

Actively Passive: The Rise of Market Volatility

Lars A. Lochstoer and Tyler Muir*

October 20, 2026

[Click here for the latest version](#)

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 toward a larger transitory component in returns. We trace the increase in volatility to two related shifts in market structure: investors increasingly trade the market index, through futures, index funds, and ETFs, and can trade over more hours of the day. We exploit changes in NYSE trading hours, and in the hours and access of index futures trading, as natural experiments. In each case, expanding the opportunity to trade brings more trading and more volatility, so volatility accrues in trading time as well as calendar time. Together, the estimates from these experiments account for over half of the rise in volatility. A model in which arbitrageurs absorb investors' demand shocks, increasingly expressed through the index and over more trading hours, accounts for these facts, including the shift in returns from continuation to reversal.

*Lochstoer and Muir are with UCLA & NBER. We thank participants at the NBER Summer Institute 2026, the ASU Winter Finance Conference 2026, UNC-Duke Asset Pricing Conference 2026, the Q Group, the Federal Reserve Board, Stanford, UC Berkeley, UCLA, Jonathan Berk, John Campbell, 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 this rise in the rolling 5-year standard deviation of daily U.S. market returns. Stock market volatility, constructed from daily returns, is a central measure in macroeconomics and finance (e.g., Bloom, 2009; Martin, 2017; Berger et al., 2020).¹ The goal of this paper is to document and understand this fact.




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

The structure of the increase already rules out some explanations. Daily volatility has risen 123%. As we cumulate returns to lower frequencies the increase shrinks, to 74% for weekly returns and 37% for monthly returns. Because a monthly (log) return is the sum of the daily returns inside it, this pattern means that the autocorrelation of returns has fallen substantially. 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 and against changes in

¹Bloom (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.

the volatility of fundamentals. Further, we show that the rise in stock market volatility has been a rise relative to the realized volatility of fundamentals (profits, GDP, industrial production).

We argue that a large part of this increase reflects increased trading activity that affects the short-horizon dynamics of returns, rather than fundamentals. In the 1950s, an investor who wanted exposure to the market had to effectively buy stocks one at a time, and only during the exchange's trading day. 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 two shifts, (1) increased trading of the market as a whole and (2) an increased ability to trade more often (i.e., overnight index futures), account for over half of the increase in stock market volatility.

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 a retail index fund in 1976, and S&P 500 futures began trading in 1982. Overnight index trading became possible when S&P 500 futures were listed on CME's electronic GLOBEX system in 1993, took hold with the E-mini introduction in 1997, and has continued increasing relative to intraday trading in S&P 500 futures until the end of the sample. The trend in volatility shown in Figure 1 appears to start in the 1970s, 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 1945 and 1985, including the closing of Saturday trading, the Wednesday closures, and six changes in the length of the trading day. Using trading hours as an instrument for volume, we estimate that a 10% increase in trading volume raises market volatility by about 14%.

Much of the rise in daily volatility since the early 1990s, the period in which overnight index trading was introduced and grew, has come from the overnight portion of market returns. Our second set of experiments uses this window to isolate index trading. Individual stocks could not trade overnight, but the index could once CME opened an overnight electronic session for S&P 500 futures in September 1993, and the E-mini made overnight index trading far easier when it arrived in September 1997. Comparing overnight-to-intraday price movements for the market versus individual stocks around each change, we find overnight market variance rises at both, by about 3% per hour of overnight index trading in 1993 and by about 50% at the E-mini, and futures volume rises with it, modestly in 1993 and sharply in 1997. The market became more volatile precisely when and where index trading expanded.

The third design turns index trading on and off within a single day. On NYSE holidays no individual stock can trade, but S&P 500 futures trade through the morning and then halt until the evening, so one day supplies a window in which only the index can be traded and a window in which nothing can, back to back. In 2014 CME moved the halt from 10:30 to noon, so the same clock window is dead on holidays before 2014 and index-only after. If prices moved only because news arrives, an hour in which the index can be traded would have the same variance as an hour in which nothing trades. Instead it has more than ten times as much. Finally, in 2003 CME moved its nightly one-hour shutdown from 10pm to 4:30pm, so on a single date one window went from shut to open and another from open to shut. Variance rose in the

hour that opened and fell in the half hour that shut.

We end by sizing up the experiments. We apply the estimates from the trading-hours and futures experiments to the actual history of trading hours, across all of the NYSE schedule changes and the changes in index futures hours and access, and ask how much of the rise in volatility they would have produced on their own. The answer is over half.

In all of these experiments, expanding the ability to trade results in both more trading activity (volume) and more 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. Our findings instead imply that expanding the ability to trade leads to more trading and larger price movements: volatility and volume track trade time rather than only depending on calendar time. Because of this, much of the rise in volatility we document in the later part of our sample (from the early 1990s 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 findings bear on the rise of passive investing. Index investing is passive in the cross-section, but trading in index products is anything but. 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 The passive literature mostly studies cross-sectional efficiency (Sammon, 2025; Haddad et al., 2025; Brightman and Harvey, 2025) and the growth of passive ownership (French, 2008; Chinco and Sammon, 2024); we study what trade in index products does to aggregate volatility.

Our findings imply that volatility reflects how trading is organized as well as economic uncertainty. This matters for measuring macroeconomic uncertainty from daily stock returns (Bloom, 2009; Martin, 2017; Berger et al., 2020), for portfolio choice as systematic risk rises (Moreira and Muir, 2017; Lochstoer and Muir, 2022), and for how market macrostructure shapes asset prices (Haddad and Muir, 2025; Harvey et al., 2025). It also raises a policy question: proposals to allow nearly 24-hour trading of individual stocks could raise volatility further.

Our analysis draws on three literatures. The first is excess volatility. Shiller (1981) showed that stock prices move too much to be justified by subsequent dividends, and Schwert (1989) established the benchmark relationship between market volatility and

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

macroeconomic variables that we seek to explain. Our model builds on the demand-shock framework of Campbell et al. (1993), the noise traders of De Long et al. (1990), the limits of arbitrage in Shleifer and Vishny (1997), and the disagreement models surveyed by Hong and Stein (2007), and adds a new source of excess volatility: index-level trading that forces arbitrageurs to absorb systematic rather than idiosyncratic demand shocks. The second is index inclusion and passive ownership. Demand curves for stocks slope down (Shleifer, 1986), stocks added to the index comove more with it (Barberis et al., 2005), ETF ownership raises stock-level volatility (Ben-David et al., 2018), and indexing is linked to negative index serial dependence (Baltussen et al., 2019); see also Da and Shive (2018); Greenwood and Thesmar (2011); Israeli et al. (2017); Basak and Pavlova (2013); Sullivan and Xiong (2012). These designs exploit cross-sectional variation in exposure to index demand, whereas we ask whether the aggregate growth of index trading has raised the level of market volatility, using discrete changes in trading time and overnight index-trading access. The third is trading and volatility. French and Roll (1986) showed that trading itself induces volatility, volume and volatility move together (Karpoff, 1987) in a way usually attributed to information arrival (Admati and Pfleiderer, 1988), and Bessembinder and Seguin (1992) find mixed effects of futures activity on stock volatility in a sample with no overnight futures trading. We add a decomposition into overnight and intraday variance that compares systematic and idiosyncratic variance within the overnight window, where the index trades via futures but individual stocks cannot (see also Lockwood and Linn, 1990; Barclay and Hendershott, 2003; Hasbrouck, 2003; Lou et al., 2019).

A natural objection is that landmark launches (index futures in 1982, the Vanguard index fund in 1976, the SPY ETF in 1993) produce no immediate jump in market volatility. This is consistent with our story: these products had negligible trading volume at launch and took years or decades to gain adoption, so in instrumental variables language a launch has a weak first stage. Our experiments have strong ones: 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. Event studies

at the 1993 opening of overnight index trading and the 1997 E-mini anchor the slower buildup, and the within-day design needs no first stage at all: it compares hours in which the index can be traded with hours in which nothing can.

The rest of the paper is organized as follows. Section 2 documents the rise in market volatility. Section 3 presents the causal evidence, from changes in NYSE trading hours, from the opening of overnight index futures trading and the E-mini, and from turning index trading on and off within a day, and adds up how much of the rise the experiments explain. Section 4 develops the model. Section 5 concludes.

2. Market Volatility: Broad Facts

We take the CRSP value-weighted market return in excess of the risk-free rate at the daily and weekly frequency. Volatility is the square root of the annualized sum of squared returns over a rolling 5-year window, expressed in percent per year.³

Table 1 regresses this series on a linear time trend over 1955 to 2025:

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

where $\tilde{t} = (t - 1)/(T - 1)$ runs from 0 at the start of the sample to 1 at the end and T is the number of observations. The constant a is the level of volatility at the start of the sample, b is the percentage point increase over the sample, $a + b$ is the level at the end, and b/a is the percentage increase. The linear trend is a way to gauge the size of the rise. We do not mean that volatility will keep rising at this rate, or that the rise from 1955 to 2025 was linear.

For daily data the trend coefficient is 11.3% on a constant of about 9%: volatility has more than doubled over the sample. 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 com-

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

pute p -values using the fixed- b critical values of Lazarus et al. (2018).⁴⁵ Section 2.3 shows the result is robust to outliers, the sample start date, stale prices and other microstructure issues, and alternative ways of measuring volatility.

Table 1: Volatility Time Trends by Return Frequency

	Daily (1)	Weekly (2)	4-Week (3)
Constant	9.18 (0.69)	10.92 (0.95)	12.86 (0.97)
Trend	11.27 (2.03)	8.05 (1.96)	4.77 (1.97)
p -value	<0.001	<0.001	0.025
N	3,704	3,704	3,704
R^2	0.508	0.355	0.164
Vol increase	123%	74%	37%

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 CRSP value-weighted market excess returns, 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 CRSP value-weighted market 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.

Column 2 cumulates to weekly returns, and the trend declines to about 8 percentage points, a 74% rise over the sample. Column 3 cumulates to 4-week returns, and the trend declines further to about 5 percentage points, a 37% rise. Volatility rises at every horizon, but the rise is concentrated in short-horizon returns. Because a compounded (log) return over a longer horizon is the sum of the daily log returns inside

⁴See 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.

⁵Inference 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 [7.8, 15.1]; with 20-year blocks (only three effective blocks) it is [8.5, 14.4].

it, a trend that differs across horizons must reflect a change in the autocorrelation of returns. If autocorrelation at all horizons were zero, the expected squared weekly return would equal the expected sum of squared daily returns within the week.

The pattern across horizons points to changes in financial markets and trading itself. Time-varying volatility of fundamentals alone does not produce a rise that differs by return horizon: with uncorrelated returns, a rise in daily variance passes one for one into weekly and monthly variance. A rise concentrated at short horizons instead requires a rise in “discount rate” volatility, the volatility of a mean-reverting component of returns, and because this mean reversion occurs over days and weeks rather than years, the natural candidates are based on trading.⁶ We return to fundamentals shortly.

2.1 Horizon Effects and Return Autocorrelations

Figure 2 plots the trend in volatility by return horizon, from 1 day out to 1 quarter (13 weeks). The point estimate is positive at every horizon and significant out to around 6 weeks. The gradual decay is consistent with short-run reversals in returns.

Figure 3 plots the ratio of 5-year return variance from daily returns to 5-year return variance from monthly returns, alongside daily volatility from Figure 1. The ratio is the variance-ratio statistic of Lo and MacKinlay (1988) and measures autocorrelation: when daily-based variance exceeds monthly-based variance the market price mean-reverts, because short-term price pressure that reverts quickly shows up in daily variance but much less in monthly variance. We report the variance ratio in logs, so that only its sign depends on which variance is on top; with daily-based variance on top, as here, a positive log ratio means reversal and a negative one continuation, and the Lo and MacKinlay (1988) ratio is the same number with the opposite sign. The two series track each other closely, and the log variance ratio moves from around -0.6 to above 0.6 , crossing zero in the late 1990s: returns went from continuation to rever-

⁶Typical discount rate volatility in traditional macro finance models (say, with time-varying risk-aversion, or the long-run risk model of Bansal and Yaron, 2004) focuses on long-run mean-reversion at the horizon of years rather than weeks.




Figure 2: Time Trend in Realized Volatility by Return Frequency

Notes: Trend coefficient from regressing 5-year rolling annualized volatility on a linear time trend, by return horizon: daily (plotted at 0.2 weeks) and 1–12 week overlapping returns, constructed as in Table 1. The shaded band is a 95% confidence interval from Newey–West standard errors with $S = 624$ lags and fixed- b critical values. Sample: 1955–2025.

sal. The variance ratio breaks in August 1999, within a year of the break in the level of volatility (Appendix A.4), and the same rise appears in 35 international markets (Appendix A.7).

Section A.8 plots the 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 declines from roughly 0.30 in the 1970s to -0.15 by the 2020s, crossing zero around 2003, while the weekly series crosses zero around 1990 and ends near -0.12 . Both trends are strongly significant under fixed- b inference. The decline in daily autocorrelation is documented by Froot and Perold (1995) and Campbell (2018); its early positive values partly reflect lead-lag effects across stocks (Lo and MacKinlay, 1990). Appendix C.2 measures and removes the effect of such lead-lag correlations on the aggregate daily return volatility over the sample. We find that these effects account for only about 1.3 percentage points of the 11.3 percentage point increase in daily volatility.




Figure 3: Variance Ratio and Volatility

Notes: Left axis (solid): log variance ratio, the log of 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 log ratio above zero indicates negative return autocorrelation (mean reversion); below zero indicates positive autocorrelation. Right axis (dashed): 5-year rolling annualized volatility as in Figure 1. Shaded regions indicate NBER recessions. Sample: 1955–2025.

2.2 Volatility, Volume, and Fundamentals

Figure 4 plots market volatility against total stock market volume as a share of market capitalization over 5-year windows from 1955 to 2025. Since the 1960s the two move together, in the low-frequency trend and in the cycle. A natural interpretation is that more trading produces more volatility, especially at short horizons, consistent with the horizon pattern above. But other factors could move both series, and high volatility could itself raise the demand for trading. Section 3.1 addresses causality with changes in trading hours as an instrument for volume and with the introduction of the E-mini futures.

Figure 5 plots stock market volatility relative to the volatility of fundamentals from 1926 on, with the variance ratio beside it. The black line is stock market volatility divided by the volatility of industrial production growth, both over ten-year rolling




Figure 4: Volatility and Turnover

Notes: Left axis (solid): 5-year rolling annualized volatility of the CRSP value-weighted market excess return, computed from daily returns as in Figure 1. Right axis (dashed): 5-year rolling annualized turnover. Daily turnover is total dollar volume divided by total market capitalization across all CRSP common stocks; it is averaged within each month, then over a 60-month rolling window, and multiplied by 252 to annualize. Shaded regions indicate NBER recessions. Sample: 1955–2025.

windows, and therefore a rise means the market has become more volatile relative to the real economy. The blue line is the variance ratio of Figure 3 on the same windows. Both are log deviations from their 1950–1974 means.

The Great Depression is the natural test. Stock volatility averaged 24% through the Depression and the war, the highest in the sample, but the volatility of the macroeconomy was also extremely high. Scaled by the volatility of fundamentals, the Depression looks like the 1950s and 1960s. The variance ratio makes the same point on a separate axis: its log was -0.30 in the Depression and -0.32 in 1950–1974. Both series trend up over the last fifty years. Since 2000 stock volatility has averaged more than eight times industrial production volatility, compared with about two in the 1930s and the 1950s, and the log variance ratio has averaged 0.40. Real GDP and corporate profits give the same picture as industrial production (Figure 16 in Appendix A.2).

This echoes a long literature starting with Schwert (1989) that struggles to tie stock




Figure 5: Stock Market Volatility Relative to Fundamental Volatility, and the Variance Ratio, 1926–2025

Notes: Solid black: annualized volatility of the daily CRSP value-weighted market excess return divided by the annualized volatility of monthly industrial-production growth (FRED INDPRO), both over 10-year rolling windows. Solid blue: the variance ratio of Figure 3, the ratio of cumulative daily to cumulative monthly realized variance, over the same 10-year windows. Both series are plotted as log deviations from their 1950–1974 means; the dotted line marks a variance ratio of one. The dashed line is a linear trend from 1955 fitted to the average of the two series. Data begin in 1926, so the first window ends in mid-1936. Industrial-production growth for March–July 2020 is excluded: the shutdown and reopening of the economy produce a swing in measured growth that would otherwise remain in every window for ten years, but it was a temporary pause rather than a change in the volatility of fundamentals.

market volatility to fundamentals. Fundamentals may also have changed in character around 2000, toward more persistent shocks (Chernov et al., 2026); a shift of that kind changes the level of volatility at every horizon, not the gap between daily and monthly volatility that the term structure in Figure 2 shows. Controlling for industrial-production volatility sharpens the point. On the full 1931–2025 sample a linear trend in stock volatility is statistically zero, because Depression-era volatility offsets the modern rise, but adding 5-year rolling industrial-production volatility as a control moves the trend to 10.5 percentage points (s.e. 4.5). On the post-1960 subsample the trend is 10.7 (s.e. 2.1) without the control and 10.6 (s.e. 2.2) with it.⁷ Together

⁷Newey-West HAC with 60 monthly lags; the appendix also reports a structural-break version with a 1960 trend break.

with the horizon pattern, this makes fundamentals a poor explanation for the rise in volatility and trading activity a promising one.

