Actively Passive: The Rise of Market Volatility

Lars A. Lochstoer and Tyler Muir*

July 20, 2026

Abstract

We show that U.S. stock market volatility has roughly doubled since the 1950s. The increase in volatility is strongest for daily returns and weaker for longer horizons, which points against changes in the volatility of fundamentals and towards explanations with transitory return shocks. We trace the increase in volatility to two related shifts in market structure: investors increasingly trade the market index — via futures, index funds, etc. — and can do so over more hours of the day. Exploiting the introduction of E-mini S&P 500 futures and historical changes in trading hours as natural experiments, we show that expanded index trading opportunities raise trading volume and market variance, and that this is large enough to account for roughly half the total rise. A model incorporating both shifts rationalizes these findings: arbitrageurs absorb larger market demand shocks, raising volatility, reducing return autocorrelation, and amplifying systematic risk.

*Lochstoer and Muir are with UCLA & NBER. We thank participants at the ASU Winter Finance Conference 2026, UNC-Duke Asset Pricing Conference 2026, UC Berkeley, UCLA, Mark Garmaise, Mark Grinblatt, Stefano Giglio, Valentin Haddad, Joren Koeter and John Shim for useful comments. A previous draft circulated under the name "(Why) Have Financial Markets Become More Volatile? The Role of Market Index Trading."

1. Introduction

We document that stock market volatility has roughly doubled over the last 70 years. Figure 1 shows the persistent nature of this increase by plotting the daily US stock market volatility from the 1950s to today using 5-year rolling windows. Yet over the same period, the economy has gotten calmer — the volatility of GDP growth, industrial production, and corporate profits has been flat or declining. Because stock market volatility, constructed from daily returns, is a central measure in macroeconomics and finance (e.g., Bloom, 2009; Martin, 2017; Berger, Dew-Becker, and Giglio, 2020), understanding this secular shift is important.[1]

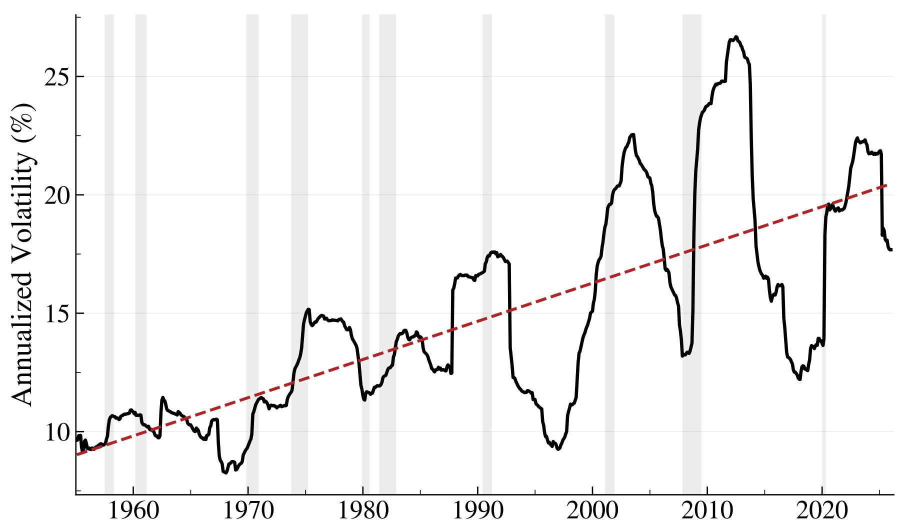

Figure 1: Stock Market Volatility. 5-year rolling standard deviation of daily CRSP value-weighted market return.

The explanation, we argue, is not about fundamentals but about both how, what, and when investors trade. In the 1950s, an investor who wanted exposure to the mar-

[1]Bloom (2009) uses realized stock market volatility to identify uncertainty shocks that drive business cycles. Martin (2017) argues for a tight link between the conditional volatility of the stock market and the conditional market risk premium. Berger, Dew-Becker, and Giglio (2020) show it is specifically realized volatility — not forward-looking implied volatility — that predicts contractions.

ket had to effectively buy stocks one at a time. Today hundreds of billions of dollars change hands around the clock in products that track the market index, including index futures, ETFs, and funds. This paper provides stylized facts and causal evidence that shifts towards (1) increased trading of the market as a whole and (2) increased ability to trade more often (i.e., overnight futures) can help account for a meaningful share of the increase in stock market volatility.

The structure of the increase reinforces this conclusion. Daily volatility has risen 127%, weekly 76%, monthly 40%, while at annual horizons, there is no trend at all. If cash flow volatility were driving the increase, volatility would rise similarly across horizons, so the steep frequency gradient points to something about how markets operate day to day.

To measure the shift toward index trading directly, we first construct a market volume beta: the slope from a daily cross-sectional regression of each stock's share of total volume on its weight in the market portfolio. If all trading were index trading, this slope would equal one. It has doubled from 0.4 to 0.8 over our sample, tracking the rise in volatility closely. Thus, investors are trading relatively more along the market dimension than they used to. This change occurs slowly over time, consistent with the gradual move towards increasing trading of a market index. This started early in the 1960s when institutional investors became more dominant concentrating trade in the nifty-fifty. Wells Fargo started an institutional index fund in 1971, Vanguard retail index fund in 1976, and S&P500 futures in 1982. Overnight trading took hold with the S&P500 E-mini introduction in 1997 and has continued increasing relative to intraday-trading in S&P500 futures until the end of the sample. The trend in volatility shown in Figure 1 really appears to start in the 1970s, which is consistent with the introduction and gradual adoption of these index trading instruments.

This correlation does not establish causation — volatility could drive trading rather than the reverse, or a third factor could drive both. We provide causal evidence from natural experiments. The first builds on French and Roll (1986), who study what happens when the New York Stock Exchange closed every Wednesday for six months in 1968. If price movements reflect information, weekly returns should be unaffected — the same news arrives whether or not the exchange is open on Wednesday. Instead, weekly volatility falls by roughly one-fifth, implying that volatility tracks trading time more than calendar time. We extend this logic to nine changes in NYSE trading hours between 1952 and 1985, including the elimination of Saturday trading, extensions and reductions in daily hours, and the Wednesday closures. Using trading hours as an instrument for volume, we estimate that a 10% increase in trading volume raises market volatility by about 12%.

The second experiment isolates index trading specifically. The introduction of E-mini S&P 500 futures in September 1997 expanded after-hours trading in the market index while leaving overnight trading in individual stocks unchanged. Comparing overnight-to-intraday price movements for the S&P 500 versus individual stocks before and after the launch, we find a clear increase in relative overnight market volatility, coinciding with a surge in futures volume. The market became more volatile precisely when and where index trading expanded.

We complement these experiments with evidence tracing the gradual buildup of overnight index trading, decomposing overnight returns for individual stocks into systematic and idiosyncratic components. Because index futures allow investors to trade the market overnight but not individual stocks, our prediction is sharp: the systematic component of overnight variance should rise over time while the idiosyncratic component should not. Using a panel of the 500 largest U.S. stocks from 1993–2024, we estimate a triple-difference specification comparing market versus idiosyncratic variance, overnight versus intraday, around three discrete expansions of CME futures trading hours. Overnight market variance rises at these events—by roughly 6% per additional hour of futures trading—while idiosyncratic overnight variance does not. The magnitudes are substantial: the rise in overnight variance accounts for roughly half of the total rise in market variance over these three decades, and taken at face value the event-study estimates imply the futures channel covers most of that overnight rise. Consistent with this channel, we provide further evidence that the opening of futures trading on NYSE holidays increased market volatility on those days, with effects that grow with the volume of futures trading on the holiday.

In all of these experiments, expanding the ability to trade results in both increased trading activity (volume) and in increased volatility. In contrast, in canonical asset pricing models, volatility is based on calendar time and not trade time. That is, if the stock market was only open one hour per week, the volatility of weekly compounded returns should be the same as in the case when the market is open 24 hours a day. In the former case, traders would collect all relevant information through the week and then move prices in large amounts during the short time when the market was open, rather than continuously as information arrived. But price moves over a calendar week would be the same. In reality, increased ability to trade leads to more trade and larger price fluctuations. We find that volatility and volume are more closely aligned with trade time than calendar time. Because of this, much of the rise in volatility we document in the later part of our sample (from 1995 onwards) comes from a rise in volatility of overnight returns as investors are now easily able to trade market indices nearly around the clock.

To interpret these patterns, we develop a simple model of index trading, along the lines of Campbell et al. (1993). In a market without index trading, investors hold concentrated portfolios (following Merton (1987)), while risk-averse arbitrageurs that trade in all stocks absorb investors' idiosyncratic demand shocks and lean against the systematic component of investors' demand shocks. When index trading is introduced, investors instead trade the market portfolio, which allows a more aggressive allocation to stocks, due to its higher Sharpe ratio relative to undiversified portfolios. Further, we let demand shocks scale with trade-time rather than calendar-time: more trading opportunities within a calendar day generate a larger total flow of demand shocks, motivated by the idea that trading and observing transaction prices generate new signals, attention, or sentiment shocks, consistent with the spirit of French and Roll (1986) and with our empirical findings. Arbitrageurs are thus required to absorb larger systematic demand shocks than before, which they only partially accommodate. This shift in trading structure raises market volatility, reduces return autocorrelation, and increases the share of variance explained by market-wide factors. The predictions of the model are borne out in the data, which show both a decline in return autocorrelations and a rise in the systematic component of stock returns over time. Further, the model matches increased volume in market index trades and the large increase in market volatility.

Related literature and implications. Our work has important implications for the rise of passive investing. While index investing may be cross-sectionally passive, the amount of trade in index products is far from passive. As David Booth puts it: "*There's been an incredible movement to indexing, and yet there's been an incredible increase in trading volume. [...] Index funds are the ideal market timing vehicle. [...] Instead of individual stock selection, it's kind of like a big gambling casino.*"[2] Much of the passive literature focuses on market efficiency in the cross-section (Sammon, 2025; Haddad et al., 2021; Brightman and Harvey, 2025); our focus is on volume in index products affecting aggregate volatility. French (2008) and Chinco and Sammon (2024) measure the growth of passive ownership, but the consequences for market-level volatility remain underexplored.

Our findings imply that volatility is no longer solely a reflection of economic uncertainty, but also of how trading is organized. This matters for measuring macroeconomic uncertainty from daily stock returns (Bloom, 2009; Martin, 2017; Berger et al., 2020), for portfolio choice given the rise in systematic risk (Moreira and Muir, 2017; Lochstoer and Muir, 2022), and for understanding how market macrostructure shapes asset prices (Haddad and Muir, 2025; Harvey et al., 2025). It also raises policy ques-

[2]Odd Lots Podcast, Eugene Fama and David Booth on the Birth of Modern Finance, minute 27.

tions: proposals to extend trading hours—for example, the NYSE's plan for 24-hour trading by the end of 2026—have the potential to increase volatility further.

Several strands of prior work inform our analysis. First, our paper contributes to a literature on excess volatility. Shiller (1981) first showed that stock prices move too much to be justified by subsequent dividends, and a large literature has sought explanations. Our paper identifies a new source of excess volatility: the shift toward index-level trading, which forces arbitrageurs to absorb systematic rather than idiosyncratic demand shocks. Our model builds on the noise trader framework of De Long et al. (1990) and the limits-of-arbitrage logic of Shleifer and Vishny (1997), applied to the specific setting of index trading.

Second, we build on a literature studying how index inclusion and passive ownership affect prices and return dynamics. Shleifer (1986) shows that demand curves for stocks slope down using S&P 500 additions, and Barberis et al. (2005) show that stocks added to the index subsequently comove more with it—exactly the mechanism our model predicts. A key paper in this literature is Ben-David et al. (2018), who use Russell index reconstitutions to show that ETF ownership increases stock-level volatility, providing clean micro-level evidence that index products transmit non-fundamental demand shocks to underlying stocks. Relatedly, Baltussen et al. (2019) link indexing to negative index serial dependence, with their strongest identification coming from cross-sectional variation in exposure to index demand. These papers show that index products affect prices, comovement, return dynamics, and stock-level volatility. Our paper addresses a complementary aggregate question: whether the secular growth of index trading—across futures, ETFs, and extended trading hours—has raised the level of *aggregate market* volatility over decades. Our time-series experiments exploit discrete changes in trading time and overnight index-trading access, which shift aggregate index-trading opportunities rather than the relative exposure of particular stocks to index demand (see also Da and Shive, 2018; Greenwood and Thesmar, 2011; Israeli et al., 2017; Basak and Pavlova, 2013; Sullivan and Xiong, 2012).

Third, our empirical strategy builds on the classic insight of French and Roll (1986) that trading itself can induce excess volatility. The positive relationship between volume and volatility is well established (Karpoff, 1987), but the canonical explanation attributes this to joint arrival of information (Admati and Pfleiderer, 1988). Bessembinder and Seguin (1992) study the contemporaneous relationship between daily futures activity and stock volatility, finding mixed results—deeper futures markets are associated with lower volatility, while unexpected surges in futures volume are associated with higher volatility. Overnight futures trading did not yet exist during their sample (1978–1989), so their analysis relies entirely on close-to-close returns and cannot isolate the overnight channel. We improve on this by decomposing variance into overnight and intraday components and comparing systematic versus idiosyncratic variance within the overnight window, since individual stocks cannot trade overnight but the index can via futures. Hasbrouck (2003) shows that E-mini futures dominate price discovery among index instruments. Our overnight decomposition extends the framework of Lockwood and Linn (1990) and Barclay and Hendershott (2003), and connects to Lou et al. (2019), who show that overnight and intraday expected returns are driven by different forces.

A natural objection is that some landmark events (the launch of index futures in 1982, the Vanguard index fund in 1976, or the SPY ETF in 1993) do not produce immediate jumps in market volatility. But this is consistent with our story: these products had negligible trading volume at launch and took years or decades to gain widespread adoption. In instrumental variables language, product launches have a weak first stage—the event barely moves index trading volume in a narrow window. Our natural experiments are chosen precisely because they have strong first stages: closing a trading day removes activity that was already happening at full volume, extending trading hours gives immediate access to people who already know how to trade, and the E-mini made an existing contract cheaper and electronic. This distinction also helps reconcile prior findings: French and Roll (1986) find large effects from Wednesday closures because the first stage is mechanical, while Bessembinder and Seguin (1992) find ambiguous results studying a period when futures trading was still building up. To identify the gradual channel, the slow buildup of index trading over decades, we use a triple-difference design that compares systematic versus idiosyncratic variance, overnight versus intraday—anchored by event-study estimates at the discrete expansions of futures trading hours and cumulated over the full sample. The cross-sectional prediction supplies the control groups, making the design well suited to a channel that accumulates slowly.

Finally, Hong and Stein (2007) reviews models in which differences in beliefs generate volume and volatility. Our model falls within this tradition, as the demand shocks driving index trading arise from heterogeneous beliefs about market prospects. Campbell et al. (1993) provides the demand-shock framework we build on, and Schwert (1989) establishes the benchmark findings on the relationship between stock market volatility and macroeconomic variables that our paper seeks to explain.

The rest of the paper is organized as follows. Section 2 carefully documents the historical rise in market volatility. Section 3 presents causal evidence from the introduction of E-mini futures and changes in trading hours. Section 4 develops the theoretical model of index trading and derives testable predictions. Section 5 concludes.

2. Market Volatility: Broad Facts

We collect data from Ken French on the CRSP value-weighted excess market return at the daily and weekly frequency over the risk-free rate. We compute the sum of squared returns on a rolling 5-year basis and normalize to produce an annualized measure of variance. We compute the square-root of this number as a measure of volatility over the same horizon, where we express volatility in percent per year.[3]

Table 1 runs regressions with a linear time trend using data from 1955 to 2025. That

[3] Computing volatility at shorter horizons, e.g., monthly, and then averaging these values over a 5-year period produces qualitatively similar results.

is, we run

$$\sigma_t = a + b\tilde{t} + \varepsilon_t \tag{1}$$

where the trend $\tilde{t} = (t-1)/(T-1)$ and T is the total number of observations. This means $\tilde{t}$ ranges from 0 at the start of the sample to 1 at the end of the sample. The coefficient a then represents the estimated level of volatility at the beginning of the sample, while b represents the percentage point increase in volatility so that the value $a + b$ is the level of volatility at the end of the sample. The percentage increase in volatility is then b/a. While we use a simple linear time trend as a way to gauge the increase in volatility, we do not necessarily mean to imply that volatility will continue to increase in the future at this linear rate, or that the increase is necessarily linear from 1955-2025.

For daily data, the coefficient on the trend is 11.4% and the constant is about 9%. This highlights that volatility has more than doubled over the sample period. Because volatility is highly persistent, conventional standard errors would overstate the significance of a trend. We report Newey-West standard errors with 260 lags (5 years of weekly observations), and we also compute p-values using the fixed-b critical values of Lazarus et al. (2018).[4][5] We show robustness of the results to outliers, time-periods, stale prices or other microstructure liquidity issues, alternative ways of measuring volatility, and so on in Section 2.2.

Column 2 uses weekly returns and finds an increase of about 8 percentage points, resulting in a total increase in volatility of about 77% over the sample. Column 3 constructs monthly data using compounded overlapping 4-week returns and finds

[4]See also Vogelsang (1998) for the general fixed-b framework. The fixed-b approach accounts for the well-known size distortions of standard HAR tests when the bandwidth is a non-negligible fraction of the sample size, as it is here.

[5]Inference is also robust to a moving-block residual bootstrap: with 10-year blocks (seven effective blocks across the sample) the 95% bootstrap confidence interval for the daily-column trend is $[+7.9, +15.2]$; with 20-year blocks (only three effective blocks) it is $[+8.6, +14.6]$.

Table 1: Volatility Time Trends by Return Frequency

	Daily (1)	Weekly (2)	4-Week (3)
Constant	9.00	10.75	12.77
	(0.68)	(0.91)	(0.98)
Trend	11.44	8.24	4.85
	(2.03)	(1.94)	(1.99)
p-value	<0.001	<0.001	0.024
N	3,705	3,704	3,704
R^2	0.517	0.370	0.168
Vol increase	127%	77%	38%

Notes: Each column regresses 5-year rolling annualized volatility (in %) on a constant and a linear time trend normalized to zero at the start and one at the end of the sample. The dependent variable is the square root of the annualized 5-year realized variance. Column (1) uses daily excess returns from the Fama–French market factor, with weekly RV computed as the sum of squared demeaned daily returns within each calendar week, and 5-year RV as the sum of 260 weekly RVs. Column (2) uses weekly Fama–French excess returns. Column (3) uses 4-week excess returns, constructed by geometric compounding of weekly gross market and risk-free returns. All regressions use weekly overlapping observations from January 1955 to December 2025. Standard errors (in parentheses) use the Newey–West estimator with Bartlett kernel and $S = 260$ lags. p-values use fixed-b critical values with $\kappa = \lfloor 3/(2b) \rfloor$ degrees of freedom, where $b = S/T$. The last row reports the percentage increase in trend volatility from the beginning to the end of the sample.

an increase of about 5 percentage points, resulting in a total increase of around 38%. Thus, while volatility increases substantially across all return horizons, we can see that it is more heavily concentrated in short-horizon returns. Because compounded (log) returns over longer-horizons are equal to the sum of log returns at a daily horizon, the difference in the trend in volatility across horizons necessarily reflects changes in the autocorrelation of returns in the sample. For example, if return autocorrelation at all horizons is zero, then the expectation of squared weekly returns is equal to the expectation of the sum of squared daily returns within a week.

Immediately, the pattern across horizons suggests differences in financial markets and trading behavior itself to explain the increase in volatility. For example, a story of time-varying volatility of fundamentals will not explain the change in volatility patterns across return horizons. One needs an increase in "discount rate" volatility – the volatility of a mean-reverting component of returns – to explain these differential patterns. Because these patterns are at relatively short horizons, this naturally suggests candidates based on trading behavior.[6] We will return to explanations based on fundamentals shortly.

2.1 Volatility, Volume, and Fundamentals

Figure 2 plots market volatility against total stock market volume as a share of total market capitalization over 5-year periods from 1931 to 2025. There is a strong correlation between trading volume and volatility since the 1960s, both in terms of low frequency trend and in the cyclical movements. A natural story is that higher trading activity leads to higher volatility, especially for shorter horizon returns. This fits well with the return-horizon pattern we've shown so far. Of course, there could be other factors affecting both volatility and volume, and in addition high volatility could lead to more demand for trading activity, a reverse causality story. In Section 3.1 we doc-

[6] Typical discount rate volatility in traditional macro finance models (say, with time-varying risk-aversion) focuses on long-run mean-reversion at the horizon of years rather than weeks.

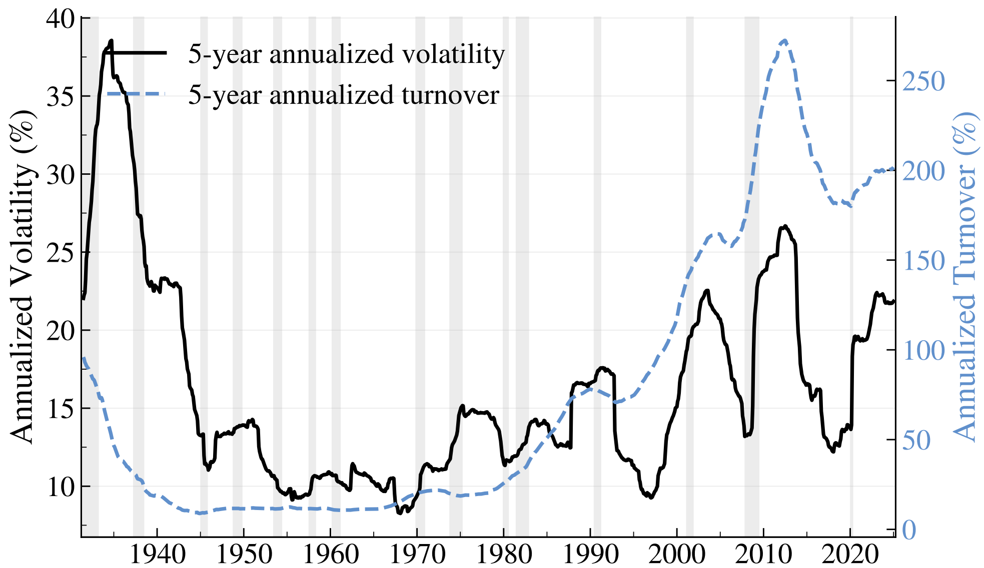

Figure 2: Volatility and Turnover

Notes: Left axis (solid): 5-year rolling annualized volatility of the Fama–French market excess return, computed from daily returns as in Figure 1. Right axis (dashed): 5-year rolling annualized turnover, defined as the cross-sectional average of daily dollar volume divided by market capitalization from CRSP, averaged over 60 months and annualized. Shaded regions indicate NBER recessions.

ument a causal relation between volume and volatility by using changes in trading hours as an instrument for volume as well as the introduction of the E-mini futures.

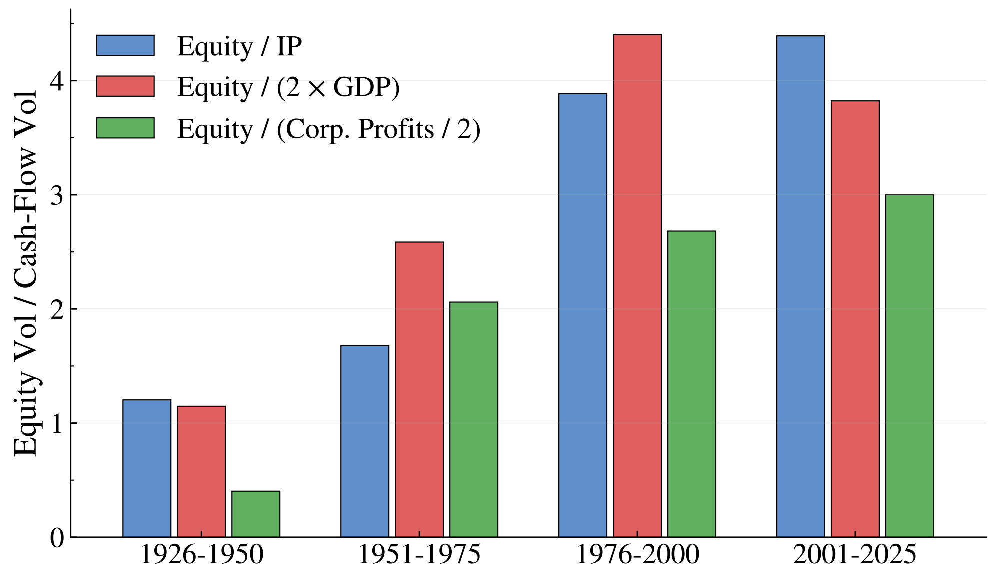

Figure 3: Equity Volatility Relative to Fundamental Volatility

Notes: Ratio of annualized equity return volatility to annualized fundamental volatility, by 25-year sub-period. Equity volatility is the annualized standard deviation of daily Fama–French market excess returns within each sub-period. Fundamental volatility measures: industrial production (IP) uses 12-month overlapping growth rates; real GDP per capita uses quarterly data post-1947 and annual data pre-1947, spliced; corporate profits uses annual real growth rates. GDP volatility is scaled by 2 and corporate profits volatility is scaled by 1/2 for visual comparability. Data sources: Fama–French daily factors, FRED (INDPRO, real GDP per capita, real corporate profits).

Figure 3 documents that the increased volatility in the post-WW2 data is not driven by more volatile macroeconomic fundamentals. Specifically, the figure shows the ratio of annualized daily stock market volatility to annualized macro volatility over the four non-overlapping 25-year periods from 1926 to 2025. We consider three underlying series to measure macro volatility: industrial production growth, real per capita GDP growth, and corporate profit growth. The first has monthly data going back to 1920 and the two other has quarterly post-WW2 data, which we splice with the annual data going back to 1929. Corporate profits are useful because they deal with changes in corporate leverage (they are a levered cash flow).

The figure shows that the ratio of stock market volatility to macroeconomic volatility was lowest in the early period, slightly increased in the 1951-1975 period, and then strongly increased in the 1976–2000 period and remained high thereafter. In other words, high volatility of fundamentals explains the high stock return volatility in the Great Depression and WWII. However, since then the volatility of macroeconomic fundamentals has declined substantially (through the "Great Moderation" period of the 1980s to mid 2000s). The macroeconomy has, if anything, looked *less* volatile over time – and hence is a poor candidate to explain the rise in volatility we've seen in the past 70 years. This echoes a long literature starting with Schwert (1989) which has difficulty tying variation in stock market volatility to fundamentals. Controlling for industrial-production volatility sharpens this point. On the full 1931–2025 sample, a simple linear trend in stock volatility is statistically zero – elevated Depression-era vol offsets the modern rise – but adding 5-year rolling industrial-production volatility as a control flips the trend to +11.0 percentage points (s.e. 4.3). On the post-1960 subsample, the trend is nearly identical with or without the IP-vol control: +10.8 (s.e. 2.2) vs. +10.9 (s.e. 2.1) percentage points on a normalized linear trend.[7] IP vol absorbs the high stock vol of the Depression era but leaves the post-1960 rise essentially unexplained. This combination – rising stock volatility against, if anything, declining fundamental volatility – is essentially unique to the post-1960 period; in earlier high-vol episodes the two moved together. Between these findings and the effects of return horizon, volatility of fundamentals is a poor explanation for return volatility, while trading activity appears a promising explanation.

