Research Statement

Eben Lazarus

Haas School of Business, University of California, Berkeley, and NBER

NOVEMBER 2026

I am a financial economist working primarily in the field of asset pricing. A central theme of my research is that asset prices are a rich source of information about the underlying economy: because prices reflect investors' beliefs, preferences, and risk exposures, they shed light on how people process information and form expectations, which risks matter in the macroeconomy, and how these forces help explain otherwise puzzling patterns observed in financial markets. Much of my work makes use of derivatives, which are particularly informative about the likelihood and pricing of concentrated risks. In addition to asset pricing, my projects often address questions in behavioral economics and macroeconomics.

While I have an empirical focus, a key methodological feature of my work is its use of economic theory to identify what should be measured and tested to answer a given question. I have also developed econometric tools that I use to conduct careful statistical inference in my empirical work.

Below, I summarize my research in three areas: beliefs and information in financial markets; risk, valuations, and the macroeconomy; and econometric methodology. I then briefly discuss my teaching and service record.

Beliefs and Information in Financial Markets

The papers in this first area study how people form beliefs, how those beliefs are reflected in asset prices, and how to reconcile seemingly conflicting patterns in market reactions to new information.

"**Overinference from Weak Signals and Underinference from Strong Signals**" **(with Ned Augenblick and Michael Thaler; *Quarterly Journal of Economics*, 2025)** studies why beliefs overreact to some pieces of information and underreact to others. A longstanding difficulty is that overreaction and underreaction have often been documented in different settings — for example, prior work has argued that firm-level stock prices often underreact to earnings news, but that the aggregate stock market exhibits overreaction — with each pattern attributed to a separate bias.[1] We instead propose a unified explanation for these patterns based on the *strength* of the information: people tend to overreact to weak (or uninformative) pieces of information and underreact to strong ones. Intuitively, if people know which direction to update in but find it difficult to determine the precise strength of a signal, they will tend to shrink toward "moderate" updates. We first formalize this hypothesis in a general model, which imposes minimal assumptions on how people estimate a signal's strength and does not rely on the Gaussian structure that is standard in related models.

[1] Similarly, recent work (e.g., Ba, Bohren, and Imas 2025) has aimed to explain these patterns by combining multiple biases. Meanwhile, we consider a single important factor that helps explain behavior in a range of common situations.

We then test our theory in a range of settings, using two abstract lab experiments, a naturalistic experiment of our design, and observational data from sports-betting and option markets. The abstract experiments use a classic belief-updating task, but with a much wider range of signal strengths than in previous work. The naturalistic experiment asks basketball fans to predict the outcome of a game given sequences of game scenarios, in which the strength of the same news (like a scored basket) changes over the course of the game. For the sports-betting and option data, we develop new empirical tests of over- and underreaction to signals of differing strengths, comparing movements in market-implied beliefs with uncertainty reduction over time. In these markets, signals are generally weak far from the resolution of uncertainty and strong close to the end of a given contract. Across all the settings we study, we find strong and consistent evidence of overinference from weak signals and underinference from strong signals.

This paper draws on the broadest range of methods and fields of my research to date, combining economic theory, lab experiments, and asset-price data to answer a core question in behavioral economics. Subsequent work has adapted our framework to help explain findings in other contexts (e.g., Lyu and Xu 2025, Bastianello 2026, De Silva et al. 2026).

The empirical tests in the paper above build on methods developed in "**Excess Movement in Option-Implied Beliefs" (with Ned Augenblick; *Journal of Finance*, forthcoming)**, which began my agenda of using derivatives to study market participants' beliefs and preferences. It is well known that stock prices vary far more than fundamentals like dividends (Shiller 1981), but this does not by itself imply excess volatility in beliefs: changes in risk aversion or the quantity of risk could also generate large price swings. We show that option prices provide a sharper way to distinguish between these explanations.[2] Specifically, we derive bounds on the rational variation in option-implied (or risk-neutral) beliefs about market returns. As risk-neutral beliefs are constructed using the relative prices of options with different strike prices but the same maturity, they strip out variation arising from time discounting and common shocks to marginal utility or the quantity of risk. Our bounds reflect that under rational expectations, movement in beliefs must correspond on average to a reduction in uncertainty. While risk preferences distort risk-neutral beliefs away from subjective beliefs, similar logic can be used to bound risk-neutral belief movement. Our bounds require no knowledge of the true data-generating process. They do require a general assumption on the stochastic discount factor (SDF); informally, the slope of the SDF is assumed to be locally constant within an option contract. This is significantly weaker than the assumption of constant discount rates used by Shiller (1981), and it is also met either exactly or approximately under many standard modeling frameworks.