The horizon pattern also says how much of the rise fundamentals could account for. Permanent news raises the volatility of returns at every horizon alike, and therefore a rise in the volatility of fundamentals would show up as much in quarterly returns as in daily ones. Instead, Figure 2 shows the trend shrinking quickly with the horizon: it is less than half as large at a month as at a day and smaller still by a quarter. Most of the increase in daily volatility is therefore specific to short horizons, which is the same statement as the rise in the variance ratio in Figure 3. A direct estimate of long-run variance, using the estimator of Dew-Becker (2017), does not change over the sample (Appendix C.4). Part of the early gap between short- and long-horizon volatility reflects the slow diffusion of common news across stocks, which Appendix C.2 measures, and its disappearance accounts for a relatively small part of the rise in daily volatility.

2.3 Robustness of the rise in volatility

Table 2 shows that the rise in volatility survives outliers, a later start date, and stale prices.

2.3.1 Outliers

Volatility is right skewed, with extreme events such as the 1987 crash. Column (1) of Table 2 therefore replaces the rolling variance with the interquartile range of weekly realized variances within each 5-year window, rescaled to annualized volatility; the trend is almost identical at 11.1%. Thus, events like the 1987 crash or Covid are not driving the rise.

2.3.2 Trend robustness

Institutions were increasingly holding and trading market portfolios from the 1960s, the first retail index fund (Vanguard) opened in 1976, and index futures and options

followed in the early 1980s, and therefore the start date could matter. Column (2) starts the regression in 1980 with the trend variable unchanged (the coefficient is still scaled as a 70-year change) and finds a slightly stronger rise at a little under 12%.

We do not take the trend literally. As in Campbell et al. (2001), it is a convenient summary of a slow-moving rise. Estimating structural breaks in the mean of volatility instead gives two breaks, around 1975 and 2000, with the mean rising substantially at each and the overall increase consistent with the trend (Appendix A.4). The trend is stable across start dates, rolling-window lengths, and 30-year rolling subsamples (Appendix A.5).

2.3.3 Concerns about stale prices

Volatility can look low when many stocks do not trade on a given day. Column (3) weights stocks by dollar volume instead of market capitalization (stocks that do not trade get zero weight), and column (4) uses only the 100 largest stocks, the most liquid and least exposed to stale prices. Both give a nearly identical trend. Campbell et al. (1993) reach the same conclusion on nonsynchronous trading by re-running their volume and autocorrelation results on the Dow Jones Industrial Average and on individual large stocks. The rise also holds under equal weighting, across size segments, and for the S&P 500 directly. Thus, it is not an artifact of value-weighting or mega-cap concentration (Appendix A.6).

2.3.4 Systematic vs. idiosyncratic volatility

The rise in market volatility is systematic, not firm-specific. We follow 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 their 49 industries, starting in 1962 to match them and extending through 2024; our implementation replicates their headline numbers within tolerance on their 1962–1997 sample.⁸ The market share of average-stock variance nearly doubles, from 0.18 in 1962–1971 to 0.31 in 2015–2024, and the total systematic share (market plus industry)

⁸See 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	8.70 (0.70)	8.79 (2.24)	12.47 (0.94)	10.28 (0.84)
Trend	11.11 (1.91)	11.74 (3.74)	12.48 (3.22)	10.04 (2.13)
p -value	<0.001	0.008	<0.001	<0.001
N	3,705	2,401	3,653	3,653
R^2	0.531	0.247	0.360	0.447
Vol increase	128%	133%	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.

risers from 0.31 to 0.45. Both trends are significant: 2.50 percentage points per decade for the market share ($t = 3.95$) and 3.25 for the systematic share ($t = 4.01$).

The variance-ratio regime change in Figure 3 above also appears in the systematic component. The market-component log variance ratio rises from -0.34 in 1961–1995 to 0.39 in 1996–2024, the industry component shifts the same way but less, and the firm-specific variance ratio shows no regime change, only a steady decline consistent with falling microstructure noise (narrower spreads, decimalization). Our 1962–1997 numbers reproduce the CLMX result on rising firm-specific volatility, and the reversal after 2001 that Campbell et al. (2023) document appears in our extended sample too. The rise in the market component persists across 1962–2024, while firm-specific volatility has fluctuated around its long-run level.

2.3.5 Implied volatility

The VIX, the most widely cited measure of equity volatility, has not trended up since 1990, and Dew-Becker and Giglio (2026) make this unsurprising: they document a large decline in the variance risk premium over the same period and attribute it to declining intermediary frictions in the options market. In our data realized volatility rises by roughly 1.8 percentage points per decade since 1990, the variance risk premium falls at an almost exactly offsetting rate, and implied volatility is flat. The implied-volatility sample is about 30 years rather than 70, so a trend test on it alone has little power, but the joint pattern across the three series is consistent with the rise we document (Appendix A.3).

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&P 500 index funds, S&P 500 futures, or similar products, it has become increasingly easy and prevalent over the past 70 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 of “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

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

where $Volume_{i,t}$ represents the dollar volume in stock i on 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 .

Panel A of 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. This suggests that today a large fraction of trading volume in the stock market may be driven by trading the market index. A slope near one is what index trading implies rather than proof of it, since any trading that is proportional to market capitalization produces the same slope.

Panel B shows a more direct measure of index investing – the asset-weighted share of funds whose returns have a correlation with the S&P 500 of at least 0.98 (or 0.99). We use fund returns and assets from the CRSP mutual fund database and focus on U.S. domestic equity funds. Both shares increase substantially after 1990, to around 35–40% by the end of the sample.

3.1 Causal Evidence: Volume and Volatility

The link between market volatility and market trading is suggestive, but it is not causal. Our claim is that the ease of trading index products drove trading in the




Figure 6: Market Trading vs Market Volatility

Notes: Panel A, 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. Panel B, left axis (solid): asset-weighted share of U.S. domestic equity funds whose monthly returns over the previous three years have a correlation with the S&P 500 of at least 0.98 or 0.99 (CRSP mutual fund data, 24-month average). Right axis in both panels (dashed): 5-year rolling annualized volatility as in Figure 1. Shaded regions indicate NBER recessions. Sample: 1955–2025 (Panel B from 1966).

market itself, and that this trading raised volatility. We test it with three sets of quasi-natural experiments. The first uses changes in the trading hours of the NYSE as an instrument for trading volume. The second uses the S&P 500 futures market. Two changes altered when and how easily the index could be traded: an overnight electronic session opened in 1993 and the E-mini arrived in 1997, while individual stocks could not trade overnight either way. The third turns index trading on and off within a single day: on NYSE holidays futures trade through the morning and then halt, and CME's schedule changes moved the boundaries between open and shut. We close by chaining the estimates together to ask how much of the rise in volatility they can account for.

An important consideration in choosing these experiments is the strength of the first stage. The most visible events in the history of index trading (Vanguard's index fund in 1976, S&P 500 futures in 1982, the SPY ETF in 1993) did not produce immediate surges in volume. They started small and took decades to catch on, 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. The within-day design goes one step further: it needs no first stage at all, because the comparison is between hours in which the index can be traded and hours in which nothing can.

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 pool every scheduled change in NYSE trading hours between 1945 and 1985 into one regression. From 1945 to 1952 the exchange shut its two-hour Saturday session each summer and re-

⁹If 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.

opened it each autumn, until it closed Saturdays for good in 1952. In 1968 it closed on Wednesdays for six months and then reopened them with a shorter trading day. Finally, between 1952 and 1985 it changed the length of the weekday session six times: a half-hour extension in 1952, and, after the closure ended in January 1969 with a 4-hour trading day, five extensions that took the trading day to the 6.5 hours of today.

We present two versions of this analysis. The cleanest test is at the daily level (Panel A of Table 3), where the Saturday and Wednesday closures move the trading hours of one day of the week and leave the other days alone. This lets us hold calendar time fixed and ask whether trading time matters for volatility above and beyond calendar time. The weekly aggregate (Panel B) asks a different question: whether the variance and volume that a closure removes from one day reappear on the other days of the week. There the comparison across days is no longer available and the test is correspondingly less sharp, but any reallocation within the week nets out by construction.

Panel A: Trading hours and volatility. Our main evidence on whether trading hours affect volatility uses daily returns, the sharpest test of our story. Each schedule change dictates which returns gain or lose trading hours, and the returns it leaves alone provide the control, either the other days of the same week when a session is closed or the same days in the surrounding weeks when the trading day is lengthened. This is where hours should show up most directly, and the comparison gives both identification and power. When the NYSE closes on Saturdays, the return from Friday's close to Monday's close still spans the same calendar hours but trading hours fall, while the Tuesday through Friday hours are unchanged. When the NYSE closes on Wednesdays, the return from Tuesday's close to Thursday's close still spans 48 hours but contains one session instead of two. When the NYSE lengthens the trading day, every return contains more trading hours and the calendar is unchanged. We therefore measure returns between consecutive closes on Monday through Friday, so that a Saturday session, when there is one, falls inside the Friday-to-Monday return, a closed Wednesday falls inside the Tuesday-to-Thursday return, and the calendar

length of each return is fixed by its day of the week. Fixed effects for the day of the week within each event and for each calendar month then hold calendar time and the level of volatility fixed. For the weekday-session changes, where every day is affected, each window is a single block and the comparison is simply the month after against the month before. Thus, the only variation left in trading hours is the schedule change itself.¹⁰

Under the null hypothesis that information arrival drives variance, a return's variance depends on the calendar time it spans and not on the trading hours inside it. Under the alternative that trading generates variance, more trading hours raise variance at fixed calendar time. Specifically, we estimate

$$\ln r_t^2 = fe + \beta \times \ln \text{trading hours}_t + \varepsilon_t \quad (3)$$

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

where r_t is the market excess return between consecutive closes, turnover is dollar volume divided by market capitalization over the sessions the return spans, and fe are the fixed effects just described. Under the information null $\beta = 0$, and if variance is proportional to trading time $\beta = 1$. Instrumenting turnover with trading hours gives the elasticity of variance to volume as the ratio β/γ . We use ± 22 trading day windows around each weekday-session change and the full closure and Saturday periods for the other two experiments, yielding 4,790 daily returns. Standard errors are from a block bootstrap that resamples calendar months within each event.

Our identifying assumptions are as follows. First, changes in trading hours affect volatility only through trading activity. Second, the timing of a schedule change does not depend on expected volume in the days around it. Demand for trading surely affects trading hours in the long run, but this itself is not a problem. We need only that a permanent change in hours was not timed to a predicted spike in volume inside the event window, and this is most plausible when the window is narrow. For example,

¹⁰The Saturday closings ran from late May to late September each year, so the weekend return also gets month-of-year effects, estimated from the five years on either side of 1945–52 in which Saturdays did not switch. Returns that contain a holiday are excluded.

the assumption is that the NYSE did not permanently lengthen the trading day on September 29, 1952 because it anticipated unusual volume in the following 22 trading days. For the Saturday and Wednesday closures the comparison is across days of the same week, so the level of volume and volatility in the window does not matter. For the weekday-session changes the comparison is before against after, and the width of the window trades identification against power. We use a month on either side at the daily level and three months at the weekly level, short relative to the years over which trading demand builds.

Panel A of Table 3 presents the results. The first stage (column 1) is very strong: a 10% increase in trading hours increases turnover by about 8%, with an F -statistic of 206 on the excluded instrument. Column 2 is the reduced form and the experiment itself. Variance increases with trading hours with an elasticity of 2.2 ($t = 3.5$), meaning a 10% increase in trading hours raises variance by about 22% and volatility by about 11%. This rejects the information null of zero. Moreover, the estimate is above one, the value at which variance would simply be proportional to trading time. Column 3 instruments turnover with trading hours and gives an elasticity of variance to volume of 2.7, significant at the 1% level, meaning a 10% increase in turnover raises variance by about 27%. The two closure experiments also answer a different null. In both, the return runs from one close to the next across the closed session (Friday to Monday for the Saturday closings, Tuesday to Thursday in 1968). If demand arrives on the calendar and a closure only changes when it is priced, for example because arbitrageurs are less willing to carry positions across a longer closure, then the same news and the same demand are priced by the following close either way and the elasticity on these spans is zero. Instead, the two closure experiments on their own give an elasticity of 2.3 (s.e. 0.6), and a return that starts the session before the 1968 closure and ends at the reopen (Monday close to Thursday close) gives 2.5 (s.e. 0.9) relative to the Monday and Friday returns of the same weeks. The variance a closed session removes is therefore not recovered at the reopen, and it did not move to the session before.

Table 3: Trading Hours, Volume, and Variance on the NYSE

Notes. Panel A uses daily data and Panel B weekly data. The dependent variable in columns (2)–(4) is the log squared market excess return, and coefficients are therefore elasticities of variance (halve them for volatility). Trading hours are the hours the schedule provides in the day or week; turnover is CRSP dollar volume divided by market capitalization. Panel A pools every scheduled change in NYSE hours, 1945–1985: the Saturday closings of 1945–52, the 1968 Wednesday closure, and six changes to the weekday session in ± 22 -trading-day windows, with fixed effects that hold calendar time fixed. Panel B uses the closure and the five weekday changes of 1969–85 in ± 13 -week windows with event fixed effects. The 1952 change is excluded from Panel B because weekly hours barely moved (the appendix reports the estimate with it). Standard errors in parentheses are from a block bootstrap of calendar months within event (1,000 draws); the weekly IV reports the bootstrap 95% interval, since its first stage is weaker. ***, **, * denote significance at the 1%, 5%, and 10% level.

Panel A: Trading hours and volatility

	(1) First Stage	(2) Reduced Form	(3) IV	(4) OLS
Dep. variable:	ln Turnover	ln r^2	ln r^2	ln r^2
Log Trading Hours	0.82*** (0.06)	2.21*** (0.63)		
Log Turnover			2.70*** (0.74)	3.13*** (0.18)
Observations	4,790	4,790	4,790	4,790
R-squared	0.77	0.13		0.22
First-stage F			206	
FE: Month, Event $\times$ Span Type, Weekend Season	Y	Y	Y	Y

Panel B: Weekly totals and reallocation across days

	(5) First Stage	(6) Reduced Form	(7) IV	(8) OLS
Dep. variable:	ln Turnover	ln r^2	ln r^2	ln r^2
Log Trading Hours	0.95*** (0.32)	2.74** (1.23)		
Log Turnover			2.90 [0.67, 9.45]	1.46*** (0.46)
Observations	167	167	167	167
R-squared	0.91	0.51		0.53
First-stage F			9	
FE: Event Window, Holiday, Week Length	Y	Y	Y	Y

Panel B: Weekly totals and reallocation across days. A concern with the daily evidence is that the volume and volatility lost on a closed Wednesday may simply move to the other days of the week. Investors who would have traded on Wednesday may trade more on Monday, Tuesday, or Friday instead, in which case the daily esti-

mate overstates how much variance the closure removed from the week. The variance should show up on Thursday, whose return spans the closed day, but it is still useful to check whether any of it migrated elsewhere in the week. The weekly test is also simpler. Once we aggregate to the week, calendar time is fixed by construction and any reallocation across days nets out, and we therefore need neither the day-of-week fixed effects nor a story about which day absorbs the trading. The cost is that the test is less sharp than comparing adjacent days, and to recover power we use longer windows of 13 weeks on either side of each change.

The dependent variable is the sum of squared daily returns within the week, the regressor is the trading hours the schedule provides in that week, and a holiday no longer moves the regressor and enters instead as a dummy for the session lost. Each change gets its own fixed effect, and we estimate (3) and (4) on weeks in place of daily returns. Without the comparison across days of the week, the Saturday closings have nothing to identify from: the summer closings recur in the same months every year and cannot be separated from a seasonal control, and the September 1952 extension of the weekday session offset the final closing, so that weekly hours barely moved. We therefore leave the Saturday closings and the 1952 extension out of Panel B. The Wednesday closure and the five later changes to the weekday session remain, which gives 167 weeks in total.

Panel B of Table 3 presents the weekly results. Trading hours increase volume one for one (column 5) – when the stock market is open longer, people trade more. If investors had made up for a closed Wednesday by trading more on the other days, weekly volume would respond to weekly hours with an elasticity well below one. It does not, and the same holds for variance. Variance increases with trading hours with an elasticity of 2.7 (column 6), close to the daily estimate of 2.2 and with twice the standard error, as we would expect from a design that no longer compares days within the week. In column 7 the IV is 2.9, with a wide interval because the weekly first stage is weaker. Using the squared weekly return in place of the sum of squared daily returns gives the same elasticity with more noise (see the appendix). This indicates that longer hours raise the variance of daily returns without changing how the days

within a week are correlated.

Overall, more trading causes more volatility in the market index. The daily design holds calendar time fixed and rejects the information null: trading hours matter in their own right, not just calendar time, and the weekly design shows that the variance a closure removes does not reappear elsewhere in the week.

3.1.2 The E-mini Futures as a Natural Experiment

The E-mini S&P 500 futures contract launched on September 9, 1997, and it did not add trading hours: the full-size S&P 500 contract had traded overnight on CME's GLOBEX system since September 1993, and the E-mini's own overnight session added at most half an hour to that (Section 3.1.3 and Table 14 in the appendix give the record). Instead, the E-mini changed access. It was a fraction of the full-size contract's size, it traded electronically around the clock, and it ran alongside the pit during regular hours, so an investor with a screen could trade the index at retail size at any hour without going through a floor broker. Volume in S&P 500 futures rose sharply within weeks of the launch. Individual stocks could not trade overnight either before or after.

We compare the market's overnight-to-intraday variance ratio with the same ratio for individual stocks (idiosyncratic vol), before and after the E-mini, since more overnight trading should raise overnight variance for the market but not for individual stocks. The stock ratio controls for news that arrives at night and moves stocks in general; the intraday denominator controls for anything that moves overall volatility in this period.