2.2 Robustness of the rise in volatility

The rise in market volatility is robust to several robustness concerns, as shown in Table 2, including outliers, the starting date of the trend, and stale prices.

[7]Newey-West HAC with 60 monthly lags. See Appendix A.2 for full specifications including the structural-break version with a 1960 trend break.

2.2.1 Outliers

It is well known that volatility is right skewed with some extreme events, such as the crash of 1987. Thus, our finding of increased volatility could potentially be driven by several extreme events. To address this, we compute weekly realized variances using daily returns, then take the interquartile range of these weekly variances within each rolling 5-year window. We rescale this to annualized volatility to get the dependent variable in the trend regression. Column (1) of Table 2 shows that the trend is almost identical at 11.3%. Thus, the rise in market volatility is not driven by outliers as events such as the 1987 crash or the Covid crisis have no effect on this alternative measure of volatility.

2.2.2 Trend robustness

Pension funds and other institutional investors were increasingly holding and trading market portfolios starting in the 1960s, with the first index fund (Vanguard) becoming available in 1976. In the early 1980s, index futures and options become available and increasingly traded. It is therefore natural to ask if the start date of the trend regressions are critical for documenting the rise in volatility. Column (2) of Table 2 shows the trend regression with a 1980 start date. The trend variable is the same, so the magnitude of the coefficient is still scaled as a 70-year change to be easily comparable to that in the full postwar sample regression. The rise in volatility with the 1980 starting date is slightly stronger than, but overall very similar to, the full sample trend at a little over 12%.

While we estimate a trend, we do not view it literally but as a convenient way to summarize the slow-moving rise in volatility over the postwar sample. This approach follows how Campbell et al. (2001) characterize idiosyncratic volatility. In Appendix A.4 we instead estimate structural breaks in the mean of market volatility and find two breaks dated around 1975 and 2000. The trend is also stable across start dates, rolling-window lengths, and 30-year rolling subsamples (Appendix A.5).

In each case, the mean of volatility increases substantially from the earlier mean and the overall increase in the mean is consistent with the trend result.

2.2.3 Concerns about stale prices

Volatility can appear low if there is no trade in a significant fraction of stocks on a given day. To address this, we weight stocks by total dollar volume on a given day, rather than weight by market capitalization. This ensures stocks that don't trade get zero weight and more liquid, heavily traded stocks get higher weight. Column (3) of Table 2 shows that volume weighting produces a similar trend. Column (4) restricts to the top 100 stocks by market capitalization, which are the most liquid and least susceptible to stale-price issues, and finds a nearly identical increase. Campbell et al. (1993) reach the same conclusion regarding nonsynchronous trading in their analysis of volume and return autocorrelations: re-running on the Dow Jones Industrial Average and on individual large stocks confirms the pattern is not driven by stale prices. The rise is also not an artifact of value-weighting or mega-cap concentration: it holds under equal weighting, across size segments, and for the S&P 500 directly (Appendix A.6).

2.2.4 Systematic vs. idiosyncratic volatility

The rise in market volatility is primarily systematic rather than firm-specific. We follow the construction of Campbell et al. (2001) (henceforth CLMX) exactly, decomposing average-stock variance into market, industry-relative-to-market, and firm-specific components on the full CRSP universe, with industry returns constructed from CRSP firms across the 49 industries used in their paper. We start the sample in 1962 to match CLMX directly and extend it through 2024; our implementation replicates their headline numbers within tolerance on their original 1962–1997 sample.[8] The market component's share of average-stock variance has nearly doubled, from 0.18 in 1962–

[8]See Appendix A.1 for the validation table, figures, and additional results.

Table 2: Volatility Trend Robustness

	IQR (1)	Post-1980 (2)	Volume-Weighted (3)	Top 100 (4)
Constant	8.51	8.45	12.47	10.28
	(0.67)	(2.25)	(0.94)	(0.84)
Trend	11.28	12.15	12.48	10.04
	(1.90)	(3.75)	(3.22)	(2.13)
p-value	$<$0.001	0.006	$<$0.001	$<$0.001
N	3,705	2,401	3,653	3,653
R^2	0.542	0.260	0.360	0.447
Vol increase	133%	93%	100%	98%

Notes: Each column regresses 5-year rolling annualized volatility (in %) on a constant and a linear time trend. The trend is normalized to zero at the start (January 1955) and one at the end of the full sample for all columns, ensuring comparability of slope coefficients. Column (1) uses the interquartile range (IQR) of weekly realized variances based on daily returns within each 5-year window, scaled by the $\chi^2(5)$ IQR to recover annualized volatility. Column (2) restricts the sample to January 1980 onward using the baseline daily-return RV measure from Table 1. Column (3) replaces the market-capitalization-weighted market portfolio with a dollar-volume-weighted portfolio constructed from individual CRSP stocks, using lagged dollar volume as weights. Column (4) uses a value-weighted portfolio of the 100 largest stocks by lagged market capitalization. Standard errors (in parentheses) use the Newey–West estimator with Bartlett kernel and $S = 260$ lags. p-values use fixed-b critical values with $\kappa = \lfloor 3/(2b) \rfloor$ degrees of freedom. The last row reports percentage increase in trend volatility from the beginning to the end of the sample.

1971 to 0.31 in 2015–2024. The total systematic share (market plus industry) has risen from 0.31 to 0.45 over the same period. Both trends are economically and statistically significant: 2.50 percentage points per decade for the market share ($t = 3.95$) and 3.25 percentage points per decade for the total systematic share ($t = 4.01$).

The variance-ratio result in Figure 5 also concentrates in the systematic component. Decomposing the 5-year variance ratio by CLMX component, the market-component variance ratio more than doubles from 0.72 in 1961–1995 to 1.51 in 1996–2024; the industry component follows a similar but smaller shift. The firm-specific variance ratio shows no regime change – it declines monotonically over the long sample, consistent with declining microstructure noise (narrower spreads, decimalization). The regime change in the variance ratio, like the rise in volatility, sits in the systematic component.

Our 1962–1997 numbers reproduce the headline CLMX result on rising firm-specific volatility. The same authors revisited the data in Campbell et al. (2023) and documented a reversal post-2001 – a pattern our extended sample confirms. The persistent feature across the full 1962–2024 sample is the rise in the market component, while firm-specific volatility has fluctuated around its long-run level.

2.2.5 Additional Results on Horizon Effects and Return Autocorrelations

Figure 4 plots the increase in volatility as a function of return horizon, going from 1 day out to 1 quarter (13 weeks). The increase is significant out to around 6 weeks, though the point estimate is positive for all return horizons. The gradual decay is consistent with short-run mean reversion or reversals in returns.

Figure 5 shows the rolling ratio of 5-year return variance based on daily returns to 5-year return variance based on monthly returns. This difference is driven by the autocorrelation pattern. If daily returns-based variance is higher than monthly returns-based variance, there is mean reversion in the price of the market. In other words, short-term price pressure that mean-reverts quickly will show up in daily return variance, but not (as much) in monthly return variance. The figure shows the variance

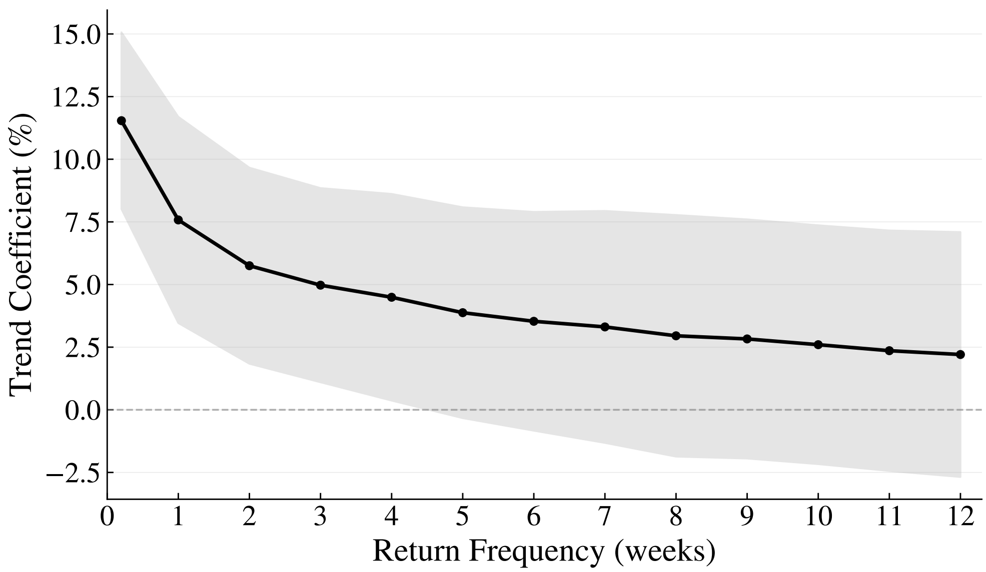

Figure 4: Time Trend in Realized Volatility by Return Frequency

Notes: Trend coefficient from regressing 5-year rolling annualized volatility on a linear time trend, as a function of return frequency. The x-axis shows the return horizon: daily (plotted at 0.2 weeks) and 1–12 week overlapping returns. Excess returns at each horizon are constructed by compounding gross market and risk-free returns separately, then differencing. The 5-year rolling window uses 1260 trading days (daily) or 260 weekly observations. Standard errors use the Newey–West estimator with Bartlett kernel and $S = 624$ lags (12×52 weeks). The shaded band shows 95% confidence intervals using fixed-b critical values with $\kappa = \lfloor 3/(2b) \rfloor$ degrees of freedom. Sample: 1955–2025.

ratio over the sample plotted against the daily return volatility from Figure 1. The two track each other quite closely, and the quantitative changes in the variance ratio are large, going from around 0.6 to over 2. Thus, the trading effects this pattern suggests are quantitatively large. The variance ratio and the level of volatility both break at March 2000 in independent structural-break tests (Appendix A.4), and the same variance-ratio rise appears in 35 of 35 international markets with very different fundamentals (Appendix A.7).

Section A.8 plots the underlying daily and weekly return autocorrelations directly and extends the volume-conditioned regression of Campbell et al. (1993) (CGW) to the modern sample. The daily autocorrelation falls from roughly $+0.30$ in the 1970s to -0.15 by the 2020s, crossing zero around 2003; the weekly series crosses zero around 1990 and settles near -0.10. Both trends are strongly significant under fixed-b inference. The CGW volume-reversal mechanism – high-volume returns more strongly reversed – holds throughout the full 1955–2025 sample and continues operating after CGW's own 1962–1987 window.

2.2.6 Implied volatility

A natural concern is that the VIX, the most widely-cited measure of equity market volatility, has not appeared to trend up over the post-1990 period. In light of Dew-Becker and Giglio (2024), this pattern is unsurprising: they document a substantial decline in the variance risk premium over this period and attribute it to declining intermediary frictions in the options market. In our data, realized volatility trends up by roughly 1.8 percentage points per decade since 1990, the variance risk premium declines at an almost exactly offsetting rate, and implied volatility shows no substantial trend up. The implied-volatility sample is short, about 30 years rather than the 70-year sample we use for realized volatility, so the power to detect a trend in implied volatility on its own is limited. Even so, the joint pattern across the three series is consistent with the rise in volatility we document. Appendix A.3 reports the corre-

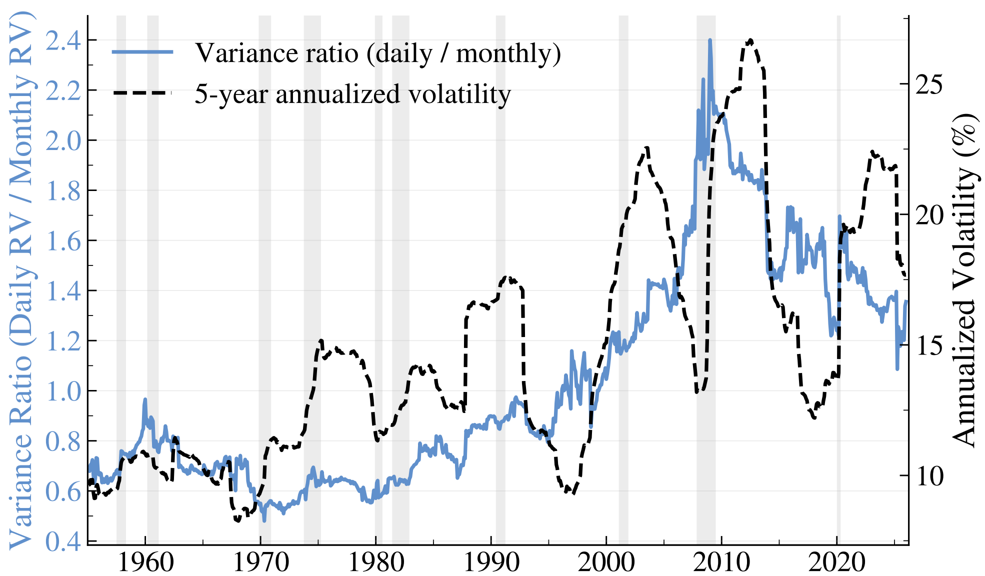

Figure 5: Variance Ratio and Volatility

Notes: Left axis (solid): variance ratio defined as the ratio of 5-year cumulative daily realized variance to 5-year cumulative monthly realized variance. Daily realized variance within each month is the sum of squared demeaned daily returns; monthly realized variance is the squared demeaned monthly return. Both are demeaned using full-sample means. A ratio above 1 indicates negative return autocorrelation (mean reversion); below 1 indicates positive autocorrelation. Right axis (dashed): 5-year rolling annualized volatility as in Figure 1. Shaded regions indicate NBER recessions. Sample: 1955–2025.

sponding figure and regressions.

3. Volume and Volatility: Stylized Facts and Causal Evidence

We propose that the market is more volatile because investors trade the overall market a lot more than they used to. That is, with the rise of S&P500 index funds, S&P500 futures, or similar products, it has become increasingly easy and prevalent over the past 50 years to trade in and out of the overall market rather than in individual stocks.

This concentration of trading volume in the overall stock market means that investors facilitating trade – e.g., arbitrageurs and liquidity providers – will have to absorb market risk, rather than idiosyncratic risk of individual stocks. Their compensation is given by the degree of short-term reversal, with higher reversal leading to more volatility of high frequency returns.

To explore the prevalence for "trading the market" we propose a simple measure that assesses how much volume for individual stocks can be explained by volume of trade in the overall stock market. More specifically, suppose all volume activity in a given day came from flows in and out of the overall stock market – for example flows in and out of market index funds. Then volume on that day for all stocks would be proportional to their weights in the market index – that is, their market capitalization.

This motivates the following cross-sectional regression. For a given time period t, we run

$$\frac{Volume_{i,t}}{Total\ Volume_t} = a_t + b_t \times MktWeight_{i,t-1} + \varepsilon_{i,t} \quad i = 1, ..., N \tag{2}$$

where $Volume_{i,t}$ represents the dollar volume in stock i in day t, $Total\ Volume$ represents the total dollar volume across all stocks, and $MktWeight_{i,t-1}$ represents stock i's weight in the market portfolio. If all trading on a given day was trading by market index funds, then we would have $b_t = 1$ because volume in every stock would be proportional to that stock's market capitalization.

We run the regression in (2) daily, collect estimates b_t, and then compute rolling averages of the time-series of b_t.

Figure 6 plots the rolling 5 year average of volume coming from trading the market along with market volatility. Market volume betas double over this period from around 0.4 in 1960 to around 0.8 today, coinciding with the rise in market volatility. Recall that a market volume beta of 1 is consistent with all trades being trades in and out of the market portfolio, which suggests that today a huge fraction of trading volume in the stock market may be driven by trading the market index.

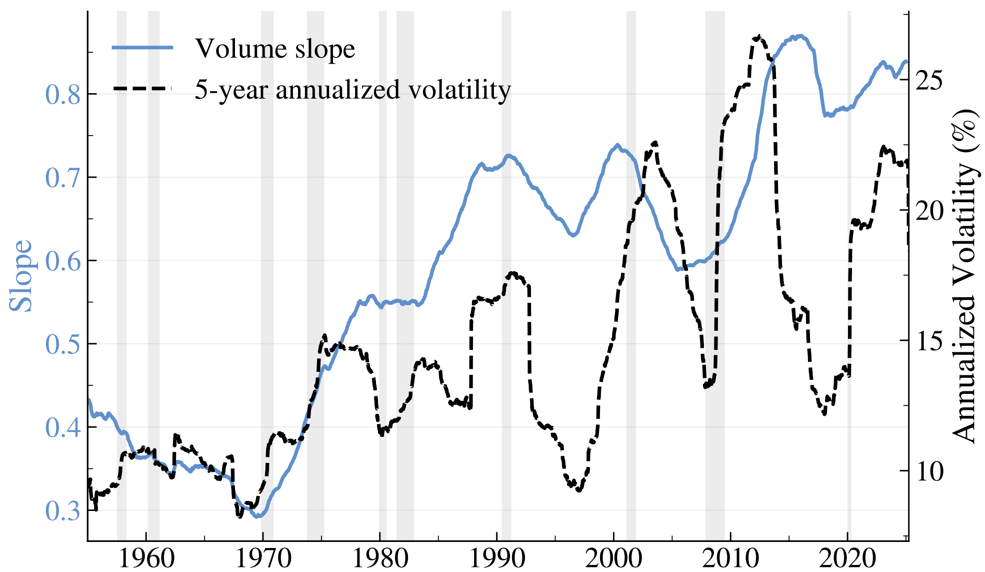

Figure 6: Market Trading vs Market Volatility

Notes: Left axis (solid): 5-year rolling average of the daily cross-sectional OLS slope from regressing volume shares on value weights within the top 500 stocks by market capitalization. A higher slope indicates that trading volume is more concentrated in stocks with larger market-cap weights, consistent with index-linked trading. Right axis (dashed): 5-year rolling annualized volatility as in Figure 1. Shaded regions indicate NBER recessions. Sample: 1955–2025.

3.1 Causal Evidence: Volume and Volatility

While the link between the volatility of the overall market and volume in the stock market is suggestive, they of course do not point to a causal channel. We argue that it is natural that the prevalence and ease of trading in market index products drove significant trading activity in the market itself which increased volatility. We now exploit quasi-natural experiments that test this view. The first uses changes in the trading hours of the NYSE as an instrument for trading volume. The second exploits the E-mini futures as a popular way to trade the market index. The third uses a triple-difference design to trace the gradual buildup of overnight index trading over decades. The fourth exploits CME's decision to open futures on NYSE holidays, a within-year comparison in which the main change is whether the index can be traded.

An important consideration in choosing these experiments is the strength of the first stage. Many of the most visible events in the history of index trading—the launch of Vanguard's index fund in 1976, S&P 500 futures in 1982, the SPY ETF in 1993—did not produce immediate surges in trading volume. These products started small and took years or decades to gain widespread adoption, so a short-window event study around their launch dates has little power. Our experiments instead exploit changes to existing trading behavior: closing a trading day that was already active, extending hours for traders who already know how to trade, or making an existing futures contract cheaper and electronic via the E-mini. In each case, the first stage is strong because the change is immediate. To capture the effects of index trading as it spread into the overnight hours, we turn to the triple-difference design in Section 3.1.3, which pairs a cross-sectional prediction—systematic versus idiosyncratic variance, overnight versus intraday—with the discrete expansions of futures trading hours, and then cumulates the estimated per-hour effect over the full sample.

3.1.1 Changes in Trading Hours as an Instrument for Volume

We exploit changes in NYSE trading hours and trading days, as instruments for trading activity. The basic idea goes back to work by French and Roll (1986). They exploited a paperwork backlog at the NYSE in 1968 which led the NYSE to close on Wednesdays from June-December 1968. The amount of information content relevant for stock prices presumably did not change because of this closure. During the closure period, the return on Thursdays is now a 2-day cumulative return – that is, the return Thursday reflects the price movements from the close on Tuesday to the close on Thursday.

Strikingly, however, French and Roll (1986) find that the *variance* of the Thursday return did not change due to the closure, when in fact it should double.[9] As a result, total weekly variance (the sum of squared returns in a week) *falls* during the closure period by around 20% (one-fifth, given one in five trading days is removed). This decrease in variance is relative to the period of January through May of 1968 (the first half of the year) before the Wednesday closure was in place. Further, the variance on other days (i.e., Tuesday) did not change during the closure period, so that it does not appear there were any general changes in the level of volatility.

We take these results and push them further. Specifically, we exploit 9 changes in trading hours on NYSE, 1952-1985. These include: the 1952 cut from 6 to 5 trading days per week (trading used to occur on Saturdays), the 1968 Wednesday closure, and the subsequent extensions of the trading day from 4 hours of trading per day in 1969 to eventually 6.5 hours today.

We present two versions of this analysis. The cleanest test is at the daily level (Panel A of Table 3), where each schedule change generates specific predictions about which days are affected and we can decompose total elapsed time into *trading hours* and *calendar hours*. This decomposition lets us ask the question: does trading time

[9] If the returns on any two days are roughly uncorrelated, variance of a 2-day return will be twice the variance of a 1-day return. Letting r_t be the log daily return, we have $var(r_t + r_{t+1}) = 2var(r)(1 + \rho)$ where ρ is the autocorrelation of returns.

matter for volatility above and beyond calendar time? The weekly aggregate (Panel B) confirms the same elasticity persists when we integrate over the full calendar week, where trading hours and calendar hours are mechanically conflated and the test is correspondingly less sharp.

Panel A: Daily IV. Each NYSE schedule change generates a specific prediction about which days are affected. The 1968 Wednesday closure means Thursdays now span two calendar days instead of one, while Tuesdays are unaffected. The elimination of Saturday trading in 1952 means Mondays now span the full weekend rather than returning from a half-day Saturday session. We exploit this within-week heterogeneity by working at the daily level, separating two distinct components of the time between market closes: *trading hours* (hours the market is open) and *calendar hours* (total hours elapsed from the previous close to the current close). These vary independently across event types. When the NYSE eliminates Saturday trading, Monday's calendar hours jump while trading hours are unchanged. When the NYSE extends closing time, trading hours increase but calendar hours remain at 24 for weekdays.

Under the null hypothesis that information arrival drives variance, daily squared returns should scale with calendar time regardless of trading hours – a two-day return should have twice the variance of a one-day return. Under the alternative that trading generates variance, more trading hours should increase variance beyond what calendar time predicts. We estimate:

$$\ln|r_t| = fe + b \times \ln \text{turnover}_t + \beta_c \times \ln \text{calendar hours}_t + \varepsilon_t \quad (3)$$

$$\ln \text{turnover}_t = fe + \gamma \times \ln \text{trading hours}_t + \gamma_c \times \ln \text{calendar hours}_t + \nu_t \quad (4)$$

where $|r_t|$ is the absolute daily market return, turnover is daily dollar volume divided by market capitalization, and log trading hours instruments for log turnover. We use log absolute returns rather than log squared returns so that coefficients are elasticities of volatility, directly comparable to Panel B. We include fixed effects fe, for each event, day-of-the-week, and individual holidays (a separate dummy for each post-holiday day, which prevents holidays from contaminating the calendar hours coefficient). We use ± 22 trading day windows (roughly one month) around each schedule change, yielding 9 events and 403 daily observations. Newey-West standard errors with 10 lags account for serial correlation.

Panel A of Table 3 presents the results. The first stage (column 3) is very strong: a 10% increase in trading hours increases turnover by about 13%, with an F-statistic of 57 on the excluded instrument. The IV estimate in column 2 implies an elasticity of daily volatility with respect to turnover of 1.2, significant at the 5% level. The OLS estimate in column 1 is similar at 1.3, suggesting minimal attenuation bias at the daily frequency. Column 4 presents the reduced form: log trading hours enters with a coefficient of 1.5, significant at the 5% level, while the coefficient on log calendar hours is 0.7 and statistically insignificant. Under the information-only null, the trading hours coefficient should be zero and the calendar hours coefficient should be 0.5 (since the dependent variable is log volatility, not log variance). We test this joint null directly: $F(2, 369) = 3.61$, $p = 0.028$, rejecting at the 5% level. The rejection is driven by the trading-hours coefficient – the marginal $\delta_T = 0$ test rejects at $p = 0.015$, while the marginal $\delta_C = 0.5$ test cannot reject ($p = 0.70$) because the individual holiday and day-of-week fixed effects absorb nearly all the variation in calendar hours. Trading hours have a large positive effect on volatility above and beyond the calendar-time scaling that information arrival alone would generate.

Panel B: Weekly aggregate. The daily analysis exploits within-week heterogeneity to identify the trading-hours channel separately from calendar time. We can also ask whether the effect persists at the weekly aggregate, where trading hours and calendar hours are mechanically conflated and the test is less sharp. We compute total trading hours per week as our instrument. The first stage runs log stock market volume (scaled by market capitalization) on log trading hours. The second stage esti-

Table 3: IV: Instrumenting volume with trading hours

Notes. Panel A uses daily data. The dependent variable is log absolute daily market return, so that coefficients are elasticities of volatility (not variance). Turnover is daily dollar volume divided by market capitalization. Trading hours and calendar hours (hours from previous close to current close) are measured for each trading day. The sample uses ± 22 trading day windows around 9 NYSE schedule changes, 1952–1985. Panel B uses weekly data. The dependent variable is log realized volatility (annualized) of the weekly value-weighted market return. Volume is weekly dollar volume divided by market capitalization, as in Panel A. Trading hours is total trading hours per week. The sample includes 6-month windows around each of the 9 changes in NYSE trading hours. Standard errors in parentheses: Newey-West with 10 lags in Panel A (with event, day-of-week, and individual holiday FEs); Newey-West with 2 lags in Panel B (with event and month-of-year FEs). ***, **, * denote significance at the 1%, 5%, and 10% level.

Panel A: Daily

	(1) OLS	(2) IV	(3) First Stage	(4) Reduced Form
Dep. variable:	$\ln \lvert r \rvert$	$\ln \lvert r \rvert$	ln Turnover	$\ln \lvert r \rvert$
Log Turnover	1.30***	1.16**		
	(0.26)	(0.47)		
Log Trading Hours			1.27***	1.47**
			(0.17)	(0.61)
Log Calendar Hours	−0.05	0.08	0.55***	0.72
	(0.57)	(0.64)	(0.13)	(0.56)
Observations	403	403	403	403
R-squared	0.37		0.92	0.32
First-stage F		57.0		
Joint F (info-only null)				3.61**
FE: Event, DOW, Holiday	Y	Y	Y	Y

Panel B: Weekly

	(5) OLS	(6) IV	(7) First Stage	(8) Reduced Form
Dep. variable:	$\ln \sigma$	$\ln \sigma$	ln VM	$\ln \sigma$
Log Volume	0.72***	1.42*		
	(0.12)	(0.82)		
Log Trading Hours			1.02***	1.45**
			(0.34)	(0.67)
Observations	326	326	326	326
R-squared	0.54	0.49	0.90	0.49
First-stage F		9.0		
FE: Event & Month	Y	Y	Y	Y

mates the relationship between log volatility and log volume. Volatility is computed as the square root of the sum of squared daily returns within the week, annualized.

$$\ln \sigma_t = a_m + a_e + b \times \ln VM_t + \varepsilon_t \tag{5}$$

$$\ln VM_t = a_m + a_e + \gamma \times \ln \text{trading hours}_t + \nu_t. \tag{6}$$

We use trading hours *only* in 6-month windows before and after each change as the instrument for weekly volume (VM). Restricting to shorter windows is important because trading hours have trended up over time along with volatility and volume; further, the levels of volume and of volatility are different around each of these changes. We include event dummies, a_e (equal to 1 in the 6 months before and after a change), and month-of-year fixed effects, a_m, to control for seasonal patterns. Newey-West standard errors with 2 lags account for serial correlation; 2 weekly lags is the same autocorrelation horizon as the 10 daily lags used in Panel A.