Implementing our bounds using S&P 500 index options, we find that there is so much movement in risk-neutral beliefs that the bounds are routinely violated. This implies either excess volatility in beliefs or large, high-frequency variation in the local price of risk, and we find little evidence that

[2] In using derivatives prices for such tests, we complement recent work including Giglio and Kelly (2018). Their framework and tests concern claims on volatility, commodities, and interest rates, whereas we focus on equity claims.

the latter is a key driver of our results.[3] We also show in simulations that well-known option pricing models produce much less excess movement in risk-neutral beliefs than in the data. Finally, we conduct a rough quantification of how these results translate to excess volatility in the underlying index, and we estimate that the variance of the S&P 500 is about 20% to 30% higher than it would be under rational expectations.

"**Forward Return Expectations**" **(with Mihir Gandhi and Niels Gormsen; *Review of Financial Studies*, forthcoming)** examines a previously unexplored aspect of investors' return expectations. The literature has typically studied *contemporaneous* expectations: what investors expect the market return to be over, for example, the coming year. We instead study *forward* return expectations: what investors today expect the expected return to be at future dates. Forward expectations matter for asset prices because today's prices embed expectations at all future horizons, not just the near term. We measure investors' forward return expectations using the difference between their long- and short-horizon return expectations (just as forward interest rates are obtained from the yield curve). We then compare the forward expectations with investors' own subsequent short-horizon return expectations, which lets us test the consistency of expectations over time without needing to measure objective expected returns. We implement this approach using both index option prices and survey data. For the option data, we derive new theoretical results (building on Gao and Martin 2021) showing that one can identify changes in forward return expectations under relatively weak conditions. We also measure forward expectations from surveys of professional forecasters, CFOs, and Vanguard retail investors, which elicit return forecasts at multiple horizons.

Across all our measures, we find that forward expectations are countercyclical, and excessively so: in bad times, investors believe that the equity risk premium will be substantially elevated in the future, but these expectations are too high relative to their own subsequent beliefs. The consistency in this pattern across investor groups contrasts with past findings on short-term return expectations, whose cyclicality appears to vary by investor type (e.g., Greenwood and Shleifer 2014, Nagel and Xu 2023). Our results help account for excess volatility in stock prices: when a crisis hits, the decline in prices partly reflects investors' overestimation of future risk premia. We estimate that this can explain a sizable share (on the order of one half) of the price declines in the 2008 and 2020 crises.

Like the *QJE* paper described above, this paper draws on a range of data and methods to help reconcile seemingly conflicting evidence from past work.

Risk, Valuations, and the Macroeconomy

Another central strand of my research aims to better understand fundamental risks in financial markets and the macroeconomy, and how those risks are reflected in asset valuations and returns.

"**Duration-Driven Returns**" **(with Niels Gormsen; *Journal of Finance*, 2023)** studies why familiar firm characteristics such as valuation ratios, profitability, and investment predict risk-adjusted

[3] These results concern the average reaction to news. As discussed above, our follow-up work finds that when we distinguish signals by strength, option markets appear to overreact to weak signals but underreact to strong ones.

stock returns. These return patterns are the basis of leading factor models in asset pricing, but their underlying economics have remained unclear because they are difficult to relate to fundamentals and to each other. We provide a unifying explanation based on the timing of firms' cash flows: the long legs of the major equity risk factors invest in firms with short-duration cash flows, and short-duration firms earn higher risk-adjusted returns than long-duration firms. We first show empirically that the value, profitability, investment, low-risk, and payout factors do in fact invest in firms with a short cash-flow duration (measured using analysts' long-term growth expectations and validated using realized growth rates), and that a single duration-based risk factor summarizes the CAPM alphas of these five factors.