We use close to open and open to close returns as measures of overnight and intraday returns, respectively.¹¹ 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). We compute the sum of squared returns within a rolling 22 day pe-

¹¹Thus, 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.

riod (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

$$\frac{\text{var}(r_{mkt}^{night})/\text{var}(r_{mkt}^{day})}{\text{var}(r_{stock,idio}^{night})/\text{var}(r_{stock,idio}^{day})},$$

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&P 500 increased overnight after the introduction of the E-mini.

However, for our story to be correct, we also need that the E-mini generated significant trading in the S&P 500 when it was introduced. Panel B adds the (log) total volume of S&P 500 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&P 500 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&P 500, consistent with our story.

Table 4 tests the increase formally. Because the dependent variables are 22-day rolling averages, a simple pre/post dummy gets the timing wrong: on September 10 only 1 of 22 observations in the rolling window is post-launch. 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 is 0.43 ($t = 7.9$), meaning the market's




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&P 500 futures volume (22-day rolling average) on the right axis.

Table 4: The E-mini, Futures Trading Hours, and Overnight Market Variance

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 for columns (1)–(2) is March 1997 to March 1998. Column (3) pools four CME schedule changes in ± 6 -month windows with event fixed effects: the GLOBEX listing of the full-size contract (September 20, 1993, 16.25 overnight hours), the E-mini (September 9, 1997), the session expansion of September 29, 2003 (0.75), and the maintenance window of November 19, 2012 (–0.25). The three hours changes enter as a linear dose in hours ($\text{Post} \times \Delta\text{Hours}$) and the E-mini as its own indicator; each event window keeps its own coefficient on the idiosyncratic control; Newey-West standard errors (44 lags) are computed within event blocks. The E-mini coefficient in column (3) differs from column (1) only because the window is ± 6 months (252 observations for that event) rather than March 1997 to March 1998. The per-event estimates are in Table 15. HAC standard errors (Newey-West, 44 lags) in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% level.

	(1) E-mini Log Market Var Ratio	(2) E-mini Log Volume	(3) Pooled, four CME changes Log Market Var Ratio
E-mini ramp (access)	0.434*** (0.055)	0.599*** (0.048)	0.430*** (0.058)
$\text{Post} \times \Delta\text{Hours}$ (per hour)			0.0301** (0.0120)
Log idio stock var ratio	0.858*** (0.169)		Y (event-specific)
Event FE			Y
Observations	273	273	1,006
R-squared	0.528	0.839	0.44

overnight variance relative to intraday rose by about 50% ($e^{0.43} \approx 1.5$) after the E-mini launch, relative to the same ratio for stocks. Column (2) shows log futures volume rose by 0.60 over the same window, an increase of roughly 80%.

3.1.3 Futures Trading Hours, Access, and Overnight Market Variance

The E-mini is one step in a larger shift. Over the past thirty years, index trading has moved steadily into the overnight hours, and overnight volatility has increased with it. Figure 8 plots the volatility of the market return using 5-year rolling windows as

in our main plots earlier, but 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 overnight returns: overnight volatility rises from about 4% in the mid-1990s to about 14% in the most recent window, and the ratio of overnight to intraday volatility goes from under 60% around 2000 to over 90% today.




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.

This rise coincided with the growth of overnight index trading: overnight activity in S&P 500 futures has grown steadily relative to daytime activity since the E-mini launch (Figure 24 in the appendix).

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 when overnight index trading expands, 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 intraday. 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; overnight and intraday returns use the same construction anchored on CRSP total returns. The object of interest is the daily triple-difference

$$TD_t = (\log V_{mkt,t}^{ovn} - \log V_{mkt,t}^{intra}) - (\log V_{idio,t}^{ovn} - \log V_{idio,t}^{intra}), \quad (5)$$

the log ratio of overnight-to-intraday market variance relative to the same ratio for idiosyncratic returns, the 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. The remainder is specific to the market's overnight window, where index futures trade and individual stocks do not.

The record. Table 14 in the appendix gives the exchange's own record of trading hours, sourced line by line. Hours available for trading S&P 500 futures outside NYSE regular hours were essentially zero until September 20, 1993, when CME listed the full-size contract on GLOBEX and overnight hours went from a quarter of an hour (the pit's tail after the NYSE close) to 16.5 per weekday. The E-mini in September 1997 added at most half an hour. The September 2003 session expansion added three quarters of an hour and the November 2012 maintenance window removed a quarter. Thus, the futures market gave the index two distinct things at two distinct dates: *hours*, in 1993, and *access*, in 1997.

Specification. We treat the two channels separately:

$$TD_t^{mkt} = a_e + \beta_H \text{Post}_t \times \Delta \text{Hours}_e + \beta_A \mathbf{1}[\text{E-mini}]_t + \gamma_e TD_t^{idio} + \varepsilon_t, \quad (6)$$

where the dependent variable is the 22-day rolling log ratio of overnight-to-intraday market variance, the control is the corresponding idiosyncratic ratio with its own coefficient in each event window, and Post_t is the ramp used in Table 4, rising from zero to one over the 22 trading days after each event. Each event gets a ± 6 -month window and an event fixed effect a_e . The three hours changes (1993, 2003, 2012) enter as a linear dose in hours, so β_H is the effect of one additional hour of overnight index trading. The E-mini gets its own indicator, so the hours channel is not asked to absorb an access event whose hours change is at most half an hour. Newey-West standard errors use 44 lags, since consecutive observations of a 22-day rolling series overlap by 21 days, and are computed within event blocks. The design mirrors Table 3: pool the events, absorb event fixed effects, identify from within-event variation, so that anything that evolves smoothly, such as growing demand for around-the-clock trading or the globalization of information flow, cannot produce a step at a CME schedule date.

Results. Column (3) of Table 4 pools the four CME schedule changes. The hours coefficient is $\beta_H = 0.030$ ($t = 2.5$): each hour of overnight index trading raises the market's overnight variance by about 3% relative to stocks, so the 16.25 hours that opened in 1993 raise it by $e^{0.49} \approx 1.6$, or 60%. The access coefficient is $\beta_A = 0.43$ ($t = 7.4$), the E-mini estimate of column (1) re-estimated inside the pooled design: access raised overnight market variance a further 50% with almost no change in hours. Both channels come with more futures trading, by 0.60 log points at the E-mini (column (2)) and by 0.17 at the GLOBEX listing (Table 15). The per-event estimates are in Table 15 in the appendix: the 1993 listing alone gives 0.48 ($t = 2.5$), while the 2003 and 2012 changes are a fraction of an hour each and imprecise.¹²

The per-hour effect is small because in 1993 the overnight session existed but was thin, and the overnight variance ratio therefore builds over the following years as volume arrives (Figure 8) rather than stepping at the hours change. Hours are neces-

¹²Two measurement choices bias these estimates toward zero. Idiosyncratic returns impose a unit market beta, and where a stock's true beta differs from one the idiosyncratic ratio retains some market exposure, which the control then absorbs. In addition, more firms moved their earnings releases to after the close over the sample, which raises idiosyncratic overnight variance, the control, and therefore lowers the estimate.

sary for overnight index trading, not sufficient. Because the 2003 and 2012 changes are fractions of an hour, the pooled dose is identified almost entirely by 1993. The appendix takes up the robustness of both large events (Table 16).

Over the full 1993–2024 sample, the market-overnight cell rises about 2.1% per year relative to the other three cells (SE 0.4%), from a stacked regression of the four log-variance cells on cell-specific trends. That is a description, not identification: over three decades the growth of overnight index trading is collinear with everything else that changed in market structure. The event evidence identifies the two channels locally, and the trend describes what they grew into.

3.1.4 Turning Index Trading On and Off Within the Day

Where the experiments so far compare periods before and after a change in the ability to trade, the last design compares hours within a single day.

Consider first the null hypothesis. If prices move because news arrives, then variance accumulates on the clock whether or not a market is open. Closing the market defers the adjustment, and the deferred price changes show up as a jump when trading resumes. Under that view the squared jump across a closure should have as much variance per hour as those same hours would have if the market were open. Under our view it should carry much less, because part of the variance is generated by trading itself. French and Roll (1986) ran this test on the NYSE’s 1968 Wednesday closures and found the jump small. We run it on the index futures market, where CME’s schedule supplies closures inside a day, and where the thing that is switched on and off is the ability to trade the index and nothing else.

Two windows on one day. On a U.S. market holiday the NYSE is closed all day, so no individual stock can trade. S&P 500 futures open the evening before, trade through the morning, and halt at midday; from the halt until the evening reopen, nothing trades at all. Thus, one holiday supplies two windows back to back. In the morning the index can be traded and its components cannot: the only claim anyone in the world can buy or sell is the market as a bundle. In the afternoon nothing can be

traded. The two windows share the same day, the same news flow, the same season, and the same state of the world, and the only thing that changes is whether a market is open. Panel A of Figure 9 shows what that does to the price. On an ordinary weekday the afternoon is the busiest part of the session; on a holiday the afternoon rate declines to a level far below the same day's morning and below an ordinary weeknight.


Moreover, the boundary moved: in 2014 CME moved the holiday halt from 10:30 to noon Chicago time. The clock window 10:30–12:00 on a holiday is therefore *dead* before 2014 and *index-only* after it, and fully open on ordinary days. Panel B of Figure 9 plots the hour-of-day profile of variance on holidays before and after the change, each scaled by its own overnight rate. The same clock window on the same kind of day is a small fraction of the overnight rate when it is shut and at or above the overnight rate when the index trades. This is the comparison the tables below turn into a number, and it rules out the objection that mornings are simply busier than afternoons.

Specification. Let $V_{d,w}$ be the variance per hour of the futures price on day d in clock window w , measured in levels from five-minute bars. A shut window has no bars, and its variance is therefore one squared increment, from the halt price to the price five minutes after the reopen, spread over the shut minutes plus five. We use five minutes as a conservative choice to avoid stale reopens. For every weekday from 2004 to 2019 we take the 10:30–12:00 window and subtract the same day's overnight rate over 20:00–07:00 Central. In that overnight window only the index trades on every kind of day, and the news is Asian and European rather than U.S., and therefore a quiet holiday is quiet in both windows and the level cancels within the day. We then estimate

$$V_{d,10:30-12:00} - V_{d,20:00-07:00} = a + \beta_I \mathbf{1}[\text{index tradeable}]_d + \beta_B \mathbf{1}[\text{ordinary business day}]_d + \varepsilon_d. \quad (7)$$

Dead is the baseline. The two indicators nest, and therefore β_I is the increment from turning the index on, identified from holidays before and after 2014, and β_B is the

Panel A: Variance by hour of day, ordinary weekdays and holidays



Panel B: The 10:30–12:00 window on holidays, before and after the 2014 halt change




Figure 9: Index Trading Against a Shut Market.

Variance per hour of the E-mini S&P 500 futures price by hour of day (Central time), from five-minute bars. Panel A: hours-weighted average over ordinary weekdays and over NYSE holidays, July 2003 to March 2020, log scale, with NYSE regular hours marked. On holidays futures halt at 10:30 (before 2014) or noon (after) and reopen at 17:00; the shaded region marks the shut afternoon, whose variance is one jump from the halt to the reopen, drawn as its average rate; the band is a 95% interval. Panel B: mean profile in 15-minute blocks on Monday–Thursday holidays halted at 10:30 (2004–2013, 57 days) and at noon (2015–2019, 29 days), and on ordinary Tuesday–Thursday weekdays, each relative to its own 20:00–07:00 overnight rate; a shut window’s variance is one increment from the halt price to the price five minutes after the reopen. In the 10:30–12:00 window the dead window is at 0.07 of the overnight rate and the index-only window at 1.77. Panel B classifies holidays by calendar year, omits 2014, and keeps September 2, 2019; Table 5 keys the change to its date, May 26, 2014, includes the 2014 holidays, and excludes September 2, 2019, so the two samples differ slightly (58 and 33 days in the table). Without September 2, 2019 the index-only window is at 0.85.

further increment from an ordinary business day, on which the cash market is open and U.S. news arrives. Under the null $\beta_I = 0$: the jump carries the same variance per hour as live trading would. The sample is every scheduled holiday session with a same-day reopen (Monday through Thursday holidays; one-off closures such as Hurricane Sandy are excluded), which gives 58 dead afternoons and 33 index-only mornings, against 4,188 ordinary days. We exclude Labor Day 2019, whose morning carried the September 1, 2019 tariff announcement and on its own accounts for half of the index-only cell's variance; with it in the estimate is larger, and the appendix reports it. Standard errors come from a bootstrap over days.

Results. Table 5 presents the results. Turning the index on multiplies variance per hour by 13 ($t = 7.2$), against a null of one: on the same holiday, an hour in which the index can be traded has thirteen times the variance of an hour in which nothing trades. With day-of-week fixed effects the ratio is 12.6. Column (3) makes the comparison within each holiday, using a panel of 15-minute buckets with time-of-day and day fixed effects, so each holiday's index-only morning is set against its own dead afternoon and the day's level of news drops out; that gives $\times 11.3$ ($t = 3.3$). Column (4) is the 2003 schedule change described next. Across the columns the index margin runs from $\times 2.4$ to $\times 13$, and in each the null that the index adds nothing is rejected.

We do not report β_B as a clean effect, because an ordinary weekday differs from a holiday by a business day's worth of U.S. news as well as by the cash market being open, and this design cannot separate the two. The coefficient we are interested in is β_I , whose two cells are both holidays, so that the absence of U.S. news and of the cash market is common to both and only the futures market's state differs.

The 2003 schedule swap. On Monday, September 29, 2003, CME moved the E-mini's nightly one-hour shutdown from 22:00–23:00 to 16:30–17:00 Central. Two clock windows changed state on the same date, in opposite directions: the late-evening hour went from shut to open, and the half hour beginning at 16:30 went from open

Table 5: Index Trading Against a Shut Market

Notes: Each cell is one plus the estimated coefficient on the index-tradeable indicator divided by the dead cell's mean of the dependent variable, i.e. the ratio of variance per hour when only the index can be traded to variance per hour when nothing trades; a ratio of one means the index adds nothing. The t -statistic in parentheses is for the coefficient itself, against zero, so it tests that null. Variance per hour is from E-mini S&P 500 five-minute bars, in levels; a shut window's variance is one squared increment, from the halt price to the price five minutes after the reopen, spread over the shut minutes plus five. Columns (1)–(2): the 10:30–12:00 CT window on every weekday 2004–2019, which is dead on holidays before the 2014 halt change (58 afternoons), index-only on holidays after it (33 mornings), and fully open on ordinary days (4,188); the same day's overnight rate (20:00–07:00 CT) is subtracted; an ordinary-business-day indicator is included, so the reported coefficient compares holidays to holidays. Column (3): every 15-minute bucket of every scheduled holiday session, with time-of-day and day fixed effects, so the comparison is the index-only morning against the same day's dead afternoon. Column (4): the two windows the September 29, 2003 schedule change moved (16:30–17:00, open then shut; 22:00–23:00, shut then open) against the same day's 18:00–20:00, with day and window fixed effects, Monday–Thursday, July 2003 to March 2004. Holiday cells use scheduled sessions with a same-day reopen (Monday–Thursday; unscheduled closures excluded); September 2, 2019 is excluded from every column and Table 18 reports the estimate with it. t -statistics from a day bootstrap (2,000 draws) in columns (1), (2), and (4); clustered by day in column (3), where the bucket count overstates the effective sample. Appendix Table 17 reports the same columns under log, Poisson, and free-control specifications, together with the same bucket panel run on all weekdays, and Table 18 reports sample cuts.

	Matched window, day level one window per day, holidays		15-min bucket panel holidays	2003 swap day $\times$ window
	(1)	(2)	(3)	(4)
Index tradeable vs. nothing trading	$\times 13.1$ (7.2)	$\times 12.6$ (5.7)	$\times 11.3$ (3.3)	$\times 2.4$ (2.9)
Same-day overnight subtracted	Y	Y		
Day-of-week FE		Y		
Time-of-day FE			Y	
Day FE			Y	Y
Window FE				Y
Observations	4,279	4,279	8,736	372

to shut.¹³ That structure rules out the usual worry about hours changes, that the exchange expanded hours because demand was growing. Figure 10 plots the evening profile before and after: prices stop moving at 16:30 and start moving at 22:00. Column (4) of Table 5 compares the two switched windows with the same day's 18:00–20:00, with day and window fixed effects, Monday through Thursday, July 2003 to

¹³A third window also changed: the afternoon reopen moved from 15:45 to 15:30, so 15:30–15:45 went from shut to open. That window falls in the slot in which companies release earnings after the NYSE close, so it measures when earnings get priced rather than whether trading generates variance, and we exclude it.

March 2004. Open versus shut is $\times 2.4$ ($t = 2.9$): the same clock window has 2.4 times the variance per hour when the E-mini trades as when it is shut. The hour that opened rises by a factor of 4.5 ($t = 4.0$); the half hour that shut falls to two thirds of its former rate ($t = 0.9$).




Figure 10: The 2003 Schedule Swap: Variance by Time of Evening Before and After CME Moved the Nightly Shutdown.

Variance per hour of the E-mini S&P 500 futures price by 15-minute block, 16:00–24:00 Central, log scale, each era divided by its own 02:00–08:00 rate, averaged over Tuesday–Thursday evenings before (July 1 to September 26, 2003, 32 days) and after (September 29, 2003 to March 26, 2004, 65 days) the September 29, 2003 schedule change, which moved the nightly one-hour shutdown from 22:00–23:00 to 16:30–17:00. A shut window’s variance is one increment from the halt price to the price five minutes after the reopen, drawn dotted. The August 14, 2003 blackout evening and the quarterly roll weeks are excluded.