Our identifying assumptions are as follows. First, increases in trading hours only affect volatility through increased trading activity. Second, the decision to increase trading hours does not depend on an expected increase in volume 6 months before / after the official change. We do not assume that increased demand for trading doesn't affect trading hours in general, but that a decision to permanently increase trading hours does not depend on a predicted spike in volume during the event window. The second assumption is most appropriate when the event window is relatively narrow. Our baseline uses 6 months because we face a tradeoff in statistical precision in using realized volatility. Results are robust at tighter 2-month windows (IV coefficient 2.18, significant at 5%) and at longer 12-month windows (IV coefficient 1.69, significant at 1%).

Panel B of Table 3 presents the weekly results. Column 3 is the first stage. Trading hours increase volume nearly one to one – when the stock market is open longer, people trade more. The first stage shows that trading hours are a strong instrument for volume. Column 2 finds a positive causal effect of volume on volatility: a 10% increase in trading hours increases volume by 10% and leads to about a 14% increase in volatility. The point estimate is significant at the 10% level. Column 1 presents OLS estimates of volatility on volume and finds a similarly strong positive relationship. Column 4 regresses volatility directly on trading hours; since the first-stage coefficient was 1, the coefficient in column 4 is about the same as that in column 2 and significant at the 5% level. The weekly elasticity of 1.4 is close to the daily Panel A estimate of 1.2, confirming the result is consistent across frequencies. The weekly IV is noisier than the daily because aggregating to the calendar-week level conflates trading time with calendar time – the very decomposition that makes the daily Panel A sharp.

Overall, the evidence strongly supports a positive causal relationship in which increased trading activity drives higher volatility of the stock market index. The daily specification sharpens identification by exploiting the specific daily predictions of each schedule change and delivers a first-stage F-statistic of 57 – well above conventional thresholds for instrument strength – and lets us reject the information-only null in favor of trading hours per se. The weekly specification confirms the elasticity at the calendar-week aggregate, where the trading-vs-calendar-hours decomposition is no longer available but the result is qualitatively unchanged.

3.1.2 The E-mini Futures as a Natural Experiment

We exploit the introduction of the E-mini futures on Sept 9, 1997, which led to a large amount of trading on the S&P500 index. Importantly, the E-mini could be traded during non-trading hours, unlike individual stocks at the time it was introduced. We exploit this differential trading by comparing the variance of the S&P500 overnight vs intraday relative to the same comparison for individual stocks in the S&P500 (idiosyncratic vol). We make this comparison before and after the E-mini is introduced. The idea is that more trading activity overnight results in more volatility for the market overnight but not for individual stocks. The relative comparison controls for overall news coming out at night that would move stocks in general, and other factors that would move the overall volatility of the market during this period (either during the day or night).

We use close to open and open to close returns as measures of overnight and intraday returns, respectively.[10] The market return is the value-weighted return of the top 500 stocks by prior-December market capitalization. For each individual stock i, we use the simple market-adjusted return $r_{i,t} - r_{m,t}$ as a proxy for the idiosyncratic component. This imposes a unit market beta on each stock, following Campbell et al. (2001), who show this avoids the noise of beta estimation while still yielding a valid variance decomposition at the aggregate level.[11] We compute the sum of squared returns within a rolling 22 day period (roughly 1 month). For individual stocks, we aggregate into a single measure of idiosyncratic volatility by value weighting the resulting overnight and intraday idiosyncratic variances at the stock level. We then compute $\left(var(r^{night}_{mkt})/var(r^{day}_{mkt})\right)$ / $\left(var(r^{night}_{stock,idio})/var(r^{day}_{stock,idio})\right)$. This represents the ratio of the two variance ratios.

Figure 7 plots the results. Panel A shows the increase in this variance ratio. The first vertical line represents the introduction of the E-mini in September 1997. Since our variance computations use 22 day rolling windows, the full effect is not felt until the second vertical line, which then only uses data since the E-mini introduction. The red dashed line estimates the means before and after using the 3-month window in the plot, with a linear transition between to reflect the transition window induced by rolling returns. It is clear from the plot that the relative variance of the S&P500 increased overnight after the introduction of the E-mini.

However, for our story to be correct, we also need that the E-mini generated sig-

[10]Thus, we also use weekend returns. Returns are constructed using CRSP's split- and dividend-adjusted daily total return combined with raw same-day open/close prices, which gives correct overnight and intraday returns through corporate actions.

[11]Section 3.1.3 below uses the identical unit-beta decomposition and return construction, so the E-mini event study and the overnight panel are directly comparable.

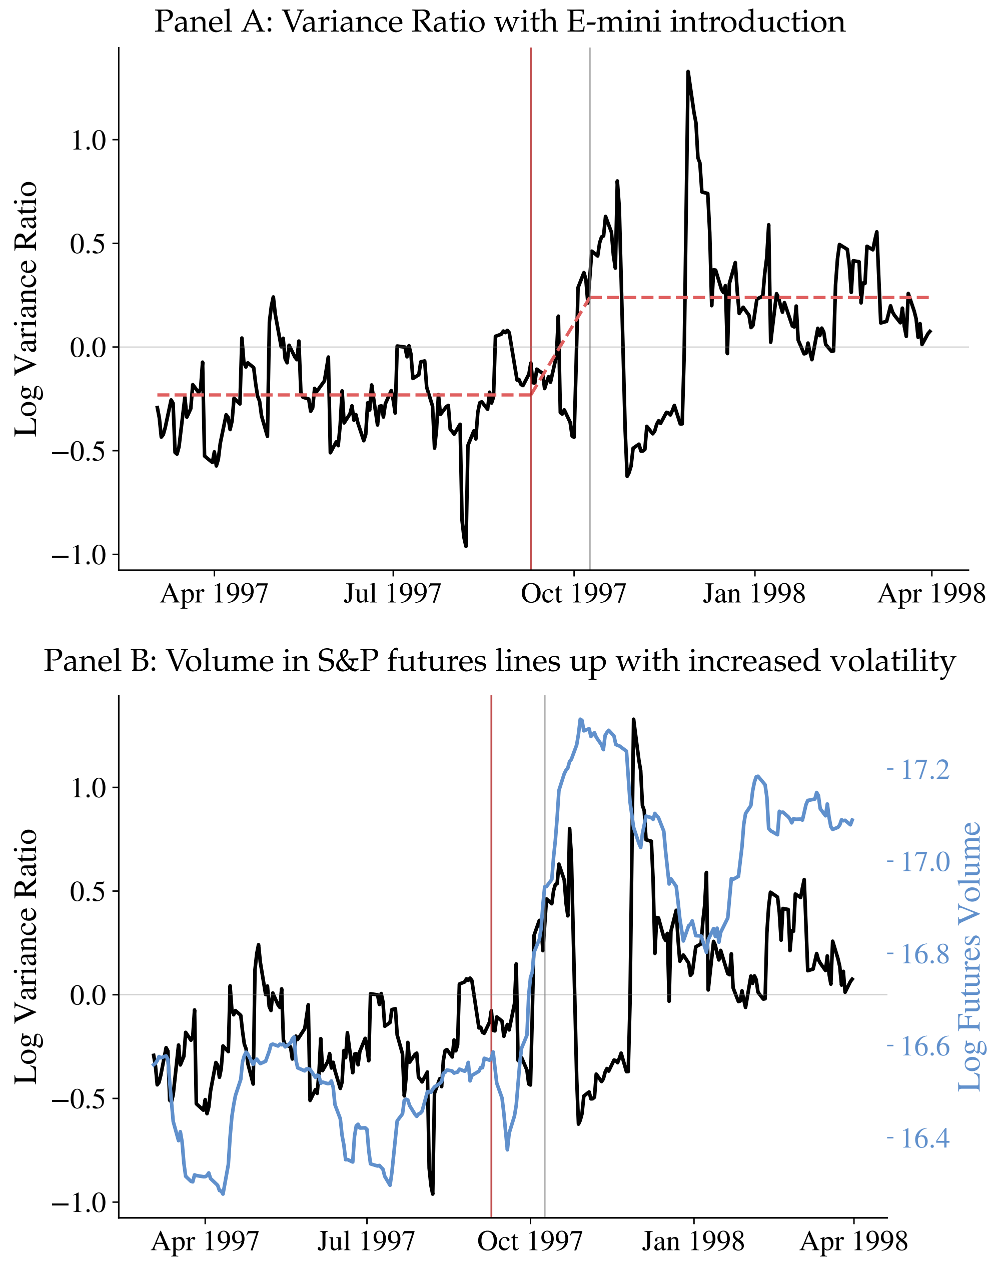

Figure 7: The Introduction of the E-mini Futures and the Effects on Volatility.

Panel A plots the log of the ratio of overnight-to-intraday variance for the market relative to individual stocks, using 22-day rolling windows around the E-mini introduction (Sept 9, 1997). The red dashed lines show pre- and post-event means with a ramp connecting them. Panel B adds the log of total S&P500 futures volume (22-day rolling average) on the right axis.

nificant trading in the S&P500 when it was introduced. Panel B adds the (log) total volume of S&P500 futures contracts and shows they indeed rise with the introduction of the E-mini. We use a 22 day rolling average of total volume in S&P500 futures contracts to match the window at which volatility is computed. Thus, we see the rise in relative volatility coinciding with a large increase in trading activity in the S&P500, consistent with our story.

Table 4 tests the increase formally. Because the dependent variables are 22-day rolling averages, a simple pre/post dummy would misrepresent the timing of the shift — on September 10 only 1 of 22 observations in the rolling window reflects the post-launch regime. We therefore use a ramp variable that equals zero before the E-mini launch, increases linearly from 0 to 1 over the next 22 trading days, and equals one thereafter. The dependent variable in column (1) is the log ratio of overnight-to-intraday variance for the value-weighted market return, controlling for the corresponding log ratio computed from idiosyncratic stock returns. The ramp coefficient of 0.43 ($t = 7.9$) implies the market variance ratio increased by a factor of $e^{0.43} \approx 1.5$ after the E-mini launch, holding stock-level patterns fixed. Column (2) confirms that log futures volume increased by 0.60 over the same window.

3.1.3 The Rise of Futures and Overnight Variance

More generally, the volume of the E-mini futures has exploded since it was introduced, leading to a huge amount of trading volume in the market index. Because futures are traded nearly 24 hours a day, this can also contribute to a general increase in volatility in overnight returns due to increased trading activity. Figure 8 plots the volatility of the market return using 5-year rolling windows as in our main plots earlier, but it splits returns into overnight (close to open) and intraday (open to close). A substantial amount of the increase in volatility over the past 30 years comes from increased volatility of overnight returns.

This rise in overnight volatility has coincided with a large increase in overnight

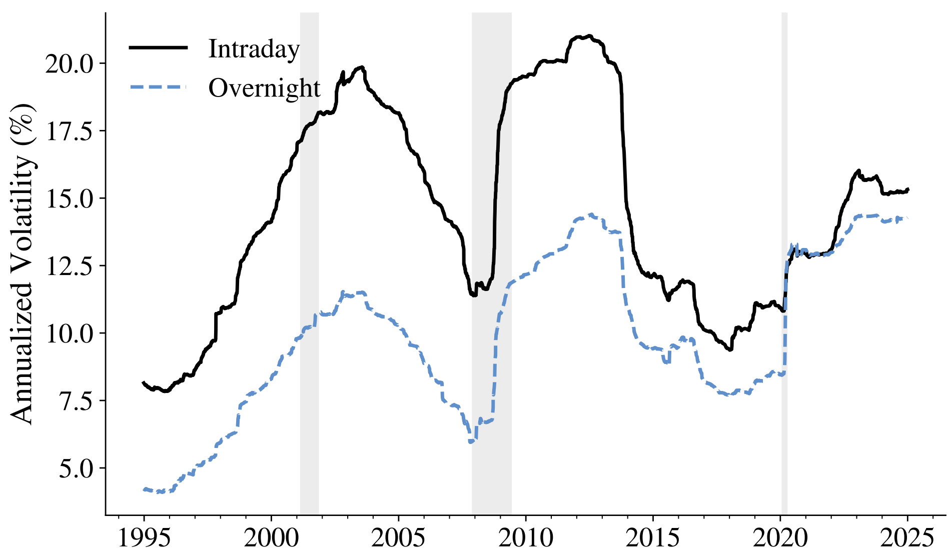

Figure 8: The Rise of Overnight Volatility on the Market.

5-year rolling standard deviation of the value-weighted market return, split into overnight (close-to-open) and intraday (open-to-close) components.

Table 4: E-mini Futures Introduction and Overnight Volatility

Notes. The dependent variable in column (1) is the log ratio of overnight-to-intraday variance for the value-weighted market return (top 500 stocks by prior-Dec market capitalization), controlling for the corresponding log ratio of idiosyncratic stock returns, value-weighted across stocks. The idiosyncratic return is defined in the text. Column (2) uses log total S&P 500 futures volume. Both dependent variables are 22-day rolling averages. The E-mini ramp variable equals 0 before September 9, 1997, increases linearly from 0 to 1 over the next 22 trading days, and equals 1 thereafter, matching the rolling window over which the effect builds. The sample window is March 1997 to March 1998. HAC standard errors (Newey-West, 44 lags) in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% level.

	(1) Log Market Var Ratio	(2) Log Volume
E-mini Ramp	0.434***	0.599***
	(0.055)	(0.048)
Log Idio Stock Var Ratio	0.858***	
	(0.169)	
Observations	273	273
R-squared	0.528	0.839

futures trading activity. Figure 9 plots the ratio of overnight to daytime trading activity in S&P 500 E-mini futures, measured by both tick count and contract volume. Overnight activity has grown steadily relative to daytime activity since the E-mini launch in 1997, consistent with growing around-the-clock demand for trading the market index.

But is this rise in overnight volatility driven by the kind of trading our theory predicts? We now decompose overnight returns for individual stocks into a systematic component (driven by the market) and an idiosyncratic component (stock-specific). Our prediction is sharp: because index futures allow investors to trade the market overnight but not individual stocks, the systematic component of overnight variance should rise over time while the idiosyncratic component should not.

Setup. We construct four daily variance series from the 500 largest U.S. common stocks by market capitalization, re-selected at the start of each year from 1993 to 2024: market overnight, market intraday, idiosyncratic overnight, and idiosyncratic intra-

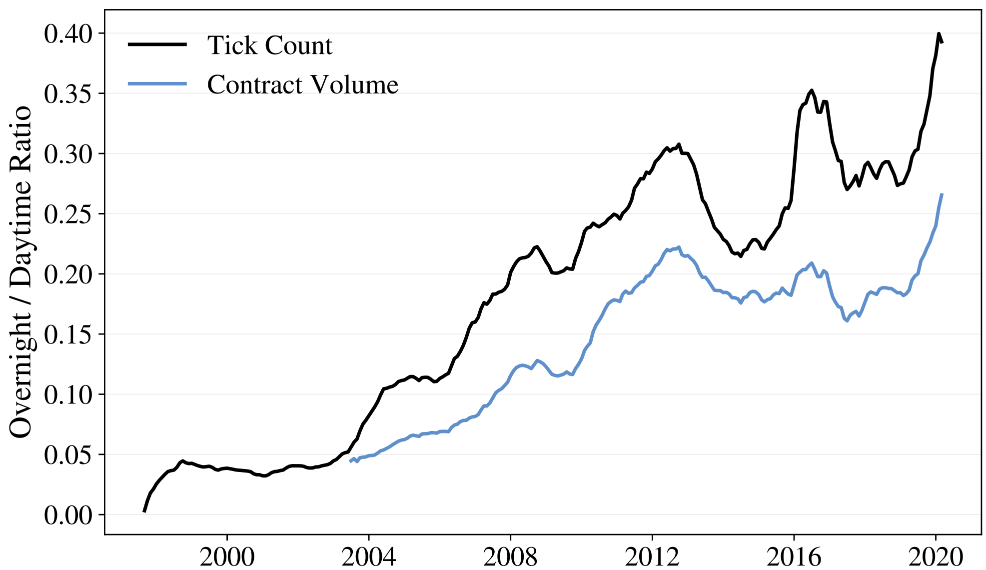

Figure 9: Overnight-to-Daytime Trading Activity in S&P 500 Futures

Notes: Ratio of overnight to daytime trading activity for S&P 500 E-mini futures, measured by tick count (black) and contract volume (blue). Overnight is defined as trading outside regular NYSE hours (9:30am–4:00pm ET). Both series are smoothed using rolling averages. The steady increase in overnight relative to daytime activity is consistent with the growing role of around-the-clock index trading following the E-mini introduction in September 1997. Tick-count data begin in 1997; contract volume data begin in 2004. Sample: 1997–2020.

day. The market return is the value-weighted return across the top 500 stocks, and idiosyncratic returns impose a unit market beta ($r_{i,t} - r_{m,t}$), exactly as in the E-mini event study above; overnight and intraday returns use the same construction anchored on CRSP total returns. For each trading date we record the log of each squared return (value-weighted across stocks for the idiosyncratic series), giving a stacked panel of four log-variance cells per date—8,054 dates and 32,216 observations.

The object of interest is the daily triple-difference

$$\mathrm{TD}_t \;=\; \big(\log V^{\mathrm{ovn}}_{\mathrm{mkt},t} - \log V^{\mathrm{intra}}_{\mathrm{mkt},t}\big) \;-\; \big(\log V^{\mathrm{ovn}}_{\mathrm{idio},t} - \log V^{\mathrm{intra}}_{\mathrm{idio},t}\big), \tag{7}$$

the log ratio of overnight-to-intraday market variance relative to the same ratio for idiosyncratic returns—the same object plotted in Figure 7. Same-day intraday variance controls for anything that moves overall volatility on a given day; the idiosyncratic ratio controls for market-wide changes in overnight information flow that affect stocks in general. What remains loads specifically on the market's overnight window—the cell where index futures trade and individual stocks do not.

Identification. We exploit three discrete changes in CME overnight futures trading hours: the launch of the E-mini in September 1997 (overnight index trading goes from zero to 8.5 hours), the expansion of the Globex electronic session in July 2003 (8.5 to 17 hours), and the introduction of a daily maintenance window in November 2012 (a small *reduction* of half an hour). The design mirrors the NYSE trading-hours analysis of Table 3: pool the events, absorb event fixed effects, and identify the average per-hour effect from within-event variation. Smoothly evolving confounds—growing demand for around-the-clock trading, globalization of information flow, the secular changes in market structure over these decades—cannot produce step responses at CME schedule dates. Parallel trends support the design: over the 24 months before the events, TD_t shows no significant pre-trend, and what slope there is runs in the wrong direction to explain the post-event jumps.

Results. Panel A of Table 5 presents the results. Columns (1)–(3) estimate an event-study ramp regression at each event separately, using the same specification as Table 4; column (1) is the E-mini launch itself and closely replicates that result. The 2003 expansion carries a positive but imprecise coefficient (0.20, $t = 1.0$); the small 2012 *reduction* in hours is correctly signed (negative) and near zero. Column (4) pools the three events, regressing TD_t on Post $\times$ ΔHours with event fixed effects: overnight market variance rises by 5.7% per additional hour of futures trading ($t = 1.9$). With only three schedule changes, conventional standard errors overstate the effective sample, so we also report permutation inference: against 5,000 placebo event-date triples, the pooled estimate has a one-sided p-value of 0.077. All three events move in the predicted direction, and the estimate is insensitive to dropping any single event. The effect is sharpest in $\pm$6-month windows and fades at longer horizons, so we read the event evidence as identifying the per-hour effect locally rather than tracing out the full transition path.

Two measurement biases run *against* finding this pattern. With unit betas imposed, any true rise in overnight market variance leaks into measured idiosyncratic overnight variance, shrinking the triple-difference; and the migration of earnings announcements into after-hours trading over the 1990s and 2000s raises true idiosyncratic overnight news, again working against the market-overnight cell. Both push the estimates toward zero.

Accounting. Panel B asks how much of the long-run rise in overnight market variance the futures-hours channel can cover, using the full 1993–2024 sample. These are descriptive regressions rather than identification: over three decades the expansion of futures hours is collinear with everything else that changed in market structure, so we do not lean on them causally. Interacting the four cells with a linear trend, market overnight variance rises 2.1% per year relative to the three control cells ($t = 5.7$); interacting instead with futures hours under year$\times$month fixed effects gives 2.0% per hour ($t = 3.5$), the same order as the event-window estimate. Taken at face value, the pooled per-hour effect of 0.057, cumulated over the roughly 17-hour expansion of overnight futures trading, implies a factor of $e^{0.94} \approx 2.6\times$ on market overnight variance—most of the observed rise in the market-overnight cell over this period. Two ways to put this in units of the headline trend: the rise in overnight variance accounts for roughly half of the total rise in market variance from the early 1990s to the recent decade (total annualized volatility rose from about 11% to 17% over this window), and the channel-implied portion amounts to roughly a quarter of the total rise, with a wide confidence band.[12]

3.1.4 Holiday Natural Experiment

The overnight panel identifies the futures-hours effect from three discrete schedule changes. We now examine the same mechanism on a different margin: holidays.

NYSE has been closed on the same holidays—MLK Day, Presidents Day, Memorial Day, Independence Day, Labor Day, Thanksgiving—throughout our sample. What changed is that CME S&P 500 futures, which were also closed on these holidays before 2004, began trading on them as electronic sessions expanded.[13] Individual stocks remain untradeable on these days regardless: NYSE is closed, and extended-hours equity trading on holidays is negligible.

The prediction is the same as before. When futures open on a holiday, the overnight return spanning that day should carry more systematic variance—index trading now generates market-level price moves during the gap. Idiosyncratic variance should be unaffected, because individual stocks cannot trade on the holiday regardless of what futures do. This margin holds the stock market fixed by construction: NYSE is always

[12]The conversion from log points to variance levels depends on the averaging convention, because daily variance is heavily right-skewed: the quarter-of-the-total figure uses arithmetic means, the convention under which variance components add up to the total.

[13]Six holidays see futures go from closed to open around 2004. We drop Good Friday (always closed, no variation), Christmas and New Year's (sporadic opens, short sessions, double-holiday clusters). MLK Day enters the sample in 1997, when NYSE first observed it as a holiday.

Table 5: Futures Trading Hours and Overnight Market Variance

Notes: For each trading date we construct four log-variance cells: market overnight, market intraday, idiosyncratic overnight, and idiosyncratic intraday. Overnight and intraday returns use the clean construction anchored on CRSP total returns, correct through splits and dividends; the market is the value-weighted top 500 stocks by prior-December market capitalization; idiosyncratic returns impose a unit market beta, $r_i - r_m$, following Campbell et al. (2001) and are value-weighted across stocks. Sample: 1993–2024. Panel A, columns (1)–(3): for each CME futures trading-hours change, the dependent variable is the 22-day rolling log ratio of overnight-to-intraday market variance, controlling for the corresponding idiosyncratic log variance ratio; the ramp variable equals 0 before the event, rises linearly to 1 over the next 22 trading days, and equals 1 thereafter; windows are $\pm$6 months; Newey-West standard errors with 44 lags are in parentheses. Column (4) pools the three events: the daily triple-difference $\text{TD}_t = (\log V^{ovn}_{mkt} - \log V^{intra}_{mkt}) - (\log V^{ovn}_{idio} - \log V^{intra}_{idio})$ is regressed on Post $\times$ ΔHours with event fixed effects over $\pm$6-month windows; standard errors are clustered by date. Panel B reports full-sample descriptive accounting regressions on the stacked four-cell panel, with 4 observations per date and standard errors clustered by date, not identification: column (1) interacts the cells with a linear year trend centered at 2008; column (2) interacts the cells with overnight CME futures trading hours and includes year$\times$month fixed effects. ***, **, * denote significance at the 1%, 5%, and 10% level.

Panel A: Event-study identification at CME trading-hours changes

	(1) E-mini launch Sep 1997	(2) Globex expansion Jul 2003	(3) Maintenance window Nov 2012	(4) Pooled (per hour)
ΔOvernight hours	+8.5	+8.5	−0.5	
Event ramp	0.432***	0.200	−0.124	
	(0.056)	(0.196)	(0.174)	
Log idio stock var ratio	0.819***	0.850**	0.628**	
	(0.221)	(0.387)	(0.293)	
Post × ΔHours				0.057*
				(0.030)
Event FE				Y
Observations	253	252	251	753
R-squared	0.46	0.21	0.15	0.02

Table 5: Futures Trading Hours and Overnight Market Variance *(continued)*

Panel B: Full-sample accounting (descriptive)

	(1) Trend	(2) Futures hours
Mkt × Ovn × Year	0.021***	
	(0.004)	
Mkt × Year	0.016***	
	(0.003)	
Ovn × Year	0.018***	
	(0.001)	
Mkt × Ovn × FutHrs		0.020***
		(0.006)
Mkt × FutHrs		0.037***
		(0.004)
Ovn × FutHrs		0.026***
		(0.001)
Year×Month FE	N	Y
Observations	32,216	32,216
R-squared	0.35	0.45

closed on both sides of the comparison, so the main change across years is whether futures trade.

Take MLK Day, a Monday holiday. Before 2004, futures are closed. The Friday-close-to-Tuesday-open return spans three fully dead days—Saturday, Sunday, Monday—with no market of any kind trading. Starting in 2004, futures trade on Monday. The same Friday-to-Tuesday return now spans two dead days and one active day. The index is being traded, but individual stocks still cannot trade.

Specification. We use the same four-cell daily panel as Table 5 and estimate the triple interaction of the market indicator, the overnight indicator, and holiday futures trading. We measure the treatment two ways: an indicator for whether CME futures traded on the holiday spanned by the gap, and a continuous dose—the log of one plus E-mini volume actually traded inside the holiday gap. The full-sample specifications include weekend-gap interactions, so holiday gaps are compared against weekend gaps, which experienced the same general expansion of overnight futures hours but no holiday-specific opening, along with holiday-type and year$\times$month fixed effects. Standard errors are clustered by date, with Driscoll–Kraay statistics reported alongside.

Results. Table 6 presents the results. The evidence is positive but, with only 163 holiday gaps in the sample (64 futures-closed, 99 open), necessarily modest in precision. The continuous-dose specification in column (1) is the sharpest: systematic overnight variance rises by 5.0% per log-contract of futures volume traded over the holiday gap ($t = 1.9$). The binary versions point the same way: within holiday gaps, systematic variance is higher in years when futures traded (column (2), 0.77, $t = 1.5$), and the full-sample triple-difference with weekend controls is 0.46 ($t = 1.5$). Taken at face value, moving from a closed holiday to a typical post-2008 holiday session of roughly 140,000 contracts implies systematic overnight variance about 80% higher ($e^{0.050\times\ln(140{,}000)} \approx 1.8$).

