-horizon noise around fundamentals. Section 4.3 addresses this question using the Valuation and macro-financial Cycle framework defined in Section 3.4.

4.3. Ex-Ante Information and the Structure of Subsequent Market Returns

Using the specification developed in Section 3.4, we now examine whether predetermined information from the Valuation and macro-financial Cycle dimensions is systematically associated with subsequent one-year market returns, and whether the resulting relation exhibits an economically coherent structure. With only 37 quarterly observation dates, the exercise is deliberately diagnostic rather than intended to establish an invariant forecasting law.

4.3.1. Baseline relation

Table 5 reports the baseline estimates and joint Wald tests for the full specification and its two conceptual blocks. The model explains 70.2% of the variation in subsequent one-year returns, with an adjusted R^2 of 65.4%.

Table 5. Baseline regression of subsequent one-year stock market returns

Overall fit							
Observations	37						
R^2	0.7019						
Adjusted R^2	0.6538						
RMSE	7.60%						
AIC	-178.65						
Wald χ^2 - All predictors	190.42						
Wald degrees of freedom - All	5						
Wald p-value - All	3.18E-39						
Wald χ^2 - Valuation	104.77						
Wald degrees of freedom - Valuation	2						
Wald p-value - Valuation	1.78E-23						
Wald χ^2 - Cycle	155.69						
Wald degrees of freedom - Cycle	3						
Wald p-value - Cycle	1.56E-33						
Variables	Coefficient	Std. error	t-statistic	p-value	95% CI lower	95% CI upper	VIF
Intercept	10.96%	1.55%	7.06	6.21E-08	7.79%	14.12%	
<i>Valuation</i>							
PE	9.66%	1.19%	8.14	3.42E-09	7.24%	12.08%	1.88
IRR	6.30%	2.21%	2.86	7.54E-03	1.81%	10.80%	1.70
<i>Cycle</i>							
Δg_B	3.51%	1.36%	2.58	1.49E-02	0.73%	6.28%	1.59
Δg_P	10.98%	2.40%	4.58	7.16E-05	6.09%	15.87%	6.09
GAP	-15.33%	2.47%	-6.20	6.93E-07	-20.37%	-10.29%	6.19

Notes: The model is estimated using standardized explanatory variables. The dependent variable is the subsequent one-year logarithmic S&P 500 price return. Coefficients measure the percentage-point change in the subsequent annual log return associated with a one-standard-deviation increase in each explanatory variable. Newey–West standard errors are computed with three lags, consistent with the one-year return horizon and quarterly sampling. The estimator uses the Bartlett kernel, no prewhitening, and the finite-sample adjustment. Reported Wald tests use the Newey–West covariance matrix. AIC is computed as $n \times \ln(RSS/n) + 2k$, where k includes the intercept.

The five predictors are jointly highly significant, and both the Valuation and Cycle blocks are jointly significant within the full specification. The baseline results therefore provide a first

affirmative answer to the question posed above: within the available sample, subsequent one-year market returns exhibit a substantial structure associated with the predetermined Valuation and macro-financial Cycle configuration.

The expected IRR provides a direct link with the preceding exercises: the valuation parameter examined in Sections 4.1 and 4.2 is also associated with subsequent market dynamics. A one-standard-deviation increase in the estimated ex-ante IRR is associated with a 6.3-percentage-point increase in the subsequent annual return. The positive sign is consistent with the aggregate IRR predictability documented by Li et al. (2013). Higher estimated ex-ante required returns are followed, on average, by higher realized returns. This result is consistent with Clère (2026a), which documents similar predictive content using an IRR proxy constructed from CAPE and P/B multiples.

The trailing P/E ratio also enters positively and with a larger partial effect. This sign is consistent with the evidence at the one-year horizon in Clère (2026a), after controlling for the IRR proxy. At this horizon, its coefficient should not be interpreted as the unconditional long-run effect of market expensiveness. Conditional on the remaining variables, P/E conveys here information about recent profitability and growth relative to the more normative level incorporated in the IRR. Its role is therefore better understood jointly with the IRR and the cycle variables than as that of an isolated valuation multiple.

The Cycle coefficients reveal a complementary dynamic. Annual price-growth acceleration, $\Delta g_{P,t}$, and the acceleration in book-equity accumulation relative to per-capita output, $\Delta g_{B,t}$, enter positively, while GAP_t enters strongly negatively. The latter relation is consistent with a subsequent correction: when trailing price growth has moved further ahead of underlying profitability, subsequent market performance tends to weaken over the following year.

The comparatively modest partial coefficient on $\Delta g_{B,t}$ warrants particular attention. Relative book-equity accumulation occupies a central position in the valuation mechanism developed above, yet its partial coefficient is smaller than those of several other predictors—a point revisited below. Moreover, some Cycle variables are strongly correlated, and the IRR may share information with both valuation and accumulation variables. These features make it necessary to distinguish a variable's partial regression coefficient from its broader informational contribution. They also motivate a broader assessment of specification robustness and of how shared information is distributed across the predictors.

4.3.2. Robustness and information structure

Table 6 first shows that the baseline relation is not associated with obvious conventional specification failures.

Residual normality is not rejected, Breusch–Godfrey tests provide no evidence of serial correlation at three lags, and neither Breusch–Pagan nor White tests detect heteroskedasticity. The inconclusive Durbin–Watson result therefore does not alter the general diagnostic assessment. Newey–West inference is nevertheless retained throughout because the overlap of annual returns provides an independent structural reason for HAC correction.

Influence and measurement tests point in the same direction. The leave-one-out exercise preserves the sign of every coefficient in all 37 re-estimations, while adjusted R^2 remains between 62.3% and 68.1%. Both conceptual blocks remain jointly significant throughout. Moreover, replacing end-of-month prices with monthly-average prices raises adjusted R^2 from 65.4% to 71.5%, improves the information criterion, and leaves all five predictors significant under Newey–West inference. The strengthening of the relation under a smoother price measure is consistent with reduced sensitivity to transitory end-of-month price fluctuations.

The more informative results concern the distribution of explanatory content across correlated predictors. The LMG–Budescu decomposition attributes 34.0% of explained variance to the Valuation variables and 66.0% to the Cycle variables. Most notably, $\Delta g_{B,t}$ accounts for 33.4% of explained variance, the largest contribution of any individual predictor. Its comparatively modest partial coefficient therefore understates its informational role, reinforcing the importance of changes in relative book-equity accumulation in the preceding analysis.

At the same time, this centrality does not make the IRR redundant: it accounts for 11.6% of explained variance, a conclusion that is robust to alternative allocations of common information. When the Valuation block is given priority, it remains jointly significant after Cycle variables are purged of their shared component. A symmetric, order-neutral orthogonalization leaves all five orthogonal dimensions significant at the 5% level. Finally, when P/E and $\Delta g_{B,t}$ are treated as the priority valuation–accumulation core, the component of the IRR orthogonal to these two variables remains significant, with a coefficient of 5.4 percentage points. Within the sample, the expected IRR thus retains distinct information about subsequent returns that is not reducible to current valuation or relative book-equity accumulation alone.