The 2003 estimate is smaller than the holiday estimates because the windows are short and fall in the least active hours of the day, and therefore the shut side is bounded below by price discreteness. The 2003 swap establishes *when* prices move: at the hour the exchange is open and not at the hour it is shut, in both directions on the same date.

Objections. The dead-zone comparison invites four objections, and each can be answered with the data in hand.

Since news is priced at the reopen, one concern is that dead afternoon variance could be undercounted. Our construction already charges the reopen's first five minutes to the closure in full, and that is more than the catch-up warrants: charging only the reopen's excess variance over an ordinary day's reopen, all the way through 19:00, leaves the holiday morning at more than twenty times the dead afternoon's rate.

10:30 Chicago is the European close; the pre-2014 afternoon is quiet because Europe went home. The 2014 move answers this: after 2014 the halt comes at noon, ninety minutes after Europe closes, so the two events separate on the same holiday. Holiday variance runs at 0.009 percent squared per hour with Europe open, *rises* to 0.021 after Europe closes at 10:30, and drops by an order of magnitude at the noon halt. Variance declines at the halt, not at the European close. Daylight saving time gives a direct estimate of Europe's share: the U.S. and E.U. switch on different Sundays, so for a few weeks a year Europe closes at 11:30 Chicago while CME's clock stays put. Switching Europe's state is worth 0.48 log points of variance per hour against a pre-2014 holiday drop of 3.14, so Europe explains at most about 15% of it. Tokyo's close explains nothing, and foreign news reaching the tape while U.S. stocks are shut is in any case the mechanism, not an alternative to it: an index-only market prices information its components cannot trade on.

Holidays are quiet, so of course the afternoon is quiet. Then the morning would be quiet too, and it is not. In matched index-only hours, midnight to 8am Chicago on both holidays and ordinary days, holidays run at about 60% of a normal day's rate: quieter, not dead. And the within-day normaliser in (7) removes the level of holiday quiet entirely.

Variance measured from transaction prices includes bid-ask bounce, and bounce requires trades. Fong and Frino (2001) raise exactly this for the drop in index-futures volatility when the cash market closes. The bounce variance implied by the first-order autocovariance of five-minute returns is about 11% of the measured live rate, which cannot produce a ratio of 13.

Relation to prior estimates. The total open-versus-shut ratio is old. French and Roll (1986) report per-hour variance during trading hours at 13 times its value across mid-week holidays and 72 times across weekends; Stoll and Whaley (1990) report about 16 for daytime versus overnight on the NYSE. Our all-weekday appendix estimate is that number on the index (Table 17). Chang et al. (1995) show on S&P 500 futures that volatility falls after the NYSE close and name the behavior of futures prices while the underlying market is shut as an open question. The new object is the index-only cell against a dead baseline. The closest magnitude comes from a cash market: Ito et al. (1998) find that yen variance over the Tokyo lunch hour doubled when lunchtime trading was allowed in 1994, with no change in public news. Our estimates for switching the index on are five to seven times that. The leading theory of periodic closure, Hong and Wang (2000), delivers a trading-to-nontrading ratio of 1.0 to 1.12 with constant information flow: a floor far below our estimates rather than an explanation of them.

3.1.5 Sizing Up the Experiments

Each experiment identifies a local effect: an elasticity to trading hours, a step at an access change, a within-day contrast. While none of them is a decomposition of the seventy-year rise, the first two can be chained. Figure 11 takes the estimated coefficients at face value, applies them to the actual history of trading hours and access, and asks how much of the rise in volatility they would have produced on their own.

The chain runs as follows. The NYSE experiments say that weekly volatility moves with weekly trading hours with an elasticity of about 1.4 (half the variance elasticity of 2.7 in Table 3, Panel B). Weekly hours went from 27 to 32.5 across the schedule changes of 1952–1985, which takes a 10% baseline to about 13% by 1985. The futures experiments say that each overnight hour of index trading raises overnight market variance by 3% and that the E-mini raised it by a further 50%; the overnight cell was about a quarter of daily market variance in the mid-1990s. Applying the 1993 hours and the 1997 access step to that cell, and holding intraday variance fixed, multiplies daily market variance by about 1.4 and takes 13% to about 15%.




Figure 11: Volatility Implied by the Experiments, Against the Actual Series.

The black line is the 5-year rolling annualized volatility of the daily market return, as in Figure 1. The red line starts at 10% in 1950 and applies two sets of estimates to the actual history of trading hours and access. NYSE hours: the elasticity of weekly volatility to weekly trading hours from Table 3, Panel B (1.37, half the variance elasticity), applied to the schedule changes of 1952–1985, over which weekly hours went from 27 to 32.5. Index futures: the per-hour effect on overnight market variance from Table 4, column (3) (0.030 per hour), applied to the record of overnight futures hours in Table 14, plus the E-mini access effect (0.43) from September 1997, both applied to the overnight cell only, which is 25% of daily market variance measured over 1993–97 (after the 1993 change; the pre-1993 window is too short to use). Intraday variance is held fixed. The implied series is shown as a 5-year moving average to match the actual.

Actual five-year rolling volatility went from about 10% in the 1950s to about 18% in the most recent window, and the implied path covers 5.2 of those 7.7 percentage points: the experiments account for over half of the rise.¹⁴ The identified channels are the right order of magnitude for the headline fact, and their timing matches the actual series.

¹⁴Against the fitted trend of Table 1, which ends near 20% because the five-year window ending in 2025 is below the trend line, the share is about half; in variance rather than volatility units it is about 60%. The 10% starting point is the level of the actual series in the mid-1950s.

4. A model of index trading and market returns

The evidence above points to two related changes in how the market is traded. Investors increasingly trade market-wide positions rather than individual stocks, and index products increasingly allow those positions to be traded throughout the day. In this section, we present a simple model with both margins. Following Merton (1987), we consider an economy where investors initially trade only a subset of available financial assets. When an index-trading vehicle is introduced, investors can express the market-level component of their demand directly through the diversified market claim. As in Campbell, Grossman, and Wang (1993) (henceforth, CGW), risk-averse arbitrageurs absorb the resulting order flow.

The second margin is trade time. Some investor demand evolves with trading opportunities rather than calendar time, so expanding the number of trading opportunities within a day can increase the amount of demand pressure that reaches prices.

The model also accounts for the change in short-run return dynamics documented above: early in the sample, market returns exhibit continuation, while modern returns exhibit reversal. We capture the early continuation by allowing common fundamental news to be incorporated gradually across stocks, while transitory demand pressure generates reversal. Index trading affects both forces: it amplifies the market-level effect of demand pressure and, in the benchmark, speeds the incorporation of common information, while more trading opportunities increase the amount of transitory demand pressure generated within a day.

4.1 Asset market without index trading

We first consider an economy without an index-trading vehicle. There are N risky assets with an $N \times 1$ vector of dividends $\mathbf{d}_t$, which follow random walks,

$$\mathbf{d}_{t+1} = \mathbf{d}_t + \boldsymbol{\varepsilon}_{t+1}^F, \quad \boldsymbol{\varepsilon}_{t+1}^F \sim N(\mathbf{0}, \sigma_F^2 C_F), \quad (8)$$

where σ_F^2 is the variance of each stock's dividend innovation and C_F is the correlation matrix with constant pairwise correlation ρ_F . Fundamental news is permanent and evolves in calendar time, independently of the number of trading opportunities within a day.

There is a unit mass of behavioral investors and a competitive, fully informed arbitrageur who can trade all assets. All agents have CARA utility with absolute risk aversion a and choose myopic, one-period portfolios. Given Gaussian payoffs, their demands take the mean-variance form. The risk-free asset has gross daily return $R_f > 1$ and is supplied elastically. We set supply $x_i = 1/N$ for each risky asset i . The arbitrageur has rational expectations, denoted by $E_t^A(\cdot)$. Since dividend levels follow random walks, $E_t^A[d_{i,t+1}] = d_{i,t}$ for all i .

Behavioral investors are evenly distributed across N stock-specialist groups, with each investor trading only one stock in addition to the risk-free asset. Their expected returns differ from those of the fully informed arbitrageur in two ways: they incorporate common fundamental news gradually, and their expectations are affected by transitory nonfundamental beliefs (animal spirits). The latter depend on the number of trading opportunities within a day, as specified below.

Let $E_t^B(\cdot)$ denote the average subjective expectation within the specialist group for stock i , and let $\hat{d}_{i,t} \equiv E_t^B[d_{i,t+1}]$ denote its average dividend forecast. Decompose dividends into common and stock-specific components, $d_{i,t} = d_{m,t} + d_{i,t}^\perp$, where $d_{m,t} \equiv N^{-1} \sum_j d_{j,t}$.

We abstract from delays in the incorporation of stock-specific fundamental news, while allowing gradual updating of the common component. The average stock-specialist forecast is

$$\hat{d}_{i,t} = d_{i,t}^\perp + \hat{d}_{m,t}^{\text{stk}}, \quad \hat{d}_{m,t}^{\text{stk}} = (1 - \alpha_{\text{stk}}) \hat{d}_{m,t-1}^{\text{stk}} + \alpha_{\text{stk}} d_{m,t}, \quad (9)$$

where $\alpha_{\text{stk}} \in (0, 1]$ governs the speed at which common fundamental news enters stock investors' expectations. When $\alpha_{\text{stk}} = 1$, their average dividend forecast equals the rational forecast. When $\alpha_{\text{stk}} < 1$, new permanent common dividend news is in-

corporated gradually. Appendix B.4 rationalizes this updating rule through a costly monitoring decision; here we take the updating rate as given.

Define the difference between stock specialists' average dividend forecast and the rational forecast as $\delta_{i,t} \equiv E_t^B[d_{i,t+1}] - E_t^A[d_{i,t+1}]$. Under our information assumptions, $\delta_{i,t} = \hat{d}_{i,t} - d_{i,t} = \hat{d}_{m,t}^{\text{stk}} - d_{m,t}$. Thus, the dividend forecast gap is identical across stocks: it reflects delayed incorporation of common fundamental news.

The second belief distortion captures transitory, nonfundamental shifts in expected returns (animal spirits). Let $z_{i,t}$ denote this distortion for investors specialized in stock i , and let $S \geq 1$ denote the number of trading opportunities within a day. At the daily frequency,

$$z_{i,t+1} = \lambda z_{i,t} + u_{i,t+1}^z, \quad 0 \leq \lambda < R_f^{-1}, \quad (10)$$

where the innovations are jointly Gaussian, serially independent, and independent of fundamental news, with

$$\begin{aligned} \text{Var}(u_{i,t+1}^z \mid S) &= \sigma_\eta^2 S^{1-\xi}, \\ \text{Corr}(u_{i,t+1}^z, u_{j,t+1}^z) &= \rho_z, \quad i \neq j, \end{aligned} \quad (11)$$

where $\xi \in [0, 1]$.

The parameter λ governs the persistence of demand pressure in calendar time, while ξ governs how new demand-pressure risk scales with trading opportunities. At $\xi = 0$, the daily innovation variance increases proportionally with S ; at $\xi = 1$, it is independent of S . Thus, lower ξ implies a stronger trade-time shock-arrival channel. The Appendix constructs an intraday market with S explicit trading opportunities and shows that sampling it at daily boundaries reproduces equations (10)–(11) exactly.

Let $P_{i,t}$ denote the end-of-day price of stock i . Its one-day excess dollar payoff, which we refer to as its return for brevity, is

$$\Pi_{i,t+1}^e = P_{i,t+1} + d_{i,t+1} - R_f P_{i,t}.$$

The average expected return among investors specialized in stock i differs from the arbitrageur's rational expectation through the two belief distortions introduced

above:

$$E_t^B(\Pi_{i,t+1}^e) = E_t^A(\Pi_{i,t+1}^e) + g_{i,t} + z_{i,t}, \quad (12)$$

where $g_{i,t}$ reflects stale fundamental information and $z_{i,t}$ captures transitory animal spirits. Because dividend news is permanent, a forecast error affects both the next dividend and its continuation value. Consequently, $g_{i,t} = \frac{R_f}{R_f - 1} \delta_{i,t}$, as derived in the Appendix. Since $\delta_{i,t}$ is common across stocks, so is $g_{i,t}$.

Define the market-level belief distortions as

$$g_{m,t}^{\text{stk}} \equiv \frac{1}{N} \sum_{i=1}^N g_{i,t}, \quad z_{m,t} \equiv \frac{1}{N} \sum_{i=1}^N z_{i,t}.$$

By symmetry, these aggregate belief distortions are the only behavioral states affecting the market price, in addition to the fundamental dividend level.

Proposition 1. *Let*

$$P_{m,t} \equiv \sum_{i=1}^N x_i P_{i,t} = \frac{1}{N} \sum_{i=1}^N P_{i,t}$$

denote the price of the market portfolio. The equilibrium market price in the economy without index trading is

$$P_{m,t} = p_S + \frac{d_{m,t}}{R_f - 1} + b_{m,S}^F g_{m,t}^{\text{stk}} + b_{m,S} z_{m,t}, \quad (13)$$

where

$$b_{m,S}^F = \frac{\psi_S}{1 + \psi_S} \frac{1}{R_f - (1 - \alpha_{\text{stk}})}, \quad b_{m,S} = \frac{\psi_S}{1 + \psi_S} \frac{1}{R_f - \lambda'}, \quad (14)$$

and $\psi_S \equiv v_{m,S}/v_S < 1$. Here p_S is a constant reflecting the equilibrium risk-premium discount, v_S is the common conditional one-day return variance of an individual stock, and $v_{m,S}$ is the corresponding market return variance. These variances are determined jointly with the equilibrium price coefficients.

The term ψ_S captures the effect of investor underdiversification. Behavioral investors trade individual stocks and therefore bear idiosyncratic risk, making their demand less sensitive to a given market-level belief distortion than that of the diversified arbitrageur. Since their respective demand sensitivities are proportional to $1/v_S$ and $1/v_{m,S}$, market clearing gives behavioral beliefs a weight of $\psi_S/(1 + \psi_S)$. Thus, un-

derdiversification attenuates the impact of both fundamental belief distortions and transitory demand pressure on the market price.

The two belief distortions have different implications for return dynamics. Define the market-level innovations in demand pressure and fundamentals as

$$u_{m,t+1}^z \equiv \frac{1}{N} \sum_{i=1}^N u_{i,t+1}^z, \quad \varepsilon_{m,t+1}^F \equiv \frac{1}{N} \sum_{i=1}^N \varepsilon_{i,t+1}^F.$$

Since demand pressure enters the market price with loading $b_{m,S}$, its contribution to the one-day excess return is

$$b_{m,S}(z_{m,t+1} - R_f z_{m,t}) = b_{m,S} u_{m,t+1}^z - (R_f - \lambda) b_{m,S} z_{m,t}.$$

New demand shocks move returns on impact, while existing demand pressure subsequently reverses. In particular, a positive $z_{m,t}$ raises the current price but predicts a lower subsequent expected excess return.

Gradual incorporation of common dividend news produces the opposite return dynamics. Since the fundamental belief distortion enters the market price with loading $b_{m,S}^F$, its contribution to the one-day excess return is

$$\begin{aligned} & b_{m,S}^F (g_{m,t+1}^{\text{stk}} - R_f g_{m,t}^{\text{stk}}) \\ &= -\frac{R_f}{R_f - 1} (1 - \alpha_{\text{stk}}) b_{m,S}^F \varepsilon_{m,t+1}^F - [R_f - (1 - \alpha_{\text{stk}})] b_{m,S}^F g_{m,t}^{\text{stk}}, \end{aligned}$$

where we use the updating rule for common fundamental news.

A positive permanent dividend innovation raises fundamental value, but also creates a negative belief distortion $g_{m,t+1}^{\text{stk}}$, so the market price initially rises by less than its full-information response. As investors gradually incorporate the news, the belief distortion disappears and prices continue to adjust upward. Thus, gradual incorporation of permanent fundamental news generates return continuation, whereas transitory demand pressure generates reversal.

To connect these return dynamics to the empirical evidence, let $V_h(S)$ denote the variance of cumulative market returns over h daily periods, with S trading opportu-

nities per day. Define the variance ratio

$$VR_{1,h}(S) = \frac{hV_1(S)}{V_h(S)}. \quad (15)$$

The ratio equals one for serially uncorrelated returns, falls below one with return continuation, and rises above one with return reversal; as in Section 2, we report its log, which is zero, negative, or positive in the three cases. Gradual incorporation of permanent fundamental news generates continuation, while transitory demand pressure generates reversal.

For $\xi < 1$, increasing S raises the variance of daily demand-pressure innovations while leaving the distribution of fundamental news unchanged. More trading opportunities therefore increase the relative importance of transitory demand pressure, pushing the variance ratio toward reversal. This trade-time channel weakens as ξ approaches one.

4.2 Asset market with index trading

We next introduce an index fund that allows investors to trade the diversified market portfolio directly. A fraction $\omega \in [0, 1]$ of behavioral investors, drawn symmetrically from the stock-specialist groups, switches from individual stocks to the index fund; the remaining fraction $1 - \omega$ continues to trade individual stocks. We refer to ω as the *index-participation share*.¹⁵ The index fund is a pass-through claim on the underlying stocks and creates no new supply of risky cash flows. Normalizing one share of the fund to hold one unit of the market portfolio, its price and dividend are $P_{m,t}$ and $d_{m,t}$, and its one-day excess payoff is $\Pi_{m,t+1}^e = P_{m,t+1} + d_{m,t+1} - R_f P_{m,t}$.