If futures trading drives the effect, it should be stronger on holidays with more actual trading. CME's decision to open on holidays applied to all six at roughly the same time, but the volume of actual trading varies across them. In the 2014–2019 period, a normal trading day sees about 1.2 million E-mini contracts ($146 billion notional). Holiday volume is a fraction of that, but a meaningful one: MLK Day averages 12% of a normal day's contracts, Presidents Day and Labor Day about 10%, and Thanksgiving 9%—roughly $17 billion in index notional on a day when individual stocks cannot trade at all. Memorial Day (5%) and Independence Day (7%) see substantially less activity.

Caveats. We are careful not to over-read the binary comparison. Because there are no futures-closed holidays after 2004, the indicator is collinear with the post-2004 era, and a permutation test that moves the switch to placebo years shows the holiday-versus-weekend differential grows with futures volume over the sample rather than stepping at the legal opening (see table notes). That pattern is what the trading channel predicts—the 2004–2007 launch years saw near-zero holiday volume, averaging 17,000 contracts per holiday against 140,000 by 2008–2019—but it means the holiday evidence is better read as a volume gradient than as a sharp switch. We therefore treat the dose specification in column (1) as the preferred one and present this margin as complementary evidence rather than a stand-alone experiment. Controlling for foreign-market returns realized during the gap, which load on holiday gaps mechanically (a Monday-holiday gap contains one more foreign trading day than the weekend gaps it is compared against), trims the full-sample coefficient by roughly a sixth.[14]

4. A model of index trading and market returns

In this section, we present a simple model for studying how expanded opportunities to trade the market portfolio affect market outcomes. The model borrows features from Merton (1987) and Campbell et al. (1993) (henceforth, CGW). Following Merton, we consider an economy where investors initially trade only a subset of available financial assets.[15] When an index-trading vehicle is introduced, investors can express the market-level component of their demand directly through the diversified market claim. Following CGW, risk-averse arbitrageurs accommodate buying and selling pressure arising from these demand shocks. Finally, we depart from both models by distinguishing trade time from calendar time by letting a subset of agents have demand shocks that scale with trade time. This allows us to study how increasing trade

[14] Good Friday serves as a placebo: CME equity futures have been closed on Good Friday throughout the entire sample, so there is no treatment. Including a Good Friday indicator in the stacked stock-level regression yields Systematic $\times$ GoodFriday $= -0.10$ (SE $= 0.14$)—statistically zero, as expected when futures do not trade.

[15] Merton writes: "The prime motivation for this assumption is the plain fact that the portfolios held by actual investors (both individual and institutional) contain only a small fraction of the thousands of traded securities available. ... Because this behavior can be derived from a variety of underlying structural assumptions, the formally derived equilibrium-pricing results are the theoretical analog to reduced-form equations."

Table 6: Futures Trading on NYSE Holidays and Overnight Market Variance

Notes: Four-cell stacked panel at the daily frequency, 1993–2024: for each trading date the dependent variable is the log squared return ($\times 10^4$) of the value-weighted market (top 500 stocks by prior-December market capitalization, re-selected annually) and of value-weighted idiosyncratic stock returns ($\beta = 1$: $r_i - r_m$), each split into overnight and intraday sessions. Returns use the clean RET-anchored construction described in the text. The reported coefficient is the triple interaction of the market indicator, the overnight indicator, and the treatment. In column (1) the treatment is $\ln(1 +$ E-mini volume traded inside the holiday gap$)$, a continuous dose defined on holiday-spanning gaps; in columns (2)–(3) it is an indicator for CME futures trading on the NYSE holiday inside the gap. Columns (1) and (3) use all trading dates with weekend-gap interactions, holiday-type fixed effects, and year–month fixed effects; column (2) conditions on holiday-spanning gaps only, with holiday-type fixed effects. The sample contains 163 holiday-spanning gaps (64 with futures closed, 99 open). Restricting column (2) to the four holidays with the highest futures volume (MLK Day, Presidents Day, Labor Day, Thanksgiving) yields 0.306 ($t = 0.50$); the cross-holiday activity gradient is instead captured by the continuous dose in column (1). Standard errors clustered by date in parentheses; Driscoll–Kraay (20 lags) t-statistics in brackets. Two caveats temper a causal reading of columns (2)–(3). First, a permutation test that moves the binary switch to placebo years (1996–2002, 2006–2018) yields a one-sided p-value of 0.67 for the column (3) coefficient, and the placebo coefficients grow steadily with the cutoff year: the holiday-versus-weekend differential tracks the growth of futures trading volume rather than stepping at the 2004 legal switch, which is why we read the continuous-dose column (1) as the preferred specification. Second, controlling for foreign-market returns realized during the gap attenuates the column (3) coefficient from 0.459 to 0.381 ($t = 1.22$), roughly 17% smaller. A placebo using Good Friday, on which futures remained closed throughout the sample, is approximately zero in the stock-level build (-0.10, SE 0.14). ***, **, * denote significance at the 1%, 5%, and 10% level.

	(1) Volume dose (full sample)	(2) All holidays (holiday gaps)	(3) All holidays (full sample)
Mkt × Ovn × ln(1+ES holiday volume)	0.0495*		
	(0.0258)		
Mkt × Ovn × FuturesOpen		0.770	0.459
		(0.510)	(0.303)
Driscoll–Kraay t-statistic	[1.94]	[1.12]	[1.53]
Weekend interactions	Y	–	Y
Holiday-type FE	Y	Y	Y
Year×Month FE	Y	–	Y
Holidays	All 6	All 6	All 6
Observations	32,216	652	32,216
R^2	0.450	0.303	0.450

time within a fixed calendar interval affects volatility, reversals, and trading volume.

4.1 Market structure and trading time

There are N risky assets. Per-share dividends are collected in the vector D_t and satisfy

$$D_t \overset{i.i.d.}{\sim} N\left(\boldsymbol{\mu}, \sigma_D^2 C_D\right), \tag{8}$$

where $\boldsymbol{\mu}$ is the mean vector with every element equal to μ, σ_D^2 is the variance of each stock's dividend, and C_D is the correlation matrix with constant pairwise correlations ρ_D. There is a unit mass of behavioral investors and a competitive rational arbitrageur that can invest in all assets. In the no-index economy, a mass $1/N$ of behavioral investors is specialized in each stock. All agents have CARA utility over next trading-period wealth with risk aversion a, so demands are myopic mean-variance. The risk-free asset has gross return $R_f > 1$ and is supplied elastically. In the main text, we set supply $x_i = 1/N$ for each asset i. The Appendix gives derivations and proofs of the results given here in the main text.

Each calendar period t is divided into S trading periods, indexed by $s = 1, \ldots, S$. Dividends arrive as subperiod payoff news,

$$D_{t,s} \overset{i.i.d.}{\sim} N\left(\boldsymbol{\mu}/S, \sigma_D^2 C_D/S\right), \qquad D_t = \sum_{s=1}^{S} D_{t,s}. \tag{9}$$

This scaling keeps the calendar-period distribution of fundamental payoff news fixed as S changes. The object $D_{t,s}$ should be read as a subperiod payoff news flow, not literally as corporate dividends paid at trading frequency. In the following, we suppress the t subscript when discussing subperiod dynamics.

The excess trading-period payoff from holding asset i is

$$\Pi_{i,s+1}^e = P_{i,s+1} + D_{i,s+1} - R_f^{1/S} P_{i,s}, \tag{10}$$

where $P_{i,s}$ and $D_{i,s}$ are the price and dividend of asset i, respectively. We will subsequently refer to excess payoffs as returns, since this is the model object most closely comparable to the empirical return. Investor i's subjective expected return on asset i is assumed to be

$$E_s^i\left(\Pi_{i,s+1}^e\right) = E_s^A\left(\Pi_{i,s+1}^e\right) + z_{i,s}, \tag{11}$$

where $E_s^A(\cdot)$ is the arbitrageur's rational expectation and $z_{i,s}$ is the subjective payoff distortion.

The key assumption is that subjective payoff distortions follow an AR(1) in trade time:

$$z_{s+1} = \lambda_S z_s + \eta_{s+1}^S, \qquad \eta_{s+1}^S \sim N(0, \Sigma_\eta^S), \tag{12}$$

where we in the following will refer to η_s^S as demand shocks and z_s as demand pressure. We set

$$\lambda_S = \lambda^{1-\xi+\xi/S}, \qquad \xi \in [0,1], \tag{13}$$

and assume $0 \leq \lambda_S R_f^{1/S} < 1$ for all S.[16] The covariance matrix Σ_η^S has variance $\sigma_{\eta,S}^2$ on the diagonal and $\rho\sigma_{\eta,S}^2$ off the diagonal. The parameter ξ controls how demand-pressure risk scales with trading frequency. When $\xi = 0$, each additional trading opportunity brings a new demand shock of the same variance, so the calendar-period variance of demand shocks increases in S. When $\xi = 1$, higher frequency is instead a pure time refinement: the variance of cumulative demand pressure over a calendar

[16]Complete notation would be

$$z_{t,s+1} = \lambda_S z_{t,s} + \eta_{t,s+1}^S,$$

with the convention that, when $s = S$,

$$z_{t+1,1} = \lambda_S z_{t,S} + \eta_{t+1,1}^S.$$

The notation treats $z_{i,s}$ as the common subjective payoff distortion of the investors specialized in stock i. Equivalently, it can be interpreted as the average distortion within that investor group.

period is invariant in S. To state this object, define

$$Z_t\left(S,\xi\right)=\sum_{s=1}^{S}z_{t,s}. \tag{14}$$

We choose the scaling

$$\sigma_{\eta,S}^2=\sigma_\eta^2\frac{1-\lambda_S^2}{1-\lambda^2}\frac{\left[\mathcal{A}_S(\lambda)\right]^{1-\xi}}{\mathcal{A}_S(\lambda_S)},\qquad \mathcal{A}_S(x)=S+2\sum_{j=1}^{S-1}(S-j)x^j, \tag{15}$$

which delivers these two benchmarks; when $\xi=0$, $\sigma_{\eta,S}^2=\sigma_\eta^2$. Intermediate values of ξ interpolate between these cases; the Appendix shows that $Var\left(Z_t\left(S,\xi\right)\right)$ is decreasing in ξ for all $S>1$. Thus lower values of ξ correspond to a stronger trade-time shock-arrival channel. We interpret $0\leq\xi<1$ as capturing the possibility that trading and observing transaction prices generate new signals, attention, or sentiment shocks, consistent with the spirit of French and Roll (1986) and with our empirical finding that volatility responds to the number of hours the market is open.

4.2 Asset market without index trading

In the no-index economy, each investor trades only one stock in addition to the risk-free asset. For notational simplicity, investor i trades stock i. Define the market-level demand pressure as

$$z_{m,s}\equiv\frac{1}{N}\sum_{i=1}^{N}z_{i,s}. \tag{16}$$

Because all stocks have equal supply and the covariance structure is symmetric, the equilibrium market price is a scalar function of this market-level signal.

Proposition 1. *Denote $P_{m,s}=\sum_{i=1}^{N}x_iP_{i,s}$ as the price of the market portfolio. The equilibrium market price in the economy without index trading is*

$$P_{m,s}=p_{m,S}+b_{m,S}z_{m,s}, \tag{17}$$

where

$$b_{m,S} = \frac{\psi_S}{1+\psi_S} \frac{1}{R_f^{1/S} - \lambda_S}. \tag{18}$$

Here

$$\psi_S = \frac{v_{m,S}}{v_S} < 1, \tag{19}$$

where v_S is the conditional return variance of a representative individual stock and $v_{m,S}$ is the corresponding conditional market return variance. These conditional variances are the objects that enter investors' mean-variance demands. The realized variances used to compute volatility, market regression R^2's, and variance ratios are unconditional return variances and are defined separately in the Appendix. See the Appendix for derivations.

The term ψ_S captures the effect of investor underdiversification. Investors trade individual stocks, not the market claim. The exposure to idiosyncratic risk makes their risky asset demand less sensitive to fluctuations in their subjective return expectations than if they could trade a diversified portfolio with no idiosyncratic risk. Thus, the market-level component of subjective beliefs, $z_{m,s}$, reaches the market portfolio with an attenuated multiplier $b_{m,S}$, where the degree of attenuation is driven by ψ_S.

Define the market return as

$$\Pi^e_{m,s+1} = P_{m,s+1} + D_{m,s+1} - R_f^{1/S} P_{m,s}, \tag{20}$$

where $D_{m,s} = \sum_{i=1}^N x_i D_{i,s}$. Substituting the equilibrium market price and the law of motion

$$z_{m,s+1} = \lambda_S z_{m,s} + \eta^S_{m,s+1} \tag{21}$$

into the market return, the component of the return that comes from the demand-pressure component is

$$b_{m,S}\eta^S_{m,s+1} - \left(R_f^{1/S} - \lambda_S\right) b_{m,S} z_{m,s}. \tag{22}$$

The first term reflects demand shocks, which add variance. The second term is the reversal of past demand pressure. A high value of $z_{m,s}$ raises the current price, but because the signal mean-reverts, it predicts a lower subsequent return. Thus, as in CGW, temporary demand pressure raises short-horizon variance and generates reversals.

The empirics consider variance ratios, and the natural model object is the variance of cumulative returns over calendar horizons. For h calendar periods, define the cumulative return from calendar period t to $t+h$ as

$$\Pi^e_{m,t,h} = \sum_{i=0}^{h-1} \sum_{s=1}^{S} \Pi^e_{m,t+i,s}, \qquad V_h(S) = \mathrm{Var}\left(\Pi^e_{m,t,h}\right). \tag{23}$$

The Appendix shows that

$$V_h(S) = hS\sigma^2_{D,m,S} + b^2_{m,S}\sigma^2_{\eta,m,S}\left(hSA_S + B_{h,S}\right), \tag{24}$$

where A_S and $B_{h,S}$ collect the demand-pressure variance and autocovariance terms. The first term is the fundamental payoff component; the second is the demand-pressure component.

We define the h-calendar period variance ratio as

$$VR_{1,h}(S) = \frac{hV_1(S)}{V_h(S)}. \tag{25}$$

With mean-reverting demand pressure, $VR_{1,h}(S) > 1$ for $h > 1$. That is, demand pressure raises short-horizon variance relative to long-horizon variance.[17] When $\xi = 0$, increasing the number of trading opportunities S increases the number of demand-pressure shocks without decreasing their variance, while the fundamental calendar-period variance is unchanged. Higher S therefore raises the relative importance of the mean-reverting demand-pressure component. As ξ increases, this trade-time shock-arrival channel is attenuated.[18]

[17] Note that our definition of the variance ratio is the inverse of the Campbell-MacKinlay-Lo definition.

[18] The model can only generate $VR \geq 1$. In the data, the early period has $VR < 1$. A straightforward

4.3 Asset market with index trading

We next consider the economy where investors can trade a diversified market claim directly. The index claim has price $P^I_{m,s}$, payoff $D_{m,s+1}$, and one-trading-period return

$$\Pi^{e,I}_{m,s+1} = P^I_{m,s+1} + D_{m,s+1} - R_f^{1/S} P^I_{m,s}. \tag{26}$$

We further assume that investors substitute from trading individual stocks to trading only this index claim.[19] The index vehicle is a fund backed by the underlying stocks with value equal to the assets it holds. Its introduction leaves the aggregate supply of risky cash flows unchanged and creates no new market-level demand signal; it changes only the asset through which investors express that signal.

Proposition 2. *The equilibrium market price in the economy with index trading is*

$$P^I_{m,s} = p^I_{m,S} + b^I_{m,S} z_{m,s}, \tag{27}$$

where

$$b^I_{m,S} = \frac{1}{2}\frac{1}{R_f^{1/S} - \lambda_S}. \tag{28}$$

See the Appendix for derivations.

The comparison with the no-index economy is immediate:

$$\frac{b^I_{m,S}}{b_{m,S}} = \frac{1+\psi_S}{2\psi_S} > 1, \tag{29}$$

since $\psi_S < 1$. In fact, $b^I_{m,S}$ is exactly $b_{m,S}$ evaluated at $\psi_S = 1$: the index claim car-

extension could generate such low VRs allowing demand pressure to diffuse gradually across stocks or investor groups. For example, shocks to one subset of stocks could raise demand pressure for related stocks with a lag. Such a mechanism would add a momentum component to returns, partially offsetting the reversal component emphasized here.

[19] This shift is a reduced-form way of capturing the introduction of a low-cost vehicle for trading the diversified market claim. One can microfound such a shift through participation costs, attention costs, or fixed costs of portfolio formation.

ries only market risk, so the attenuation from underdiversification disappears. Index trading amplifies the market-price response to the same market-level demand signal because investors now express the signal through the diversified market claim rather than through concentrated individual-stock positions. The price-pressure decomposition is the same as in the no-index economy (Equation (22)), except that $b_{m,S}$ is replaced by $b^I_{m,S}$. Hence index trading increases the demand-pressure share of calendar-horizon variance. In the calibration, this substantially raises the variance ratio relative to the no-index economy.

The same comparison has a useful price-pressure implication. We can separate the non-fundamental demand-pressure component of prices into market and idiosyncratic parts. Let

$$\widetilde{z}_{i,s} = z_{i,s} - z_{m,s}$$

denote the stock-specific component of demand pressure. In the no-index economy, individual prices can be written as

$$P_{i,s} = p_S + b_{m,S} z_{m,s} + b_{\perp,S} \widetilde{z}_{i,s}.$$

Here $b_{\perp,S}$ is the loading of individual prices on stock-specific demand pressure. In the index economy, individual prices instead take the form

$$P^I_{i,s} = p^I_{i,S} + b^I_{m,S} z_{m,s}.$$

Thus, moving from the no-index economy to the index economy eliminates the loading of relative prices on stock-specific demand pressure. In the no-index economy, relative prices load on $\widetilde{z}_{i,s}$ through $b_{\perp,S}$. In the index economy, this loading is zero.

This implication is useful because it distinguishes aggregate prices from relative prices. Index trading amplifies the effect of market-level demand pressure on aggregate prices because investors can express the market signal through a diversified claim. At the same time, it removes stock-specific demand pressure from relative prices because investors no longer trade individual stocks in the index economy. Thus index trading can increase reversals at the market level while reducing reversals in the idiosyncratic component of individual stock returns.

4.4 Market regression R^2's and volume betas

The model produces counterparts to the empirical R^2 of daily stock returns regressed on market returns, as well as the volume-beta regressions. The R^2 is simply calculated within the model as the market regression of daily stock returns on the daily market return. Trading volume is

$$dv_s = \left| \left(n_s^A - n_{s-1}^A \right) \odot P_s \right|, \tag{30}$$

where $\odot$ denotes element-by-element multiplication. By market clearing, the arbitrageur's position changes mirror investors' trades, so dv_s equally measures investor volume. We report daily cross-sectional volume regressions within the model, mimicking those in the empirical section.

4.5 Calibration and comparative statics

We calibrate the model to illustrate the quantitative size of the two mechanisms emphasized in the paper: direct trading of the market claim and expanded opportunities to trade within a fixed calendar period. The calibration is deliberately simple. We discipline a set of primitive parameters using external moments, choose a preferred benchmark, and then study how the model's implications vary with the strength of the trade-time channel and the persistence of demand pressure.

One calendar period in the model is one trading day. All volatilities reported below are annualized using 252 trading days. The baseline calibration is summarized in Table 7.

Table 7: Baseline model calibration

Parameter	Value	Description
N	500	Number of risky assets
ρ_D	0.25	Pairwise correlation of fundamental payoff news
ρ	0.25	Pairwise correlation of demand-pressure innovations
R_f^{ann}	1.02	Annual gross risk-free return
$\sigma_{D,m}^{ann}$	7.7%	Annualized market fundamental-news volatility
σ_D^{ann}	15.35%	Implied annualized firm-level fundamental-news volatility
σ_D^{day}	0.967%	Implied daily firm-level fundamental-news volatility
λ	0.98	Demand-pressure persistence
σ_η	0.000603	Demand-shock scale
ξ	0.1	Trade-time scaling parameter
S_L	1	Low-trading-opportunity benchmark
S_H	4	High-trading-opportunity benchmark

Parameters calibrated to standard moments. We set the number of assets to $N = 500$, so the index can be interpreted as a stylized S&P 500 index. Historical evidence suggests very limited diversification among direct stockholders. Blume et al. (1974) and Blume and Friend (1975) document highly concentrated individual stock portfolios in the 1960s and early 1970s. In their samples, most individual investors held the overwhelming majority of their stock portfolio in a single security.

We set the average correlation of fundamental payoff news and the average correlation of demand-pressure innovations to $\rho_D = \rho = 0.25$. The annual gross risk-free return is $R_f^{ann} = 1.02$. The one-subperiod gross risk-free return when there are S trading opportunities per day is then $R_S = 1.02^{1/(252S)}$.

We calibrate fundamental payoff news to match annualized market dividend volatility over our sample of $\sigma_{D,m}^{ann} = 7.7\%$. Under the equicorrelation structure, the model implies

$$\sigma_{D,m}^{ann} = \sigma_D^{ann} \sqrt{\rho_D + \frac{1-\rho_D}{N}}, \tag{31}$$

where σ_D^{ann} is the annualized firm-level fundamental-news volatility. Given $N = 500$

and $\rho_D = 0.25$, this implies $\sigma_D^{ann} \simeq 15.35\%$.

Demand-pressure persistence and shock scale. The key remaining primitive is the persistence of demand pressure, λ. The persistence parameter governs the term structure of reversals. Low values make demand pressure short-lived and generate strong short-horizon reversals. High values make demand pressure more persistent, raising price-level effects while spreading reversals over longer horizons.

Our baseline value, $\lambda = 0.98$, is disciplined by the variance-ratio implication of the model. Conditional on matching the low-trading-opportunity no-index volatility target, lower values of λ produce counterfactually sharp daily reversals. A high value of λ ensures the one-month variance ratio at the beginning-of-sample benchmark is low. In the sensitivity analysis below, we vary λ over a grid and, for each value, recalibrate σ_η so that the no-index economy with trade time equal to calendar time ($S = 1$) matches the beginning-of-sample annualized market volatility of 10%.

Trade-time scaling and open hours. The parameter ξ governs how much of the additional trade-time demand-pressure variance survives at the calendar-period level. The pure trade-time shock-arrival benchmark is $\xi = 0$, while $\xi = 1$ is the pure time-refinement benchmark in which the variance of cumulative demand pressure over a calendar period is invariant in S.

The market-closure evidence in French and Roll (1986), together with our empirical evidence on expanded trading hours, points toward a low value of ξ. French and Roll show that return variance accumulates much more slowly when markets are closed than when they are open: a two-calendar-day return with only one trading day has variance only modestly above that of a normal one-day return. By contrast, when markets are open on both days, the variance of a two-day return is approximately twice the one-day variance. This contrast suggests that a substantial part of return variance is tied to trading and price discovery rather than to elapsed calendar time alone. We therefore set $\xi = 0.1$ as a baseline in which trade time strongly affects the cumulation of demand shocks over a calendar day.

Trading hours increased from roughly 5 hours per day at the beginning of the sample to about 23 hours per day at the end of the sample. This corresponds to an expansion of trade time by a factor greater than four. We normalize the low-trading-opportunity benchmark to $S = 1$ and use $S = 4$ as a conservative high-trading-opportunity benchmark.

Baseline calibration results. Figure 10 reports the baseline calibration for four cases: no index with $S = 1$, index with $S = 1$, no index with $S = 4$, and index with $S = 4$. The comparison isolates the effect of the index vehicle, the effect of expanded trade time, and the combined effect of the two.

Panel A shows that index trading and expanded trade time both raise daily market volatility. The no-index, low-S benchmark is calibrated to have annualized market volatility of 10.0%. Introducing the index vehicle at $S = 1$ raises volatility to 14.3%. Expanding trading opportunities to $S = 4$ raises no-index volatility to 19.0%. Combining the index vehicle and high trade time raises volatility to 22.8%.

Panel B shows that the same forces increase short-horizon reversals. The variance ratio rises from $VR_{1,21} = 1.08$ in the no-index, low-S benchmark to 1.14 when the index is introduced, 1.66 when trade time expands, and 1.73 when both mechanisms are present. Because the model's demand-pressure component is mean-reverting, it pushes $VR_{1,21}$ above one.

Panel C shows that the model also raises the market component of individual stock returns. The market R^2 rises from 0.36 in the no-index, low-S benchmark to 0.54 with index trading, 0.68 with expanded trade time, and 0.75 when both mechanisms are active. The volume-beta result is most stark in the index economy. When investors trade the index claim, constituent volume is mechanically proportional to index weights, so the volume beta is one. In the equal-supply no-index benchmark, market weights do not vary across stocks, so the corresponding cross-sectional volume-beta regression

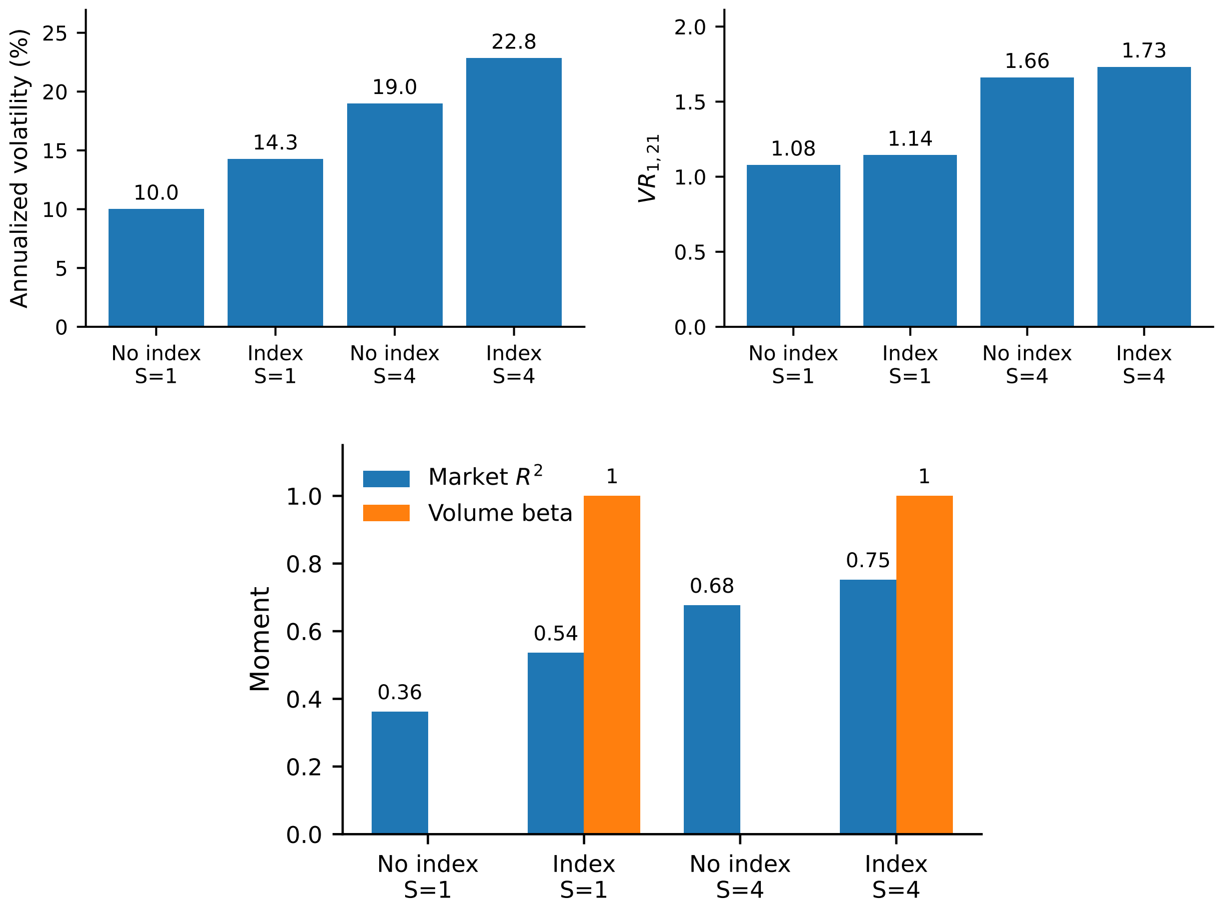

Figure 10: Baseline calibration. Panel A reports annualized market volatility. Panel B reports the variance ratio $VR_{1,21} = 21V_1/V_{21}$, where values above one indicate short-horizon reversal. Panel C reports the individual-stock market R^2 in all four cases and the volume-beta implication in the index cases. In the equal-supply benchmark, market-cap weights are identical in the no-index economy, so the cross-sectional no-index volume-beta regression is not separately identified. The index economy has volume beta equal to one because index trading generates constituent volume in proportion to index weights.

is not separately identified. If instead one solves the model with an unequal supply vector chosen to approximate the S&P 500 size distribution, the no-index volume beta becomes positive but remains below one. In sum, the model captures the broad empirical patterns documented in the earlier section.