Table 6. Robustness, Shared Information, and Principal-Component Evidence

Level of validation	Test or alternative	Main result	Interpretation
I. Standard diagnostics and influence	Residual normality	Shapiro–Wilk $p=0.798$; D’Agostino–Pearson $p=0.846$	Normality is not rejected
	Serial correlation	Durbin–Watson =1.388, inconclusive; BG(3) $p=0.247$ –0.337	No evidence of residual serial correlation from the more general BG test
	Heteroskedasticity	Breusch–Pagan $p=0.844$; White $p=0.840$	No evidence of heteroskedasticity
	Leave-one-out estimation	All coefficient signs preserved in 37 estimations; adjusted $R^2=62.3\%$ –68.1%; V and C Wald tests significant in 37 cases	Results are not driven by an individual observation
	Monthly-average prices	Adjusted R^2 : 65.4% $\rightarrow$ 71.5%; AIC: $-178.7 \rightarrow -192.2$; all five predictors remain significant under NW(3)	The relation strengthens under a smoother price measure, consistent with reduced sensitivity to transitory end-of-month price fluctuations
II. Information structure and shared variance	LMG-Budescu decomposition	V: 34.0%; C: 66.0%. Δg_B accounts for 33.4% of explained R^2 , followed by GAP (23.0%), P/E (22.3%), IRR (11.6%), and Δg_P (9.6%)	Explanatory content is broadly distributed; the partial coefficient on Δg_B understates its central informational role
	V-oriented orthogonalization	Priority V: Wald $\chi^2=40.83$, $p<0.001$; residual C: $\chi^2=155.69$, $p<0.001$	Valuation and cycle retain distinct explanatory content after removing cross-block shared variance
	Symmetric non-oriented orthogonalization	All five mutually orthogonal dimensions remain significant at 5%; model fit unchanged	Results do not depend on predictor ordering or on collinearity among the original regressors
	{P/E, Δg_B }-oriented orthogonalization	Priority core: Wald $\chi^2=63.62$, $p<0.001$; residual block: $\chi^2=62.77$, $p<0.001$; residual IRR coefficient =5.42 pp, $p=0.008$	IRR retains significant information beyond current valuation and capital accumulation
III. Principal-component evidence	Dominant common component (PC1)	52.3% of predictor variance; $R^2=0.8\%$; NW $p=0.597$	The main common movement in the predictors does not explain subsequent returns
	Capital-accumulation component (PC2)	20.8% of predictor variance; Δg_B loading =0.983; $R^2=38.3\%$; NW $t=5.95$, $p<0.001$	Confirms that capital accumulation carries a major orthogonal information channel despite its modest partial coefficient
	Remaining orthogonal information	PC3 significant ($p=0.002$); PC4 not significant; PC5 explains 15.4% of returns despite only 1.6% of predictor variance ($p=0.044$). PCs 1–4 explain 98.4% of predictor variance but only 54.8% of returns; adding PC5 restores $R^2=70.2\%$	Predictive information is distributed across economically specific directions and cannot be identified from predictor variance alone

Notes: The baseline specification uses 37 quarterly observations and five standardized explanatory variables; the dependent variable is the subsequent one-year logarithmic S&P 500 price return. Newey–West inference uses three lags to account for overlapping annual returns. LMG-Budescu contributions are expressed as shares of the full-model explained R^2 . Oriented orthogonalizations preserve the designated priority block and residualize the remaining variables with respect to it; the symmetric transformation is order-neutral. Principal-component results report the share of predictor variance captured by each component together with its separate explanatory power for subsequent returns. The monthly-average-price specification replaces end-of-month prices with Shiller monthly averages.

Principal-component analysis provides a complementary view of this information structure by orthogonalizing the five predictors and identifying where their explanatory content resides. The

decomposition is particularly informative because the component that dominates the joint variation of the predictors is not the one associated with subsequent performance. PC1 accounts for 52.3% of total predictor variance and mainly combines positive movements in Δg_P , GAP, and P/E with a negative loading on the IRR, while Δg_B plays only a minor role. This configuration largely isolates a common price-related valuation dimension: price enters directly or relative to profitability in the first three variables, while, conditional on expected cash flows, a higher price is associated with a lower implied IRR. Yet this dominant common movement has virtually no explanatory power for the subsequent one-year return ($R^2 = 0.8\%$; Newey–West $p = 0.597$). This result also informs the interpretation of the Monte Carlo exercise: although the expected IRR is partly inferred from contemporaneous prices, the PCA indicates that the empirical information carried by the valuation variables is not reducible to their dominant common price-related movement. This does not exclude the transmission of price-related variation through the IRR in the simulation, but it argues against interpreting the broader valuation structure as a simple repackaging of contemporaneous price movements.

PC2, which accounts for a further 20.8% of predictor variance, is instead overwhelmingly aligned with Δg_B : its weight is 0.951 and its loading is 0.983. This component alone explains 38.3% of subsequent return variation and is strongly significant under Newey–West inference ($t = 5.95, p < 0.001$). This result complements the preceding evidence on the comparatively modest partial coefficient of Δg_B . Although relative book-equity accumulation has a relatively small coefficient in the multivariate regression, the orthogonal dimension predominantly associated with its acceleration contains substantial information about subsequent performance.

The remaining components reveal that this information is not exhausted by capital accumulation. PC3, accounting for 15.9% of predictor variance, mainly contrasts P/E with Δg_P , GAP, and the IRR; it is also significantly related to subsequent returns ($p = 0.002$). PC4 is primarily a residual valuation dimension combining the IRR and P/E and is not individually significant. PC5 is particularly revealing. It represents only 1.6% of predictor variance and is dominated by an opposition between GAP and Δg_P , yet it remains significant at the 5% level ($p = 0.044$) and accounts for 15.4% of return variation. Thus, even a direction that is almost negligible for describing the overall variance of the predictors can contain economically relevant information for subsequent returns.

The cumulative regressions reinforce this interpretation. The first three components capture 88.9% of predictor variance but explain only 48.9% of return variation, while the first four capture 98.4% of predictor variance but explain 54.8%. Adding PC5 raises R^2 to 70.2%,

necessarily reproducing the fit of the original five-variable specification because the complete set of principal components spans the same predictor space. The informative result is therefore not the final equality of fit, but the highly uneven distribution of explanatory content across components. In particular, the dominant common factor carries almost none of it, while economically more specific orthogonal directions—including the dimension dominated by Δg_B —carry substantial information. The exercise therefore cautions against dimensionality reduction based solely on the proportion of predictor variance explained.

