

The uneasy case for active management

Dirk Broeders & Kristy Jansen

ABSTRACT¹

Passive investing is compelling, but incomplete. Real-world investors face changing market regimes, fat-tailed risks and path-dependent wealth outcomes that are not fully captured by standard models. These features create a role for active management focused on adaptation and robustness: dynamic risk control, regime awareness and protection against severe losses. But they do not make active management easy. Markets remain hard to forecast, costs matter and apparent skill may reflect hidden tail risk. The case for active management is therefore real, but uneasy.

INTRODUCTION

The central economic problem in investment management is simple to state but difficult to solve: how should investors allocate capital when the future distribution of returns is unknown? One classic answer is passive investing. A passive investor holds every security in a pre-determined market in proportion to its market capitalisation (Sharpe, 1991). This view grew out of mean-variance optimisation in the 1950s, the Capital Asset Pricing Model in the 1960s and the efficient market hypothesis as formalised around 1970. Classical mean-variance analysis evaluates portfolio choices in terms of expected returns, variances and covariances. It has been enormously useful: it has improved diversification, lowered costs and challenged the

performance of high-fee active strategies (Broeders, van Oord and Rijsbergen, 2019).

There is also another powerful arithmetic behind the passive case. Active investors as a group hold the market. Before costs, the average actively managed euro must therefore earn the market return; after fees and trading costs, it must underperform the average passively managed euro (Sharpe, 1991). This arithmetic does not imply that every active manager must underperform. It implies that active management is a zero-sum game before costs and a negative-sum game after costs. For active management to be worthwhile, some investors must possess skill, access, discipline or risk-bearing capacity that others lack. The traditional Fundamental Law of Active Management, associated with Grinold and Kahn (1999), gives the active manager's version of this logic:

$$IR = IC \times \sqrt{\text{breadth}}$$

The information ratio, (*IR*), rises with forecasting skill, measured by the information coefficient, (*IC*) and with the number of independent active bets, measured by breadth. The formula is elegant and intuitive. If a manager has a positive edge and can apply it across many independent opportunities, active management can add value. But this formula presumes that

Dirk Broeders

Professor of Finance at the School of Business and Economics of Maastricht University and Senior Lead Financial Risk Expert at the European Central Bank



Kristy Jansen

Assistant Professor of Finance and Business Economics at the University of Southern California Marshall School of Business and a Research Affiliate at CEPR



the environment is sufficiently stable for skill to be estimated, that active bets are genuinely independent, and that risk can be adequately summarised by variance.

Financial markets do not consistently satisfy these conditions. Return distributions are non-stationary: parameters change, regimes shift and correlations rise when diversification is most needed. Returns are also fat-tailed: crashes and large dislocations occur more often than a normal distribution would imply, and a small number of extreme events can significantly affect realised wealth. Wealth dynamics are non-ergodic: investors experience one path through time, not the average of all possible paths. A severe loss early in the horizon can permanently impair the ability to meet future objectives, even if the long-run average return looks attractive.

The classical passive prescription assumes that investors share the same opportunity set and differ mainly in their mix of the market portfolio and the risk-free asset. In practice, they have different liabilities, different funding and liquidity constraints, and different non-tradable exposures. Nor can anyone hold the complete market portfolio, which includes inaccessible or illiquid assets. Strategic asset allocation should therefore reflect each investor's liabilities, endowments and feasible opportunity set. This limitation of CAPM supports tailored, liability-aware investing without implying that active security selection generates alpha.

This is especially relevant for institutional investors such as pension funds, insurers, sovereign wealth funds and endowments. Their objective is not merely to earn a high risk-adjusted return. It is to meet future liabilities, preserve solvency, manage drawdowns and avoid losses from which recovery may be impossible. For such investors, risk is not just volatility around an expected return. It is the possibility that the realised path of returns prevents the investor from achieving the investment objective. Further, institutional investors often also face challenges in balancing financial returns with social impact objectives. Through deliberative processes, investor democracy enables stakeholders to navigate these trade-offs collectively (Bauer et al., 2026).