Sensitivity to the trade-time scaling parameter. Figure 11 varies ξ holding $S = 4$ and $\lambda = 0.98$ fixed. The vertical line marks the preferred value $\xi = 0.1$.

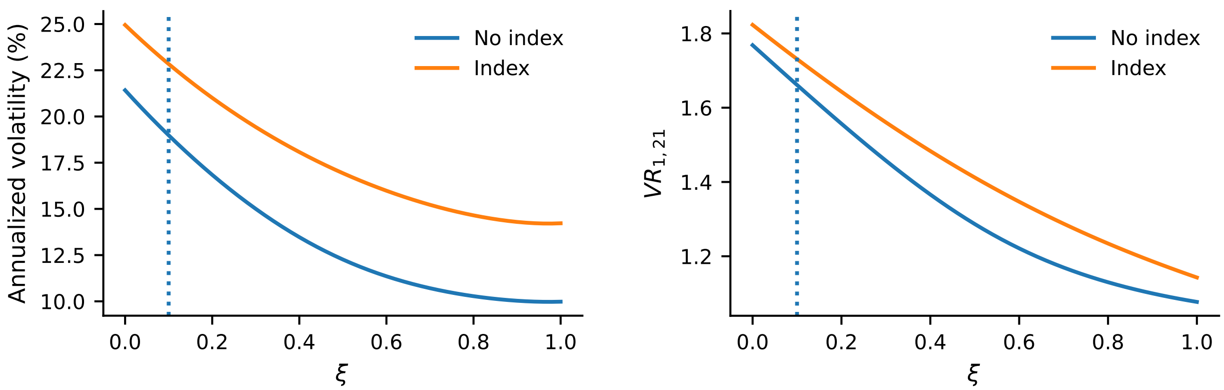

Figure 11: Sensitivity to the trade-time scaling parameter ξ, holding $S = 4$ and $\lambda = 0.98$ fixed. Panel A reports annualized market volatility. Panel B reports $VR_{1,21}$. The vertical line marks the preferred value $\xi = 0.1$.

The trade-time channel is strongest when ξ is low. At $\xi = 0$, additional trading opportunities bring additional demand-pressure shocks almost one-for-one, so both no-index and index volatility are high. As ξ rises toward one, more of the increase in trading frequency is time refinement rather than additional calendar-day demand-pressure variance. Volatility therefore falls in both economies. The index economy remains above the no-index economy throughout the range because index trading increases the market-price loading on the common demand signal.

The variance-ratio pattern is similar. Low ξ produces more demand-pressure variance within the calendar day, and because this demand pressure is mean-reverting, it raises short-horizon variance relative to long-horizon variance. As ξ rises, the demand-pressure contribution to calendar-day variance falls and the variance ratio moves back toward one. The baseline value $\xi = 0.1$ lies near the trade-time shock-arrival benchmark and therefore allows expanded trading hours to have economically large effects.

Sensitivity to demand-pressure persistence. Figure 12 varies λ holding $S = 4$ and $\xi = 0.1$ fixed. For each value of λ, the demand-shock scale σ_η is recalibrated so that the no-index economy with $S = 1$ has annualized market volatility of 10%.

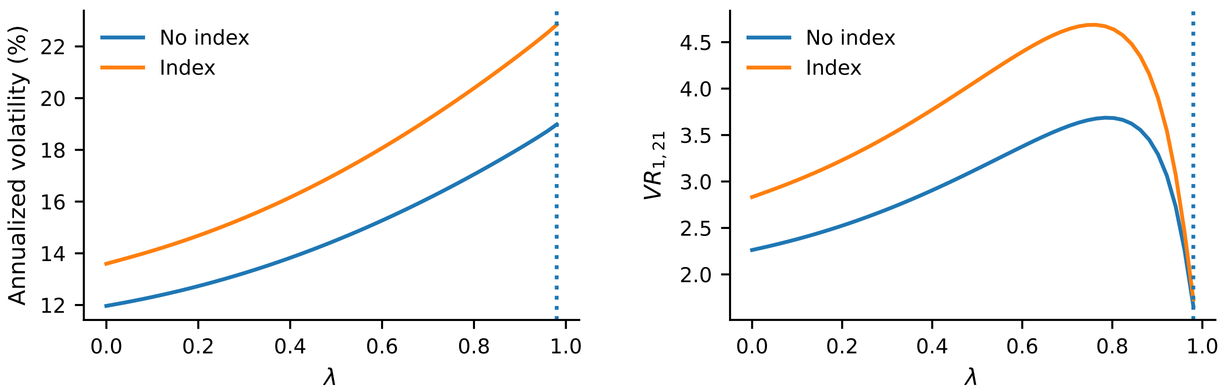

Figure 12: Sensitivity to demand-pressure persistence λ, holding $S = 4$ and $\xi = 0.1$ fixed. For each λ, σ_η is recalibrated so that the no-index economy with $S = 1$ has annualized market volatility of 10%. Panel A reports annualized market volatility. Panel B reports $VR_{1,21}$.

Persistence primarily governs the reversal implications. Low values of λ make demand pressure short-lived. Conditional on matching the low-S volatility target, this produces sharp short-horizon reversals and large variance ratios. Higher values of λ make demand pressure more persistent. This spreads reversals over longer horizons and lowers $VR_{1,21}$. The preferred value $\lambda = 0.98$ therefore balances these two forces, generating a large increase in volatility when trade time expands while avoiding implausibly large daily reversals.

Summary of model implications. Taken together, Figures 10–12 show that the two mechanisms emphasized in the paper can jointly generate a large increase in market volatility, a higher market component of individual stock returns, and a stronger link between index trading and constituent volume. The model remains intentionally stylized. In particular, the mean-reverting demand-pressure component pushes short-horizon variance ratios above one. A richer model with delayed diffusion of demand pressure across stocks or investor groups could add a momentum component and help match positive short-horizon autocorrelation early in the sample. We abstract from that force here to keep the model as parsimonious as possible.

Overall, the calibration requires a low value of ξ: demand shocks must scale substantially with trade time rather than only with calendar time. This is the key discipline delivered by the open-hours evidence. Conditional on that discipline, both index trading and expanded trade time contribute materially to the model-implied increases in volatility, variance ratios, market R^2's, and index-related volume.

5. Conclusion

The stock market has transformed over the past 60 years, with the proliferation of index products (index futures, index ETFs, and passive mutual fund products). Far from being "passive", the volume of trade in these products is enormous. We provide causal evidence that this volume has contributed substantially to increased volatility of the overall stock market.

References

Admati, Anat R, and Paul Pfleiderer, 1988, Selling and trading on information in financial markets, *The American Economic Review* 78, 96–103.

Baltussen, Guido, Sjoerd van Bekkum, and Zhi Da, 2019, Indexing and stock market serial dependence around the world, *Journal of Financial Economics* 132, 26–48.

Barberis, Nicholas, Andrei Shleifer, and Jeffrey Wurgler, 2005, Comovement, *Journal of Financial Economics* 75, 283–317.

Barclay, Michael J., and Terrence Hendershott, 2003, Price discovery and trading after hours, *Review of Financial Studies* 16, 1041–1073.

Basak, Suleyman, and Anna Pavlova, 2013, Asset prices and institutional investors, *American Economic Review* 103, 1728–1758.

Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi, 2018, Do etfs increase volatility?, *The Journal of Finance* 73, 2471–2535.

Berger, David, Ian Dew-Becker, and Stefano Giglio, 2020, Uncertainty and economic activity: A multicountry perspective, *Review of Financial Studies* 33, 3393–3445.

Bessembinder, Hendrik, and Paul J. Seguin, 1992, Futures-trading activity and stock price volatility, *Journal of Finance* 47, 2015–2034.

Bloom, Nicholas, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

Blume, Marshall E, Jean Crockett, and Irwin Friend, 1974, Stockownership in the united states: Characteristics and trends, *Survey of Current Business* 54, 16–40.

Blume, Marshall E, and Irwin Friend, 1975, The asset structure of individual portfolios and some implications for utility functions, *The Journal of Finance* 30, 585–603.

Brightman, Chris, and Campbell R Harvey, 2025, Passive aggressive: The risks of passive investing dominance, *Available at SSRN 5259427* .

Campbell, John Y, Sanford J Grossman, and Jiang Wang, 1993, Trading volume and serial correlation in stock returns, *The Quarterly Journal of Economics* 108, 905–939.

Campbell, John Y, Martin Lettau, Burton G Malkiel, and Yexiao Xu, 2001, Have individual stocks become more volatile? an empirical exploration of idiosyncratic risk, *The journal of finance* 56, 1–43.

Campbell, John Y, Martin Lettau, Burton G Malkiel, and Yexiao Xu, 2023, Idiosyncratic equity risk two decades later, *Critical Finance Review* 12, 203–223.

Chiah, Mardy, Philip Gharghori, and Angel Zhong, 2020, Has idiosyncratic volatility increased? not in recent times, *Not in Recent Times (September 26, 2020). Critical Finance Review, Forthcoming* .

Chinco, Alexander, and Marco Sammon, 2024, The passive-ownership share is double what you think it is, *Journal of Finance* Forthcoming.

Da, Zhi, and Sophie Shive, 2018, Exchange traded funds and asset return correlations, *European Financial Management* 24, 136–168.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703–738.

Dew-Becker, Ian, and Stefano Giglio, 2024, The decline of the variance risk premium: Evidence from traded and synthetic options, Working paper, Chicago Booth and Yale.

French, Kenneth R., 2008, Presidential address: The cost of active investing, *Journal of Finance* 63, 1537–1573.

French, Kenneth R, and Richard Roll, 1986, Stock return variances: The arrival of information and the reaction of traders, *Journal of financial economics* 17, 5–26.

Greenwood, Robin, and David Thesmar, 2011, Stock price fragility, *Journal of Financial Economics* 102, 471–490.

Haddad, Valentin, Paul Huebner, and Erik Loualiche, 2021, How competitive is the stock market, *Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing. SSRN* .

Haddad, Valentin, and Tyler Muir, 2025, Market macrostructure: Institutions and asset prices, *Annual Review of Financial Economics* 17.

Harvey, Campbell R, Michele G Mazzoleni, and Alessandro Melone, 2025, The unintended consequences of rebalancing, Working paper, National Bureau of Economic Research.

Hasbrouck, Joel, 2003, Intraday price formation in u.s. equity index markets, *Journal of Finance* 58, 2375–2399.

Hong, Harrison, and Jeremy C Stein, 2007, Disagreement and the stock market, *Journal of Economic perspectives* 21, 109–128.

Israeli, Doron, Charles M.C. Lee, and Suhas A. Sridharan, 2017, Is there a dark side to exchange traded funds? an information perspective, *Review of Accounting Studies* 22, 1048–1083.

Karpoff, Jonathan M., 1987, The relation between price changes and trading volume: A survey, *Journal of Financial and Quantitative Analysis* 22, 109–126.

Lazarus, Eben, Daniel J Lewis, James H Stock, and Mark W Watson, 2018, HAR inference: Recommendations for practice, *Journal of Business & Economic Statistics* 36, 541–559.

Lochstoer, Lars A., and Tyler Muir, 2022, Volatility expectations and returns, *Journal of Finance* 77, 1055–1096.

Lockwood, Larry J., and Scott C. Linn, 1990, An examination of stock market return volatility during overnight and intraday periods, 1964–1989, *Journal of Finance* 45, 591–601.

Lou, Dong, Christopher Polk, and Spyros Skouras, 2019, A tug of war: Overnight versus intraday expected returns, *Journal of Financial Economics* 134, 192–213.

Martin, Ian, 2017, What is the expected return on the market?, *Quarterly Journal of Economics* 132, 367–433.

Merton, Robert C, 1987, A simple model of capital market equilibrium with incomplete information, *Journal of Finance* 42, 483–510.

Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *Journal of Finance* 72, 1611–1644.

Sammon, Marco, 2025, Passive ownership and price informativeness, *Management Science* 71, 4582–4598.

Schwert, G William, 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1153.

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421–436.

Shleifer, Andrei, 1986, Do demand curves for stocks slope down?, *Journal of Finance* 41, 579–590.

Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35–55.

Sullivan, Rodney N., and James X. Xiong, 2012, How index trading increases market vulnerability, *Financial Analysts Journal* 68, 70–84.

Vogelsang, Timothy J, 1998, Trend function hypothesis testing in the presence of serial correlation, *Econometrica* 123–148.

A. Appendix

A.1 Systematic vs. Idiosyncratic Volatility

This appendix gives the construction behind the decomposition in Section 2.2.4 and reports the validation against Campbell et al. (2001) (CLMX), the variance-ratio cut, and the figures. We follow CLMX literally. Each stock return splits into a market component, an industry-relative-to-market component, and a firm-specific component through their unit-beta accounting identity: the industry residual is $R_{it} - R_{mt}$ and the firm residual is $R_{jit} - R_{i(j)t}$, with no betas estimated. We build industry returns by value-weighting CRSP firms within each of the 49 Fama–French industries; firm variances aggregate to the industry and market level by lagged market-cap weights; monthly variances are the within-month sum of squared daily residuals. The universe is all CRSP common stocks (share codes 10 and 11) from 1925 to 2024, roughly 82 million stock-days. We report the sample from 1962 to line up with CLMX and extend it through 2024.

Our construction reproduces their original numbers. Table 8 compares our implementation to their Table III over the 1962–1997 value-weighted daily sample they study. The firm-specific mean and trend – the moments CLMX build their paper around – match to within 3 and 8 percent, and our firm trend of 0.887 sits well inside their reported 90% confidence interval of $[0.55, 1.47]$.

Table 8: Validation against Campbell et al. (2001) Table III

	Our implementation	CLMX Table III	Δ
MKT mean $\times 100$	1.658	1.542	$+7.5\%$
IND mean $\times 100$	0.957	1.032	-7.3%
FIRM mean $\times 100$	6.628	6.436	$+3.0\%$
MKT trend $\times 10^5$	0.164	0.156	$+5.4\%$
IND trend $\times 10^5$	0.094	0.062	$+51.7\%$
FIRM trend $\times 10^5$	0.887	0.965	-8.1%

Notes: Mean and linear-trend coefficient of the market (MKT), industry-relative-to-market (IND), and firm-specific (FIRM) variance components, value-weighted, daily, over the 1962-07 to 1997-12 sample of Campbell et al. (2001). Means are average monthly variance $\times 100$; trends are the coefficient on a linear time index $\times 10^5$, matching their Table III units. CLMX report a 90% confidence interval of $[0.55, 1.47]$ for the FIRM trend.

The decomposition recovers both the original CLMX result and its later reversal. Firm-specific volatility rose from 16% in the postwar quiet to 26% by 2000–2010, the increase CLMX documented. It then fell back to 19% over 2011–2024 – the reversal Campbell et al. (2023) report, also studied by Chiah et al. (2020). The market component moved the other way: its share of average-stock variance reached a trough in the late 1990s, then rebounded.

Table 9 reports the share of average-stock variance carried by the market and by the systematic (market plus industry) component. The market share nearly doubles, from 0.18 in 1962–1971 to 0.31 in 2015–2024; the systematic share rises from 0.31 to 0.45. Both trends are strongly significant – +2.50 percentage points per decade for the market share ($t = 3.95$) and +3.25 for the systematic share ($t = 4.01$), under Newey–West HAC inference with 60 monthly lags.

Table 9: Share of average-stock variance in the systematic component

Period	Market share (= literal R^2)	Systematic share (MKT + IND) / Total
1962–1971	0.18	0.31
1988–1997	0.18	0.28
2000–2009	0.22	0.38
2015–2024	0.31	0.45
Trend (pp/decade)	+2.50	+3.25
	($t = 3.95$)	($t = 4.01$)

Notes: Market share is value-weighted aggregate market variance divided by value-weighted aggregate total variance across all CRSP common stocks – the share-of-variance interpretation of average R^2 in Campbell et al. (2001). Systematic share adds the industry-relative-to-market component. Period rows are means of the 60-month rolling share; the trend row is the coefficient on a normalized linear time index, in percentage points per decade, with Newey–West HAC t-statistics (60 monthly lags).

The variance-ratio result in Section 2.2 sits in the same component. Decomposing the five-year variance ratio by CLMX component, the market-component ratio more than doubles, from 0.72 in 1961–1995 to 1.51 in 1996–2024, and the industry component shifts the same way by a smaller amount. The firm-specific ratio shows no regime change: it has always been above one – the mechanical signature of bid-ask bounce – and declines monotonically over the full sample as spreads narrow and microstructure noise dissipates. Equal-weighting the top 500 stocks gives the same market-component shift ($0.62 \rightarrow 1.42$), so this is not a mega-cap artifact. Whatever changed the character of returns around 2000 changed the systematic component, not firm-specific behavior.

A.2 Macroeconomic Fundamentals Do Not Explain the Rise

The rise in stock volatility does not trace to more volatile fundamentals. This appendix gives the industrial-production-volatility control behind Section 2.2 and the

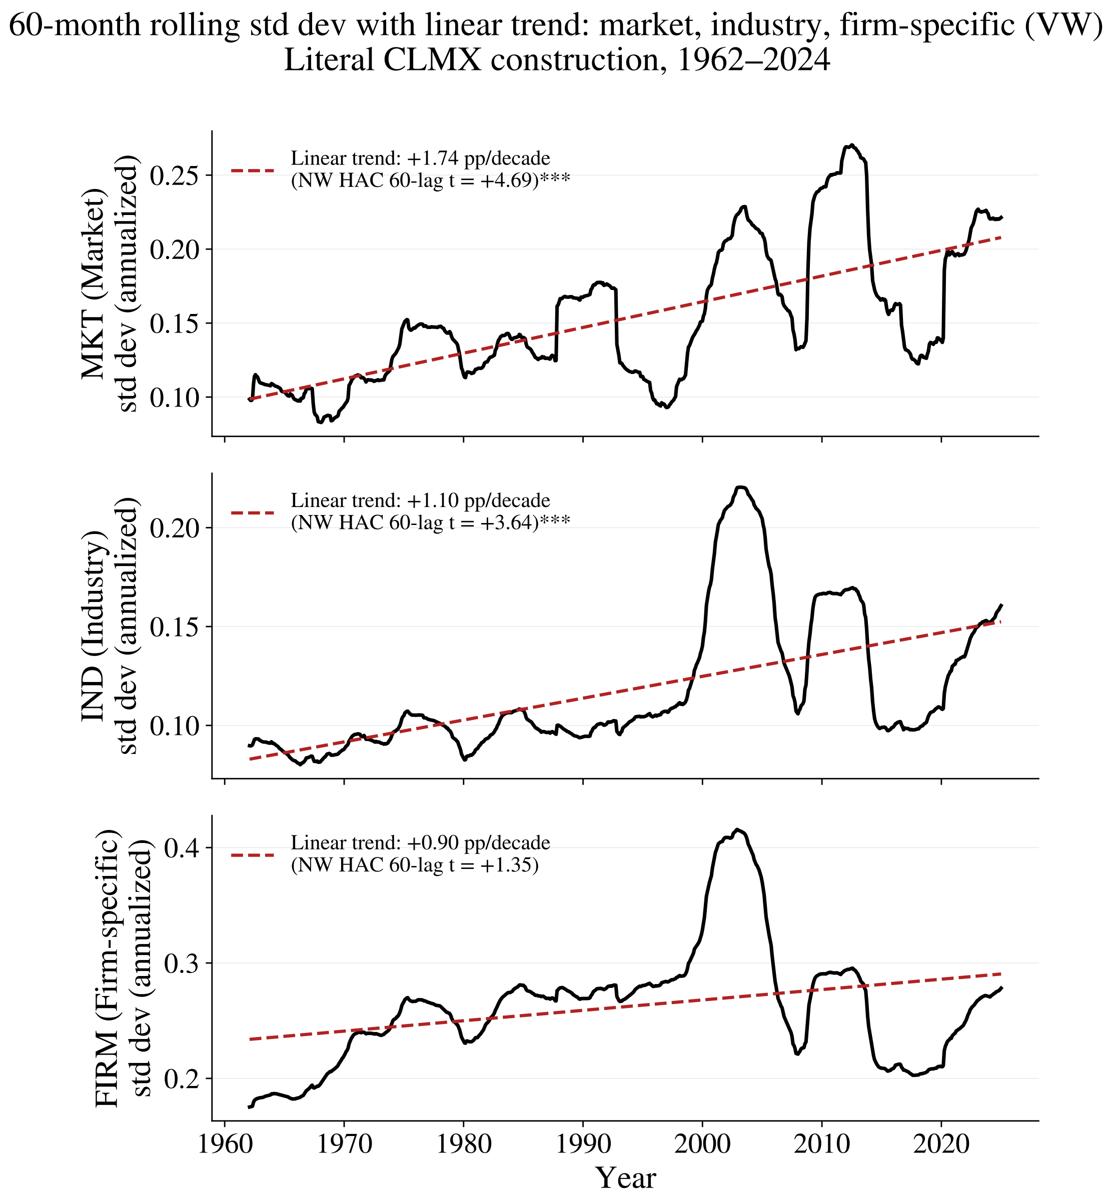

Figure 13: Market, industry, and firm-specific volatility, 1962–2024

Notes: Annualized standard deviation of the market (MKT), industry-relative-to-market (IND), and firm-specific (FIRM) components of average-stock variance, five-year rolling, value-weighted, 1962–2024, with fitted linear trends. The market and industry components trend up significantly; the firm-specific linear trend is positive but insignificant over the full sample because of the post-2001 reversal.

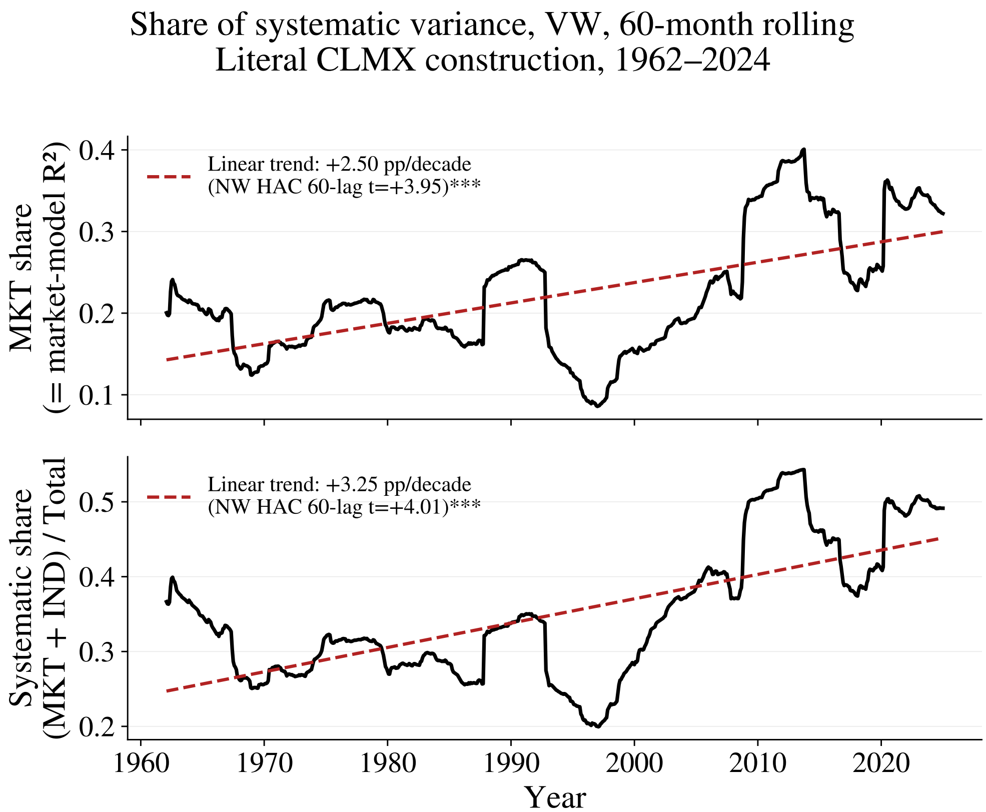

Figure 14: Systematic share of average-stock variance, 1962–2024

Notes: Market share (value-weighted aggregate market variance over total variance) and systematic share (market plus industry), 60-month rolling, 1962–2024. The late-1990s trough in the market share mirrors the firm-specific volatility peak in Campbell et al. (2001).