4.3.3. Out-of-Sample Persistence of the Valuation Configuration

The preceding results establish that the in-sample relation is robust to conventional specification concerns and that its informational content is distributed across several dimensions rather than concentrated in a single common factor. A stricter test is whether this structure retains predictive content when each forecast is generated using only information available before the forecast date. We therefore implement an expanding-window one-year out-of-sample exercise. The first forecast is formed in March 2018 from an initial estimation sample containing one-year returns fully realized by December 2017. A one-quarter gap is maintained between the end of the last return used for estimation and each forecast origin, so that the retained accounting inputs predate the forecast origin. The estimation window is then expanded by one quarterly observation at each step. This procedure produces 24 one-year-ahead forecasts through December 2023.

For this exercise, Δg_B is excluded from the forecasting specification because its construction incorporates real per-capita GDP growth, whereas the available macroeconomic series contain subsequent revisions. The OOS specification consequently retains four variables: the IRR and P/E for Valuation, and Δg_P and GAP for the macro-financial Cycle. This exclusion reflects information-set discipline rather than weak predictive content. Exploratory sensitivity calculations indicate that the predictive contribution of the book-equity accumulation dimension is sensitive to the horizon over which its acceleration is measured. Modestly longer windows can substantially improve OOS fit relative to the four-variable specification, while also slightly improving the corresponding in-sample fit. This is consistent with the prominent role of the accumulation dimension in the in-sample analysis, although the alternative timing was identified ex post and does not resolve the real-time measurement issue associated with revised per-capita output data.

Forecast performance is evaluated against an expanding-window historical mean constructed from the same information set and using the same one-quarter gap. The standard out-of-sample coefficient of determination is:

$$R_{OOS}^2 \equiv 1 - \frac{\sum_t (R_{t,t+4} - \hat{R}_{t,t+4})^2}{\sum_t (R_{t,t+4} - \hat{R}_{t,t+4}^\mu)^2},$$

so that a positive value indicates a reduction in mean squared forecast error relative to the historical mean. To examine where the predictive content resides, four economically defined specifications are considered: the IRR alone; the Valuation block $V = (k, PE)$; the Cycle block $C = (\Delta g_P, GAP)$; and the combined $V + C$ configuration. This set of specifications directly parallels the informational structure examined in Sections 4.3.1 and 4.3.2.

The results are sharply differentiated across specifications, as shown in **Table 7**. The IRR alone fails to outperform the historical mean, with $R_{OOS}^2 = -11.8\%$, while adding P/E leaves the Valuation block essentially at the benchmark, with $R_{OOS}^2 = -0.9\%$ ($CW_{HAC} = 0.90, p = 0.188$). By contrast, the Cycle block produces an R_{OOS}^2 of 22.7% and significantly outperforms the expanding-window historical mean ($CW_{HAC} = 2.46$, one-sided $p = 0.011$).

Table 7. Out-of-Sample Forecast Performance at the One-Year Horizon

Forecast specification	SSE	R_{OOS}^2	CW-HAC(3)	p-value
IRR only: k	0.79	-11.8%	—	—
Valuation: $V = (k, PE)$	0.72	-0.9%	0.90	0.188
Cycle: $C = (\Delta g_P, GAP)$	0.55	22.7%	2.46	0.011
Valuation + Cycle: $V + C$	0.47	34.0%	2.74	0.006
Historical mean	0.71	—	—	—
$V + C$ relative to C	—	14.7%	0.49	0.310

Notes: Expanding-window forecasts, March 2018–December 2023 ($n = 24$), with a one-quarter estimation gap. R_{OOS}^2 is relative to the expanding-window historical-mean benchmark, except for the last row, which uses the Cycle forecast as the benchmark. Clark–West tests use Newey–West(3) standard errors and one-sided p-values. $V = (k, PE)$; $C = (\Delta g_P, GAP)$.

The combined four-variable V–C specification performs best among the economically defined models, with an R_{OOS}^2 of 34.0% and a significant improvement over the expanding-window historical mean ($CW_{HAC} = 2.74, p = 0.006$). Measured instead against the realized variance of one-year returns over the forecast period, the model achieves a positive fit of 26.8%, whereas the expanding-window historical mean yields -10.9%. Relative to the Cycle specification alone, the combined model reduces squared forecast error by 14.7%, although this incremental improvement is not statistically significant ($CW_{HAC} = 0.49, p = 0.31$).

Forecasts based simply on the immediately preceding one-year return perform markedly worse, indicating that the result does not reflect mechanical return persistence.

The OOS pattern reinforces the configuration-based interpretation. Predictive content is not concentrated in the market IRR or in Valuation alone: the Cycle block significantly outperforms the expanding-window historical mean, while the combined V–C specification achieves the lowest forecast error and highest R_{OOS}^2 . Its incremental gain over Cycle is not precisely estimated over the short forecasting sample, but the broader result is consistent with the in-sample PCA and orthogonalization evidence: the required return is one component of the valuation process rather than a sufficient statistic for the information associated with subsequent one-year returns. With only 24 overlapping forecasts, the exercise remains exploratory and cannot establish an invariant forecasting relation, but the significant OOS performance of the Cycle and combined specifications provides a more demanding test of the identified structure than the in-sample evidence alone.

The magnitude of the out-of-sample gains is substantial, and the fact that they are obtained at the one-year horizon is especially noteworthy. Li, Ng, and Swaminathan (2013) report OOS predictability from aggregate IRR, with an R_{OOS}^2 of 1.22% at the one-month horizon. This is already below the performance reported in Clère (2026a) at the same horizon, although that study relies on observable market multiples, their dynamics, and macro-financial information rather than on a directly estimated implied required return. The contrast is all the more notable because the two monthly OOS periods substantially overlap, while Clère (2026a) employs a fixed rolling estimation window and Li et al. an expanding sample.

The present results extend this comparison. OOS performance remains of the same order of magnitude as in Clère (2026a), but is now obtained at the one-year rather than the one-month horizon, with a directly estimated market IRR incorporated into the broader Valuation–Cycle configuration. The evidence therefore raises the question whether predictive information resides less in the implied required return considered in isolation than in the broader valuation configuration in which it is embedded. This interpretation must remain tentative, since return definitions, predictor construction, forecast horizons, and estimation procedures differ across the studies. A common-sample, common-horizon comparison would be required to distinguish a configuration effect from differences attributable to measurement, horizon, or estimation design.

4.3.4. Converging evidence

The Monte Carlo analysis asks whether the joint dynamics of the market IRR, book-equity accumulation, profitability, and expected growth can generate a credible level of annual market volatility,

$$(k, g_B, ROE, g') \rightarrow \sigma^* \stackrel{?}{\approx} \sigma_{(realized)}.$$

The required-return exercise then asks whether accumulation, distributed income, and recurrent volatility imply an economically plausible required-return reference (RRR),

$$(g_B, DY, \sigma^*) \rightarrow k^* \stackrel{?}{\approx} k.$$

Finally, the return analysis turns from levels to dynamics by testing whether prevailing Valuation and Cycle conditions contain information about the subsequent one-year returns through which realized market volatility unfolds, both in sample and out of sample,

$$(V, C)_t \stackrel{?}{\rightarrow} R_{t,t+4} \quad \text{and} \quad (V, C)_t \stackrel{?}{\rightarrow} \hat{R}_{t,t+4}^{OOS}.$$

What emerges across the three exercises is a common valuation structure that is informative about the magnitude of market volatility, the economic level of the required return, and the dynamics of subsequent market returns.