Write each stock's transitory belief distortion as $z_{i,t} = z_{m,t} + z_{i,t}^\perp$, where $z_{i,t}^\perp \equiv z_{i,t} - z_{m,t}$ is the stock-specific component. These stock-specific components cancel in the diversified market portfolio. Investors who switch to the index fund act on the common distortion $z_{m,t}$. Index trading thus allows the same market-level demand signal to be expressed directly through a diversified claim rather than indirectly through

¹⁵We treat ω as exogenous; limited participation can reflect direct or nonpecuniary costs of switching.

individual stocks.

Let $E_t^{B,\text{idx}}(\cdot)$ denote the average subjective expectation among index investors. As with stock specialists, their expected excess return differs from the arbitrageur's rational expectation through fundamental and nonfundamental belief distortions:

$$E_t^{B,\text{idx}}(\Pi_{m,t+1}^e) = E_t^A(\Pi_{m,t+1}^e) + g_{m,t}^{\text{idx}} + z_{m,t}, \quad (16)$$

where $g_{m,t}^{\text{idx}}$ captures delayed incorporation of common fundamental news, while $z_{m,t}$ is the same market-level demand distortion introduced above.

Let $\hat{d}_{m,t}^{\text{idx}} \equiv E_t^{B,\text{idx}}[d_{m,t+1}]$ denote the average dividend forecast among index investors. Since market dividends follow a random walk, the arbitrageur's rational forecast is $E_t^A[d_{m,t+1}] = d_{m,t}$. Index investors incorporate common fundamental news at rate α_{idx} :

$$\hat{d}_{m,t}^{\text{idx}} = (1 - \alpha_{\text{idx}})\hat{d}_{m,t-1}^{\text{idx}} + \alpha_{\text{idx}}d_{m,t}, \quad (17)$$

where $\alpha_{\text{idx}} \in (0, 1]$. The resulting dividend forecast gap is $\delta_{m,t}^{\text{idx}} \equiv \hat{d}_{m,t}^{\text{idx}} - d_{m,t}$. As with stock specialists, this gap generates an expected-return distortion $g_{m,t}^{\text{idx}} = \frac{R_f}{R_f - 1} \delta_{m,t}^{\text{idx}}$.

We allow $\alpha_{\text{idx}} > \alpha_{\text{stk}}$. Both groups process the same common dividend news, but index investors can act on it through a diversified claim without bearing idiosyncratic stock risk. This makes common information more valuable to monitor through the index. Appendix B.4 shows how this diversification advantage can rationalize the faster updating rate of index investors.

For compactness, define $\Delta_S(\omega) \equiv 1 + \omega + (1 - \omega)\psi_S$. The three terms reflect the market-level demand sensitivities of the arbitrageur, index investors, and stock specialists, respectively. The equilibrium return variances v_S and $v_{m,S}$, and hence ψ_S , depend on ω ; we suppress this dependence in the notation.

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

$$P_{m,t} = p_S + \frac{d_{m,t}}{R_f - 1} + b_{m,S}^{F,\text{stk}}(\omega)g_{m,t}^{\text{stk}} + b_{m,S}^{F,\text{idx}}(\omega)g_{m,t}^{\text{idx}} + b_{m,S}(\omega)z_{m,t}, \quad (18)$$

where

$$\begin{aligned}
b_{m,S}^{F,\text{stk}}(\omega) &= \frac{(1-\omega)\psi_S}{\Delta_S(\omega)} \frac{1}{R_f - (1-\alpha_{\text{stk}})}, \\
b_{m,S}^{F,\text{idx}}(\omega) &= \frac{\omega}{\Delta_S(\omega)} \frac{1}{R_f - (1-\alpha_{\text{idx}})}, \\
b_{m,S}(\omega) &= \frac{\omega + (1-\omega)\psi_S}{\Delta_S(\omega)} \frac{1}{R_f - \lambda}.
\end{aligned} \tag{19}$$

The no-index economy is nested as the special case $\omega = 0$. The equilibrium constant p_S can also depend on ω ; we suppress this dependence in the notation.

The index mechanism is easiest to see in the loading on transitory demand pressure, $b_{m,S}(\omega)$. When behavioral investors trade individual stocks, they must bear idiosyncratic risk to act on market-level beliefs, limiting the size of their positions and attenuating their price impact. Trading the diversified index fund removes this idiosyncratic risk, allowing investors to take larger positions in response to the same market-level demand distortion. As ω rises, the impact of this distortion on market prices therefore increases.

In the limiting case in which all behavioral investors trade the index fund, their demand sensitivity equals that of the diversified arbitrageur, giving

$$b_{m,S}(1) = \frac{1}{2} \frac{1}{R_f - \lambda}, \quad \frac{b_{m,S}(1)}{b_{m,S}(0)} = \frac{1 + \psi_S(0)}{2\psi_S(0)} > 1, \tag{20}$$

where $\psi_S(0)$ denotes the no-index value of ψ_S . Thus, index trading amplifies the price impact of existing market-level demand pressure without increasing the underlying demand shock.

In our benchmark, index trading also accelerates the incorporation of common fundamental news into market prices. Because index investors update more quickly than stock specialists ($\alpha_{\text{idx}} > \alpha_{\text{stk}}$), higher ω shifts price determination toward their more current beliefs. This reduces the delayed response of prices to permanent news.

The two changes in market structure reinforce one another in generating transitory market volatility. A higher ω increases the price impact of a given market-level demand distortion, while, for $\xi < 1$, a higher S increases the variance of new demand-pressure shocks generated within a day. Additional trading opportunities therefore

have a larger effect on market volatility when index participation is high.

4.3 Stock-level implications

The model also has implications for stock-level return comovement and trading activity. We measure market comovement by the R^2 from regressing individual stock returns on the market return. We measure intermediated dollar trading at daily boundaries by the change in the arbitrageur's stock position,

$$V_{i,t}^{\text{model}} = P_{i,t} \left| n_{i,t}^A - n_{i,t-1}^A \right|. \quad (21)$$

By market clearing, $\Delta n_{i,t}^A = -\Delta n_{i,t}^{B,\text{net}}$. As ω increases, more behavioral demand is expressed through the diversified claim, increasing the common component of stock-level order flow.

Together with the market-level results above, the model delivers three empirical predictions. First, more trading opportunities raise short-horizon volatility and push variance ratios toward reversal, as examined in the trading-hours IVs, the E-mini and futures-hours experiments, and the within-day holiday design. Second, expanded index trading opportunities raise systematic variance more than idiosyncratic variance, since investors can trade the market portfolio directly. Third, the volatility effect of additional trading opportunities is larger when index participation is high. The Appendix examines a further implication: when index products were concentrated in the S&P 500, index membership should have accelerated the incorporation of common information.

4.4 Benchmark parameterization and comparative statics

We next ask whether plausible changes in index participation and trading opportunities can reproduce the historical rise in market volatility and the shift from return continuation to reversal. We use empirical evidence to guide parameter choices, calibrate the scale of demand shocks to early-sample volatility, and choose stylized his-

torical paths for ω and S . The exercise assesses the quantitative plausibility of the mechanisms rather than attempting to identify all model parameters.

Benchmark parameters. We set $N = 500$, interpreting the index as a stylized S&P 500. The pairwise correlations of fundamental innovations and demand-pressure innovations are $\rho_F = \rho_z = 0.25$, and the annual gross risk-free return is $R_f^{\text{ann}} = 1.02$. We set the annualized volatility of market-level fundamental news to 12.8%, corresponding to approximately 25% for individual stocks given $N = 500$ and $\rho_F = 0.25$.

The scale of transitory demand-pressure innovations, σ_η , is calibrated to target early-sample daily market volatility and held fixed thereafter. Thus, the model's subsequent rise in volatility is driven by changes in index participation and trading opportunities, rather than changes in shock magnitudes.

For fundamental information, we set $\alpha_{\text{stk}} = .30$, implying a roughly two-day half-life for stock specialists' forecast errors, and $\alpha_{\text{idx}} = 1$, so index investors incorporate common fundamental news without delay at the daily frequency. Appendix B.4 rationalizes these updating rates through a monitoring problem. At the benchmark $\psi_S \simeq 0.21$, the same common information is about five times as valuable when acted on through the diversified index. Appendix C provides complementary evidence of delayed common-information propagation across stocks and of faster information incorporation associated with S&P 500 membership when index products were concentrated in that index.

Demand-pressure persistence. The persistence of demand pressure determines how quickly its price impact reverses. We set $\lambda = .90$, corresponding to a half-life of approximately 6.6 trading days. Demand-driven price pressure therefore dissipates over days and weeks, consistent with the short-horizon return reversal observed in the modern data.

Trade-time scaling and market structure. The parameter ζ determines how strongly demand-pressure risk scales with trading opportunities. Motivated by our trading-hours evidence, we set $\zeta = .1$, so the variance of new demand shocks increases nearly




Figure 12: Market structure and transitory demand. Panel A plots the annualized contribution of the transitory z component to market volatility as the index-participation share ω varies, for $S = 1, 2, 4$ trading opportunities. Panel B plots the corresponding 21-trading-day log variance ratio. All other parameters are held at their benchmark values.

proportionally with the number of trading opportunities.

Trading opportunities have expanded from the regular cash-market day to nearly around-the-clock index trading. We normalize the early trading environment to $S = 1$ and use the stylized historical paths

$$S = (1, 1, 2, 4), \quad \omega = (.35, .70, .72, .90),$$

for 1950–70, 1971–90, 1991–2004, and 2005–24, respectively. The path for ω reflects the broad rise in market-directed trading documented in the data, rather than estimates of its exact level in each era.

Quantitative implications. Figure 12 illustrates the two mechanisms. Greater index participation increases the price impact of market-level demand pressure, while more trading opportunities increase the variance of new demand shocks. Both raise the transitory component of daily market volatility (Panel A) and push variance ratios toward reversal (Panel B). The two mechanisms reinforce one another: additional trading opportunities have a larger effect when index participation is high.




Figure 13: Historical volatility and return dynamics. Panel A compares broad-market daily volatility with the model; error bars show two standard errors from a year-block bootstrap. Panel B compares the model with the S&P 500 and Dow using the log of a horizon-matched variance-ratio estimator. The shaded region spans the two index benchmarks; the horizontal line marks zero. The S&P 500 first point is 1957–70; the final era is 2005–24.

Figure 13 compares the model with the historical data, using the S&P 500 as the primary variance-ratio benchmark and the Dow as a second large-stock index benchmark.

The parameterization captures both historical patterns that motivate the model. Daily market volatility rises from about 10% to about 20%, while return dynamics move from continuation to roughly unpredictable returns and then to reversal. Between the first two eras, S remains unchanged, and greater index participation moves the variance ratio toward one. The subsequent expansion of trading opportunities raises volatility and strengthens reversal. The model’s stock-level R^2 , the systematic share of a stock’s return variance, roughly doubles from the early to the modern economy, from 0.21 to 0.44, while in the data the systematic share of average-stock variance in Section 2.3.4 rises somewhat less than twofold, from 0.18 to 0.31. The model slightly overshoots, in part because idiosyncratic variance also increased somewhat in the data over this period.

Transitory demand pressure rises from 23% of daily market variance in the early

economy to 61% in the modern economy and accounts for 82% of the modeled increase in variance. Permanent fundamental news remains important for the level of market risk but does not drive the secular increase. Together, greater index participation and expanded trading opportunities reproduce the broad historical evolution of market volatility and return dynamics without requiring an increase in the volatility of fundamental news.

5. Conclusion

Stock market volatility has roughly doubled over the past 70 years. The increase is largest at short horizons and reflects a change in the autocorrelation of returns, from continuation to reversal. Over the same period the stock market transformed, with the proliferation of index products (index futures, index ETFs, and passive mutual fund products) and with substantially expanded trading hours for these products. We provide causal evidence that trading in these products has contributed substantially to the increase in volatility: wherever the opportunity to trade the index expands, trading and volatility rise with it, and together the experiments account for over half of the rise. A model in which arbitrageurs absorb investors' demand shocks reproduces both the rise in volatility and the shift from continuation to reversal.

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 Bakkum, and Zhi Da, 2019, Indexing and stock market serial dependence around the world, *Journal of Financial Economics* 132, 26–48.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- 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 shocks as second-moment news shocks, *Review of Economic Studies* 87, 40–76.
- 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.
- Brightman, Chris, and Campbell R Harvey, 2025, Passive aggressive: The risks of passive investing dominance, *Available at SSRN* 5259427 .
- Campbell, John Y., 2018, *Financial Decisions and Markets: A Course in Asset Pricing* (Princeton University Press, Princeton).
- 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.
- Chang, Eric C., Prem C. Jain, and Peter R. Locke, 1995, S&P 500 index futures volatility and price changes around the New York Stock Exchange close, *Journal of Business* 68, 61–84.

- Chernov, Mikhail, Lars A. Lochstoer, and Dongho Song, 2026, The real channel for nominal bond-stock puzzles, *Journal of Finance* forthcoming.
- Chiah, Mardy, Philip Gharghori, and Angel Zhong, 2023, Has idiosyncratic volatility increased? Not in recent times, *Critical Finance Review* 12, 125–170.
- Chinco, Alex, and Marco Sammon, 2024, The passive ownership share is double what you think it is, *Journal of Financial Economics* 157, 103860.
- 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, 2017, How risky is consumption in the long-run? Benchmark estimates from a robust estimator, *Review of Financial Studies* 30, 631–666.
- Dew-Becker, Ian, and Stefano Giglio, 2026, The decline of the S&P 500 variance risk premium, Conditionally accepted, *Critical Finance Review*.
- Fong, Kingsley, and Alex Frino, 2001, Stock market closure and intraday stock index futures market volatility: “contagion”, bid–ask bias or both?, *Pacific-Basin Finance Journal* 9, 219–232.
- 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.
- Froot, Kenneth A., and André F. Perold, 1995, New trading practices and short-run market efficiency, *Journal of Futures Markets* 15, 731–765.
- Greenwood, Robin, and David Thesmar, 2011, Stock price fragility, *Journal of Financial Economics* 102, 471–490.
- Haddad, Valentin, Paul Huebner, and Erik Loualiche, 2025, How competitive is the stock market? Theory, evidence from portfolios, and implications for the rise of passive investing, *American Economic Review* 115, 975–1018.
- Haddad, Valentin, and Tyler Muir, 2025, Market macrostructure: Institutions and asset prices, *Annual Review of Financial Economics* 17, 133–150.
- Harvey, Campbell R, Michele G Mazzoleni, and Alessandro Melone, 2025, The unintended consequences of rebalancing, Working paper 33554, National Bureau of Economic Research.

- Hasbrouck, Joel, 2003, Intraday price formation in U.S. equity index markets, *Journal of Finance* 58, 2375–2400.
- Hong, Harrison, and Jeremy C Stein, 2007, Disagreement and the stock market, *Journal of Economic Perspectives* 21, 109–128.
- Hong, Harrison, and Jiang Wang, 2000, Trading and returns under periodic market closures, *Journal of Finance* 55, 297–354.
- 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.
- Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin, 1998, Is there private information in the FX market? the Tokyo experiment, *Journal of Finance* 53, 1111–1130.
- 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.
- Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies* 1, 41–66.
- Lo, Andrew W., and A. Craig MacKinlay, 1990, When are contrarian profits due to stock market overreaction?, *Review of Financial Studies* 3, 175–205.
- 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.
- Stoll, Hans R., and Robert E. Whaley, 1990, Stock market structure and volatility, *Review of Financial Studies* 3, 37–71.
- 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* 66, 123–148.

A. Appendix

A.1 Systematic vs. Idiosyncratic Volatility

This appendix gives the construction behind the decomposition in Section 2.3.4 and its validation against Campbell et al. (2001) (CLMX). 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. Industry returns value-weight CRSP firms within each of the 49 Fama–French industries, firm variances aggregate to the industry and market level by lagged market-cap weights, and 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; we report the sample from 1962 to line up with CLMX and extend it through 2024.

Table 6 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 is well inside their reported 90% confidence interval of $[0.55, 1.47]$.

Table 6: 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.

Figure 14 plots the three components over 1962–2024. Firm-specific volatility increased from 16% in the postwar quiet to 26% by 2000–2010, the increase CLMX documented, and then declined to 19% over 2011–2024, the reversal Campbell et al. (2023) report (see also Chiah et al. (2023)). The market component moved the other way, with its share of average-stock variance at a trough in the late 1990s and increasing since. Figure 15 decomposes the five-year log variance ratio by component. The market-component log ratio rises from -0.34 in 1961–1995 to 0.39 in 1996–2024, the industry component shifts the same way by a smaller amount, and the firm-specific ratio shows no regime change: its log has been positive in every five-year window

but one, the signature of bid-ask bounce, and declines over the full sample as spreads narrow. Equal-weighting the top 500 stocks gives the same market-component shift ($-0.49 \rightarrow 0.33$), and therefore the shift is not a mega-cap artifact.

60-month rolling std dev with linear trend: market, industry, firm-specific (VW)
 Literal CLMX construction, 1962–2024




Figure 14: 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 15: Log variance ratio by CLMX component, 1931–2024

Notes: Log of the 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.