PASSIVE IS POWERFUL

This article argues that these facts do not invalidate passive investing, nor do they vindicate active management in general. They change the question. The issue is not whether active managers can reliably predict next period's returns. It is whether investors need adaptive decision-making to survive in a world where return distributions move, extremes dominate and paths matter. In that world, the more defensible role for active management is not heroic stock-picking, but regime awareness, dynamic risk control, explicit tail-risk management and robust

position sizing. At the same time, it remains an uneasy case: the same features that create the need for active management also make it difficult to execute successfully.

The rest of the article first explains why non-stationarity, fat tails and path dependence make fixed or mechanically maintained allocations incomplete. It then shows how these same features create a role for adaptive risk management, before explaining why that role remains difficult to implement in practice.

WHY PASSIVE INVESTING LEAVES IMPORTANT PORTFOLIO CHOICES UNRESOLVED

In this article, we use "passive investing" as an umbrella term for rules-based investment approaches that do not seek to exploit discretionary views about mispricing or changing investment opportunities. This includes index-tracking strategies, such as cap-weighted ETFs, whose security weights change mechanically with prices, corporate actions and index composition. It also includes buy-and-hold strategies, in which portfolio weights drift with realised returns after the initial allocation. These approaches differ from one another: cap-weighted indexing is not a fixed portfolio, and buy-and-hold is not the same as benchmark replication. The issue is therefore not passive investing as such, but the stronger assumption that a fixed or mechanically maintained allocation is sufficient when return distributions are non-stationary, fat-tailed and path-dependent.

The case for low-cost passive investing is compelling. It does not, however, determine how much market risk an investor should bear, which liabilities should be hedged or how much liquidity should be maintained. These decisions depend on the investor's objectives and constraints and may change as market conditions evolve. Non-stationarity, fat tails and path dependence therefore warrant closer examination of asset allocation and risk management.

NON-STATIONARITY: PARAMETERS CHANGE

Non-stationarity means that the data-generating process for asset returns changes over time. Expected returns, volatilities, correlations, factor loadings and higher-order co-moments may drift, jump or vary across regimes (Ang and Bekaert, 2002; Broeders and Jansen, 2020). In such an environment, long historical samples are not necessarily reliable guides to future return distributions. This is a practical manifestation of the Lucas Critique in financial markets. Financial markets are prone to boom-bust cycles, in which periods of apparent stability encourage leverage and risk-taking until a Minsky moment exposes the fragility that had accumulated beneath the surface (Kocken, 2012).

This creates a challenge for portfolio allocations calibrated using historical estimates of expected returns, variances and covariances. Mean-variance optimization can incorporate updated estimates, but maintaining allocations based on outdated inputs may leave portfolios poorly aligned with prevailing risks. In practice, key parameters are moving targets. A covariance matrix estimated over several decades may combine inflationary and disinflationary

regimes, financial repression and liberalisation, crises and calm markets. The estimate may be statistically significant and yet misrepresent the risks that matter today. Long historical samples can nevertheless remain informative when structural breaks are explicitly modelled (Pástor and Stambaugh, 2001).

Expected returns and factor exposures are similarly conditional on the prevailing environment. Risk premia embedded in equity valuations, bond yields and currency exchange rates may compress or widen as macroeconomic conditions, policy regimes and market structure change. Broad diversification strategies that are beneficial from the perspective of individual institutions may become a source of systemic risk when many institutions hold similar exposures and market stress prompts them to de-risk simultaneously (“diversification paradox”, Beale et al., 2011). Factor models estimated under the assumption of time invariance may misstate the contemporary risks embedded in traded portfolios.

FAT TAILS: EXTREMES DOMINATE LONG-RUN OUTCOMES

Fat tails mean that asset returns exhibit extreme observations more frequently than a normal distribution would predict (Gabaix et al., 2006). Large price moves, abrupt regime changes and volatility spikes are not rare anomalies; they are recurring features of financial markets. Crashes occur more often, and with greater severity, than standard Gaussian models imply.

The consequence is that long-run wealth can be shaped disproportionately by a small number of extreme events. Ordinary daily fluctuations matter less than large losses and, in some cases, very large gains. This concentration also appears across securities: a small minority of individual stocks accounts for a disproportionate share of the aggregate stock market’s long-run wealth creation, while most stocks underperform Treasury bills (Bessembinder, 2018).

Standard deviation of returns (“volatility”) compresses risk into an average measure of dispersion. In fat-tailed markets this can be misleading, because the main concern is not ordinary fluctuation but the probability and size of severe losses (or gains). A portfolio with acceptable volatility may still expose an investor to drawdowns large enough to force asset sales, breach constraints or permanently impair wealth (Goldstein and Taleb, 2007).