The first two exercises contain a limited degree of endogenous information because the market-implied IRR contributes to the simulated volatility subsequently used in the RRR calculation. This dependence does not, however, make the result mechanical. Simulated volatility is an output rather than a calibration target, and replacing it with realized market volatility leaves the resulting required-return reference materially unchanged.

More importantly, the two required-return measures have different economic status. The market IRR is an implied estimate recovered from market prices and expected cash flows; it does not by itself explain why the required return should lie at that level. By contrast, the RRR, k^* , provides a direct economic rationale for that level by relating it to expectations of reinvested and distributed income and to the perception of risk (capital accumulation, distributed income, and recurrent volatility). Their proximity is therefore best interpreted as a conditional consistency result: an economically motivated required-return reference lies close to a market-implied estimate obtained from a different valuation procedure, without implying an accounting identity or an independently identified structural relation.

The broader implication concerns the nature of the valuation process itself. The return evidence indicates that the relevant state is multidimensional rather than reducible to the market IRR alone, while the persistence of its predictive content out of sample suggests sufficient local continuity for current valuation conditions to remain informative about subsequent market movements. The empirical picture is therefore one of structured but revisable valuation: market fluctuations are neither mechanically determined by the initial state nor indistinguishable from autonomous noise. This coexistence of structure, local continuity, and revision motivates the interpretative framework developed in the following section.

5. Discussion

The preceding evidence leaves an interpretative question: how can valuation remain sufficiently structured to generate persistent empirical relations while being repeatedly revised as information and economic conditions change? Clère (2026a) develops a Bayesian valuation-convention framework designed to address precisely this coexistence of structure and instability. Rather than restating that framework in full, this section recalls only the elements needed to interpret the present evidence and to assess how the three empirical exercises bear upon it.

The starting point is fundamental uncertainty. Future economic conditions cannot be observed at the valuation date, nor need all relevant contingencies be describable through an objectively known probability distribution. Valuation must nevertheless take place. Investors therefore rely on workable procedures for forming expectations and assigning value, consistently with Simon's (1978) notion of procedural rationality. When sufficiently shared, such procedures can perform the coordinating role described by Orléan (1999) as a market convention. Persistence does not imply invariance: consistent with Lo's (2004) Adaptive Markets Hypothesis, valuation relationships may remain locally coherent while adjusting as information and economic conditions change.

Following Clère (2026a), let θ denote one conceivable representation of future economic and financial conditions at date t , and let Θ_t denote the universe of such representations considered possible at that date. Each representation jointly incorporates expectations about future cash flows, the required return, and relevant macro-financial conditions. Market participants assign subjective weights $b_t(\theta)$ to these alternative representations. The resulting probabilities are therefore working beliefs over conceivable representations, rather than objective probabilities

governing an exhaustively specified future. This distinction allows probabilistic learning to coexist with Knightian uncertainty.

5.1. Bayesian Updating and Valuation Revision

Within such a locally persistent but revisable convention, and conditional on a sufficiently stable universe Θ_t , incoming information can modify the relative weights assigned to the representations it contains. As developed in Clère (2026a), this updating process can be represented in reduced form as:

$$b_t(\theta) = \frac{p(X_t | \theta) b_{t-1}(\theta)}{\int_{\Theta_t} p(X_t | \tilde{\theta}) b_{t-1}(\tilde{\theta}) d\tilde{\theta}} \quad (13)$$

where X_t denotes newly observed information and $p(X_t | \theta)$ its likelihood under representation θ . **Equation (13)** is not intended to imply that investors literally perform Bayesian calculations, share identical beliefs, or know the objective likelihood of future outcomes. It provides a parsimonious representation of sequential learning: current assessments depend both on new information and on beliefs inherited from earlier periods. The Bayesian step applies within a sufficiently stable Θ_t ; fundamental uncertainty remains relevant because the universe of conceivable representations may itself change and is not assumed to follow a known probabilistic transition law.

Each representation θ is mapped through a valuation function $V(\cdot)$ into a conditional intrinsic value $V(\theta)$. The prevailing market valuation can therefore be represented as the belief-weighted aggregation of these conditional values:

$$P_t = \int_{\Theta_t} V(\theta) b_t(\theta) d\theta \quad (14)$$

This formulation preserves the intrinsic-valuation logic of Williams (1938): prices remain grounded in expected cash flows and required returns. What is relaxed is the assumption that one uniquely identifiable valuation state is known to be correct at date t . The convention consists instead of the representations regarded as relevant, the valuation procedure applied to them, and the relative weights assigned to them. The convention can therefore evolve through changes in $b_t(\theta)$ even when Θ_t remains locally stable, while more substantial historical change may also transform Θ_t itself.

The distinction between valuation under the prevailing convention and subsequent revisions to that convention extends naturally to realized returns. Conceptually,

$$R_{t,t+h} = E(R_{t,t+h}|\Theta_t) + \Gamma_{t,t+h}^{rev} \quad (15)$$

where the first term denotes the return conditionally implied by the valuation convention prevailing at t , while $\Gamma_{t,t+h}^{rev}$ collects subsequent revisions in conditional valuation. These revisions may result from confirmation or invalidation of previously considered outcomes, changes in their belief weights, reassessment of known uncertainty, or developments that were not represented in the initial universe of possibilities. The first three cases can be represented as Bayesian revisions of $b_t(\theta)$ within a locally stable Θ_t ; the last may instead require a change in Θ_t itself and therefore lies outside the Bayesian updating rule in **equation (13)**. In this respect, **equation (15)** is interpretative rather than an empirically identified decomposition; its purpose is to distinguish variation conditional on an initial valuation configuration from subsequent changes in that configuration.

5.2. Empirical Implications and Limits of the Valuation-Convention Interpretation

The empirical results provide several ways of confronting the valuation-convention interpretation with observable market outcomes. The first concerns the distinction between recurrent variability within a prevailing regime and revision of the regime itself. The Monte Carlo simulation generates returns while profitability, expected growth, capital accumulation, and the market IRR vary within calibrated distributions and dependence structures, but the valuation rule and the probabilistic structure governing those variations remain unchanged. The resulting volatility can therefore be interpreted as recurrent variability conditional on a given valuation regime. A persistent displacement in a central parameter is conceptually different because it may alter the configuration itself rather than merely generate another realization within it. The decline in the expected IRR over 2014–2022 is consistent with such a revision, and the counterfactual exercise shows that its valuation effect is economically substantial. The empirical distinction between recurrent volatility within a convention and revision of the convention thus maps naturally onto the two components of equation (15), even though neither is directly observable.