A.2 Macroeconomic Fundamentals Do Not Explain the Rise

Figure 16 plots the ratio of annualized daily stock market volatility to annualized macro volatility over four non-overlapping 25-year periods from 1926 to 2025, for industrial production growth, real per capita GDP growth, and corporate profit growth. Industrial production is monthly from 1920, while the other two are quarterly after WWII, spliced with annual data back to 1929. Corporate profits are a levered cash flow, so they absorb changes in corporate leverage. For all three measures the ratio was lowest in the first period, increased slightly in 1951–1975, increased strongly in 1976–2000, and stayed high.




Figure 16: Equity Volatility Relative to Fundamental Volatility, by 25-Year Period

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 CRSP value-weighted 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: CRSP, FRED (INDPRO, real GDP per capita, real corporate profits).

Table 7 gives the industrial-production-volatility control behind Section 2.3. On the full 1931–2025 sample a linear trend in stock volatility is essentially zero, because Depression-era volatility offsets the modern rise, and adding five-year rolling industrial-production volatility as a control moves the trend to 10.5 percentage points on the normalized index, close to the 1955–2025 baseline (Panel A). IP volatility declined over the sample, and therefore partialling it out uncovers the rise it was masking. On the post-1960 subsample the trend is nearly identical with or without the control, 10.58 vs. 10.72 (Panel B). Allowing a slope break at 1960 (column 3b) puts the change in slope at 30 points on the normalized trend, about 3 pp/decade ($p = 0.012$).

Table 7: 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	17.99*** (4.18)	3.42 (3.16)	6.98* (3.76)
Trend $\tilde{t}$	−2.98 (6.32)	10.53** (4.48)	−24.28 (16.17)
σ_{IP}		1.60*** (0.34)	1.63*** (0.28)
Post-1960 $\times \tilde{t}$			30.14** (12.03)
R^2	0.016	0.478	0.573

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

	(4a) Trend only	(4b) + σ_{IP}
Constant	9.84*** (0.77)	8.34*** (1.19)
Trend (1960→end)	10.72*** (2.12)	10.58*** (2.22)
σ_{IP}		0.52* (0.27)
R^2	0.467	0.506

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 8 updates the Schwert (1989) regression of stock volatility on the volatilities of industrial-production growth, inflation, unemployment, and the short rate. On 1947–2024 the R^2 is 0.26, close to the 0.39 we recover on his original window, and roughly three-quarters of the variation in stock volatility is unexplained by fundamentals. Since the 1980s macroeconomic volatility has declined to historic lows while stock volatility has increased.

A.3 Implied Volatility and the Variance Risk Premium

The VIX shows little trend over its post-1990 sample because the variance risk premium has declined at almost exactly the rate realized volatility has increased. Table 9 splices the VXO (1986–1989) to the VIX (1990 onward) and measures realized

Table 8: Schwert (1989) regression, updated sample

Sample	N	σ_{IP}	σ_{CPI}	σ_{UNR}	σ_{RF}	R^2
Schwert original (1947–1987)	468	−0.29*	2.35***	0.00	0.13***	0.39
Updated full (1947–2024)	912	−0.92***	6.00**	0.03***	0.05	0.26
Pre-1985	432	−0.35***	2.53***	0.01	0.13***	0.48
Post-1985	480	1.58	6.62**	−0.01	0.11	0.39
Post-2000	300	1.31	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 $p < 0.10, 0.05, 0.01$.

volatility, implied volatility, and their difference as five-year rolling averages. Realized volatility trends up at 1.82 pp/decade, the variance risk premium declines at −2.00 pp/decade, and implied volatility is flat at −0.17 pp/decade, indistinguishable from zero. The declining premium is the central finding of Dew-Becker and Giglio (2026), who attribute it to easing intermediary frictions in the options market. Thus, the flat VIX reflects a shrinking premium rather than the absence of a rise in volatility.

Table 9: 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)	6.38 (3.35)	1.82	0.086	0.155
Implied vol (IV)	−0.61 (2.93)	−0.17	0.84	0.003
Variance risk premium (VRP)	−6.99 (0.85)	−2.00	9×10^{-6}	0.621

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. Figure 17 instead estimates Bai–Perron breaks in the mean of the daily five-year volatility series of Figure 1, allowing two breaks. The breaks are dated October 1973 and March 2000, with regime means of 10.3%, 13.8%, and 19.3%, and the increase in mean volatility over the sample is about the same as in the trend specification.

The variance ratio breaks at almost the same date. Its log was near −0.29 from 1931 through the mid-1990s – identical in the Depression and the quiet postwar period




Figure 17: Structural Breaks in 5-Year Rolling Market Volatility

Notes: Annualized realized volatility of the CRSP value-weighted 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 10.3%, 13.8%, and 19.3%. Shaded regions indicate NBER recessions. Sample: 1955–2025.

despite volatility levels that differed by a factor of two and a half – and then it rose by about 0.7. A break test on the five-year log variance ratio places a single break at August 1999, significant at the 10% level, with segment means of -0.29 before and 0.41 after (Figure 18). Two moments of returns shift within a year of each other, and this is 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 does not depend on the sample window or the rolling-window length. Figure 19 estimates the trend on $\tilde{r}$ separately for every start year from 1931 to 2010, with the sample ending in 2025 and 90% fixed- b confidence intervals. The trend is positive at every post-WWII start year, and the interval excludes zero for start years from about 1945 through 1993 (daily) and 1987 (weekly). Earlier start dates absorb Depression-era volatility, and start dates after the mid-1990s cut off the rise itself. Figure 20 varies the rolling-window length from 1 month to 10 years on the 1955–2025 sample. The trend increases with window length and levels off around 7 years, and even at 1-month windows it is 8.3 pp (daily) and 5.5 pp (weekly), with intervals excluding zero. Figure 21 estimates the trend on rolling 30-year windows. Every window ending after 1980 is positive, 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 estimate. Windows ending in the 1960s are negative because they still contain the Depression-to-postwar decline.




Figure 18: Structural break in the log variance ratio, 1931–2025

Notes: Log of the five-year rolling variance ratio (daily-implied over monthly-implied annualized variance) with a single Bai–Perron break, dated August 1999 (sup- F 8.24 against a 5% critical value of 8.58, so significant at 10%). The segment means are -0.29 before and 0.41 after, a shift that comes within a year of the March 2000 break in the level of volatility.

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

Table 10 reports the trend for alternative market constructions. Equal-weighting gives a steeper trend than value-weighting (2.16 vs. 1.59 pp/decade), and by size the top-decile trend (1.44) is below the full-market trend while the mid- and small-cap segments are steeper (2.18 and 2.35). If a few large stocks drove the result, the largest stocks would show the steepest trend; they show the shallowest. Measuring volatility from the S&P 500 index directly, the index traded through futures and ETFs, gives 1.54 pp/decade vs. 1.67 for the value-weighted market over the matched 1967–2024 sample, with a correlation of 0.99.

A.7 The Variance-Ratio Rise Is Global

We build daily value-weighted market returns for 35 countries from Compustat Global firm-day data and estimate the trend in each country’s five-year log variance ratio, with Newey–West HAC (60 monthly lags) and fixed- b critical values. Figure 22 plots the coefficients. The log variance ratio increases in all 35 countries, 30 of them significant at the 5% level, with a mean trend of 0.205 per decade, and every major developed market is significantly positive (United Kingdom 0.25, France 0.26, Germany 0.14, Japan 0.26, vs. 0.10 for the United States). Volatility levels are mixed across countries, since emerging markets enter the sample at their 1990s crisis highs and decline through it, but the variance ratio increases everywhere. A common, daily-frequency




Figure 19: 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 $[\text{start year}, 2025]$, where σ_t is the 5-year rolling annualized volatility from daily (top panel) or weekly (bottom panel) CRSP value-weighted 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.

channel of the kind index trading produces is consistent with this pattern, and a fundamentals story is not.

A.8 Return Autocorrelations and Volume

The horizon decay in Table 1 is a statement about daily and weekly return autocorrelation. Figures 23a and 23b plot the rolling 5-year first-order autocorrelation of daily and weekly CRSP value-weighted market excess returns, estimated as the slope in $r_{t+1} = \alpha + \beta r_t + \varepsilon$ over the prior 60 calendar months, with ± 1.96 HC1 standard er-




Figure 20: Trend coefficient by rolling-window length

Notes: Each point is the trend coefficient b from the regression $\sigma_t = a + b\tilde{\tau} + \varepsilon_t$ on the 1955–2025 sample, with σ_t constructed as the rolling annualized volatility of daily (top panel) or weekly (bottom panel) CRSP value-weighted 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.

rors. Daily autocorrelation peaks at roughly 0.30 in the mid-1970s, declines steadily, and crosses zero around 2003, while the weekly series crosses zero around 1990 and settles near -0.12 . The linear trend on $\tilde{\tau} \in [0, 1]$ over 1955–2025 is -0.39 (NW HAC standard error 0.08) for the daily series and -0.29 (0.03) for the weekly, with fixed- b p -values below 10^{-3} in both cases. On the 1962–1974 sample of Campbell et al. (1993) (henceforth CGW) we recover the daily autocorrelation of 0.28 they report.




Figure 21: 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.

Table 10: Volatility trend by market construction and size segment

Construction	Trend (pp/decade)
Value-weighted market	1.59***
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 CRSP value-weighted market; 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.

Table 11 extends the CGW Table II regression

$$r_{t+1} = \alpha + \beta r_t + \gamma (V_t \cdot r_t) + \sum_{d=1}^4 \delta_d (D_d \cdot r_t) + \varepsilon_{t+1}, \quad (22)$$

to the modern sample, 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 dum-






Figure 22: Log variance-ratio trend coefficients by country

Notes: Per-country trend coefficient of the log five-year rolling variance ratio (daily-implied over monthly-implied annualized variance, per decade) 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), with Newey–West HAC standard-error bars (60 monthly lags). All 35 bars are positive; the United States is in the lower third of the distribution.

mies interacted with r_t , with HC1 standard errors. The CGW prediction $\gamma < 0$, that high-volume returns are more strongly reversed, holds on the full 1955–2025 sample under all three V_t specifications, and our estimates on their samples match theirs (their γ on Sample A was -0.328 with their detrended measure, vs. our -0.351). On the post-CGW 1988–2025 sample γ remains strongly negative (-0.36 , $t = -3.2$ for the 1-yr detrended specification). The 1980–2000 sub-period attenuates toward zero, the period when the autocorrelation itself was near zero, and the post-2000 sub-period



(a) Daily returns



(b) Weekly returns

Figure 23: First-order return autocorrelation of the market, 5-year rolling

Notes: First-order autocorrelation of CRSP value-weighted 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).

returns to strongly negative estimates as the autocorrelation turns negative.

Table 11: Return autocorrelation and volume

Notes: Estimates of γ from Equation (22): $r_{t+1} = \alpha + \beta r_t + \gamma(V_t \cdot r_t) + \sum_d \delta_d(D_d \cdot r_t) + \varepsilon_{t+1}$, where V_t is log aggregate market turnover constructed from CRSP daily data following the construction underlying Figure 4. 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)	–0.129*** (0.015)	–0.289*** (0.088)	–0.228*** (0.065)	17,619
CGW Sample A (1962-87)	–0.116*** (0.024)	–0.351*** (0.063)	–0.304*** (0.049)	6,344
CGW Sample B (1962-74)	–0.173** (0.076)	–0.449*** (0.112)	–0.291*** (0.079)	3,122
CGW Sample C (1975-87)	–0.097*** (0.035)	–0.229*** (0.076)	–0.280*** (0.074)	3,221
1988-2025 (post-CGW)	–0.197*** (0.049)	–0.361*** (0.112)	–0.257*** (0.082)	9,321
Pre-1980	–0.144*** (0.047)	–0.356*** (0.093)	–0.307*** (0.072)	6,274
1980-2000	–0.067 (0.065)	–0.042 (0.110)	–0.014 (0.099)	5,055
2000-2025	–0.230*** (0.083)	–0.353*** (0.123)	–0.235*** (0.088)	6,288

A.9 NYSE Trading Hours Changes

Table 3 exploits changes in NYSE trading schedules as instruments for trading volume. Table 12 lists the nine schedule changes used in the analysis. Until 1952 the NYSE operated six days per week, with a Saturday session from 10:00 AM to 12:00 PM; from 1945 the exchange closed that session each summer, from late May to late September, and from May 31, 1952 it closed Saturdays for good. All subsequent changes affected weekday hours only, adjusting the closing time while keeping the open at 10:00 AM, until the open moved to 9:30 AM in September 1985. The 1968 Wednesday closure reduced the trading week to four days for six months. Panel A of Table 3 uses all nine changes. Panel B leaves out the two 1952 changes and uses the remaining seven, because neither can be identified once we aggregate to the week. The Saturday closings were identified from the comparison of the weekend return with the other days of the same week, which a weekly design no longer has, and because they recur in the same summer months every year a weekly comparison cannot separate them

from the seasonal pattern in trading. The September 1952 extension of the weekday session replaced the lost Saturday hours and took effect at the end of September, when Saturdays had reopened in earlier years, so that weekly hours barely moved relative to prior years and the before-and-after comparison falls on the same seasonal turn. Table 13 reports the robustness of each panel.

Table 12: NYSE Trading Hours Schedule Changes

Notes: All times Eastern. Dates are the first trading day under the new schedule, from the NYSE's own list of special closings and its 1971 chronology; where the exchange's list gives only the month (October 1974), the day is from contemporaneous accounts. The schedule shown is the weekday session, Monday to Friday; "Daily Hours" are weekday hours; the Saturday session was 2.0 hours when open, and was closed each summer from 1945 before it ended for good. "Weekly Hours" counts the trading days in the week and the shorter Saturday session. Two temporary 2:00 PM closes (August 8–18, 1967 and January 22–March 1, 1968) are absorbed by their own fixed effects and are not used as events.

#	Effective Date	Description	Schedule	Daily Hours	Weekly Hours
	<i>Pre-change</i>	<i>6-day week, Mon–Sat</i>	<i>10:00–3:00 (Sat 10–12)</i>	<i>5.0</i>	<i>27.0</i>
1	May 31, 1952	Saturday trading ended	10:00–3:00	5.0	25.0
2	Sep 29, 1952	Close extended 30 min	10:00–3:30	5.5	27.5
	<i>Stable: Sep 1952 – Jun 1968</i>				
3	Jun 12, 1968	Wednesday closure (paperwork)	10:00–3:30, no Wed	5.5	22.0
4	Jan 2, 1969	Wed reopens; hours shortened	10:00–2:00	4.0	20.0
5	Jul 7, 1969	Close extended 30 min	10:00–2:30	4.5	22.5
6	Sep 29, 1969	Close extended 30 min	10:00–3:00	5.0	25.0
7	May 4, 1970	Close extended 30 min	10:00–3:30	5.5	27.5
8	Oct 1, 1974	Close extended 30 min	10:00–4:00	6.0	30.0
9	Sep 30, 1985	Open moved 30 min earlier	9:30–4:00	6.5	32.5

A.10 S&P 500 Futures Trading Hours

Table 14 gives the exchange's record of the hours in which S&P 500 index futures could be traded outside regular equity-market hours, compiled from CME rule filings and notices, Federal Register publications, and contemporaneous accounts, with every statement labeled by its evidence status. The full-size contract was listed on GLOBEX for overnight trading in September 1993, and CME's archived calendars show holiday sessions continuously from late 1998. The four events in the table are the ones pooled in Table 4, column (3), and estimated one at a time in Table 15. Two later schedule changes are used in Section 3.1.4: the September 29, 2003 relocation of the nightly

Table 13: Trading Hours, Volume, and Variance on the NYSE: Robustness

Notes. Each row changes one thing relative to Table 3 and reports the elasticity of variance to trading hours, with the same block bootstrap standard error. In Panel A the first three rows drop one of the three experiments, the seasonal control is the month-of-year effect for the Friday-to-Monday return, and the holiday row keeps days containing a holiday, with a dummy, instead of dropping them. In Panel B the autocorrelation row estimates the lag-one autocorrelation of daily returns inside the windows, with one slope on log trading hours, and combines it with the estimate on the sum of squared daily returns to obtain the elasticity of the squared weekly return under an AR(1); the 1952 row adds the end of Saturday trading and the September 1952 extension of the trading day as a single window.

Panel A: Daily

	Variance on hours	N
As in Table 3	2.21 (0.63)	4,790
Without the Saturday closings	2.30 (0.82)	532
Without the Wednesday closure	2.01 (0.97)	4,520
Without the weekday-session changes	2.25 (0.62)	4,528
No month-of-year control for the weekend return	1.38 (0.56)	4,790
Saturday period 1926–60 instead of 1940–57	1.95 (0.57)	8,839
Days containing a holiday kept, with a dummy	2.13 (0.61)	5,022
Poisson instead of log OLS	2.50 (0.89)	4,790
Winsorised at the 95th percentile	2.07 (0.61)	4,790

Panel B: Weekly

	Variance on hours	N
As in Table 3	2.74 (1.23)	167
Squared weekly return as the outcome	2.52 (2.48)	167
Squared weekly return implied by one autocorrelation parameter	4.02 (1.95)	167
With the two 1952 changes as one window	3.03 (1.09)	211
Wednesday closure only	1.96 (2.10)	49
Weekday-session changes only	4.46 (1.88)	118
Weeks containing a holiday dropped	3.41 (1.42)	142
Poisson instead of log OLS	3.71 (1.92)	167
Windows ± 8 weeks	1.63 (1.42)	109

one-hour shutdown from 22:00–23:00 to 16:30–17:00 Central, and the 2014 move of the holiday-session halt from 10:30 to noon Central. Both are documented in the same record and both are visible in the bar data as a change in where prints stop. Figure 24 plots overnight relative to daytime trading activity in the E-mini from 1997, and the ratio increased steadily after the launch.