NON-ERGODICITY: A SINGLE PATH MATTERS MORE THAN THE ENSEMBLE

Non-ergodicity means that an investor experiences one realised path through time, not the average across all possible paths. In an ergodic system, the ensemble average and the time average coincide: what happens on average across many realisations is also a good guide to what one investor experiences over time. Financial wealth, however, does not behave in this way. It accumulates multiplicatively and is exposed to absorbing states such as default, forced selling, liquidation or breach of regulatory constraints. A sufficiently large adverse shock can permanently alter, or even terminate, the investment path. A common mistake in investment analysis is to treat the average outcome across many possible paths as equivalent to the realised

path of a single investor. This is called the ergodic illusion. For institutions with liabilities, leverage or solvency constraints, a portfolio that looks attractive in expectation may still be unacceptable if it creates a material risk of forced selling, deleveraging or irreversible loss (Shleifer and Vishny, 1997).

SURVIVAL MATTERS MORE THAN AVERAGES

This is also why the sequence of returns matters. For investors with cash in and out flows, liabilities, leverage or binding constraints, two return paths can have the same arithmetic average return and volatility, yet lead to very different outcomes if losses occur early rather than late. A large early loss reduces the capital base available to meet withdrawals, absorb further shocks or satisfy collateral and solvency requirements. The advice to “ride out” short-term volatility assumes an idealised horizon and the absence of binding constraints. While conventional wisdom may suggest that stocks are less risky for long-term investors, Bodie (1995) demonstrates that the cost of insuring against stock shortfalls relative to the risk-free return actually increases with the length of the investment horizon. In practice, investors can benefit from long-run risk premia only if interim losses do not force them to sell assets, breach constraints or abandon the strategy (Liu and Longstaff, 2004).

HOW THIS STRENGTHENS THE CASE FOR ACTIVE MANAGEMENT

If return distributions are non-stationarity, fat tailed and non-ergodic a passive approach may be incomplete. This argument should not be read as a claim that these features alone guarantee excess returns. In information-efficient markets, successful active management still requires an informational, analytical, behavioral, or structural edge over the marginal investor. Non-stationarity, for instance, does not remove this requirement. Nor does it imply that pure cap-weighted indexing is necessarily naive: cap-weighted indices adjust automatically as prices change and index composition evolves according to rules. To the extent that markets are informationally efficient, such an index may be less exposed to errors from naive historical estimation than an active strategy calibrated on past returns. Even when prices efficiently reflect available information, investors must still make active choices about asset-class exposures, benchmarks, liquidity, hedging, rebalancing, and tail-risk protection. The case made here is therefore less a claim about “alpha” and more an argument for adaptive portfolio construction and risk management under changing market conditions and depending on investor-specific risk aversion and constraints. Horizon, liabilities, liquidity needs, solvency requirements, and mandate restrictions on tracking error, leverage, derivatives, or cash holdings determine how much flexibility an investor has to adapt exposures as conditions change. We explore four implications.

TIME-VARYING RISK MANAGEMENT

Risk management relies on three broad tools whose effectiveness varies over time: diversification, hedging and risk transfer.

Diversification mitigates idiosyncratic risk by spreading exposures across assets, sectors, regions or risk factors whose returns are not perfectly correlated. Hedging uses symmetric instruments, including futures and swaps, to reduce or neutralise specific exposures. Risk transfer uses asymmetric instruments, including options and insurance contracts, to protect against risks that cannot be reliably diversified or hedged. In exchange for a premium, it transfers the financial consequences of low-probability, high-severity events to a third party. The effectiveness of all three depends on assumptions about correlations, hedge relationships, liquidity, and counterparty capacity – relationships that can change across regimes and may weaken during market stress.

Market conditions are not constant. Correlations, valuations and risk premia can shift across regimes, so a portfolio that is robust in one environment may be fragile in another. As Box 1 illustrates, in a regime-switching world portfolio risk comes not only from volatility within a given regime, but also from uncertainty about which regime will prevail. This strengthens the case for dynamic asset allocation: adjusting exposures to equities, bonds, cash and other asset classes as conditions change, rather than relying mechanically on a fixed mix calibrated to historical averages. Moreira and Muir (2017) find that scaling down risk exposure when volatility is high improves Sharpe ratios across several equity factors and in the currency carry trade.