These two facts, however, do not establish that cash-flow horizon is the *driver* of the cross-sectional alphas rather than being a proxy for some other firm-level risk. To isolate this timing effect, we use novel derivatives data on single-stock dividend futures — claims on the dividends that an individual firm will pay in a given year — to compare the pricing of cash flows across horizons within the same firm, holding all other firm characteristics fixed. We find that risk-adjusted returns decrease with cash-flow horizon within a firm, while firm characteristics do not predict returns once cash-flow horizon is held fixed. These results imply that many major return patterns, including the value premium, arise in large part because they sort firms by duration. Identifying a driver of return patterns in this way is fairly unique in cross-sectional asset pricing, which often can go no further than estimating equilibrium covariances. We also show that an alpha on near-term cash flows is consistent with greater exposure to macroeconomic risk than suggested by market beta.

This paper builds on prior work relating duration to stock returns, and to the value premium in particular (Lettau and Wachter 2007, Gonçalves 2021). We show that this insight helps explain a much broader set of return patterns, and we use single-stock dividend futures data to pin down the role of cash-flow timing.[4]

In "**Interest Rates and Equity Valuations**" **(with Niels Gormsen; revise and resubmit,** ***American Economic Review*****)**, we study duration and asset pricing from an aggregate perspective, asking how changes in interest rates affect stock-market valuations. Long-term interest rates declined significantly across countries in recent decades. Given that stocks are long-term assets with high sensitivity to discount rates, recent work has often assumed that the decline in interest rates mechanically raised stock valuations over this period. Empirically, however, there is no clear relation between changes in equity valuations and interest rates, and the correlation between stock and bond returns is weak and often negative (e.g., Campbell, Sunderam, and Viceira 2017). This disconnect arises because interest rates are endogenous and may change for multiple structural reasons, each of which affects equity differently. A decline in rates driven by lower expected growth, for example, also lowers equities' expected future cash flows, muting the effect on valuations. We develop a general decomposition that formally separates changes in real rates into three drivers: expected growth, uncertainty, and "pure discounting" (akin to a change in the rate of time

[4]In so doing, we also build on previous work studying index-level, rather than single-stock, dividend futures, including van Binsbergen and Koijen (2017).

preference). Only pure discounting shocks should pass through fully to equity valuations.

To implement the decomposition, we measure trend real rates and expected growth using long-horizon survey forecasts across G7 countries. We augment these with option-based measures of uncertainty. After stripping out the role of growth-rate and uncertainty changes, the remaining real-rate change is our estimate of the pure discounting shock. We find that pure discounting shocks indeed transmit one-for-one to equity valuations, and they explain a large share (80%) of cross-country valuation changes since 1990. While pure discounting shocks are important for stocks, they account for only part of the decline in rates: in the U.S., only 35% of the decline is attributable to pure discounting, and the remaining 65% to the other components. So while equities have benefited somewhat from falling U.S. rates, assuming full pass-through would significantly overstate the role of rate changes. We use the decomposition to measure rate exposures in the cross-section of stocks, estimate the ex ante equity premium, assess the contribution of declining rates to wealth inequality, and study market responses to monetary-policy shocks.

I am also conducting early-stage work on exchange rates in "**Long-Run Currency Shocks**" **(with Niels Gormsen and Markus Ibert)**, which builds on the analysis in the previous paper. Higher relative interest rates are generally associated with a stronger currency, but large changes in exchange rates often go in the opposite direction. In 2025, for instance, the dollar fell significantly even as U.S. rates increased, and similar dynamics describe a number of past policy-related shocks. We study such episodes using a model fit to long-horizon survey data on professional forecasters' expected exchange rates and interest rates. A model allowing for perceived permanent changes in exchange rates fits the forecast data better than a stationary model and helps make sense of large currency moves: we find that the majority of dollar variation during these episodes is perceived as permanent, helping explain the disconnect between exchange rates and interest rates.