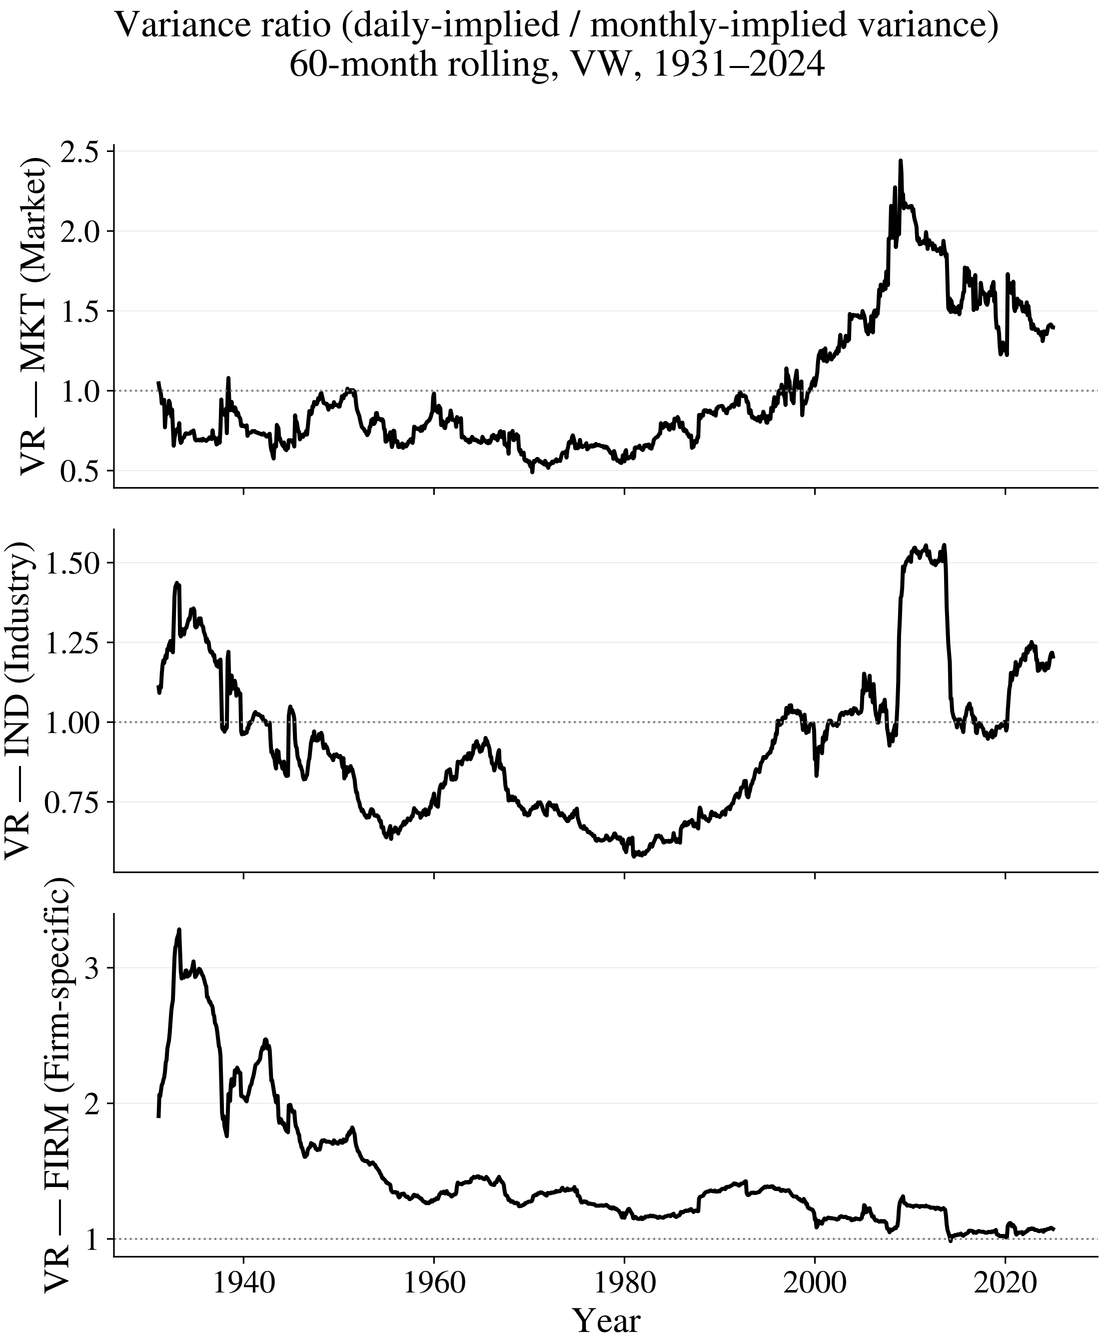

Figure 15: Variance ratio by CLMX component, 1931–2024

Notes: Five-year rolling variance ratio (daily-implied annualized variance over monthly-implied annualized variance) for the market, industry-relative-to-market, and firm-specific components. Note the differing vertical scales: the market and industry ratios shift upward through a regime change around 2000, while the firm-specific ratio declines monotonically with no regime change.

updated Schwert (1989) regression.

On the full 1931–2025 sample a simple linear trend in stock volatility is essentially zero – elevated Depression-era volatility offsets the modern rise. Adding five-year rolling industrial-production volatility as a control flips the trend to +11.0 percentage points on the normalized index, matching the paper's 1955–2025 baseline (Table 10, Panel A). The reason is that IP volatility trended *down* over the sample, so partialling it out uncovers the upward trend in stock volatility it was masking. On the post-1960 subsample the trend is nearly identical with or without the control, +10.78 versus +10.91 (Panel B): industrial-production volatility absorbs essentially none of the modern rise. Allowing a slope break at 1960 (column 3b) places the change in slope at +28 points on the normalized trend, about +3 pp/decade ($p = 0.013$).

The same conclusion holds in the classic Schwert (1989) regression of stock volatility on the volatilities of industrial-production growth, inflation, unemployment, and the short rate. Updating it to 1947–2024 gives an R^2 of 0.26 – close to the 0.38 we recover on his original window, and leaving roughly three-quarters of the variation in stock volatility unexplained by fundamentals (Table 11). The relationship has decoupled since the 1980s: macroeconomic volatility has fallen to historic lows while stock volatility has risen.

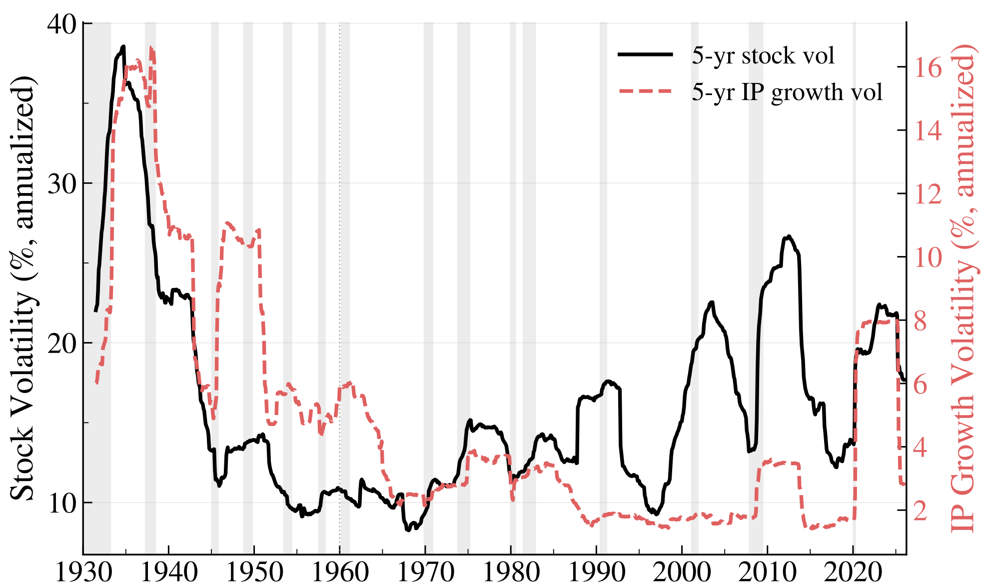

Figure 16: Stock volatility and industrial-production volatility, 1931–2025

Notes: Five-year rolling annualized stock volatility (left axis) and industrial-production-growth volatility (right axis). The two track each other through the Depression and early postwar period, then diverge after the mid-1980s: IP volatility falls to historic lows while stock volatility rises.

Table 10: Stock volatility and industrial-production volatility

Panel A: Full sample, 1931–2025 ($N = 1{,}135$)

	(1) Trend only	(2) + σ_{IP}	(3b) Slope break
Constant	17.21***	3.21	6.56*
	(3.98)	(3.02)	(3.58)
Trend $\tilde{t}$	-1.97	11.01**	-21.64
	(6.05)	(4.31)	(15.32)
σ_{IP}		1.54***	1.57***
		(0.33)	(0.26)
Post-1960 $\times\tilde{t}$			28.27**
			(11.37)
R^2	0.008	0.474	0.565

Panel B: Post-1960 subsample ($N = 792$)

	(4a) Trend only	(4b) + σ_{IP}
Constant	9.65***	8.16***
	(0.74)	(1.17)
Trend (1960$\rightarrow$end)	10.91***	10.78***
	(2.11)	(2.20)
σ_{IP}		0.51*
		(0.27)
R^2	0.477	0.515

Notes: Regressions of five-year rolling annualized stock volatility (%) on a normalized linear trend $\tilde{t}$ and the five-year rolling annualized volatility of monthly industrial-production growth, σ_{IP}. Panel A is the full 1931–2025 sample; column (3b) adds a slope change interacted with a post-1960 indicator. Panel B restricts to post-1960 with the trend renormalized to zero in 1960. Newey–West HAC standard errors (60 monthly lags) in parentheses. *, **, *** denote $p < 0.10, 0.05, 0.01$.

Table 11: Schwert (1989) regression, updated sample

Sample	N	σ_{IP}	σ_{CPI}	σ_{UNR}	σ_{RF}	R^2
Schwert original (1947–1987)	468	-0.29^{*}	2.21^{***}	0.00	0.13^{***}	0.38
Updated full (1947–2024)	912	-0.93^{***}	5.94^{**}	0.03^{***}	0.05	0.26
Pre-1985	432	-0.34^{**}	2.39^{***}	0.01	0.13^{***}	0.46
Post-1985	480	1.58	6.63^{**}	-0.01	0.11	0.39
Post-2000	300	1.30	6.08^{**}	-0.01	0.26^{**}	0.46

Notes: Joint regression of 12-month rolling annualized stock volatility on the 12-month rolling annualized volatilities of industrial-production growth, CPI inflation, unemployment-rate changes, and risk-free-rate changes, following Schwert (1989). Newey–West HAC (12 monthly lags). *, **, *** denote $p < 0.10, 0.05, 0.01$.

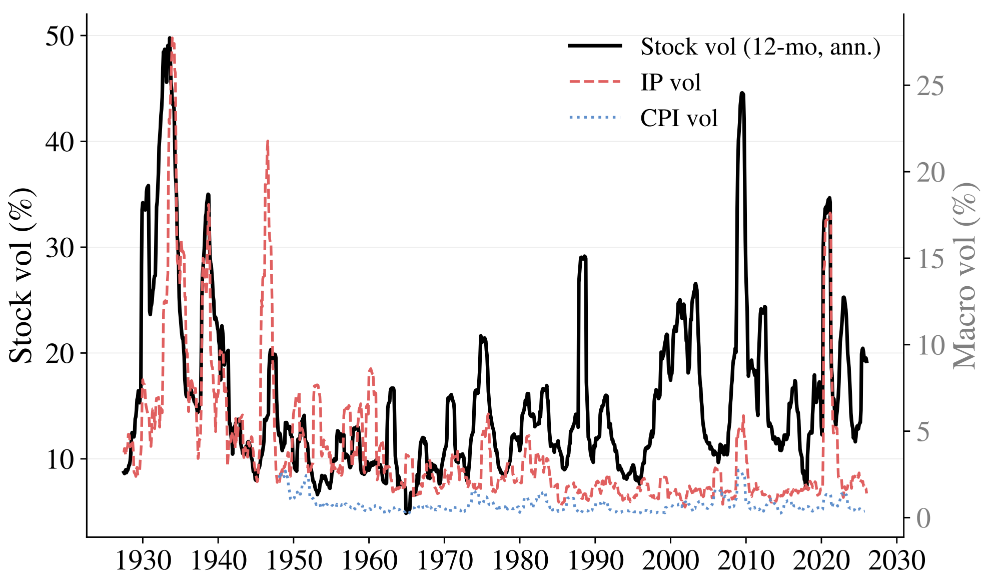

Figure 17: Stock volatility versus macroeconomic volatility, 1928–2025

Notes: Stock volatility (left axis) against industrial-production-growth and CPI-inflation volatility (right axis), all 12-month rolling annualized. Macroeconomic volatility was elevated in the Depression and the 1970s but has been near historic lows since the mid-1990s, even as stock volatility has risen.

A.3 Implied Volatility and the Variance Risk Premium

A natural objection is that the VIX, the most-cited measure of market volatility, shows little trend over its post-1990 sample. The reason is that the variance risk premium has declined at almost exactly the rate realized volatility has risen. Splicing the VXO (1986–1989) to the VIX (1990 onward) and measuring realized volatility, implied volatility, and their difference as five-year rolling averages, realized volatility trends up at $+1.84$ pp/decade while the variance risk premium falls at -2.01 pp/decade, leaving implied volatility flat at -0.17 pp/decade, indistinguishable from zero (Table 12). The declining premium is the central finding of Dew-Becker and Giglio (2024), who attribute it to easing intermediary frictions in the options market. The flat VIX reflects a shrinking premium, not an absence of rising volatility.

Table 12: Trends in realized volatility, implied volatility, and the variance risk premium

Series	Trend (pp on $\tilde{t}$)	Trend (pp/decade)	p-value	R^2
Realized vol (RV)	$+6.44$ (3.35)	$+1.84$	0.084	0.158
Implied vol (IV)	-0.61 (2.93)	-0.17	0.84	0.003
Variance risk premium (VRP)	-7.05 (0.86)	-2.01	9×10^{-6}	0.624

Notes: Trend regressions of five-year rolling realized volatility, implied volatility (VXO 1986–1989 spliced to VIX 1990+), and their difference (VRP = IV − RV) on a normalized linear trend, monthly, 1990–2025 ($N = 421$). Newey–West HAC standard errors (60 monthly lags) in parentheses; p-values use fixed-b critical values matching Table 1.

A.4 Structural Breaks in the Level and Character of Volatility

A linear trend is a convenient summary, not a literal model. As an alternative, we estimate structural breaks in the mean of the daily five-year volatility series in Figure 1. Figure 20 allows two breaks and dates them to around 1975 and 2000, with the increase in mean volatility over the sample about the same as in the trend specification.

Extending the break analysis back to 1931 sharpens the picture. With four breaks the daily volatility series splits into a Depression peak (mean 33%), a late-Depression and wartime regime (23%), the postwar low (11%), a step up around 1973 (14%), and the modern era (19%), recovering breaks in 1937, 1943, 1973, and 2000 (Figure 21). The Depression had higher volatility than today, but it came with extraordinarily volatile fundamentals; the modern rise does not.

The level of volatility is not the only moment that breaks. The variance ratio – the ratio of daily-implied to monthly-implied variance, a measure of short-horizon

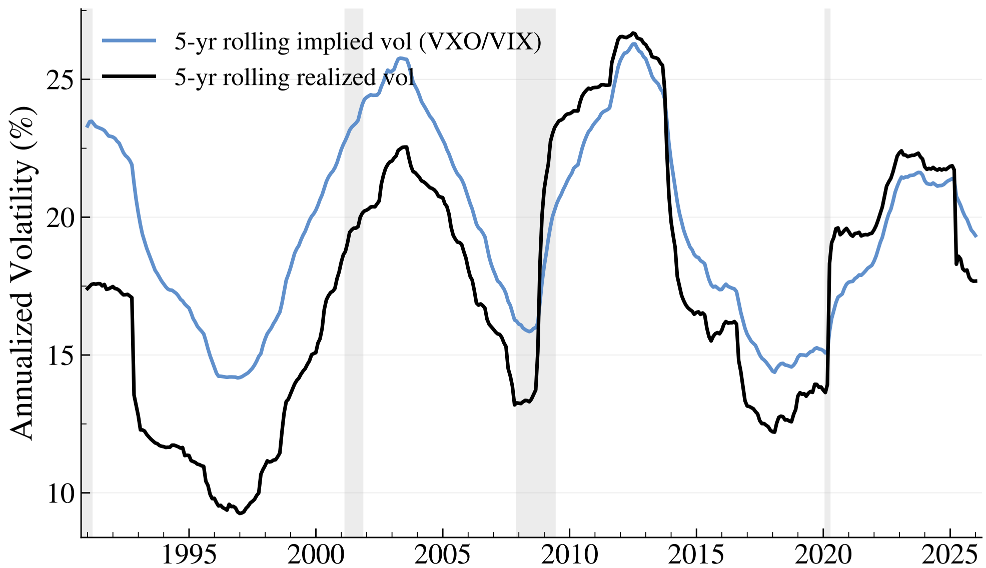

Figure 18: Implied and realized volatility, 1991–2025

Notes: Five-year rolling realized volatility and implied volatility (VXO spliced to VIX). Realized volatility trends up underneath a flat implied volatility; the gap between them – the variance risk premium – closes from roughly six points in the early 1990s to near zero by the late 2010s.

mean reversion – sat near 0.74 from 1931 through the mid-1990s, identical in the high-volatility Depression and the quiet postwar period despite volatility levels differing by a factor of two and a half. It then doubles. A break test on the variance ratio places a single significant break at March 2000 (Figures 22 and 23), the same date as the modern break in the level of volatility. Two independent moments of returns shift at the same moment – harder to attribute to fundamentals than to a change in how the market trades.

A.5 Stability of the trend across samples and rolling windows

The trend coefficient in Table 1 is robust to the choice of sample window and rolling-window length. Figure 24 plots the trend coefficient on $\tilde{t}$ estimated separately for every start year from 1931 to 2010 (sample ends in 2025), with 90% fixed-b confidence intervals. The trend is positive at every post-WWII start year. For the daily column, the 90% confidence interval excludes zero for any start year from approximately 1945 through 1993; for the weekly column, through 1992. Start dates pre-1945 absorb Depression-era volatility, which pulls the trend toward zero or negative. Start dates after the mid-1990s cut off the rise itself, so the trend in a now-elevated regime is mechanically smaller and less precise.

Figure 25 plots the trend coefficient by rolling-window length, from 1 month (4

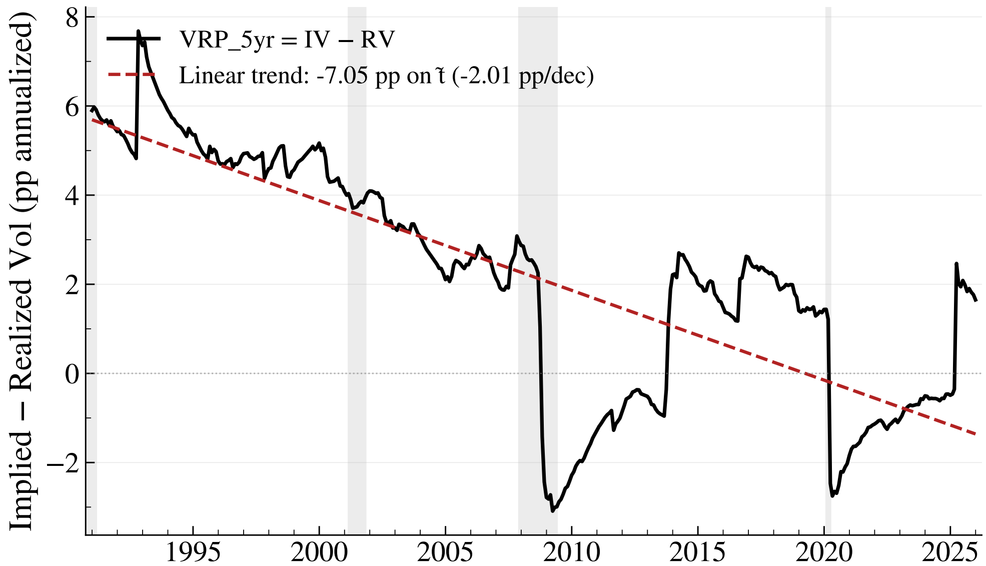

Figure 19: The variance risk premium, 1991–2025

Notes: Five-year rolling variance risk premium (implied minus realized volatility) with fitted linear trend. The premium falls from about six points in the early 1990s to near or below zero from 2018 onward, a decline of roughly two points per decade.

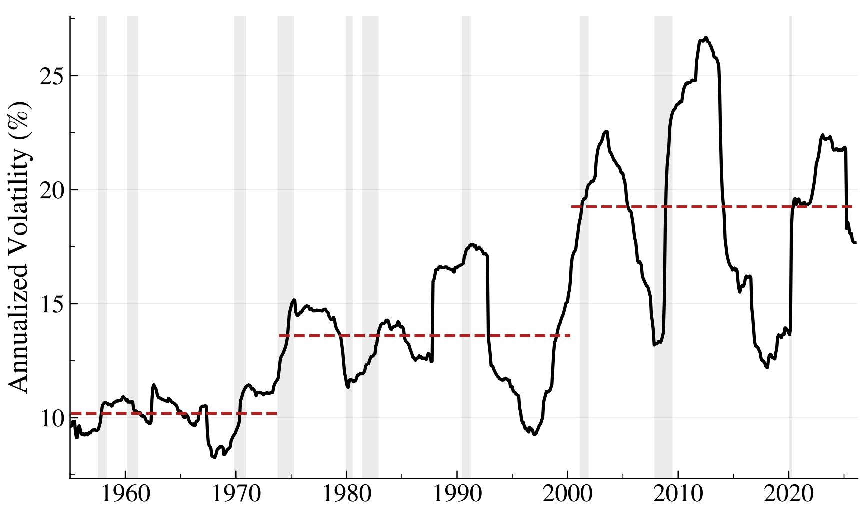

Figure 20: Structural Breaks in 5-Year Rolling Market Volatility

Notes: Annualized realized volatility of the Fama–French market excess return (same series as Figure 1). Dashed horizontal lines show estimated segment means from a Bai–Perron test for two structural breaks in the mean. The optimal break dates are October 1973 and March 2000, yielding three regimes with means of 10.2%, 13.6%, and 19.3%. Shaded regions indicate NBER recessions. Sample: 1955–2025.

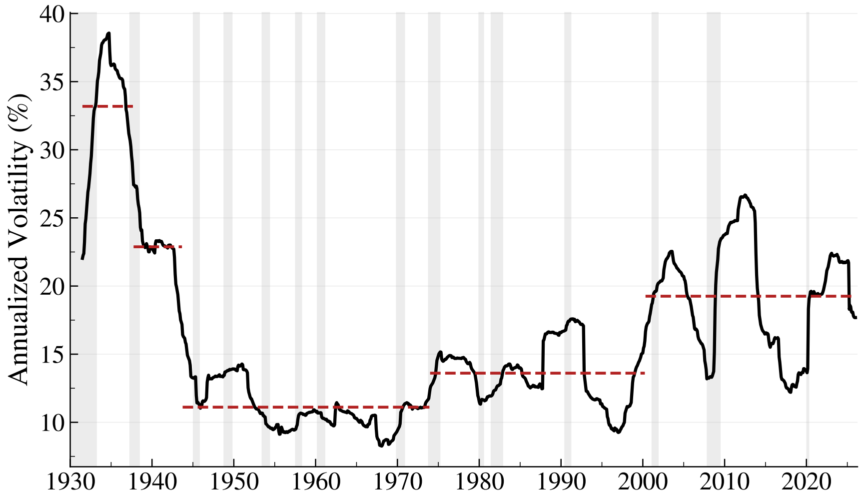

Figure 21: Structural Breaks in Volatility over the Long Sample, 1931–2025

Notes: Five-year rolling annualized volatility of the Fama–French market excess return, 1931–2025. Dashed horizontal lines show segment means from a Bai–Perron test with four breaks, dated August 1937, August 1943, November 1973, and March 2000, with regime means of 33.2%, 22.9%, 11.1%, 13.6%, and 19.3%. Shaded regions indicate NBER recessions.

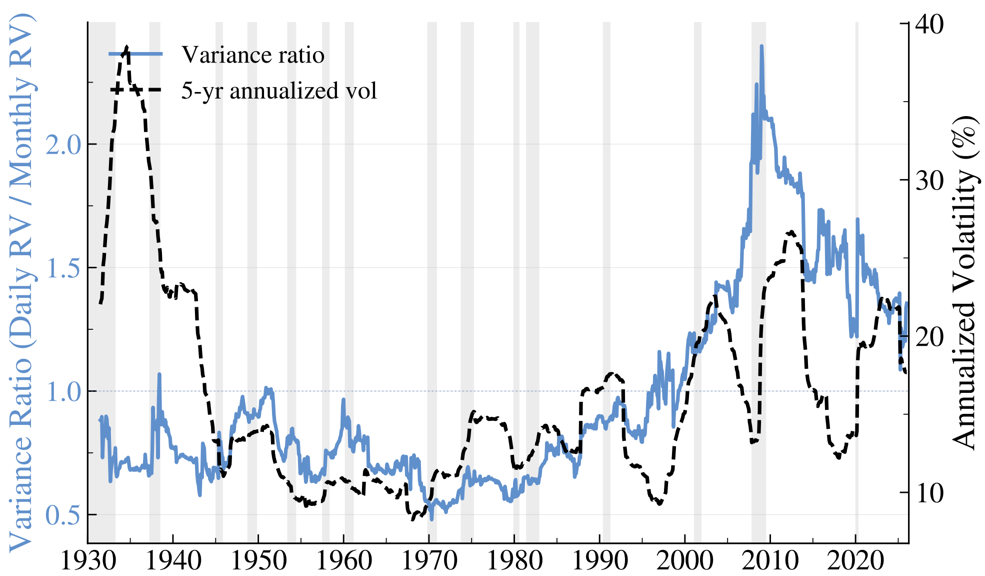

Figure 22: Variance ratio and volatility over the long sample, 1931–2025

Notes: Five-year rolling variance ratio (daily-implied over monthly-implied annualized variance, left axis) and annualized volatility (right axis), 1931–2025. The variance ratio is near 0.74 in both the Depression and the postwar period despite very different volatility levels, and rises only after the mid-1990s.

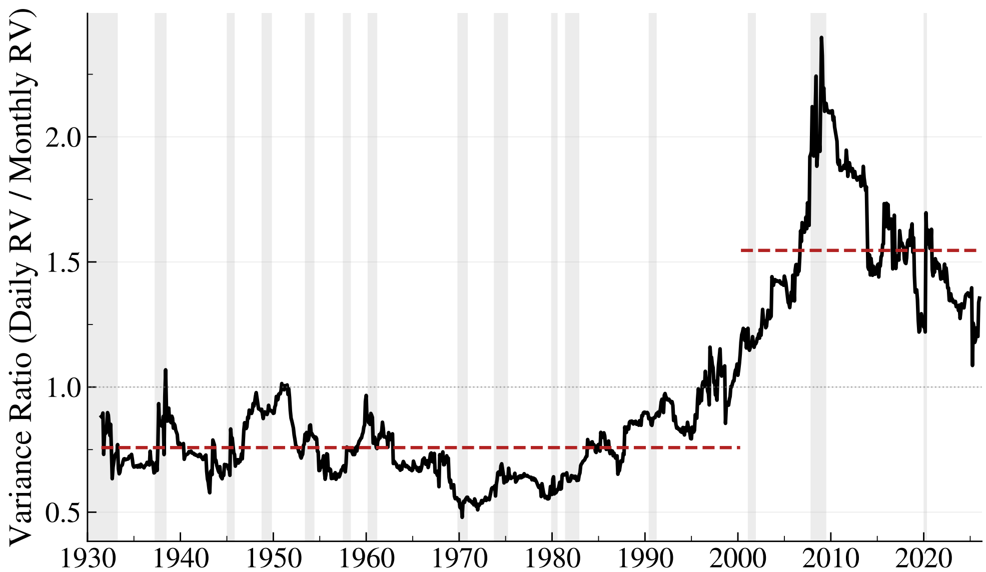

Figure 23: Structural break in the variance ratio, 1931–2025

Notes: Five-year rolling variance ratio with a single Bai–Perron break, dated March 2000. The segment means are 0.76 before and 1.55 after, doubling at the same date as the modern break in the level of volatility.

weeks) up to 10 years (520 weeks), on the 1955–2025 sample. The trend increases monotonically with window length and asymptotes around 7 years; the paper's 5-year window sits in the middle of this plateau. Even at 1-month windows the daily trend is +8.5 pp and the weekly trend is +5.6 pp, both with 90% confidence intervals excluding zero. The 5-year window choice is not a smoothing artifact.