The stock-bond correlation is a clear example. During the roughly two decades following the turn of the millennium, government bond returns typically moved inversely to equity returns, offering a diversification cushion during stock market downturns. More recently, stock-bond co-movement has increased, with global bond-stock betas turning positive in 2023-2025 (Campbell, Pflueger and Viceira, 2025). Since the diversification value of bonds depends on their covariance with equities, such shifts can materially alter the optimal stock-bond allocation (Campbell and Viceira, 2002; Campbell, Pflueger and Viceira, 2020; Pflueger, 2025). This is not primarily an argument about informational advantage or alpha generation, but a hedging-demand argument: when the state variables that determine investment opportunities change, the portfolio that best serves an investor's risk-management objective may also change.

REGIME SHIFTS AND MISPRICING

Market regimes are shaped by the economic, policy and institutional environment. Monetary policy, technology, regulation, investor behavior and financial conditions evolve, sometimes gradually and sometimes abruptly. Prices may adjust only gradually as investors interpret new information and reassess its implications. Information may diffuse slowly, investors may extrapolate the recent past, and institutional constraints may delay portfolio adjustment. Institutional

mandates, liquidity constraints, and risk limits prevent large funds from reallocating portfolios instantly (Duffie, 2010).

This creates the possibility of temporary mispricing. Credit spreads can become too tight before a default cycle, sectors can become crowded after strong performance, and neglected assets can trade below fundamental value. In principle, active managers can exploit such opportunities through timing, rotation or security selection if they can identify regime shifts in real time. Active management adds value if judgment improves decisions after costs and errors.

TAIL-RISK HEDGING

Fat-tailed markets make severe losses more likely than standard models imply. For investors with finite horizons, liabilities or solvency constraints, protection against large drawdowns may therefore be worth accepting, even at the cost of lower expected returns in normal times. Tail-risk management can take several forms: protective options, dynamic hedging, lower exposure to risky assets during periods of stress, or option-based strategies designed to pay off in severe equity or credit sell-offs. Modelling techniques that allow for time-varying volatility and fat tails, such as models from the GARCH family, can provide additional information by showing how portfolio losses behave in the tail of the return distribution. Such analysis may support precautionary adjustments to portfolio weights or hedging exposures. None is free. Options carry premia, dynamic hedging can fail when prices gap and liquidity disappears, and crisis strategies may disappoint if the next stress episode differs from the scenario for which they were designed.

These choices are inherently active. They require a judgment about which risks to insure, how much protection is worth paying for, and when protection is most valuable. A passive portfolio that ignores tail risk is not neutral. It simply accepts the full downside exposure embedded in the market.

POSITION SIZING

The size of a position matters as much as its expected return. An attractive risk premium can still be unacceptable if the exposure is large enough to create ruin, forced selling or an unrecoverable drawdown. Conversely, excessive caution may leave too little growth to meet long-term objectives.

The Kelly criterion formalises this trade-off by linking position size to the maximisation of long-term geometric wealth (Kelly, 1956). In principle, it favours larger allocations when expected returns are high and risks are low. In practice, however, full Kelly assumes a degree of knowledge investors rarely possess. Because the approach is highly sensitive to parameter estimation risk, overestimating expectations or understating tail risk triggers aggressive overbetting.

This is why fractional Kelly rules, such as half-Kelly or quarter-Kelly, are often more robust. They sacrifice some return in exchange for materially lower drawdowns and greater protection against model error. The broader lesson is that active management

is not only about selecting securities or timing markets. It is also about sizing exposures so that the investor survives. In a world where distributions move, paths matter and extremes dominate, survival is a precondition for compounding.

WHY THIS DOESN'T AUTOMATICALLY VINDICATE ACTIVE MANAGERS

The fact that non-stationarity, path dependency and fat tails create a role for active management does not mean that active managers can reliably add value. Active management still faces substantial challenges for three reasons.

First, forecasting is difficult. The same features that create potential opportunities also make those opportunities difficult to exploit. What looks like a new environment may be noise; what looks like noise may be the beginning of a new regime. Regime shifts, changing parameters and extreme events increase estimation error, and historical relationships may break down just when they appear most useful. A valuation signal may identify a genuine opportunity, or it may be the first sign of a deteriorating regime. A hedge may protect against one scenario while losses arise from another; alternatively, it may become too expensive precisely when it is most needed.