A.11 Robustness of the 1993 and 1997 Events

Table 15 reports each of the four CME schedule changes on its own, and Table 16 re-estimates the two large events under one change at a time, with both events run

Table 14: S&P 500 Futures Trading Hours Outside NYSE Regular Hours

Notes: Hours available for trading S&P 500 index futures outside NYSE regular trading hours (9:30am–4:00pm ET), per weekday, taking the maximum availability across the pit tail, the big-contract (SP) GLOBEX session, and the E-mini (ES), net of documented halts and maintenance windows. Times are Central. Sources are CME rule filings and notices, Federal Register publications, and contemporaneous accounts; the label under each event gives its evidence status (documented, double- or single-attested, or not established). The 1997 E-mini row is an *access* event rather than an hours event: its change in outside-RTH hours is at most half an hour and is not established, while its documented innovation—simultaneous electronic and pit trading *during* regular hours, with a 30-contract limit on the electronic book—is invisible to a measure defined outside those hours.

Event	Date	Before (hours)	After (hours)	Δ (hours)
SP listed on GLOBEX <i>Month double-attested; day single-source</i>	Sep 20, 1993	0.25	16.50	+16.25
E-mini launch <i>Launch-day clock not established; $\Delta \in [0, 0.5]$</i>	Sep 9, 1997	16.75	17.25	+0.50
Globex session expansion <i>Documented</i>	Sep 29, 2003	16.00	16.75	+0.75
Maintenance window <i>Documented</i>	Nov 19, 2012	16.75	16.50	−0.25

through identical code. The Newey–West lag choice matters for the 1993 event, whose t -statistic runs from 7.1 with no correction to 3.0 with 132 lags. We report 44 lags, twice the rolling window, throughout. The 1993 estimate is stable across windows from ± 3 to ± 24 months, while the E-mini estimate is confined to ± 6 months and is zero at ± 12 . The placebo is clean for 1993 but does not discriminate for the E-mini, which returns a significant coefficient at a placebo date one year after the launch. Two further probes cut the other way for 1993. On a quarterly grid of placebo dates over 1993–96, a no-event date in January 1996 returns 0.60 ($t = 4.3$), larger than the event itself, and dropping the ten largest overnight market moves inside the window reduces the 1993 coefficient to 0.21 ($t = 1.0$), while the E-mini estimate is unchanged by the same probe (0.38, $t = 4.8$).

A.12 Index Trading Against a Shut Market: Estimators and Sample Cuts

Table 17 reports every column of Table 5 under each estimator we considered, together with the same 15-minute bucket panel run on all weekdays rather than on holidays alone, and Table 18 re-estimates column (1) on alternative samples. The paper’s specification is ordinary least squares on variance per hour in levels. A shut window’s variance is one squared increment, and the log of a single draw has an expectation about 1.27 log points below the log of its expectation, and therefore a log-ratio regression that does not correct for this overstates the contrast by a factor of three or more.




Figure 24: 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. Tick-count data begin in 1997; contract volume data begin in 2004. CME changed how trades are bundled into reported prints in 2011 and again in 2015, which affects the tick-count series but not volume; the volume series is the cleaner measure where the two are both available. Sample: 1997–2020.

Moreover, a window whose price is unchanged across the closure has zero variance and no log at all (ten of the 58 dead holiday afternoons and 47 of the 124 shut windows in the 2003 sample); dropping those windows biases the ratio down and flooring them biases it up. Levels need neither a correction nor a zero rule, and Poisson pseudo-maximum likelihood, which estimates the same multiplicative model without logs and keeps every zero, gives nearly the same answer in every column. The sample cuts show where the day-level estimate comes from. Once Labor Day 2019 is out, no single remaining holiday moves it much (without MLK 2008, the largest day in the dead cell, it is $\times 11.0$), and it declines to $\times 3.8$ only if unscheduled closures are added back – a decline that is entirely December 5, 2018, a national day of mourning on which the Huawei arrest broke during the closure.

Table 15: Index Futures Trading Hours, Access, and Overnight Market Variance

Notes: 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 market is the value-weighted top 500 stocks by prior-December market capitalization; overnight and intraday returns use the CRSP-total-return-anchored construction, correct through splits and dividends; idiosyncratic returns impose a unit market beta following Campbell et al. (2001). The event ramp equals 0 before the event, rises linearly to 1 over the next 22 trading days, and equals 1 thereafter. The log futures volume row re-runs the same event-ramp regression with log S&P 500 futures volume (22-day rolling average) as the dependent variable. Windows are ± 6 months; Newey-West standard errors with 44 lags in parentheses. Column (5) pools all four events with event fixed effects, entering the three hours events through a linear $\text{Post} \times \Delta\text{Hours}$ dose and the E-mini through a separate access indicator, so that the hours channel is not asked to absorb an access event whose hours change is at most half an hour. Each event window keeps its own coefficient on the idiosyncratic control, so the access indicator in column (5) is the same object as the event ramp in column (2), re-estimated jointly, and reproduces it exactly. Because the dependent variable is a 22-day rolling average, consecutive observations share 21 of 22 days; the Newey-West correction is what handles that overlap, and Table 16 reports the lag sensitivity. In column (5) the Newey-West sum is computed within event blocks, so no autocovariance is taken across the multi-year gaps between events. Hours are measured in levels rather than logs: the two differ by whether the pre-1993 base is the 6.5-hour regular session or the 15-minute pit tail, and the level specification is invariant to that choice. Sample 1993–2024; the panel cannot begin earlier because CRSP opening prices are unavailable before mid-1992. ***, **, * denote significance at the 1%, 5%, and 10% level.

	(1) GLOBEX listing Sep 1993	(2) E-mini launch Sep 1997	(3) Globex expansion Sep 2003	(4) Maintenance Nov 2012	(5) Pooled
Channel	hours	<i>access</i>	hours	hours	
Δ Overnight hours	16.25	0.50	0.75	−0.25	
Event ramp	0.478** (0.195)	0.430*** (0.058)	0.199 (0.154)	−0.118 (0.185)	
Log futures volume	0.169* (0.088)	0.597*** (0.054)	0.030 (0.058)	0.009 (0.097)	
$\text{Post} \times \Delta\text{Hours}$ (per hour)					0.0301** (0.0120)
E-mini access indicator (= col. 2)					0.430*** (0.058)
Log idio. var. ratio	0.887 (0.783)	0.819*** (0.222)	0.758** (0.355)	0.611* (0.320)	Y
Event FE					Y
Observations	253	252	252	249	1006
R-squared	0.15	0.46	0.14	0.14	0.44

Table 16: Robustness of the Two Large Events

Notes: Each row re-estimates the event-ramp specification of Table 15 under a single change to the baseline (top row), which uses Newey-West standard errors with 44 lags, a ± 6 -month window, and a 22-day rolling variance. Both events are run through identical code. *t*-statistics in parentheses. The final row reports the range of the ramp coefficient across a leave-one-day-out jackknife in which each raw trading day in turn is removed and the rolling variances recomputed, so that no single day can be responsible for the estimate. The one-year-prior placebo is unavailable for the 1993 event because the panel begins in January 1993.

	GLOBEX listing Sep 1993	E-mini launch Sep 1997
<i>Baseline</i>	0.478** (2.46)	0.430*** (7.48)
Newey-West lags = 0	0.478*** (7.14)	0.430*** (10.55)
Newey-West lags = 132	0.478*** (3.03)	0.430*** (10.03)
Window ± 3 months	0.823*** (4.13)	0.313*** (3.75)
Window ± 12 months	0.603*** (3.74)	-0.003 (-0.02)
Window ± 24 months	0.497*** (3.63)	-0.121 (-0.87)
Rolling window = 10 days	0.427** (2.10)	0.583*** (9.34)
Rolling window = 44 days	0.607*** (4.03)	0.245*** (2.76)
No idiosyncratic control	0.422** (1.97)	0.504*** (4.82)
Step indicator, not ramp	0.580*** (3.20)	0.386*** (5.35)
Placebo: event +1 year	-0.192 (-0.77)	0.410** (2.04)
Placebo: event +2 year	0.279 (1.49)	0.058 (0.32)
Leave-one-day-out range	[0.392, 0.625]	[0.344, 0.502]

Table 17: Index Trading Against a Shut Market: Estimators

Notes: Columns (1), (2), (4), and (7) are columns (1)–(4) of Table 5, same samples and construction; column (3) adds year-month fixed effects to column (2). Columns (5)–(6) run the 15-minute bucket panel of column (4) on all weekdays 2004–2019, with time-of-day fixed effects alone and with time-of-day and day fixed effects. There the dead time is mostly Friday evening and the weekend, set against the same clock hours on Monday through Thursday, so the comparison is the weekend comparison of French and Roll (1986) rerun on the index; it is reported here as a benchmark and not in the main table. Each cell is an implied ratio with the coefficient's t -statistic in parentheses. LEVELS, OLS is the specification in Table 5. Poisson PML estimates the same multiplicative model without taking logs and keeps every zero cell. The two LOG rows differ only in the zero rule (a window whose price is unchanged across the closure has no log): dropping biases the ratio down, a floor biases it up. The second block is the further increment from an ordinary business day, a benchmark rather than a clean effect (Section 3.1.4); it is not defined in the holidays-only panel or the 2003 design. Appendix A.12 discusses the estimators.

	Day level: 10:30–12:00 CT vs. same-day overnight			Day $\times$ 15-min buckets		2003 swap
	(1)	(2)	(3)	holidays	all weekdays	
<i>Index tradeable vs. nothing trading</i>						
LEVELS, OLS (Table 5)	$\times 13.1$ (7.2)	$\times 12.6$ (5.7)	$\times 35.5$ (8.3)	$\times 11.3$ (3.3)	$\times 3.2$ (12.7)	$\times 4.0$ (2.8)
LEVELS, Poisson PML	$\times 12.0$ (5.9)	$\times 12.4$ (5.8)	$\times 16.6$ (6.3)	$\times 13.2$ (6.7)	$\times 4.7$ (10.8)	$\times 4.7$ (12.7)
LOG, zeros dropped	$\times 4.9$ (5.9)	$\times 4.9$ (5.8)	$\times 5.7$ (6.4)	$\times 4.9$ (7.7)	$\times 7.9$ (27.4)	$\times 6.5$ (31.3)
LOG, zeros floored at tick ² / 3	$\times 8.0$ (7.4)	$\times 7.9$ (7.4)	$\times 9.5$ (7.9)	$\times 7.0$ (8.7)	$\times 8.8$ (27.5)	$\times 7.4$ (31.8)
<i>Cash market open, on top of the index (benchmark)</i>						
LEVELS, OLS	$\times 7.0$ (7.6)	$\times 6.5$ (6.9)	$\times 3.8$ (4.2)	—	$\times 4.5$ (19.2)	$\times 4.0$ (17.2)
LEVELS, Poisson PML	$\times 3.5$ (3.8)	$\times 3.3$ (3.6)	$\times 2.8$ (3.4)	—	$\times 7.2$ (11.3)	$\times 3.9$ (9.8)
LOG, zeros dropped	$\times 3.6$ (8.2)	$\times 3.3$ (7.5)	$\times 2.9$ (6.7)	—	$\times 4.8$ (11.1)	$\times 3.7$ (20.9)
Same-day overnight subtracted	Y	Y	Y			
Day-of-week FE		Y	Y			
Year-month FE			Y			
Time-of-day FE				Y	Y	Y
Day FE				Y		Y
Window FE						
Cells	58/33/4,188 days	same	same	8,736	411,236	same
Zero-variance cells	10	same	same	—	15,800	same
						372
						49

Table 18: Index Trading Against a Shut Market: Sample Cuts

Notes: Column (1) of Table 5 (LEVELS, OLS, same-day overnight subtracted) re-estimated on alternative samples. “No overnight adjustment” drops the subtraction of the same-day overnight rate. Unscheduled closures are the two Hurricane Sandy days (October 29–30, 2012), the December 5, 2018 national day of mourning, and MLK Day 2006, on which the session in the bar data does not match the published schedule.

<i>Column (1), LEVELS OLS</i>	Ratio	<i>t</i>
As in Table 5 (58 dead afternoons, 33 index-only mornings)	×13.1	7.2
no overnight adjustment at all	×5.8	4.8
September 2, 2019 (Labor Day) back in	×19.4	6.7
unscheduled closures back in (Hurricane Sandy, December 5, 2018, MLK 2006)	×3.8	1.7
... all 62 dead days minus December 5, 2018 alone	×11.6	6.7
without MLK 2008 (world markets fell with the U.S. shut)	×11.0	5.5

A.13 Reversal Strategies in S&P 500 Futures

A variance ratio above one reflects negative serial dependence in returns. We assess its economic magnitude using three simple reversal strategies in S&P 500 futures. The data cover April 1982 through December 2019, using the full-size contract through 1997 and the E-mini thereafter. We adjust the continuous front-contract series for estimated roll gaps using changes in the futures–cash basis. The resulting returns approximate trading individual contracts; an alternative carry-based adjustment produces similar results.

The position of strategy h on day t is

$$\omega_t^{(h)} = -\kappa_h \left[\sum_{j=0}^{h-1} r_{t-j}^F - \mathbf{1}_{\{h=21\}} h \bar{r}_t^F \right], \quad h \in \{1, 5, 21\},$$

where r_t^F is the roll-adjusted daily futures log return. Here, $\omega_t^{(h)}$ is signed futures notional per unit of fixed reference capital. The trailing mean $\bar{r}_t^F$ uses up to 1,260 daily log returns, with at least 60 observations required initially. For each strategy, the positive constant κ_h normalizes full-sample average absolute exposure to one. This constant scaling leaves gross and net Sharpe ratios unchanged under the proportional trading costs used here. Signals are formed at the cash-market close, and positions are updated at the first available five-minute bar ending at least five minutes later. We report performance before costs and after charging half a tick per executed side, corresponding to an assumed one-tick bid–ask spread.

Table 19 reports the log variance ratio $\log \text{VR}(21)$ of daily futures and S&P 500 index returns and annualized Sharpe ratios. The five-day strategy earns a full-sample Sharpe ratio of 0.36 before costs and 0.34 after costs. Its net Sharpe ratios are 0.27 in 1982–1997 and 0.50 in 1998–2019. The final row reports full-sample performance excluding October 1987.

Table 19: Reversal Strategies in S&P 500 Futures

Period	log VR(21)		Signal horizon		
	Futures	Index	1-day	5-day	1-month
<i>A. Gross Sharpe ratio</i>					
1982–1997	0.40	0.12	0.00	0.28	0.32
1998–2019	0.30	0.31	0.52	0.53	0.36
1982–2019	0.34	0.24	0.18	0.36	0.33
1982–2019, excluding October 1987	0.29	0.25	0.48	0.49	0.36
<i>B. Sharpe ratio net of bid–ask spread costs</i>					
1982–1997			−0.02	0.27	0.31
1998–2019			0.44	0.50	0.34
1982–2019			0.14	0.34	0.32
1982–2019, excluding October 1987			0.40	0.45	0.34

Notes: The table reports log VR(21), where $\text{VR}(21) = 21 \text{Var}(r_t) / \text{Var}(\sum_{j=0}^{20} r_{t-j})$, computed from underlying returns, not from strategy P&L: the roll-adjusted daily futures log returns r_t^F , and daily log total returns on the S&P 500 index, including dividends, on the same days. It uses the available one-day returns and the complete 21-trading-day windows within each period, and it is shown once because the same VRs apply to both panels. The futures log VR exceeds the index log VR in 1982–1997; the two are similar in 1998–2019. Annualized Sharpe ratios equal the mean daily strategy P&L divided by its standard deviation, multiplied by $\sqrt{252}$. The signal-horizon columns show the return horizon used to form each signal; all strategies rebalance daily. The monthly signal uses a trailing 1,260-day mean, with expanding history initially and at least 60 observations. All strategies are evaluated over the same dates after this initial period. Net performance deducts half a contemporaneous tick per executed side, including both legs of contract rolls. The exclusion row omits October 1987 P&L observations; signal histories are otherwise unchanged. For VR(21), it omits October 1987 daily returns and every 21-day window that contains one.

B. Model details

This appendix derives Propositions 1 and 2, gives the intraday implementation of the daily model, rationalizes the updating rates through information acquisition, and reports the benchmark variance decomposition.

B.1 Equilibrium without index trading: Proposition 1

We use the primitives of Section 4.1. Let $\mathbf{1}$ denote an N -vector of ones and $q \equiv \mathbf{1}/N$ the vector of supplies and market-portfolio weights.