A second implication concerns the economic anchoring of the convention. A valuation convention formed under uncertainty need not be arbitrary or purely self-referential. Section 4.2 shows that expectations concerning reinvested and distributed income and the perception

of risk provide an economically motivated required-return reference (RRR) close to the market-implied IRR observed over the recent period, while the longer-run evidence suggests that inflation may constrain its nominal level. This does not establish structural determination, but it indicates that a central parameter of the convention can remain connected to observable economic conditions rather than evolving independently of fundamentals.

The return evidence provides additional empirical support for two properties already present in the framework developed in Clère (2026a). First, the valuation state is represented there as multidimensional; the present results show that its dimensions are also empirically differentiated. Information about subsequent returns is not exhausted either by the IRR or by the dominant common movement among the predictors, implying that distinct components of the configuration need not carry the same information. Second, local persistence was already part of the earlier framework and received shorter-horizon empirical support. The present expanding-window exercise extends that evidence to the one-year horizon: information embedded in the prevailing Valuation–Cycle configuration remains useful after temporal separation between estimation and subsequent return realization. This does not identify the belief weights in equation (13) or establish the duration of persistence, but it provides further empirical support for local continuity without requiring the underlying relations to be globally invariant.

This interpretation remains compatible with semi-strong informational efficiency, since predictability need not imply delayed incorporation of public information when expected returns or other state-dependent compensation for risk vary over time (Fama, 1970, 1991; Cochrane, 2011). The evidence therefore does not by itself reject conventional explanations based on time-varying expected returns. The valuation-convention perspective adds a different possibility: information may already be reflected in current prices while remaining relevant because the valuation configuration through which it is interpreted persists locally and is subsequently revised. The empirical results cannot determine whether this interpretation dominates standard asset-pricing explanations, but they show that structured valuation and subsequent predictability need not be treated as mutually exclusive.

The evidence nevertheless remains conditional on a short and unusually constrained sample. The annual specification contains only 37 quarterly observation dates and the expanding-window exercise contains only 24 overlapping forecasts; the IRR is model-based; some in-sample macroeconomic inputs rely on subsequently revised series; and the Monte Carlo calibration uses the same broad regime whose volatility it seeks to reproduce. The latter should

therefore be read as a conditional consistency exercise rather than as an independent causal reconstruction. More generally, neither the belief distribution nor its updating process is observed directly, and the counterfactual IRR exercise identifies valuation sensitivity rather than a structurally estimated causal effect. The framework should accordingly be regarded as an interpretation that organizes a coherent set of empirical regularities, not as a fully identified structural model.

These limitations point naturally to tests over longer histories, alternative required-return measures, and other market environments.

6. Conclusion

This paper develops a valuation-based interpretation of stock market volatility, required returns, and subsequent market performance by building directly on two results established in the implied-cost-of-capital literature. Pástor et al. (2008) show that an implied cost of capital inferred from prices and expected cash flows can provide informative variation about the latent conditional expected return. Li et al. (2013) subsequently show that an aggregate implied required return contains predictive information about future market returns. These results provide the starting point for the present analysis, which shifts attention from the informational content of the internal rate of return (IRR) to the broader valuation process in which it operates. It asks whether the joint dynamics of required returns and fundamentals can generate a credible level of market volatility, whether the economic level of the required return can be motivated without first being inferred from the market price, and whether information about subsequent returns is concentrated in the IRR or distributed across the broader valuation configuration.

The Monte Carlo analysis calibrates the joint behavior of the valuation inputs from the observed regime and lets annual market volatility emerge from the valuation mechanism. The resulting volatility is of the same order of magnitude as that subsequently observed, although realized volatility is not used as a calibration target. In this respect, the direction of inquiry is the reverse of that in Pástor et al. (2008): rather than specifying expected-return and volatility properties upstream and asking whether the IRR recovers the latent expected return, the present exercise takes the market IRR as one valuation input and asks what level of volatility emerges from the joint valuation dynamics. The second exercise introduces a further reversal of the usual implied-return logic. Rather than solving from prices and expected cash flows for an implied discount rate, it constructs a required-return reference (RRR) directly from expectations of reinvested and distributed income and the perception of risk. Its proximity to the market-implied IRR is

therefore economically informative: the IRR is recovered from market prices and expected cash flows, while RRR provides a rationale for the level of return investors require. Although IRR variation partly contributes to the simulated volatility entering this calculation, the result is not mechanical; substituting realized volatility leaves the RRR materially unchanged. The correspondence is therefore one of conditional quantitative consistency, not an accounting identity. Longer-run evidence further suggests that inflation may condition the nominal level of required returns.

The return analysis adds the dynamic dimension. Information about subsequent one-year market movements is not concentrated in the IRR or in a single common factor, but distributed across distinct Valuation and Cycle dimensions. More importantly, this structure survives expanding-window OOS evaluation: predictive content persists at the one-year horizon and the combined Valuation–Cycle specification produces substantial forecast gains. Their magnitude is markedly larger than the OOS improvements reported for aggregate IRR by Li et al. (2013), although differences in samples, horizons, and predictor construction preclude a direct performance comparison. The contrast nonetheless raises a substantive question left unresolved by the present short sample: whether the stronger predictability reflects the broader multidimensional valuation configuration or features specific to the recent period and annual horizon.

The evidence therefore points to a valuation process that is economically structured but revisable. Recurrent volatility can arise while profitability, growth, accumulation, and required returns vary within a prevailing valuation regime; more persistent changes in central valuation parameters may instead alter the regime through which future cash flows are capitalized. This distinction offers an alternative reading of excess volatility: market fluctuations need not consist solely of deviations from a uniquely identifiable fundamental value, but may reflect both variability within an economically grounded valuation configuration and revisions of that configuration. The valuation-convention framework developed in Clère (2026a) provides one possible interpretation of this combination of multidimensionality, local continuity, and revision. The present evidence does not identify that framework structurally, but it is consistent with a valuation process that remains anchored in economic quantities while adapting under fundamental uncertainty.

A longer history of market-implied IRR estimates would provide the most useful next test. It would strengthen the annual OOS analysis, permit predictive persistence to be examined across horizons, and allow RRR to be constructed over comparable historical periods and confronted

with the market-implied IRR across different valuation and nominal regimes. It would also make possible a stricter test of the volatility mechanism by calibrating it on one regime and evaluating it on another. Evidence from alternative implied-return measures and other developed markets would then indicate whether the structure identified here is specific to the recent U.S. experience or reflects a more general feature of market valuation under fundamental uncertainty.