Second, costs remain decisive. Active management is not evaluated before costs, but after fees, trading costs and market impact. Even if some managers identify opportunities, those opportunities must be large enough to survive implementation. The arithmetic is unforgiving: before costs, active investors in aggregate hold the market; after costs, the average active investor must lag a low-cost passive alternative with comparable exposures.

Third, complexity is not the same as opportunity. Markets may be adaptive and shaped by frictions, constraints and investor flows, but that does not make them easy to trade. Apparent mispricings may reflect compensation for hidden risks, crowding, leverage, illiquidity or exposure to rare crash states. A strategy that looks skilful in normal times may simply be earning a premium for bearing risks that are revealed only in stress.

The bar for active management is therefore high. A credible strategy needs a clear economic rationale, evidence that the opportunity can be exploited after costs, robust implementation and disciplined risk management. Without these, the existence of opportunity in theory will not translate into reliable value for investors.

DEMAND-DRIVEN PRICING, INVESTOR FLOWS AND BENCHMARK SELECTION

A further, relatively recent, insight is that investor behavior also matters. Prices are not shaped only by fundamentals and macroeconomic shocks, but also by mandates, flows and benchmark-driven investor behavior demand. This insight is central to demand-based asset pricing. Traditional financial theory assumes highly elastic markets, where prices primarily

Box 1: Two Sources of Risk in a Regime-Switching World

Investors face two distinct sources of risk. The first is risk within a regime: even if the prevailing economic environment (*S*) were known, asset returns (*R*) would remain uncertain because of shocks specific to that environment. The second is regime uncertainty: investors may not know which economic environment will prevail, and different regimes can imply different expected returns and volatilities. Total variance can be decomposed into these two sources using the law of total variance:

$$\text{Var}(R) = \underbrace{\mathbb{E}[\text{Var}(R/S)]}_{\text{risk within regimes}} + \underbrace{\text{Var}(\mathbb{E}[R/S])}_{\text{risk from regime uncertainty}}$$

Regime uncertainty becomes particularly important when expected returns differ substantially across regimes, even if returns are relatively stable within each regime. In the case of two possible regimes, *{H, L}*, portfolio variance can be written as:

$$\text{Var}(R) = \underbrace{p\sigma_H^2 + (1 - p)\sigma_L^2}_{\text{risk within regimes}} + \underbrace{p(1 - p)(\mu_H - \mu_L)^2}_{\text{risk from regime uncertainty}}$$

where *p* is the probability of regime *H*, *μ_H* and *μ_L* are the portfolio's expected returns in the two regimes, and *σ_H* and *σ_L* are the corresponding standard deviations. Intuitively, regime uncertainty is greatest when both regimes are equally plausible (*p* = 1/2) and the portfolio would perform very differently in each.

reflect fundamental information. The demand system approach challenges this view by showing that prices can be materially affected by the mandates, constraints and latent demand shocks of different investor groups (Kojien and Yogo, 2019).

FLOWS MOVE MARKETS, TOO

Because aggregate investor demand is inelastic, large-scale capital reallocations – driven by institutional mandates, shifting preferences or behavioral sentiment – can have a substantial impact on the prices of the underlying securities (Jansen, 2025). Gabaix and Kojien (2021) formalise this mechanism in the inelastic markets hypothesis: when aggregate demand is insufficiently elastic, investor flows can move market prices materially. In this framework, a significant share of cross-sectional stock return variation and macroeconomic volatility is generated not by changes in cash-flow fundamentals alone, but also by who owns the assets and how they trade.

For active managers, information about demand and flows may therefore be valuable, especially when such flows are partly predictable and lead to systematic buying or selling. For example,

Pavlova and Sikorskaya (2023) show that increases in benchmarking intensity – that is, the extent to which investors manage portfolios closely against benchmark indices – generate inelastic demand, raise prices and reduce subsequent stock returns over horizons of up to five years.

This also complicates the meaning of passive investing. A passive strategy may minimise tracking error to an index, but the choice of that index is itself an active decision. Broeders and De Haan (2020) argue that a benchmark is not a neutral representation of the market. It reflects policy choices about objectives, risk exposures and eligible assets. Benchmark selection therefore shapes long-term outcomes and is one of the most important active decisions an investor makes.