A third cut holds the trend stable. Estimating the trend on rolling 30-year windows, every window ending after 1980 has a positive coefficient, and the windows ending in the 2010s are the steepest in the sample at about +2.8 pp/decade – above the +1.6 pp/decade full-sample headline (Figure 26). Windows ending in the 1960s are negative because they still contain the Depression-to-postwar decline. If anything, the linear-trend summary understates the recent steepening.

A.6 The Rise Is Not a Weighting or Concentration Artifact

The rise is broad-based across the size distribution, not a product of value-weighting or a handful of mega-caps. Table 13 reports trend coefficients for alternative market constructions. Equal-weighting the market gives a steeper trend than value-weighting (+2.16 versus +1.61 pp/decade), and decomposing by size, the top-decile trend (+1.44) sits below the full-market trend while the mid- and small-cap segments run steeper (+2.18 and +2.35). If a few large stocks drove the result, the largest stocks would show

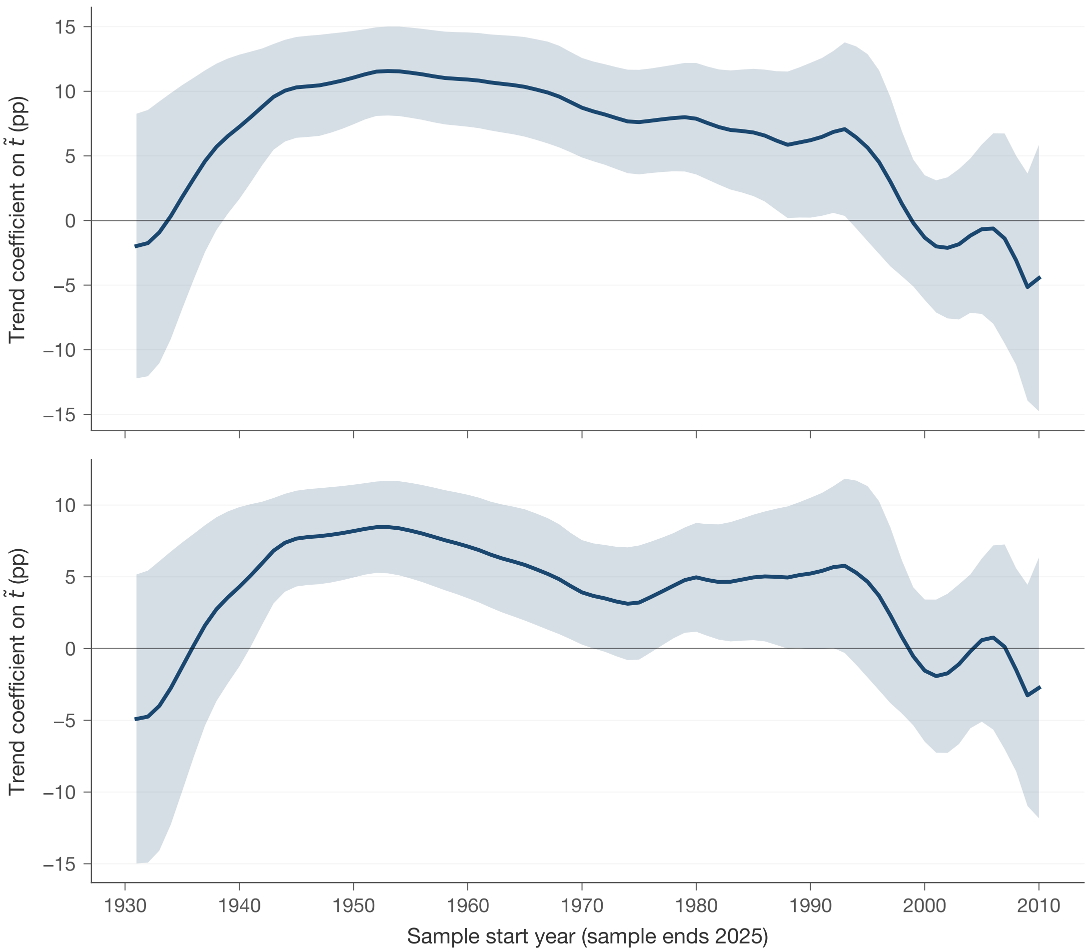

Figure 24: Trend coefficient by sample start year (sample ends 2025)

Notes: Each point is the trend coefficient b from the regression $\sigma_t = a + b\tilde{t} + \varepsilon_t$ estimated over [start year, 2025], where σ_t is the 5-year rolling annualized volatility from daily (top panel) or weekly (bottom panel) Fama–French market excess returns, and $\tilde{t} \in [0, 1]$ over the sample window. The shaded band is the 90% confidence interval under Newey–West HAC with Bartlett kernel and $S = 260$ weekly lags, fixed-b critical values with $\kappa = \lfloor 3/(2b) \rfloor$ degrees of freedom.

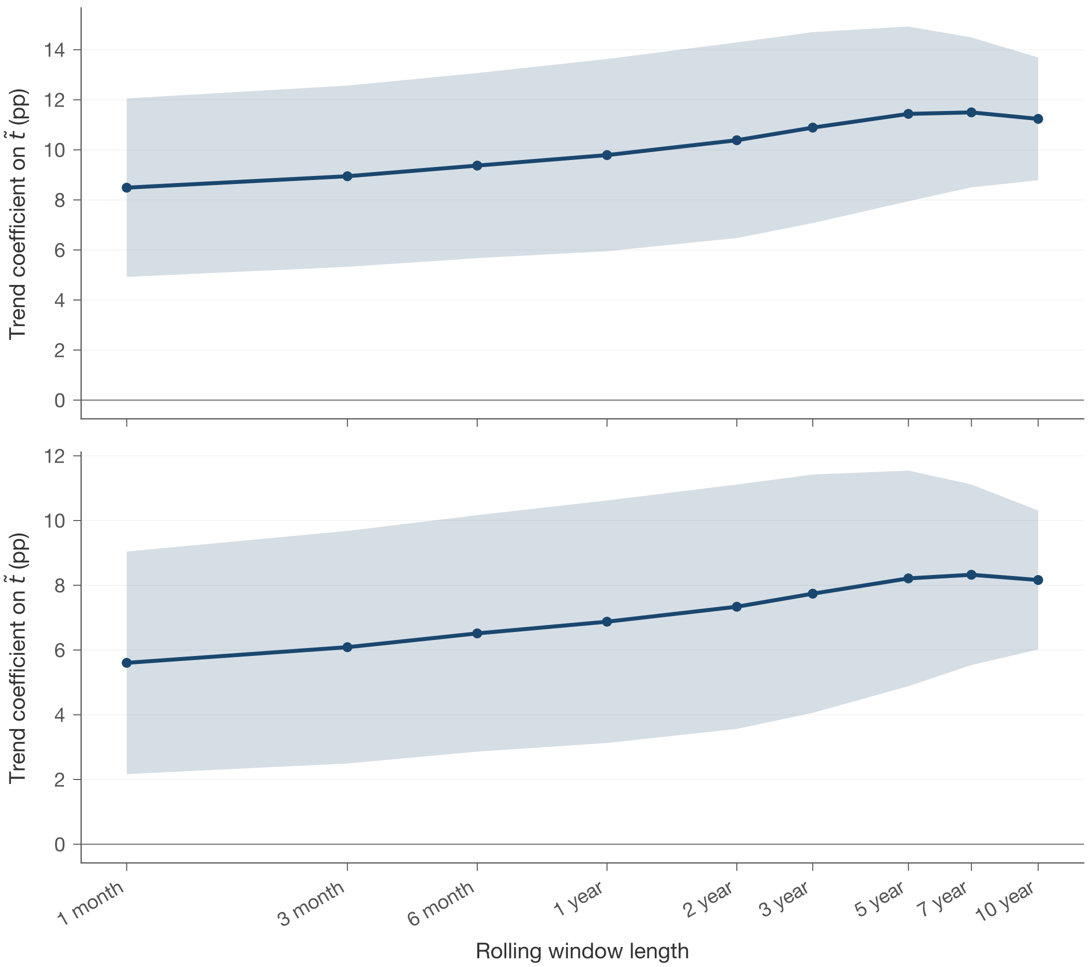

Figure 25: Trend coefficient by rolling-window length

Notes: Each point is the trend coefficient b from the regression $\sigma_t = a + b\tilde{t} + \varepsilon_t$ on the 1955–2025 sample, with σ_t constructed as the rolling annualized volatility of daily (top panel) or weekly (bottom panel) Fama–French market excess returns over the window indicated on the x-axis. The shaded band is the 90% confidence interval under Newey–West HAC with Bartlett kernel and $S = 260$ weekly lags, fixed-b critical values. X-axis is on a log scale.

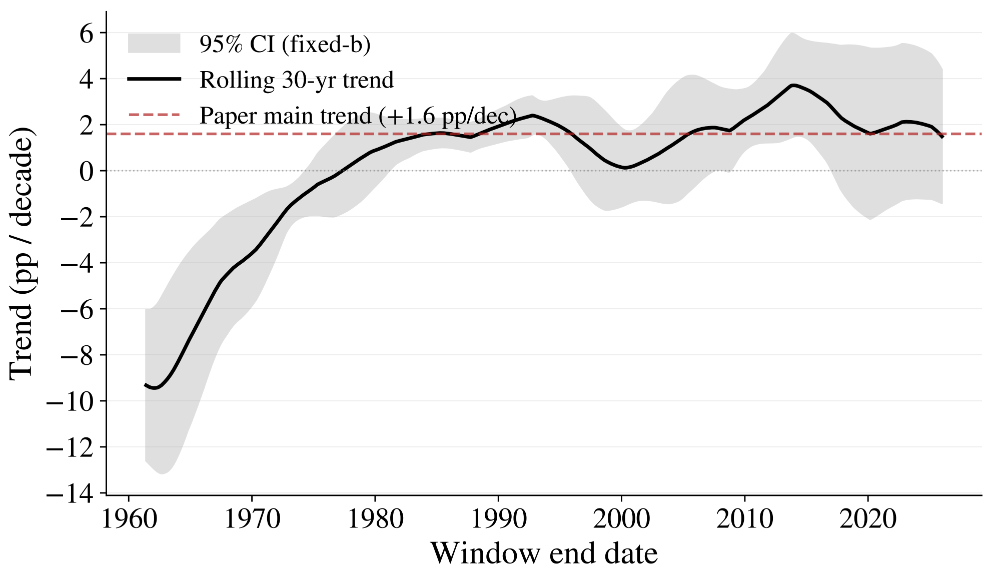

Figure 26: Trend coefficient on rolling 30-year windows

Notes: Trend coefficient (pp/decade) from estimating $\sigma_t = a + b\tilde{t} + \varepsilon_t$ on each rolling 30-year window, dated at window end, with 95% fixed-b confidence band. The dashed line marks the $+1.6$ pp/decade full-sample estimate from Table 1. Windows ending in the 1960s are negative because they still contain the Depression-to-postwar decline.

the steepest trend; they show the shallowest. Measuring volatility from the S&P 500 index directly – the index actually traded through futures and ETFs – gives +1.54 pp/decade against +1.71 for the value-weighted market over the matched 1967–2024 sample, with a correlation of 0.98.

Table 13: Volatility trend by market construction and size segment

Construction	Trend (pp/decade)
Value-weighted market	$+1.61^{***}$
Equal-weighted market	$+2.16^{***}$
Top decile	$+1.44^{***}$
Mid 40%	$+2.18^{***}$
Bottom 30%	$+2.35^{***}$

Notes: Trend coefficient (pp/decade) from regressing five-year rolling annualized volatility on a normalized linear trend, 1955–2025. The value-weighted market is the Fama–French market factor; the equal-weighted market averages Ken French equal-weighted size deciles; size segments are Ken French capitalization portfolios. Newey–West HAC (60 monthly lags). *** denotes $p < 0.01$ under fixed-b critical values.

A.7 The Variance-Ratio Rise Is Global

The variance-ratio rise is not specific to the United States. Building daily value-weighted market returns for 35 countries from Compustat Global firm-day data, the five-year variance ratio rises in all 35, with 29 significant at the 5% level under fixed-b inference and a mean trend of +0.189 per decade (Table 14). Every major developed market is significantly positive. Volatility *levels* are mixed across countries – emerging markets enter the sample at their 1990s crisis highs and decline through it – but the variance ratio, harder to pin on country-specific fundamentals, rises everywhere. A common, daily-frequency channel of the kind index trading produces fits this pattern; a fundamentals story does not.

A.8 Return Autocorrelations and Volume

The horizon decay in Table 1 is mechanically a statement about daily and weekly return autocorrelation. We show the underlying series and link them to volume, following Campbell et al. (1993) (henceforth CGW).

Figures 31a and 31b plot the rolling 5-year first-order autocorrelation of daily and weekly Fama–French market excess returns. Each window estimates $r_{t+1} = \alpha + \beta r_t + \varepsilon$ over the prior 60 calendar months; the shaded band shows ± 1.96 heteroskedasticity-consistent (HC1) standard errors. Daily autocorrelation peaks at roughly +0.30 in the

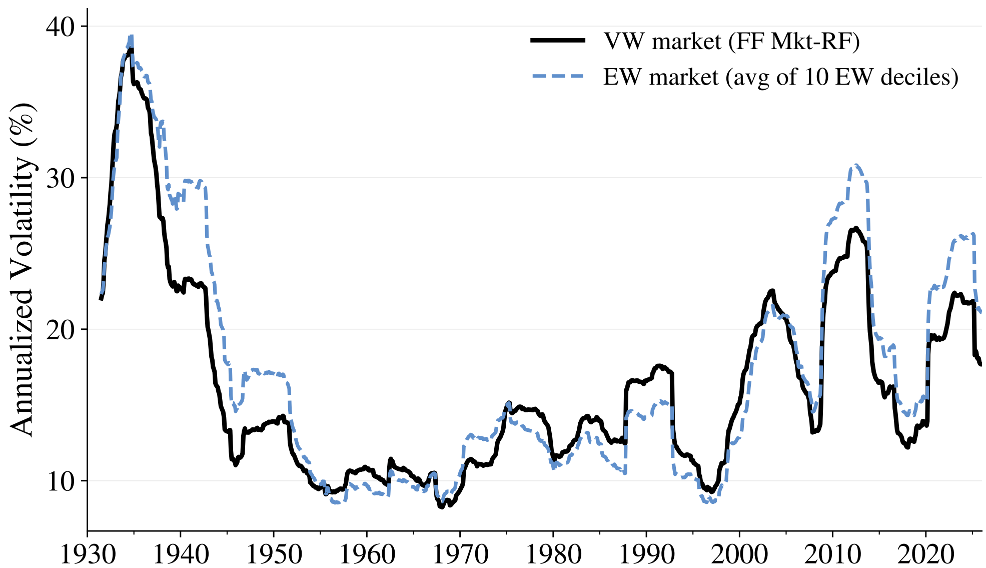

Figure 27: Equal-weighted and value-weighted market volatility, 1932–2025

Notes: Five-year rolling annualized volatility of the value-weighted and equal-weighted market, 1932–2025. The two track closely; the equal-weighted series runs slightly higher and trends slightly steeper in the modern era.

Table 14: Variance-ratio trends across countries

	Variance-ratio trend
Countries with a positive trend	35 of 35
Significant at 5% (fixed-b)	29 of 35
Mean across countries (per decade)	+0.189
United States	$+0.10^{***}$
United Kingdom	$+0.22^{***}$
France	$+0.29^{**}$
Germany	$+0.15^{**}$
Japan	$+0.29^{***}$

Notes: Five-year rolling variance ratio (daily-implied over monthly-implied annualized variance) for daily value-weighted market returns built from Compustat Global firm-day data, 35 countries, samples of 17–35 years (95 years for the US). Trends are the coefficient on a normalized linear time index, per decade, with Newey–West HAC (60 monthly lags) and fixed-b critical values. **, *** denote $p < 0.05, 0.01$.

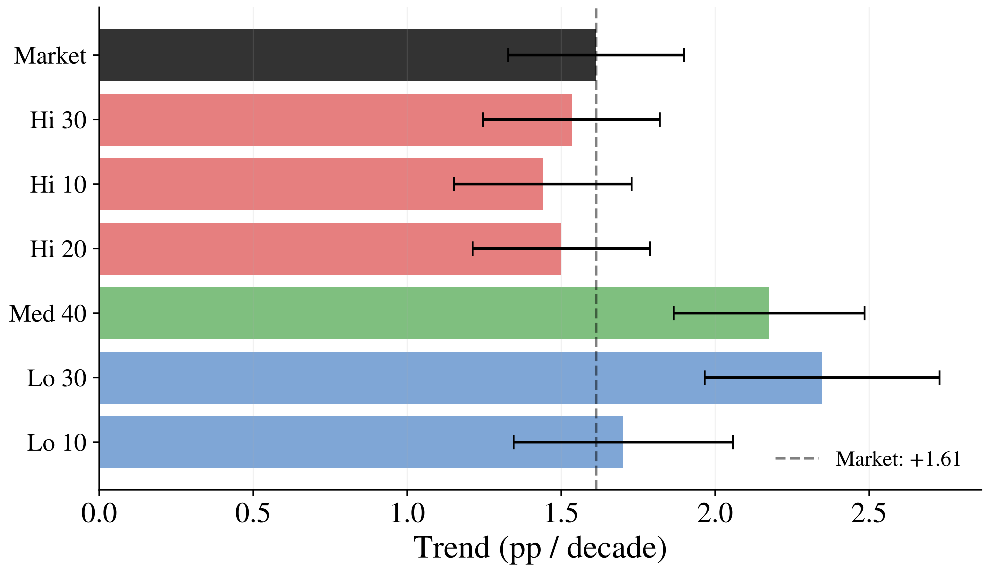

Figure 28: Volatility trend by size segment

Notes: Trend coefficient (pp/decade) of five-year rolling volatility by capitalization segment, 1955–2025, with 95% confidence intervals. The dashed line marks the full-market trend. The largest segments sit at or below the market trend; mid- and small-cap segments sit above it.

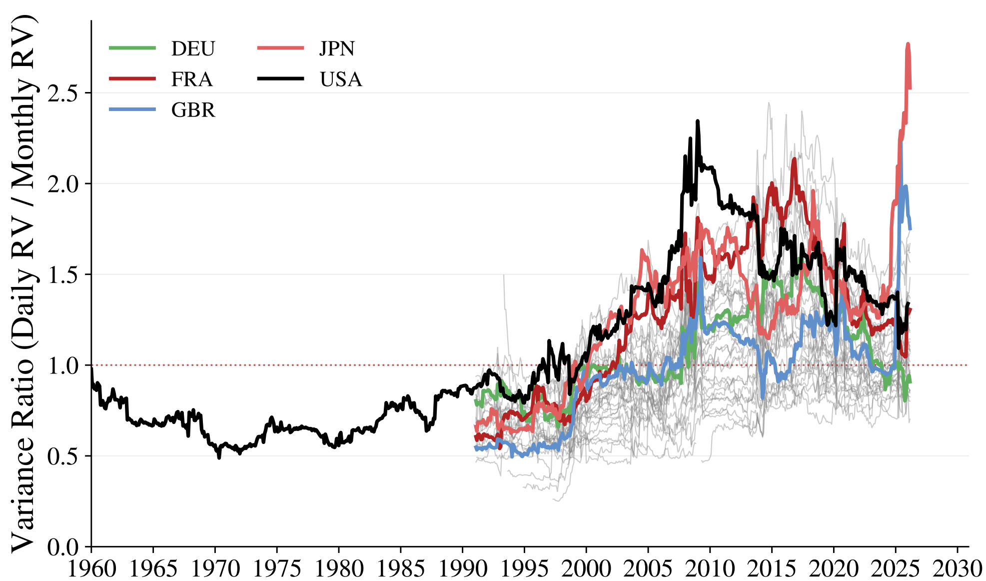

Figure 29: Variance ratio over time, 35 countries

Notes: Five-year rolling variance ratio for 35 countries built from Compustat Global firm-day data; the United States (longest sample) in black. By 2010 nearly every country has a variance ratio above one.

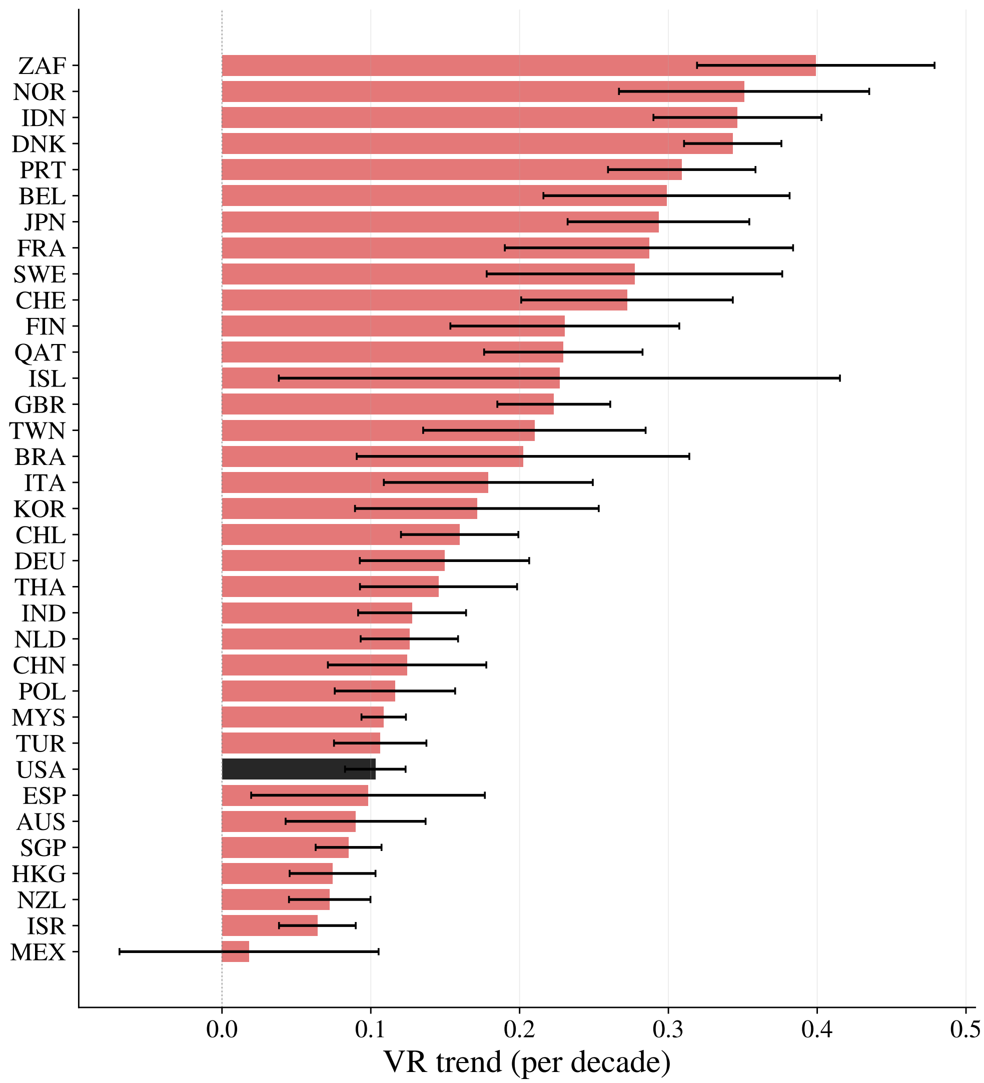

Figure 30: Variance-ratio trend coefficients by country

Notes: Per-country trend coefficient of the five-year rolling variance ratio (per decade), with Newey–West HAC standard-error bars. All 35 bars are positive; the United States sits in the middle of the distribution.

mid-1970s, declines steadily, and crosses zero around 2003. The weekly series crosses zero earlier (around 1990) and settles near -0.10. The linear time-trend coefficient on $\tilde{t} \in [0, 1]$ over 1955–2025 is -0.39 (NW HAC standard error 0.08) for the daily series and -0.30 (0.03) for the weekly, with fixed-b p-values below 10^{-3} in both cases.

CGW's own decade-by-decade evidence (their Table VI) already showed that the daily index autocorrelation moved substantially across regimes: 0.015 in the 1930s, 0.11 in the 1940s, 0.13 in the 1950s, rising to 0.28 in 1962–1974 and declining to 0.17 in 1975–1987. The series here extends their evidence by nearly four decades; the trajectory continues, with the autocorrelation turning negative in the 2000s and reaching -0.16 in the 2020s. Replicating their original Sample B (1962–1974), we recover the same first AC of 0.28 that they reported.

Table 15 runs the CGW Table II regression

$$r_{t+1} = \alpha + \beta r_t + \gamma \left(V_t \cdot r_t\right) + \sum_{d=1}^{4} \delta_d (D_d \cdot r_t) + \varepsilon_{t+1}, \tag{32}$$

where V_t is log aggregate market turnover (raw, or detrended by a backward 1-yr or 5-yr moving average as in CGW), and D_d are day-of-week dummies interacted with r_t. Standard errors are heteroskedasticity-consistent (HC1). The CGW prediction $\gamma < 0$ (high-volume returns more strongly reversed) holds on the full 1955–2025 sample under all three V_t specifications. On the CGW samples our estimates match theirs closely (their γ on Sample A was -0.328 with their detrended measure; ours is -0.349). On the post-CGW 1988–2025 sample γ remains strongly negative (-0.36, $t = -3.2$ for the 1-yr detrended specification), so the volume-reversal mechanism continued operating after CGW's sample. The 1980–2000 sub-period attenuates toward zero – the period when the autocorrelation itself was near zero, leaving little to interact with volume – and the post-2000 sub-period returns to strongly negative estimates as the autocorrelation moves into negative territory.

Figure 32 plots γ estimated on rolling 10-year windows. The interaction has been negative throughout, deepest in the 1970s and in the post-2008 / COVID period, and remains strongly negative at the end of the sample.

A.9 NYSE Trading Hours Changes

Table 3 exploits changes in NYSE trading schedules as instruments for trading volume. Table 16 lists all schedule changes used in the analysis. Before June 1952, the NYSE operated six days per week, with a shortened Saturday session (10:00 AM–12:00 PM). All subsequent changes affected weekday hours only, adjusting the closing time while keeping the open at 10:00 AM (until 1985, when it moved to 9:30 AM). The 1968 Wednesday closure temporarily reduced the trading week to four days.

Both specifications use all 9 events. The daily specification (Panel A) exploits the distinct within-week predictions of each schedule change; the weekly specification (Panel B) aggregates to calendar weeks.