References

- Bancel, F., Mittoo, U.R., 2014. The gap between the theory and practice of corporate valuation: Survey of European experts. *Journal of Applied Corporate Finance* 26, 106-117.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637-654.
- Campbell, J.Y., Shiller, R.J., 1988. Stock prices, earnings, and expected dividends. *Journal of Finance* 43, 661-676.
- Claus, J., Thomas, J., 2001. Equity premia as low as three percent? Evidence from analysts' earnings forecasts for domestic and international stock markets. *Journal of Finance* 56, 1629-1666.
- Clère, R., 2026a. Aggregate U.S. stock market returns across horizons: Structure, dynamics, and predictability. Working paper. SSRN 7353080.
- Clère, R., 2026b. Expected equity risk premium, stock market performance, and fundamental uncertainty. In: Bourghelle, D., Grandin, P., Jawadi, F., Rozin, P. (Eds.), *Financial Market Dynamics: Essays in the Honor of Pascal Alphonse*. World Scientific, forthcoming.
- Cochrane, J.H., 2011. Presidential address: Discount rates. *Journal of Finance* 66, 1047-1108.
- Easton, P.D., 2004. PE ratios, PEG ratios, and estimating the implied expected rate of return on equity capital. *Accounting Review* 79, 73-95.
- Fama, E.F., 1970. Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25, 383-417.
- Fama, E.F., 1991. Efficient capital markets: II. *Journal of Finance* 46, 1575-1617.
- Fama, E.F., French, K.R., 1988. Permanent and temporary components of stock prices. *Journal of Political Economy* 96, 246-273.
- Fernandez, P., 2003. Equivalence of ten different methods for valuing companies by cash flow discounting. Unpublished working paper. IESE Business School, University of Navarra.
- French, K.R., 1980. Stock returns and the weekend effect. *Journal of Financial Economics* 8, 55-69.
- French, K.R., Roll, R., 1986. Stock return variances: The arrival of information and the reaction of traders. *Journal of Financial Economics* 17, 5-26.
- Gebhardt, W.R., Lee, C.M.C., Swaminathan, B., 2001. Toward an implied cost of capital. *Journal of Accounting Research* 39, 135-176.

- Gordon, M.J., Shapiro, E., 1956. Capital equipment analysis: The required rate of profit. *Management Science* 3, 102-110.
- Han, Y., Lu, Y.J., Zhou, G., 2026. Macro financial trends and market expected returns. *Review of Asset Pricing Studies* 16, 241-282.
- Hoang, K., Cannavan, D., Huang, R., Peng, X., 2021. Predicting stock returns with implied cost of capital: A partial least squares approach. *Journal of Financial Markets* 53, 100576.
- Jagannathan, R., McGrattan, E.R., Scherbina, A., 2000. The declining U.S. equity premium. *Federal Reserve Bank of Minneapolis Quarterly Review* 24, 3-19.
- Jawadi, F., Ben Ameer, H., Bigou, S., Flageollet, A., 2022. Does the real business cycle help forecast the financial cycle? *Computational Economics* 60, 1529-1546.
- Ke, R., Liu, J., 2014. Estimating the equity risk premium using implied cost of capital. Unpublished working paper. Southern Methodist University and Cheung Kong Graduate School of Business.
- Keynes, J.M., 1921. *A Treatise on Probability*. Macmillan, London.
- Keynes, J.M., 1936. *The General Theory of Employment, Interest and Money*. Macmillan, London.
- Knight, F.H., 1921. *Risk, Uncertainty and Profit*. Houghton Mifflin, Boston, MA.
- LeRoy, S.F., Porter, R.D., 1981. The present-value relation: Tests based on implied variance bounds. *Econometrica* 49, 555-574.
- Li, Y., Ng, D.T., Swaminathan, B., 2013. Predicting market returns using aggregate implied cost of capital. *Journal of Financial Economics* 110, 419-436.
- Lintner, J., 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47, 13-37.
- Lo, A.W., 2004. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management* 30, 15-29.
- Mehra, R., Prescott, E.C., 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15, 145-161.
- Merton, R.C., 1974. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance* 29, 449-470.
- Mishra, D.R., O'Brien, T.J., 2019. Fama-French, CAPM, and implied cost of equity. *Journal of Economics and Business* 101, 73-85.
- Ohlson, J.A., 1995. Earnings, book values, and dividends in equity valuation. *Contemporary Accounting Research* 11, 661-687.
- Orléan, A., 1999. *Le pouvoir de la finance*. Odile Jacob, Paris.
- Pástor, L., Sinha, M., Swaminathan, B., 2008. Estimating the intertemporal risk-return tradeoff using the implied cost of capital. *Journal of Finance* 63, 2859-2897.
- Poterba, J.M., Summers, L.H., 1988. Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics* 22, 27-59.
- Roll, R., 1984. Orange juice and weather. *American Economic Review* 74, 861-880.

- Sharpe, W.F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19, 425-442.
- Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421-436.
- Simon, H.A., 1978. Rationality as process and as product of thought. *American Economic Review* 68, 1-16.
- Williams, J.B., 1938. *The Theory of Investment Value*. Harvard University Press, Cambridge, MA.

Appendix A. Monte Carlo Calibration and Simulation Procedure

The Monte Carlo exercise reported in Sections 3.3 and 4.1 generates one-year S&P 500 price changes from the valuation framework defined by **equations (1)–(3)**. The purpose of this appendix is to document the marginal distributions, dependence assumptions, simulation procedure, and principal distributional diagnostics underlying that exercise. The simulation is calibrated retrospectively on the recent historical regime; it should therefore be interpreted as a conditional exercise based on valuation inputs that are observable or estimable *ex ante*, rather than as a real-time forecast of the realized return path.

A.1. Calibration of stochastic inputs

Each simulation generates two successive valuation states, $t - 1$ and t . At each date, the price-to-book multiple is determined by normalized profitability, expected long-run growth, and the expected market IRR through **equation (2)**. Book equity evolves between the two states according to a simulated annual accumulation rate.

Table 8. Marginal Calibration of Valuation Inputs

Variable	Definition / construction	Distribution	Calibration
Normalized profitability, $\overline{ROE}$	Ten-year rolling average of annual aggregate S&P 500 return on equity	N(13.8%, 0.9%)	2014–2022
Expected market IRR, k	Fairness Finance aggregate market-implied IRR, adjusted for analyst-forecast optimism	N(6.4%, 0.5%)	2014–2022
Expected long-run growth, g'	75% of the IMF five-year forecast of nominal U.S. GDP growth	N(3.0%, 0.14%)	2014–2022

Notes: The second parameter of each normal distribution is the standard deviation. Normalized profitability is a ten-year trailing measure intended to capture a recurrent rather than contemporaneous level of market profitability; the 2014–2022 calibration window refers to the observation dates, with the first 2014 value averaging annual ROE from 2005 through 2014 inclusive. The expected IRR is the optimism-adjusted market measure described in

Section 3.2. The growth calibration follows the empirical rule used in the original valuation exercise to reproduce observed S&P 500 price-to-book multiples.

The marginal distributions of the three valuation inputs are calibrated primarily over 2014–2022, the period for which this study has a consistently constructed aggregate S&P 500 expected-IRR series. **Table 8** summarizes their construction and calibration.