"**Conditional Excess Volatility**" **(with Christian Jensen; working paper)** turns from the long-run perspective above to business-cycle variation in risk and expected returns. Standard theory implies that investors require compensation for bearing market risk, so that when the market's variance rises, its expected return should increase at least in proportion. (If risk aversion rises in recessions, the increase should be more than proportional.) In bad times, however, expected returns increase far less than variance. At the height of the financial crisis, for example, annualized market variance reached roughly $(80\%)^2$, about 30 times its pre-crisis level of $(15\%)^2$. Even with low risk aversion, this variance would imply an expected return of roughly 70% per year, far above plausible estimates at the time. We show that the puzzle can be resolved if much of the increase in market variance in recessions is uncompensated, in that it is unrelated to investors' marginal utility.[5] We call this uncompensated component *conditional* excess volatility, as it arises in bad times specifically. Identifying such conditional excess volatility requires weaker assumptions than in past volatility tests: using the risk–return relation in good times as a benchmark, we obtain a lower bound for excess volatility in bad times using only the assumption that risk aversion does not

[5] Equivalently, the exposure of the SDF to the market (i.e., $\beta_{M\to R,t} = \mathrm{Cov}_t(M_{t+1}, R_{m,t+1})/\sigma_t^2$) *decreases* in bad times.

decrease in those periods. Measuring the bound using three different approaches (from realized returns, estimated expected returns and variance, and options), we find in our baseline that the unpriced component of variance accounts for at least half of market variance in bad times. As a result, even though the market Sharpe ratio (excess returns relative to standard deviation) increases in bad times, investors' optimal equity positions do not. We also show that this helps explain why volatility-managed portfolios earn high alphas (Moreira and Muir 2017) and why Sharpe ratios are higher on claims to near-term dividends than on more distant claims. Finally, high conditional excess volatility predicts lower corporate investment even after controlling for other aggregate conditions, indicating that it has eventual real effects despite being unpriced.

In "**Horizon-Dependent Risk Pricing: Evidence from Short-Dated Options**" **(working paper)**, I show that the price of risk over the value of the S&P 500 increases as the investment horizon becomes shorter. I use index option prices to infer the prices of gambles on the future index return that hold *fixed* the quantity of risk by horizon. In the data, I find that the risk premium on such a gamble robustly declines with the horizon. This paper does not rule out either the quantity of risk or belief-related explanations as drivers of the alpha on near-term cash flows discussed above, but it shows that the price of risk must be at least part of the story.

"**Option-Implied Event Risk**" **(with Richard Stanton and Johan Walden; working paper)** also separates distinct sources of return risk, now focusing on discrete anticipated events such as earnings announcements, policy announcements, and court decisions, with known timing but unknown effects on stock prices. Standard option-implied measures of risk combine the event-specific risk with ordinary price risk before and after it. We develop a method to isolate the market's perceived distribution of event-specific price changes, using option prices at different maturities.[6] We first study three illustrative cases: Elon Musk's takeover bid for Twitter, the AT&T–Time Warner antitrust case, and the 2016 Brexit referendum. In the first two cases, the recovered distributions assign substantial probability to both the agreed takeover prices and the possibility that the deals would collapse. For the Brexit referendum, we use options on an exchange-traded fund tracking the British pound, and the recovered probability of a large GBP depreciation tracks the probability of a "Leave" vote implied by betting markets at the time quite well. We then apply the method to a large sample of earnings announcements. The recovered event distributions predict the variance and tail risk of realized announcement-day returns well, outperforming conventional measures. We have completed an early version of the paper and continue to develop it with additional applications.

Econometric Methodology

The two main agendas above reflect my model-agnostic approach to research: while a behavioral lens is appropriate for many of the questions I examine, I am open to a range of frameworks and use the one best suited to the specific question and data. This emphasis on careful empirical work also informs a secondary line of research on econometric inference in time-series settings.

[6] This paper connects to both main strands of my research, as we use option prices to back out market-implied beliefs.