AN UPDATED VIEW OF THE FUNDAMENTAL LAW OF ACTIVE MANAGEMENT

The traditional Fundamental Law of Active Management is elegant and useful, but it becomes less complete once we move beyond the assumptions of stable distributions, independent bets and variance-based risk. A robust active-management framework would need to recognise three additional points.

First, forecasting skill is regime-dependent. A signal is not valuable in isolation; its value depends on the environment in which it is applied. Buying high-yield bonds after spreads widen can be profitable during a normal slowdown, but dangerous if the economy is heading into a deep recession or banking crisis. The skill is therefore not merely observing that spreads are wide. It is understanding what kind of environment produced the widening, and whether the historical relationship on which the signal relies is still relevant.

Second, breadth should be measured economically, not mechanically. A large number of positions does not necessarily imply a large number of independent bets. A manager may invest in 100 different stocks, but if these are all small-cap growth companies in the same part of the US technology sector, the portfolio may in economic terms represent only one dominant exposure. Raw breadth can be large while effective breadth is small. What matters is not the number of securities, signals or trades, but the number of genuinely independent sources of return.

Third, apparent skill must be assessed after allowing for tail risk. Strategies that generate steady returns in normal times may be earning compensation for exposure to rare but severe losses. Currency carry illustrates the point. The currency carry trade – borrowing in low-yielding currencies and investing in high-yielding ones – can produce attractive returns in calm markets but suffer sharp losses when investors rush into safe-haven currencies during crises. Such losses reflect exposure to rare currency crashes that may be underrepresented in historical samples. What looks like persistent skill may partly be compensation for bearing tail risk.

This does not overturn the logic of the Fundamental Law. Rather, it changes how its ingredients should be interpreted in

real-world markets. Skill should be understood as regime-aware skill, breadth as effective and economically independent breadth, and risk as including exposure to severe drawdowns, not only ordinary volatility. In this view, active management is not simply about scaling many small forecasting edges. It is about building adaptive and robust strategies that can survive regime shifts, control tail risk and preserve capital.

ADAPT, HEDGE, SURVIVE

Advances in artificial intelligence may help implement this more adaptive approach. Kojien and Levy (2026) find that optimised agentic AI systems raise the R^2 for explaining contemporaneous stock returns around earnings announcements from 8% for standard benchmarks to nearly 20%, using only information available at announcement time. Relatedly, Gabaix et al. (2025) show that high-dimensional asset representations learned from portfolio holdings explain stock-return comovement better than a benchmark based on standard firm characteristics. These tools could help managers identify common exposures and assess whether diversification is economically meaningful. More broadly, better information processing and exposure measurement could support regime awareness and the assessment of emerging tail risks. These broader applications remain to be established, however: explaining contemporaneous returns and their comovement does not itself demonstrate an ability to anticipate regime shifts or extreme losses.

CONCLUSION: THE UNEASY CASE FOR ACTIVE MANAGEMENT

If return distributions shift over time, if wealth paths matter, and if extreme events dominate long-horizon outcomes, then a fixed buy-and-hold allocation to the market portfolio is an incomplete solution to the portfolio problem. Investors do not live in the long-run average. They experience a particular sequence of returns, with liabilities, constraints and finite horizons.

This provides an economic rationale for active, adaptive management: dynamic risk control, regime awareness, tail-risk management and careful position sizing. But it is not an unconditional endorsement of traditional active management. The fact that markets create scope for skill does not mean that most managers possess it, or can deploy it successfully after costs, estimation error and implementation frictions.

The conclusion is therefore modest. Passive investing remains a powerful benchmark and a useful core building block. But permanently holding a capitalisation-weighted index while ignoring changing risks is hard to justify. A stronger approach combines low-cost passive exposure with disciplined adaptation, whether through systematic allocation rules, risk overlays or carefully selected active managers to create robust investment strategies. That is the uneasy case for active management.