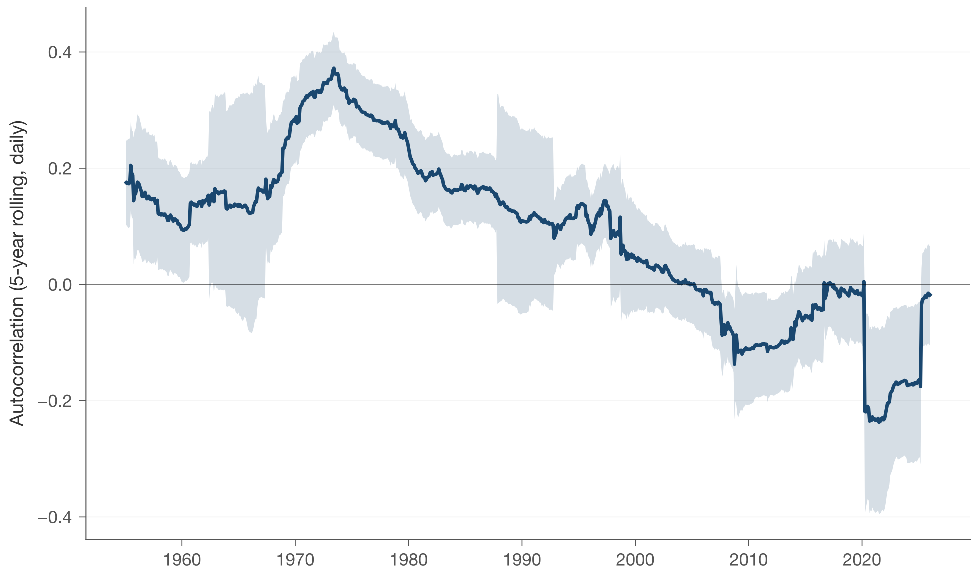

(a) Daily returns

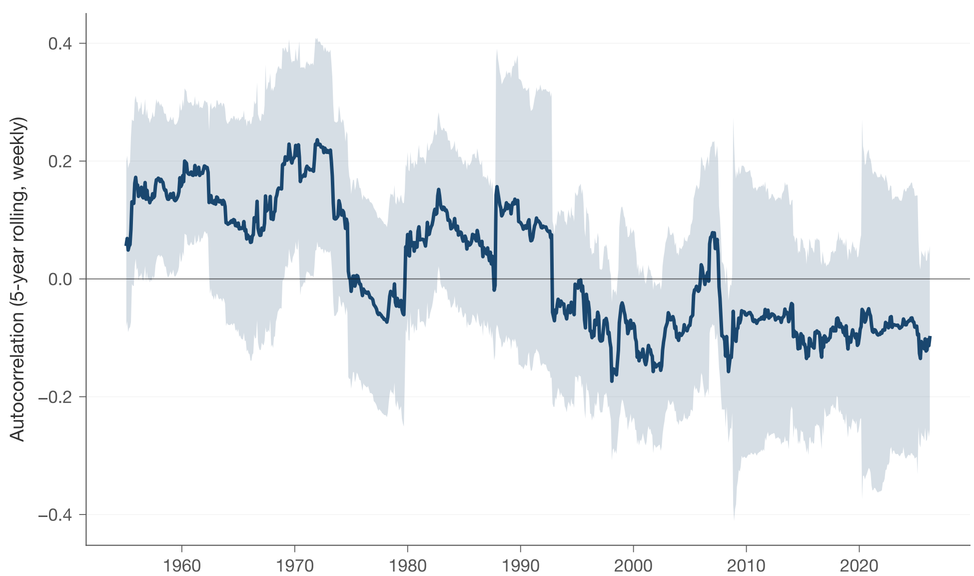

(b) Weekly returns

Figure 31: First-order return autocorrelation of the market, 5-year rolling

Notes: First-order autocorrelation of Fama–French market excess returns, estimated as the OLS slope β in $r_{t+1} = \alpha + \beta r_t + \varepsilon$ over each 5-year rolling window (60 calendar months). The shaded band is ± 1.96 HC1 standard errors estimated within each window. Panel (a) uses daily returns; panel (b) uses weekly returns. Sample: 1955–2025 (rolling window so first estimate uses data from 1950).

Table 15: Return autocorrelation and volume

Notes: Estimates of γ from Equation (32): $r_{t+1} = \alpha + \beta r_t + \gamma(V_t \cdot r_t) + \sum_d \delta_d(D_d \cdot r_t) + \varepsilon_{t+1}$, where V_t is log aggregate market turnover constructed from CRSP daily data following the construction underlying Figure 2. Detrended versions subtract a backward moving average from $\log V_t$. Standard errors in parentheses are heteroskedasticity-consistent (HC1). Stars: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$ (two-sided, normal asymptotics). Day-of-week interactions with r_t included.

Sample	Raw $\log V_t$	$\log V_t -$ 1-yr MA (CGW)	$\log V_t -$ 5-yr MA	N
Full (1955-2025)	-0.131^{***}	-0.289^{***}	-0.226^{***}	17,619
	(0.015)	(0.089)	(0.065)	
CGW Sample A (1962-87)	-0.114^{***}	-0.349^{***}	-0.299^{***}	6,344
	(0.024)	(0.063)	(0.049)	
CGW Sample B (1962-74)	-0.171^{**}	-0.452^{***}	-0.294^{***}	3,122
	(0.076)	(0.112)	(0.079)	
CGW Sample C (1975-87)	-0.104^{***}	-0.226^{***}	-0.276^{***}	3,221
	(0.035)	(0.077)	(0.076)	
1988-2025 (post-CGW)	-0.197^{***}	-0.361^{***}	-0.257^{***}	9,321
	(0.049)	(0.112)	(0.082)	
Pre-1980	-0.134^{***}	-0.356^{***}	-0.304^{***}	6,274
	(0.047)	(0.094)	(0.073)	
1980-2000	-0.074	-0.043	-0.013	5,055
	(0.065)	(0.110)	(0.100)	
2000-2025	-0.229^{***}	-0.354^{***}	-0.235^{***}	6,288
	(0.083)	(0.123)	(0.088)	

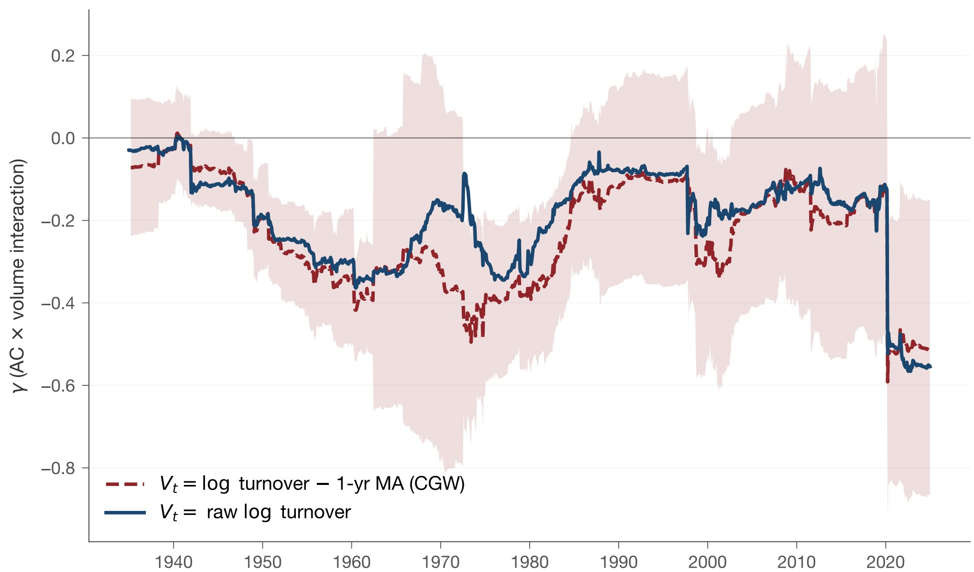

Figure 32: Rolling AC-volume interaction γ

Notes: Estimate of γ from Equation (32) on rolling 10-year windows, stepped monthly. Shaded band shows $\pm$1.96 HC1 standard errors for the detrended specification. Solid line uses raw log V_t; dashed line uses CGW's 1-year backward MA detrending.

Table 16: NYSE Trading Hours Schedule Changes

Notes: All times Eastern. Before June 1952, the NYSE was open six days per week (Monday–Saturday), with Saturday sessions from 10:00 AM to 12:00 PM. The "Daily Hours" column refers to weekday hours; Saturday hours were always 2.0 when applicable. "Weekly Hours" accounts for both the number of trading days and Saturday's shorter session.

#	Effective Date	Description	Schedule	Daily Hours	Weekly Hours
	Pre-change	*6-day week, Mon–Sat*	*10:00–3:00 (Sat 10–12)*	*5.0*	*27.0*
1	Jun 1952	Saturday trading eliminated	10:00–3:00, Mon–Fri	5.0	25.0
2	Sep 1952	Closing time extended 30 min	10:00–3:30, Mon–Fri	5.5	27.5
	Stable: Sep 1952 – Jun 1968				
3	Jun 12, 1968	Wednesday closure (paperwork)	10:00–3:30, 4 days/wk	5.5	22.0
4	Jan 2, 1969	Wed reopens; hours shortened	10:00–2:00, Mon–Fri	4.0	20.0
5	Jul 3, 1969	Closing extended 30 min	10:00–2:30, Mon–Fri	4.5	22.5
6	Sep 26, 1969	Closing extended 30 min	10:00–3:00, Mon–Fri	5.0	25.0
7	May 1, 1970	Closing extended 30 min	10:00–3:30, Mon–Fri	5.5	27.5
8	Oct 1, 1974	Closing extended 30 min	10:00–4:00, Mon–Fri	6.0	30.0
9	Jan 2, 1985	Open moved 30 min earlier	9:30–4:00, Mon–Fri	6.5	32.5

A.10 S&P 500 Futures Trading Hours

Several tests in this paper require knowledge of when S&P 500 index futures could be traded, particularly outside regular equity-market hours. We construct a daily schedule of futures trading hours from 5-minute tick data (TickData) for both the "big" S&P 500 futures contract (April 1982–March 2020) and the E-mini S&P 500 contract (September 1997–March 2020). We classify each 30-minute window as active if any trades were recorded on more than half of trading days in the relevant period, and identify regime changes as dates when new windows begin appearing consistently.

Table 17 documents the five distinct trading-hour regimes we identify. The big contract launched in April 1982 as a pit-traded instrument with no after-hours session. When the CME moved the open from 9:00 to 8:30 AM CT on September 30, 1985, total hours increased slightly. The major shift came with the E-mini launch on September 9, 1997, which added an overnight electronic session from midnight to the 3:15 PM pit close — roughly doubling available trading hours. The first overnight session appears in the tick data on September 10, which is the regime start date used in the table. On July 1, 2003, the CME expanded the session to include the afternoon and evening, bringing total weekday hours to approximately 24 and introducing Sunday-evening trading for the first time. Beginning November 16, 2012, a 30-minute daily maintenance window (4:30–5:00 PM CT) was introduced.

Table 18 documents changes in CME holiday trading policy. Before 2004, S&P 500 futures were closed on all major U.S. holidays, with one exception: the CME did not observe Martin Luther King Jr. Day as a holiday from 1986 to 1996, so futures traded normally on those days. In 1997, the CME began observing MLK Day, closing the market. The 2003 session expansion coincided with the first holiday opening (Labor Day 2003, abbreviated session). Beginning in 2004, most major holidays switched to open with near-full electronic sessions of 17–19 hours. Christmas and New Year's Day were the last holdouts, only becoming consistently traded (in short 7-hour sessions) around 2016–2017. Good Friday remains closed throughout the sample.

B. Model Derivations

This Appendix gives the derivations for the equal-supply benchmark used in Section 4. Throughout,

$$x_i = \frac{1}{N}, \qquad R_S = R_f^{1/S}, \qquad \delta_S = R_S - \lambda_S, \qquad A_S = 1 + \frac{\delta_S^2}{1-\lambda_S^2}.$$

We distinguish conditional variances from realized unconditional variances. Unbarred objects, such as $v_{m,S}$ and v_S, denote conditional return variances and are the objects that enter investors' mean-variance demands. Barred objects, such as $\bar{v}_{m,S}$ and $\bar{v}_S$, denote unconditional realized return variances and are used to compute volatility,

Table 17: S&P 500 Futures Trading Hours by Regime

Notes: All times Central Time (CT). Trading hours inferred from 5-minute tick data (TickData). A 30-minute window is classified as regularly active if trades appear on more than 50% of trading days in the period. "Overnight hours" are hours outside NYSE regular trading hours (8:30–3:15 CT). Transition dates identified from the first trading day on which new windows appear consistently. Friday closes at 3:30 PM CT in regimes 1–3 and 4:30 PM CT in regimes 4–5, with no evening session. Sunday sessions run from 5:00 PM to midnight CT.

#	Start Date	Day	Schedule (CT)	O/N hrs	Total hrs
1	Apr 21, 1982	Mon–Fri	9:00–3:30	0	6.5
	Big S&P launches (pit only)				
2	Sep 30, 1985	Mon–Fri	8:30–3:30	0	7.0
	Pit open moves 30 min earlier				
3	Sep 10, 1997	Mon–Fri	0:00–3:30	8.5	15.5
	E-mini launches on Globex	Sun	*Closed*	—	—
4	Jul 1, 2003	Mon–Thu	0:00–24:00	17.0	24.0
	Session expansion;	Fri	0:00–3:30	8.5	15.5
	Sunday trading begins	Sun	5:00p–12:00a	7.0	7.0
5	Nov 16, 2012	Mon–Thu	0:00–4:30p, 5:00p–24:00	16.5	23.5
	Maintenance window	Fri	0:00–4:30p	9.5	16.5
	(4:30–5:00p CT)	Sun	5:00p–12:00a	7.0	7.0

Table 18: Changes in CME Holiday Trading Policy for S&P 500 Futures

Notes: Holiday trading status inferred from 5-minute tick data. "First traded" is the first year the holiday has a non-zero tick count. "Consistent" means traded every year the holiday falls on a weekday. Hours shown are typical post-2014 values; earlier years averaged ∼17–18 hours.

Holiday	Pre-2004	First Traded	Consistent From	Typical Hours
MLK Day	Open 1986–96; Closed 1997–03	Jan 2004	2004	19.5
Presidents' Day	Closed	Feb 2004	2004	19.5
Good Friday	Closed	—	*Never*	—
Memorial Day	Closed	May 2004	2004	19.5
Independence Day	Closed	Jul 2004	2004	17–19.5
Labor Day	Closed	Sep 2003	2003	19.5
Thanksgiving	Closed	Nov 2004	2004	19.5
Christmas	Closed	Dec 2005 (one-off)	∼2016	7.0
New Year's Day	Closed	Sporadic	∼2017	7.0

market regression R^2's, and variance ratios. Because the equilibrium is symmetric, all individual stocks have the same conditional variance v_S and the same unconditional variance $\bar{v}_S$.

B.1 Variance components and frequency scaling

Define the market and zero-sum components

$$z_{m,s} = \frac{1}{N}\sum_{i=1}^{N} z_{i,s}, \quad \widetilde{z}_{i,s} = z_{i,s} - z_{m,s}, \qquad D_{m,t,s} = \frac{1}{N}\sum_{i=1}^{N} D_{i,t,s}, \quad \widetilde{D}_{i,t,s} = D_{i,t,s} - D_{m,t,s}.$$

The equicorrelation assumptions imply the market and zero-sum eigenvariances

$$\sigma^2_{\eta,m,S} = \sigma^2_{\eta,S}\left[\rho + \frac{1-\rho}{N}\right], \qquad \sigma^2_{\eta,\perp,S} = \sigma^2_{\eta,S}(1-\rho),$$

$$\sigma^2_{D,m,S} = \frac{\sigma^2_D}{S}\left[\rho_D + \frac{1-\rho_D}{N}\right], \qquad \sigma^2_{D,\perp,S} = \frac{\sigma^2_D}{S}(1-\rho_D).$$

Thus the variance of a representative stock's zero-sum innovation is $(N-1)\sigma^2_{\eta,\perp,S}/N$, and analogously for dividends. The calendar-period fundamental market variance is

$$S\sigma^2_{D,m,S} = \sigma^2_D\left[\rho_D + \frac{1-\rho_D}{N}\right],$$

which is invariant to S. For a general supply vector X, the market variances are obtained by replacing the equal-weight expressions with $\sigma^2_{\eta,m,S} = X'\Sigma^S_\eta X$ and $\sigma^2_{D,m,S} = (\sigma^2_D/S)X'C_D X$.

The persistence parameter and innovation variance are

$$\lambda_S = \lambda^{1-\xi+\xi/S}, \qquad \sigma^2_{\eta,S} = \sigma^2_\eta \frac{1-\lambda_S^2}{1-\lambda^2} \frac{[\mathcal{A}_S(\lambda)]^{1-\xi}}{\mathcal{A}_S(\lambda_S)},$$

where

$$\mathcal{A}_S(x) = S + 2\sum_{j=1}^{S-1}(S-j)x^j.$$

For the stationary AR(1), with $Z_t(S,\xi) = \sum_{s=1}^S z_{t,s}$,

$$Var\left(Z_t(S,\xi)\right) = \frac{\sigma^2_{\eta,S}}{1-\lambda_S^2}\mathcal{A}_S(\lambda_S) = \frac{\sigma^2_\eta}{1-\lambda^2}\left[\mathcal{A}_S(\lambda)\right]^{1-\xi}.$$

Hence $\xi = 0$ is the trade-time shock-arrival benchmark and $\xi = 1$ is the pure time-refinement benchmark. Moreover,

$$\frac{\partial}{\partial\xi}\log Var\left(Z_t(S,\xi)\right) = -\log\mathcal{A}_S(\lambda),$$

so $Var(Z_t(S,\xi))$ is decreasing in ξ for all $S > 1$.

B.2 Equilibrium price loadings

In the no-index economy, conjecture

$$P_{i,s} = p_S + b_{m,S}z_{m,s} + b_{\perp,S}\tilde{z}_{i,s}.$$

Up to constants,

$$E^A_s(\Pi^e_{i,s+1}) = -\delta_S b_{m,S}z_{m,s} - \delta_S b_{\perp,S}\tilde{z}_{i,s}.$$

Conditional on the current demand-pressure state, the market return variance, zero-sum return eigenvariance, and common individual-stock variance are

$$v_{m,S} = \sigma^2_{D,m,S} + b^2_{m,S}\sigma^2_{\eta,m,S}, \quad q_S = \sigma^2_{D,\perp,S} + b^2_{\perp,S}\sigma^2_{\eta,\perp,S}, \quad v_S = v_{m,S} + \frac{N-1}{N}q_S.$$

Define the conditional market variance share

$$\psi_S = \frac{v_{m,S}}{v_S} < 1.$$

Market clearing in the market and zero-sum directions gives

$$\frac{1-\delta_S b_{m,S}}{v_S} = \frac{\delta_S b_{m,S}}{v_{m,S}}, \qquad \frac{1-\delta_S b_{\perp,S}}{N v_S} = \frac{\delta_S b_{\perp,S}}{q_S}.$$

Solving,

$$b_{m,S} = \frac{\psi_S}{1+\psi_S}\frac{1}{R_S-\lambda_S}, \qquad b_{\perp,S} = \frac{1-\psi_S}{N-\psi_S}\frac{1}{R_S-\lambda_S}.$$

The no-index equilibrium is therefore characterized by the scalar fixed point

$$q_S(\psi_S) = \frac{N(1-\psi_S)}{(N-1)\psi_S} v_{m,S}(\psi_S),$$

where

$$v_{m,S}(\psi) = \sigma^2_{D,m,S} + \left[\frac{\psi}{1+\psi}\frac{1}{R_S-\lambda_S}\right]^2 \sigma^2_{\eta,m,S},$$

$$q_S(\psi) = \sigma^2_{D,\perp,S} + \left[\frac{1-\psi}{N-\psi}\frac{1}{R_S-\lambda_S}\right]^2 \sigma^2_{\eta,\perp,S}.$$

The unconditional realized variances implied by these loadings are

$$\bar{v}_{m,S} = \sigma^2_{D,m,S} + b^2_{m,S}\sigma^2_{\eta,m,S}A_S, \quad \bar{q}_S = \sigma^2_{D,\perp,S} + b^2_{\perp,S}\sigma^2_{\eta,\perp,S}A_S, \quad \bar{v}_S = \bar{v}_{m,S} + \frac{N-1}{N}\bar{q}_S.$$

The model-implied market regression R^2 in the no-index economy is

$$R^2_{N,S} = \frac{\bar{v}_{m,S}}{\bar{v}_S}.$$

In general, $R^2_{N,S}$ is not identical to the conditional variance share ψ_S, because the realized variances include the variance of conditional expected returns through A_S. Numerically the distinction is small in the calibration because A_S is close to one.

In the index economy, behavioral investors trade the market claim directly. With $P^I_{m,s} = p^I_{m,S} + b^I_{m,S} z_{m,s}$, the conditional market variance

$$v^I_{m,S} = \sigma^2_{D,m,S} + \left(b^I_{m,S}\right)^2 \sigma^2_{\eta,m,S}$$

cancels from the market-clearing condition,

$$\frac{1-\delta_S b^I_{m,S}}{v^I_{m,S}} = \frac{\delta_S b^I_{m,S}}{v^I_{m,S}}.$$

Hence

$$b^I_{m,S} = \frac{1}{2}\frac{1}{R_S-\lambda_S}.$$

The zero-sum idiosyncratic component has no supply in the market basket and is not traded by behavioral investors; the arbitrageur therefore chooses zero exposure to zero-sum portfolios. Individual prices can be written as $P^I_{i,s} = p^I_{i,S} + b^I_{m,S} z_{m,s}$. The unconditional realized variances are

$$\bar{v}^I_{m,S} = \sigma^2_{D,m,S} + \left(b^I_{m,S}\right)^2 \sigma^2_{\eta,m,S} A_S, \qquad \bar{v}^I_S = \bar{v}^I_{m,S} + \frac{N-1}{N}\sigma^2_{D,\perp,S}.$$

Therefore

$$R^2_{I,S} = \frac{\bar{v}^I_{m,S}}{\bar{v}^I_S} = \frac{\bar{v}^I_{m,S}}{\bar{v}^I_{m,S} + \frac{N-1}{N}\sigma^2_{D,\perp,S}}.$$

Finally,

$$\frac{b^I_{m,S}}{b_{m,S}} = \frac{1+\psi_S}{2\psi_S} > 1$$

whenever $\psi_S < 1$. Thus index trading raises the market-price loading on aggregate demand pressure whenever individual stocks contain idiosyncratic risk.

B.3 Return autocovariances and variance ratios

The return formulas apply with $b = b_{m,S}$ in the no-index economy and $b = b^I_{m,S}$ in the index economy. Let $P_{m,s} = p + bz_{m,s}$. Ignoring constants,

$$\Pi^e_{m,s+1} = D_{m,s+1} + y_{s+1}, \qquad y_{s+1} = b\eta^S_{m,s+1} - \delta_S b z_{m,s}.$$

Conditional on $z_{m,s}$, the return variance is $\sigma^2_{D,m,S} + b^2\sigma^2_{\eta,m,S}$. The unconditional realized variance also includes variation in the conditional mean. Since dividend news is serially independent and independent of demand pressure,

$$Var(\Pi^e_{m,s+1}) = \sigma^2_{D,m,S} + b^2\sigma^2_{\eta,m,S} A_S,$$

and, for $r \geq 1$,

$$Cov(\Pi^e_{m,s+r+1}, \Pi^e_{m,s+1}) = b^2\sigma^2_{\eta,m,S} C_S \lambda_S^{r-1}, \qquad C_S = -\delta_S + \frac{\lambda_S\delta_S^2}{1-\lambda_S^2}.$$

Equivalently,

$$C_S = \delta_S \frac{\lambda_S R_S - 1}{1-\lambda_S^2} < 0$$

under $\lambda_S R_S < 1$. Demand pressure therefore generates negative market-return autocovariances.

For h calendar periods, define $\Pi^e_{m,s,h} = \sum_{\ell=1}^{hS} \Pi^e_{m,s+\ell}$ and $V_h(S;b) = Var(\Pi^e_{m,s,h})$. Also define

$$\mathcal{L}_n(\lambda_S) = \sum_{r=1}^{n-1}(n-r)\lambda_S^{r-1} = \frac{(n-1) - n\lambda_S + \lambda_S^n}{(1-\lambda_S)^2}.$$

Then
$$V_h(S;b) = hS\sigma^2_{D,m,S} + b^2\sigma^2_{\eta,m,S}\left[hSA_S + 2C_S\mathcal{L}_{hS}(\lambda_S)\right].$$
The calendar-horizon variance ratio is
$$VR_{1,h}(S;b) = \frac{hV_1(S;b)}{V_h(S;b)},$$
with $b = b_{m,S}$ or $b = b^I_{m,S}$.

To see why demand pressure raises variance ratios above one, isolate its contribution,
$$Z_h(S;b) = b^2\sigma^2_{\eta,m,S}\left[hSA_S + 2C_S\mathcal{L}_{hS}(\lambda_S)\right].$$
Then
$$hZ_1(S;b) - Z_h(S;b) = 2b^2\sigma^2_{\eta,m,S}C_S\left[h\mathcal{L}_S(\lambda_S) - \mathcal{L}_{hS}(\lambda_S)\right].$$
For $h > 1$, $S \geq 1$, and $0 \leq \lambda_S < 1$, $\mathcal{L}_{hS}(\lambda_S) > h\mathcal{L}_S(\lambda_S)$: if $S = 1$, this is immediate because $\mathcal{L}_1 = 0$ and $\mathcal{L}_h > 0$; if $S > 1$, split $\mathcal{L}_{hS}$ at lag S and note that $hS - r > h(S - r)$ for $r = 1, \ldots, S-1$. Since $C_S < 0$,
$$hZ_1(S;b) - Z_h(S;b) > 0.$$
Thus the demand-pressure component has $hZ_1(S;b)/Z_h(S;b) > 1$. The fundamental component is serially uncorrelated, so adding the demand-pressure component pushes the total variance ratio above one.

B.4 Price-pressure and volume implications

The aggregate price-pressure implication follows from the loading ratio
$$\frac{b^I_{m,S}}{b_{m,S}} = \frac{1+\psi_S}{2\psi_S}.$$
For relative prices, the no-index economy has
$$P_{i,s} - P_{j,s} = b_{\perp,S}(\tilde{z}_{i,s} - \tilde{z}_{j,s})$$
up to constants. In the index economy, $P^I_{i,s} = p^I_{i,S} + b^I_{m,S}z_{m,s}$, so relative prices contain no idiosyncratic demand-pressure component.

In the polar index economy, all investor trading is in the market claim. If $\Delta n^I_{m,s}$ is the change in the index position, dollar volume in stock i is
$$dv^I_{i,s} = |\Delta n^I_{m,s}|\frac{P_{i,s}}{N}, \qquad \sum_{j=1}^{N} dv^I_{j,s} = |\Delta n^I_{m,s}|P^I_{m,s}.$$

Therefore
$$\frac{dv^I_{i,s}}{\sum_{j=1}^N dv^I_{j,s}} = \frac{P_{i,s}/N}{P^I_{m,s}}.$$
The cross-sectional regression
$$dv^I_{i,s} = a_s + b_s \frac{P_{i,s}/N}{P^I_{m,s}} \sum_{j=1}^N dv^I_{j,s} + \varepsilon_{i,s}$$
has slope $b_s = 1$, which is the volume-beta implication used in the calibration.