Book-equity accumulation requires a separate calibration because both its central tendency and its dispersion depend materially on the historical window considered. Three configurations are therefore used. They distinguish assumptions about the expected rate of accumulation from assumptions about recurrent variability around that rate.

The optimistic-prudent configuration combines the 5.0% mean accumulation rate observed over 2014–2019 with the 7.3% standard deviation observed over 2004–2022. The pessimistic-myopic configuration uses both the 3.7% mean and the 2.9% standard deviation observed over 2014–2022. The pessimistic-prudent configuration retains the 2014–2022 mean while combining it with the broader 2004–2022 dispersion.

A.2. Contemporaneous dependence and intertemporal persistence

The stochastic inputs are not simulated independently. Book-equity accumulation, normalized profitability, expected growth, and the expected IRR exhibit economically meaningful contemporaneous dependence.

Table 9. Contemporaneous Correlations among Monte Carlo Inputs

	g_B	$\overline{ROE}$	k	g'
g_B	1.00	-0.27	0.15	0.55
$\overline{ROE}$	-0.27	1.00	-0.55	0.09
k	0.15	-0.55	1.00	-0.07
g'	0.55	0.09	-0.07	1.00

Notes: Correlations are estimated from the calibration data used in the original simulation. g_B denotes book-equity accumulation, $\overline{ROE}$ normalized profitability, k the optimism-adjusted expected market IRR, and g' expected long-run growth.

Independent draws would therefore generate combinations of valuation parameters that differ materially from the configurations observed in the calibration data. The contemporaneous correlation matrix imposed in the simulation is reported in **Table 9**.

The simulation also preserves the persistence observed between successive valuation states. For normalized profitability, the expected IRR, and expected growth, the correlations between $t - 1$ and t are imposed directly. These parameters are reported in **Table 10**.

Table 10. Intertemporal Persistence and Price-to-Book Validation

Variable	Correlation between successive states	Treatment in simulation
$\overline{ROE}_{t-1}, \overline{ROE}_t$	0.71	Imposed
k_{t-1}, k_t	0.56	Imposed
g'_{t-1}, g'_t	-0.08	Imposed
$(P/B)_{t-1}, (P/B)_t$ observed	≈ 0.86	Not imposed
$(P/B)_{t-1}, (P/B)_t$ simulated	≈ 0.90	Endogenous simulation outcome

Notes: The persistence of the price-to-book ratio is not imposed directly. The two multiples are reconstructed independently from the simulated underlying valuation inputs using **equation (2)**. The resulting correlation of approximately 0.90 is therefore an endogenous consequence of the dependence structure and is close to the approximately 0.86 correlation observed in the calibration data.

This distinction is important. The simulation does not assume an autoregressive process for the market multiple itself. Persistence in P/B arises from persistence and contemporaneous dependence in the economic variables from which the multiple is reconstructed.

A.3. Simulation procedure

For each of the three book-equity configurations, 20,000 joint simulations are generated. Each simulation produces two successive valuation states.

At each state $s \in \{t-1, t\}$, normalized profitability, expected long-run growth, and the expected IRR are generated subject to the marginal distributions reported in **Table 8**. Book-equity accumulation over the annual transition follows one of the three configurations described in Section A.1. Joint dependence is governed by the contemporaneous correlation structure reported in **Table 9**, and the intertemporal correlations reported in **Table 10**. The price-to-book multiple at each date is then reconstructed using **equation (2)**:

$$\frac{P_s}{B_s} = 1 + \frac{\overline{ROE}_s - k_s}{k_s - g'_s}.$$

Book equity evolves over the annual transition according to the selected accumulation configuration, with

$$g_B \equiv \ln\left(\frac{B_t}{B_{t-1}}\right).$$

The simulated annual logarithmic price return then follows the decomposition already established in **equation (3)**:

$$\ln\left(\frac{P_t}{P_{t-1}}\right) = \ln\left(\frac{B_t}{B_{t-1}}\right) + \ln\left(\frac{P_t/B_t}{P_{t-1}/B_{t-1}}\right).$$

The first component therefore reflects capital accumulation, while the second reflects the change in the valuation multiple generated by changes in normalized profitability, expected growth, and the market IRR. Because the two components are jointly simulated, they can reinforce or offset one another.

Dividend income is excluded from this simulated price return. It is introduced separately in Section 4.2.1, where expected capital accumulation, distributed income, and the recurrent-volatility adjustment are combined to construct the required-return reference (RRR).

Two distinct return statistics must also be distinguished. Simulated volatility is measured as the standard deviation of the 20,000 annual logarithmic price returns:

$$\sigma \equiv SD\left[\ln\left(\frac{P_t}{P_{t-1}}\right)\right].$$

By contrast, the “simulated log price return” reported in **Table 2** of the main text is calculated from the expected simulated future price:

$$R^L \equiv \ln\left[\frac{E(P_t)}{P_{t-1}}\right].$$

It therefore differs from the mean of the simulated logarithmic returns because: $\ln E(P_t) \neq E[\ln(P_t)]$. The corresponding simple annual price return is defined as

$$R \equiv \frac{P_t}{P_{t-1}} - 1.$$

A.4. Distributional diagnostics

The original implementation reports detailed distributional statistics for the optimistic-prudent configuration, as shown in **Table 11**. This configuration is particularly informative because its broader dispersion provides the recurrent-volatility estimate subsequently used in Section 4.2. The headline results for all three configurations are reported in **Table 2** of the main text. The table clarifies the distinction between the different return measures.

Table 11. Distributional Statistics for the Optimistic-Prudent Simulation

Statistic	$\ln\left(\frac{B_t}{B_{t-1}}\right)$	$\Delta \ln(P/B)$	$\ln\left(\frac{P_t}{P_{t-1}}\right)$	P/B_t	P/B_{t-1}	Simple return R
Draws	20,000	20,000	20,000	20,000	20,000	20,000
Mean	5.03%	0.24%	5.28%	3.35	3.33	6.42%
Median	5.04%	-0.07%	5.04%	3.22	3.22	5.17%
Standard deviation	7.29%	10.24%	13.71%	0.81	0.73	14.78%
Variance	0.53%	1.05%	1.88%	0.66	0.54	2.18%
Skewness	-0.01	0.17	0.11	1.29	1.16	0.55
Kurtosis	3.01	3.08	3.07	6.75	6.03	3.62
Coefficient of variation	1.45	42.00	2.60	0.24	0.22	2.30
Minimum	-26.18%	-36.93%	-46.66%	1.35	1.52	-37.29%
Maximum	38.79%	45.06%	66.08%	11.69	10.26	93.64%
Range	64.97%	81.99%	112.75%	10.34	8.74	130.93%
Standard error of mean	0.052%	0.072%	0.097%	0.006	0.005	0.104%

Notes: $\Delta \ln(P/B) \equiv \ln[(P_t/B_t)/(P_{t-1}/B_{t-1})]$. The coefficient of variation for this variable is mechanically large because its simulated mean is close to zero and has no substantive interpretation here.