References

- Ang, A., & G. Bekaert (2002). International asset allocation with regime shifts. *Review of Financial Studies*, 15(4), 1137–1187.
- Bauer, R., E. Cooper, B. van der Kroft, & P. Smeets (2026). Investor democracy. MIT Sloan Research Paper No. 7357-26.
- Beale, N., et al. (2011). Individual versus systemic risk and the Regulator’s Dilemma. *Proceedings of the National Academy of Sciences*, 108(31), 12647–12652.
- Bessembinder, H. (2018). Do stocks outperform Treasury bills? *Journal of Financial Economics*, 129(3), 440–457.
- Bodie, Z. (1995). On the risk of stocks in the long run. *Financial Analysts Journal*, 51(3), 18–22.
- Broeders, D., A. van Oord, & D. Rijsbergen (2019). Does it pay to pay performance fees? Empirical evidence from Dutch pension funds. *Journal of International Money and Finance* 93, 299–312.
- Broeders, D., & L. de Haan (2020). Benchmark selection and performance. *Journal of Pension Economics and Finance* 19(4), 511–531.
- Broeders, D., & K. Jansen (2020). Pension funds and drivers of heterogeneous investment strategies. Available at SSRN: <https://ssrn.com/abstract=3446239>.
- Campbell, J. Y., & L. M. Viceira (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.
- Campbell, J. Y., C. Pflueger, & L. M. Viceira (2020). Macroeconomic Drivers of Bond and Equity Risks. *Journal of Political Economy*, 128(8), 3148–3185.
- Campbell, J. Y., C. Pflueger, & L. M. Viceira (2025). Bond-Stock Comovements. NBER Working Paper 34323, revised June 2026.
- Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *The Journal of Finance*, 65(4), 1237–1267.
- Gabaix, X., P. Gopikrishnan, V. Plerou, & H. E. Stanley (2006). Institutional investors and stock market volatility. *Quarterly Journal of Economics*, 121(2), 461–504.
- Gabaix, X., & Koijen, R. S. J. (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. NBER Working Paper No. 28967.
- Gabaix, X., Koijen, R. S. J., Richmond, R. J., & Yogo, M. (2025). Asset embeddings. Becker Friedman Institute Working Paper No. 2025-51.
- Goldstein, D. G., & Taleb, N. N. (2007). We don’t quite know what we are talking about when we talk about volatility. *Journal of Portfolio Management*, 33(4), 84–86.
- Grinold, R. C., & Kahn, R. N. (1999). *Active portfolio management: A quantitative approach for producing superior returns and controlling risk* (2nd ed.). McGraw-Hill.
- Jansen, K.A.E. (2025). Long-term Investors, Demand Shifts, and Yields. *Review of Financial Studies*, 38(1), 114–157.
- Kelly, J. L., Jr. (1956). A new interpretation of information rate. *Bell System Technical Journal*, 35(4), 917–926.
- Kocken, T. (2012). Endogene instabiliteit; je moet er niet aan denken! Inaugural Lecture VU.
- Koijen, R. S. J., & Levy, B. (2026). Assessing the benefits of optimized agentic AI systems for asset pricing. SSRN Working Paper 6474601.
- Koijen, R. S. J., & Yogo, M. (2019). A demand system approach to asset pricing. *Journal of Political Economy*, 127(4), 1475–1515.
- Liu, J., & F. A. Longstaff (2004). Losing money on arbitrage: Optimal dynamic portfolio choice in markets with arbitrage opportunities. *Review of Financial Studies*, 17(3), 611–641.
- Moreira, A., & T. Muir (2017). Volatility-managed portfolios. *Journal of Finance*, 72(4), 1611–1644.
- Pástor, L., & R. F. Stambaugh (2001). The equity premium and structural breaks. *Journal of Finance*, 56(4), 1207–1239.
- Pavlova, A., & Sikorskaya, T. (2023). Benchmarking intensity. *The Review of Financial Studies*, 36(3), 859–903.
- Pflueger, C. (2025). Back to the 1980s or Not? The Drivers of Inflation and Real Risks in Treasury Bonds. *Journal of Financial Economics*, 167, 104027.
- Sharpe, W. F. (1991). The Arithmetic of Active Management. *Financial Analysts Journal*, 47(1), 7–9.
- Shleifer, A., & R. W. Vishny (1997). The limits of arbitrage. *Journal of Finance*, 52(1), 35–55.

Note

1 We thank Fernando Monar, Ken Nyholm and Luca Bortolussi for their comments. The views expressed are those of the authors and do not necessarily reflect those of the ECB or the Eurosystem.