In such settings, serial correlation and heteroskedasticity complicate hypothesis testing, and the usual approach is to use heteroskedasticity- and autocorrelation-robust (HAR) standard errors. Conventional tests based on Newey and West (1987) standard errors, however, can reject true hypotheses far too often: a nominal 5% test can falsely reject the null around 20% of the time even with moderate data persistence. These false rejections arise because conventional tests use relatively few lags to estimate the standard errors. Using more lags reduces the false rejection rate, but it also generates substantial variability in the standard error estimate. Accounting for this variability requires larger critical values, at a cost in statistical power. And while a theoretical literature had proposed alternatives to Newey–West, it remained unclear which tests achieved the best possible tradeoff between false rejection rates and statistical power. Applied researchers in empirical finance and macroeconomics also lacked clear guidance about which estimator, bandwidth, and critical values to use. My two papers in this area address these problems, and I continue to use their recommended procedures in my applied work.

"**The Size-Power Tradeoff in HAR Inference**" **(with Daniel Lewis and James Stock; *Econometrica*, 2021)** is the central paper in this line of work. We characterize the fundamental tradeoff between false rejections (size distortions) and power across a broad class of HAR tests. These results provide an unambiguous ranking of tests, based on each test's power for a given false rejection rate. The optimal testing procedure achieves the size–power frontier, which we characterize analytically. The frontier is derived from higher-order expansions for tests of the mean of a Gaussian time series, but a large-scale simulation analysis suggests that the main findings extend to higher-dimensional time-series regressions. We also show that simple tests with familiar t and F critical values come quite close to the frontier. Implementing these tests still requires choosing a bandwidth parameter, which amounts to selecting a point on the size–power frontier. "**HAR Inference: Recommendations for Practice**" **(with Daniel Lewis, James Stock, and Mark Watson; *Journal of Business & Economic Statistics*, 2018)** takes up this practical question, providing a simple bandwidth rule and explicit test recommendations that researchers can apply in one step. Relative to textbook Newey–West, our proposed tests use substantially more lags (for example, with 200 observations, we suggest using 19 Newey–West lags rather than 5, or using an alternative test altogether), as well as critical values that account for the estimation uncertainty in the standard error. In the moderate-persistence case above, our recommendations reduce the false rejection rate from 20% to about 6%, close to the nominal 5% level, at only a modest cost in power. Earlier this year, Daniel Lewis and I released a Stata package (`harreg`; see `https://github.com/ebenlazarus/harreg`) implementing our recommendations, so that applied researchers can adopt them directly in their own work.

Teaching and Service

TEACHING. Each year at Haas, I have taught MFE 230A (Investments and Derivatives), the core asset pricing and derivatives course in the first term of the Master of Financial Engineering (MFE) program, and co-taught two sections of UGBA 103 (Introduction to Finance), the introductory finance course for undergraduate business majors. In each year, I have received the Berkeley

Haas "Club 6" teaching recognition for teaching effectiveness ratings of 6 or above out of 7, and in my most recent teaching evaluations, my trimmed mean ratings were 6.7 in the MFE course (the highest among spring MFE classes) and 6.2 in the undergraduate course. At MIT Sloan, I taught introductory finance to undergraduates and MBA students and co-taught PhD courses, including the second-year asset pricing course. I have chaired PhD dissertation committees at both MIT and Berkeley, and I have co-led a PhD reading group in behavioral economics and finance at Haas.

SERVICE. I served as Chair of the Finance Group's faculty search committee in 2024–25, and I have continued to serve as a member of the search committee for 2025–27. I serve on the Haas Undergraduate Program Committee and the Finance PhD Improvement Committee (and previously the MFE Curriculum Committee), and I am a co-organizer of the weekly finance seminar. In November 2025, I was invited to give a presentation to UC Investments, which manages the UC system's endowment and pension assets, on current equity valuations and what they imply for long-run returns and crash risk. I regularly serve as a referee for numerous journals, including the *Quarterly Journal of Economics*, *American Economic Review*, *Econometrica*, *Journal of Finance*, *Review of Financial Studies*, and *Journal of Financial Economics*. I have also served on the program committees for SFS Cavalcade North America and the Western Finance Association Annual Meetings, among others. I have served as a conference discussant for papers in multiple subfields adjacent to my work, including at the NBER Summer Institute Asset Pricing Meeting, the Chicago Booth Asset Pricing Conference, the NBER Behavioral Finance Meeting, the Cowles Foundation Conference on Macroeconomics, and the American Finance Association Annual Meetings.