The mean simulated logarithmic price return is 5.28%, whereas the mean simple price return is 6.42%. The latter implies a return based on the expected simulated future price of

$$R^L = \ln(1 + 0.0642) \approx 6.22\%, \quad (16)$$

which corresponds, after rounding, to the 6.2% simulated log price return reported for the optimistic-prudent configuration in **Table 2**. The 13.71% standard deviation of annual logarithmic returns is the simulated recurrent volatility subsequently used in the RRR calculation in Section 4.2.1. Realized S&P 500 return volatility is not imposed as a calibration target, and the expected-IRR series was constructed for valuation purposes rather than fitted to reproduce market volatility. The simulated volatility therefore emerges from the valuation system and remains close to realized volatility, although the exercise should be understood as a conditional consistency test rather than as an independent validation. The output correlations reported in **Table 12** provide an additional check on the internal structure of the simulation. Two features are noteworthy. First, capital accumulation and the change in the valuation multiple are positively correlated at 0.20 rather than independent. The annual logarithmic price return consequently reflects both dimensions, correlating at 0.68 with book-equity accumulation and at 0.85 with the change in the multiple. Second, the simulated correlation of 0.90 between successive price-to-book ratios emerges without imposing persistence directly on the multiple, providing an internal check on the dependence structure described in Section A.2.

Table 12. Correlations among Simulated Outcomes, Optimistic-Prudent Configuration

	$\ln\left(\frac{B_t}{B_{t-1}}\right)$	$\Delta \ln(P/B)$	$\ln\left(\frac{P_t}{P_{t-1}}\right)$	P/B_t	P/B_{t-1}	R
$\ln(B_t/B_{t-1})$	1.00	0.20	0.68	0.12	0.04	0.68
$\Delta \ln(P/B)$	0.20	1.00	0.85	0.39	-0.05	0.85
$\ln(P_t/P_{t-1})$	0.68	0.85	1.00	0.35	-0.02	1.00
P/B_t	0.12	0.39	0.35	1.00	0.90	0.36
P/B_{t-1}	0.04	-0.05	-0.02	0.90	1.00	-0.01
R	0.68	0.85	1.00	0.36	-0.01	1.00

Notes: Values are pairwise correlations across the 20,000 simulated observations. The simple return is an exact monotonic transformation of the logarithmic return; their sample correlation rounds to 1.00 in the simulation.

The simulations were originally implemented using Oracle Crystal Ball within Microsoft Excel. The software choice is not part of the economic specification. The exercise can be reproduced in any simulation environment capable of generating correlated draws subject to the marginal distributions, contemporaneous correlations, and intertemporal dependence parameters reported above.

Appendix B. Glossary of Key Terms and Notation

The following glossary summarizes the main concepts and notation specific to the analysis. Standard financial and econometric terms are omitted.

B-A. Valuation and required-return concepts

Term / notation	Meaning and role in the paper
Latent required return	Unobservable ex-ante return required by equity investors. It is the underlying economic object approached through the market IRR and the RRR by two distinct empirical routes.
Market IRR, k_t	Market-implied ex-ante required return obtained by solving a valuation relation from observed prices and expected future fundamentals (cash flows or excess returns). IRR is used generically for implied required-return measures, including measures commonly referred to as implied costs of capital (ICC).
Required-return reference (RRR), k_t^{L*}	Economically motivated reference for the latent required return derived from expected capital accumulation, distributed income, and recurrent market volatility. k_t^* denotes its continuously compounded form. RRR is not an equilibrium condition and need not coincide mechanically with the market IRR.
Book-equity accumulation, $g_{B,t}$	Growth in the aggregate book-equity base. It provides the accumulation component of market-price growth and plays a central role in both the valuation decomposition and the construction of RRR.
Recurrent market volatility, σ_t^*	Anticipated market volatility associated with the prevailing valuation configuration. In the Monte Carlo exercise, it emerges from the joint dynamics of the valuation inputs rather than being imposed as a volatility target.
IRR–inflation spread	Difference between the estimated market IRR and consumer-price inflation. Its relative stability over recent decades is examined as evidence that inflation may condition the nominal level of the required return, not as a structural inflation-pricing rule.

Term / notation	Meaning and role in the paper
Conditional consistency	Evidential status assigned to quantitative coherence between separately constructed elements of the valuation framework without claiming independent structural identification or causality. The proximity of RRR and the market IRR is interpreted in this sense.

B-B. Valuation dynamics and empirical configuration

Term / notation	Meaning and role in the paper
Valuation configuration	Multidimensional state formed by the economically relevant valuation variables and their relations at a given date. It need not be reducible to a single required-return or valuation measure.
Valuation–Cycle configuration (V–C)	Empirical representation of the valuation state combining Valuation information, including the market IRR and valuation multiples, with Cycle information describing recent market, profitability, and accumulation dynamics.
Valuation regime	Relatively persistent economic and probabilistic environment within which valuation variables fluctuate. Recurrent variability occurs within a regime; sufficiently persistent changes in central parameters may instead indicate a change of regime.
Recurrent variability	Market variability generated while the valuation rule and the probabilistic structure of the prevailing regime remain unchanged. It is distinguished from persistent revisions capable of altering the regime itself.
Local persistence	Temporary continuity of the valuation configuration or convention that allows current valuation conditions to remain informative about subsequent market outcomes without requiring globally invariant relations.
Valuation revision	Change in conditional valuation resulting from revised beliefs, reassessment of known uncertainty, or changes in the set of conceivable representations. It is conceptually distinguished from recurrent variability within the prevailing regime.

B-C. Valuation-convention framework under fundamental uncertainty

Term / notation	Meaning and role in the paper
Valuation convention	Locally persistent but revisable set of representations, valuation procedures, and belief weights through which market participants organize valuation under fundamental uncertainty.
θ	One conceivable representation of future economic and financial conditions, incorporating expectations about future cash flows, required returns, and relevant macro-financial conditions.
Θ_t	Universe of representations regarded as possible at date t . Under fundamental uncertainty, this set need not be fixed and may itself change as new situations become conceivable.
$b_t(\theta)$	Subjective weight assigned at date t to representation θ . These weights represent working beliefs rather than objective probabilities over an exhaustively known state space.
$V(\theta)$	Conditional intrinsic value associated with representation θ . Prevailing market valuation is represented as the belief-weighted aggregation of these conditional values.
Bayesian updating	Parsimonious representation of revisions in the relative belief weights $b_t(\theta)$ as new information arrives, conditional on a sufficiently stable Θ_t . It represents sequential learning rather than literal Bayesian calculation by investors.
Fundamental uncertainty	Uncertainty for which future states, their probabilities, or even the relevant set of conceivable representations are not fully known in advance. It allows Θ_t itself to change over time.