References

AUGENBLICK, N. AND E. LAZARUS (2026): "Excess Movement in Option-Implied Beliefs," *Forthcoming, Journal of Finance*.

AUGENBLICK, N., E. LAZARUS, AND M. THALER (2025): "Overinference from Weak Signals and Underinference from Strong Signals," *Quarterly Journal of Economics*, 140, 335–401.

BA, C., J. A. BOHREN, AND A. IMAS (2025): "Over- and Underreaction to Information: Belief Updating with Cognitive Constraints," *Working Paper*.

BASTIANELLO, F. (2026): "Bias and Predictability in Analysts' Beliefs," *Working Paper*.

VAN BINSBERGEN, J. H. AND R. S. J. KOIJEN (2017): "The Term Structure of Returns: Facts and Theory," *Journal of Financial Economics*, 124, 1–21.

CAMPBELL, J. Y., A. SUNDERAM, AND L. M. VICEIRA (2017): "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds," *Critical Finance Review*, 6, 263–301.

DE SILVA, T., E. LARSEN-HALLOCK, A. REJ, AND D. THESMAR (2026): "Expectations Formation with Fat-Tailed Processes: Evidence and Theory," *Working Paper*.

GANDHI, M., N. J. GORMSEN, AND E. LAZARUS (2026): "Forward Return Expectations," *Forthcoming, Review of Financial Studies*.

GAO, C. AND I. W. R. MARTIN (2021): "Volatility, Valuation Ratios, and Bubbles: An Empirical Measure of Market Sentiment," *Journal of Finance*, 76, 3211–3254.

GIGLIO, S. AND B. KELLY (2018): "Excess Volatility: Beyond Discount Rates," *Quarterly Journal of Economics*, 133, 71–127.

GONÇALVES, A. S. (2021): "The Short Duration Premium," *Journal of Financial Economics*, 141, 919–945.

GORMSEN, N. J., M. IBERT, AND E. LAZARUS (2026): "Long-Run Currency Shocks," *Working Paper*.

GORMSEN, N. J. AND E. LAZARUS (2023): "Duration-Driven Returns," *Journal of Finance*, 78, 1393–1447.

——— (2026): "Interest Rates and Equity Valuations," *NBER Working Paper 34814*.

GREENWOOD, R. AND A. SHLEIFER (2014): "Expectations of Returns and Expected Returns," *Review of Financial Studies*, 27, 714–746.

JENSEN, C. S. AND E. LAZARUS (2026): "Conditional Excess Volatility," *Working Paper*.

LAZARUS, E. (2022): "Horizon-Dependent Risk Pricing: Evidence from Short-Dated Options," *Working Paper*.

LAZARUS, E., D. J. LEWIS, AND J. H. STOCK (2021): "The Size-Power Tradeoff in HAR Inference," *Econometrica*, 89, 2497–2516.

LAZARUS, E., D. J. LEWIS, J. H. STOCK, AND M. W. WATSON (2018): "HAR Inference: Recommendations for Practice," *Journal of Business & Economic Statistics*, 36, 541–559.

LAZARUS, E., R. STANTON, AND J. WALDEN (2026): "Option-Implied Event Risk," *Working Paper*.

LETTAU, M. AND J. A. WACHTER (2007): "Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium," *Journal of Finance*, 62, 55–92.

LYU, T. AND Z. W. XU (2025): "Taking the Road Less Traveled? Market Misreaction and Firm Innovation Directions," *Working Paper*.

MOREIRA, A. AND T. MUIR (2017): "Volatility-Managed Portfolios," *Journal of Finance*, 72, 1611–1644.

NAGEL, S. AND Z. XU (2023): "Dynamics of Subjective Risk Premia," *Journal of Financial Economics*, 150, 103713.

NEWEY, W. K. AND K. D. WEST (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

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