B.1.1 Fundamental forecast gaps

Conjecture that prices have the full-information dividend loading $1/(R_f - 1)$, as verified below. A dividend forecast gap therefore changes the expected next dividend and its continuation value, giving $g_{i,t} = \frac{R_f}{R_f - 1} \delta_{i,t}$. Since only common dividend news is in-

corporated gradually, $\delta_{i,t} = \hat{d}_{m,t}^{\text{stk}} - d_{m,t}$ and $g_{i,t} = g_{m,t}^{\text{stk}}$ for every stock. The updating rule implies

$$g_{m,t+1}^{\text{stk}} = (1 - \alpha_{\text{stk}})g_{m,t}^{\text{stk}} - \frac{R_f}{R_f - 1}(1 - \alpha_{\text{stk}})\varepsilon_{m,t+1}^F. \quad (23)$$

B.1.2 Demands and market clearing

Write $\boldsymbol{\mu}_t \equiv E_t^A[\boldsymbol{\Pi}_{t+1}^e]$ and $\mathbf{z}_t \equiv (z_{1,t}, \dots, z_{N,t})'$. Symmetry implies the conditional covariance matrix $\Sigma_S = (v_S - c_S)I_N + c_S\mathbf{1}\mathbf{1}'$, where v_S is the variance of each stock and c_S the covariance of distinct stocks. Set $v_{m,S} \equiv q'\Sigma_S q$ and, for any vector $\mathbf{y}$, define $y_m \equiv q'\mathbf{y}$ and $\mathbf{y}^\perp \equiv \mathbf{y} - y_m\mathbf{1}$. Then

$$\Sigma_S\mathbf{1} = Nv_{m,S}\mathbf{1}, \quad \Sigma_S\mathbf{y}^\perp = (v_S - c_S)\mathbf{y}^\perp, \quad v_{m,S} = \frac{v_S + (N-1)c_S}{N}. \quad (24)$$

CARA-normal demands, after aggregating each specialist group of mass $1/N$, are

$$\mathbf{n}_t^A = \frac{1}{a}\Sigma_S^{-1}\boldsymbol{\mu}_t, \quad \mathbf{n}_t^{B,\text{stk}} = \frac{1}{Nav_S}(\boldsymbol{\mu}_t + g_{m,t}^{\text{stk}}\mathbf{1} + \mathbf{z}_t). \quad (25)$$

Market clearing requires $\mathbf{n}_t^A + \mathbf{n}_t^{B,\text{stk}} = q$. Projecting this condition onto the market direction gives

$$\mu_{m,t} = \bar{\mu}_m - \frac{\psi_S}{1 + \psi_S}(g_{m,t}^{\text{stk}} + z_{m,t}), \quad \bar{\mu}_m \equiv \frac{av_{m,S}}{1 + \psi_S}, \quad \psi_S \equiv \frac{v_{m,S}}{v_S}. \quad (26)$$

The orthogonal projection similarly yields

$$\boldsymbol{\mu}_t^\perp = -\chi_S(0)\mathbf{z}_t^\perp, \quad \chi_S(0) \equiv \frac{v_S - c_S}{Nv_S + (v_S - c_S)}. \quad (27)$$

B.1.3 Prices and equilibrium variances

Conjecture the market price in equation (13). Its rational expected excess payoff is

$$\mu_{m,t} = -(R_f - 1)p_S - [R_f - (1 - \alpha_{\text{stk}})]b_{m,S}^F g_{m,t}^{\text{stk}} - (R_f - \lambda)b_{m,S} z_{m,t},$$

where the dividend level cancels. Matching coefficients with equation (26) gives

$$p_S = -\frac{\bar{\mu}_m}{R_f - 1}, \quad b_{m,S}^F = \frac{\psi_S}{1 + \psi_S} \frac{1}{R_f - (1 - \alpha_{\text{stk}})}, \quad b_{m,S} = \frac{\psi_S}{1 + \psi_S} \frac{1}{R_f - \lambda}. \quad (28)$$

The two return components then follow from the state laws. In particular, the demand-pressure component has negative first-order return autocovariance: its sign is that of $(R_f - \lambda)(R_f - 1) < 0$.

With $\mathbf{d}_t^\perp \equiv \mathbf{d}_t - d_{m,t}\mathbf{1}$, the relative-price solution implied by (27) is

$$\mathbf{P}_t^\perp = \frac{\mathbf{d}_t^\perp}{R_f - 1} + b_{z,S}^\perp \mathbf{z}_t^\perp, \quad b_{z,S}^\perp \equiv \frac{\chi_S(0)}{R_f - \lambda}. \quad (29)$$

Together, $\mathbf{P}_t = P_{m,t}\mathbf{1} + \mathbf{P}_t^\perp$ clears all N asset markets and proves Proposition 1. Since $\chi_S(0) = O(1/N)$, stock-specific demand pressure has little price impact in a large market.

Equilibrium variances. The variances in investors' demands must equal those generated by these prices:

$$v_{m,S} = \left(\frac{R_f}{R_f - 1} \right)^2 [1 - (1 - \alpha_{\text{stk}}) b_{m,S}^F]^2 \text{Var}(\varepsilon_{m,t+1}^F) + b_{m,S}^2 \text{Var}(u_{m,t+1}^z), \quad (30)$$

$$v_S - c_S = \left(\frac{R_f}{R_f - 1} \right)^2 \sigma_F^2 (1 - \rho_F) + (b_{z,S}^\perp)^2 \sigma_\eta^2 S^{1-\xi} (1 - \rho_z). \quad (31)$$

Here $\text{Var}(\varepsilon_{m,t+1}^F) = \sigma_F^2 [\rho_F + (1 - \rho_F)/N]$ and $\text{Var}(u_{m,t+1}^z) = \sigma_\eta^2 S^{1-\xi} [\rho_z + (1 - \rho_z)/N]$. Because the loadings depend on the variances, these are two scalar fixed-point conditions in $(v_{m,S}, v_S - c_S)$, solved numerically. Although dividend levels are nonstationary, they cancel from excess payoffs; conditional payoff risk depends only on stationary innovations and belief states.

B.2 Equilibrium with index trading: Proposition 2

A fraction ω of behavioral investors trades the index, leaving $(1 - \omega)/N$ specialists per stock. The index is a pass-through claim on $q = \mathbf{1}/N$ and creates no new risky-asset supply.

B.2.1 Beliefs and market clearing

As in Section B.1.1, the index investors' fundamental distortion is $g_{m,t}^{\text{idx}} = \frac{R_f}{R_f - 1} (\hat{d}_{m,t}^{\text{idx}} - d_{m,t})$. Their updating rule implies

$$g_{m,t+1}^{\text{idx}} = (1 - \alpha_{\text{idx}}) g_{m,t}^{\text{idx}} - \frac{R_f}{R_f - 1} (1 - \alpha_{\text{idx}}) \varepsilon_{m,t+1}^F. \quad (32)$$

Index investors face market return variance $v_{m,S}$ and act on the common demand distortion $z_{m,t}$. Their aggregate underlying-stock demand is

$$\mathbf{n}_t^{B,\text{idx}} = \frac{\omega}{av_{m,S}} (\mu_{m,t} + g_{m,t}^{\text{idx}} + z_{m,t}) q. \quad (33)$$

Together with the arbitrageur and the remaining stock specialists, clearing requires

$$\frac{1}{a} \Sigma_S^{-1} \mu_t + \frac{1 - \omega}{Nav_S} (\mu_t + g_{m,t}^{\text{stk}} \mathbf{1} + \mathbf{z}_t) + \mathbf{n}_t^{B,\text{idx}} = q. \quad (34)$$

Projecting onto the market direction using (24) gives

$$\begin{aligned} \mu_{m,t} &= \bar{\mu}_m - \frac{(1 - \omega) \psi_S}{\Delta_S(\omega)} g_{m,t}^{\text{stk}} - \frac{\omega}{\Delta_S(\omega)} g_{m,t}^{\text{idx}} \\ &\quad - \frac{\omega + (1 - \omega) \psi_S}{\Delta_S(\omega)} z_{m,t}, \quad \bar{\mu}_m \equiv \frac{av_{m,S}}{\Delta_S(\omega)}. \end{aligned} \quad (35)$$

In the relative-price direction, index demand drops out, leaving

$$\mu_t^\perp = -\chi_S(\omega) \mathbf{z}_t^\perp, \quad \chi_S(\omega) \equiv \frac{(1 - \omega)(v_S - c_S)}{Nv_S + (1 - \omega)(v_S - c_S)}. \quad (36)$$

B.2.2 Prices and equilibrium variances

Compute the rational expected excess payoff from the market-price conjecture (18). As in Section B.1.3, matching its constant and state coefficients with (35) gives $p_S = -\bar{\mu}_m / (R_f - 1)$ and exactly the three price loadings in Proposition 2. The relative-price solution is $\mathbf{P}_t^\perp = \mathbf{d}_t^\perp / (R_f - 1) + [\chi_S(\omega) / (R_f - \lambda)] \mathbf{z}_t^\perp$. Together with the market price, it satisfies all N clearing conditions; setting $\omega = 0$ recovers Proposition 1.

Equilibrium variances satisfy the two fixed-point conditions (30)–(31), replacing $b_{m,S}$ and $b_{z,S}^\perp$ by their ω -dependent values and the common-news impact factor by

$$\kappa_F(\omega) \equiv 1 - (1 - \alpha_{\text{stk}}) b_{m,S}^{F,\text{stk}}(\omega) - (1 - \alpha_{\text{idx}}) b_{m,S}^{F,\text{idx}}(\omega). \quad (37)$$

Thus $\frac{R_f}{R_f - 1} \kappa_F(\omega)$ is the loading of a permanent market-news innovation in the one-day excess payoff. The two equilibrium variances are again solved jointly.

Stock-level objects. Symmetry makes relative returns orthogonal to the market return, so each stock has market beta one and $R_i^2 = v_{m,S} / v_S$. Arbitrageur holdings decompose as

$$n_{i,t}^A = \frac{1}{a} \left[\frac{\mu_{m,t}}{N v_{m,S}} + \frac{\mu_{i,t}^\perp}{v_S - c_S} \right].$$

The stock-specific component is of order $1/N$, while the common component responds to the market states in (35), providing the model counterpart to market-directed trading.

B.3 An intraday implementation

The daily equilibrium admits an implementation with S trading opportunities between consecutive daily boundaries. At each opportunity investors choose the myopic CARA-normal position they would hold to the next daily boundary, without accounting for possible subsequent rebalancing. They may revise that position when another opportunity arrives.

Timing and state laws. Let $h = 0, \dots, S$ index equally spaced points from boundary t to $t + 1$; trading occurs at $h = 0, \dots, S - 1$. The gross risk-free return remaining to the next boundary is $R_{h,S} \equiv R_f^{1-h/S}$, so $S = 1$ reproduces the daily trading schedule. Fundamental news and belief updating occur only at daily boundaries; $g_{m,t}^{\text{stk}}$ and $g_{m,t}^{\text{idx}}$ are therefore fixed between boundaries.

Demand pressure evolves in trade time. Set $\mathbf{z}_{t,0} = \mathbf{z}_t$, $\lambda_S \equiv \lambda^{1/S}$ and $B_S \equiv \sum_{k=0}^{S-1} \lambda_S^{2k}$. Between h and $h + 1$,

$$z_{i,t,h+1} = \lambda_S z_{i,t,h} + \eta_{i,t,h+1}^S, \quad \text{Var}(\eta_{i,t,h+1}^S) = \frac{\sigma_\eta^2 S^{1-\xi}}{B_S}, \quad (38)$$

for $h = 0, \dots, S - 1$. The innovations are jointly Gaussian, independent across opportunities and days, orthogonal to fundamental news, and have cross-stock correlation ρ_z . Iterating to the boundary gives

$$z_{i,t+1} = \lambda z_{i,t} + u_{i,t+1}^z, \quad u_{i,t+1}^z = \sum_{h=1}^S \lambda_S^{S-h} \eta_{i,t,h}^S, \quad \text{Var}(u_{i,t+1}^z | S) = \sigma_\eta^2 S^{1-\xi}. \quad (39)$$

The daily innovation also has cross-stock correlation ρ_z . Thus the daily-boundary law is exactly (10)–(11) for every S .

Intraday demands and prices. Let $\mathbf{Y}_{t+1} \equiv \mathbf{P}_{t+1} + \mathbf{d}_{t+1}$ be the cum-dividend payoff at the next boundary. The excess payoff from buying at opportunity h and holding to that boundary is $\mathbf{Y}_{t+1} - R_{h,S} \mathbf{P}_{t,h}$. Define $\boldsymbol{\mu}_{t,h} \equiv E_{t,h}^A[\mathbf{Y}_{t+1}] - R_{h,S} \mathbf{P}_{t,h}$ and $\Sigma_{h,S} \equiv \text{Var}_{t,h}^A(\mathbf{Y}_{t+1})$. By symmetry, $\Sigma_{h,S} = (v_{h,S} - c_{h,S}) I_N + c_{h,S} \mathbf{1}\mathbf{1}'$ and $v_{m,h,S} = q' \Sigma_{h,S} q$. CARA-normal demands and market clearing give

$$\begin{aligned} \frac{1}{a} \Sigma_{h,S}^{-1} \boldsymbol{\mu}_{t,h} + \frac{1-\omega}{N a v_{h,S}} (\boldsymbol{\mu}_{t,h} + g_{m,t}^{\text{stk}} \mathbf{1} + \mathbf{z}_{t,h}) \\ + \frac{\omega}{a v_{m,h,S}} (\mu_{m,t,h} + g_{m,t}^{\text{idx}} + z_{m,t,h}) q = q. \end{aligned} \quad (40)$$

The coefficient matrix in $\boldsymbol{\mu}_{t,h}$ is positive definite, so clearing uniquely determines the expected excess payoff and hence

$$\mathbf{P}_{t,h} = \frac{E_{t,h}^A[\mathbf{Y}_{t+1}] - \boldsymbol{\mu}_{t,h}}{R_{h,S}}. \quad (41)$$

As h increases, fewer demand shocks remain before the boundary, while uncertainty about the next fundamental innovation is unchanged. The conditional distribution of $\mathbf{Y}_{t+1}$ does not depend on the current intraday price, so no additional intraday price fixed point is needed; the market and relative-price decomposition of Section B.2.1 applies at each opportunity.

Exact daily-boundary nesting. At $h = 0$, $R_{0,S} = R_f$ and $\mathbf{z}_{t,0} = \mathbf{z}_t$. Equation (39) and the calendar-time laws for fundamentals and beliefs give the same conditional distribution of $\mathbf{Y}_{t+1}$ as in the daily model. In particular, $\Sigma_{0,S} = \Sigma_S$, with the same individual-stock and market variances. Consequently (40) reduces to (34), yielding

$$\mathbf{P}_{t,0} = \mathbf{P}_t. \quad (42)$$

The states update at $h = S$ according to the daily laws, and the argument repeats at the next boundary. Thus prices, returns, variances, and variance ratios at daily boundaries are identical to the daily model; S counts literal trading opportunities and ξ governs their effect on daily demand-shock variance.

B.4 Information acquisition and the updating rates

We rationalize the benchmark updating rates with a standing monitoring choice for the common fundamental $d_{m,t}$. We abstract from delays in incorporating stock-specific fundamental news.¹⁶ An investor chooses a monitoring probability $\alpha \in (0, 1]$ at daily cost $C(\alpha)$. Independently across investors and days, monitoring reveals the current common fundamental before trading; otherwise the investor retains the last forecast. This represents processing the market-wide implications of news, not access to public releases.

Aggregation. With a continuum of investors, a fraction α refreshes each day, so the group-average forecast follows

$$\hat{d}_{m,t} = (1 - \alpha)\hat{d}_{m,t-1} + \alpha d_{m,t}, \quad (43)$$

exactly as in the main model. The age K of an investor's information has stationary distribution $\Pr(K = k) = \alpha(1 - \alpha)^k$, hence $E[K] = (1 - \alpha)/\alpha$.

The value of monitoring. Taking equilibrium payoff variance v as given,¹⁷ using a stale expected payoff m' rather than the informed mean m incurs certainty-equivalent loss $(m - m')^2 / (2av)$. Write $\sigma_{F,m}^2 \equiv \sigma_F^2[\rho_F + (1 - \rho_F)/N]$ for the variance of common fundamental news. Since a k -day-old forecast misses k independent permanent innovations, its expected loss is kB_j , with

$$B_{\text{stk}} = \frac{(R_f / (R_f - 1))^2 \sigma_{F,m}^2}{2av_S}, \quad B_{\text{idx}} = \frac{(R_f / (R_f - 1))^2 \sigma_{F,m}^2}{2av_{m,S}}.$$

Both groups monitor the same common news, but index investors act on it through a less risky claim. Thus

$$\frac{B_{\text{idx}}}{B_{\text{stk}}} = \frac{v_S}{v_{m,S}} = \frac{1}{\psi_S}. \quad (44)$$

Optimal intensity. Expected staleness loss is $B(1 - \alpha)/\alpha$. With the common monitoring cost $C(\alpha) = c\alpha^\gamma$, $0 < \gamma \leq 1$, each investor minimizes

$$c\alpha^\gamma + B \frac{1 - \alpha}{\alpha}, \quad \alpha \in (0, 1].$$

Because α^2 times the derivative is $c\gamma\alpha^{1+\gamma} - B$, which is strictly increasing, the unique optimum is

$$\alpha^*(B) = \min \left\{ 1, \left(\frac{B}{\gamma c} \right)^{1/(1+\gamma)} \right\}. \quad (45)$$

¹⁶A separate stock-specific belief gap would not affect the market direction. Under symmetry it enters only relative prices, with loading $\chi_S(\omega) = O(1/N)$, negligible at $N = 500$.

¹⁷As elsewhere in the model, investors use objective conditional payoff variance in their demand.

With $e \equiv 1/(1 + \gamma)$ and an interior stock-monitoring rate,

$$\alpha_{\text{idx}} = \min \left\{ 1, \alpha_{\text{stk}} \left(\frac{1}{\psi_S} \right)^e \right\}. \quad (46)$$

The benchmark rates. We calibrate the common cost scale c to give $\alpha_{\text{stk}} = .30$ in the early-sample benchmark, where $\psi_S \simeq .211$. The index rate reaches its cap when

$$.30(1/.211)^e \geq 1 \iff e \gtrsim .77 \iff \gamma \lesssim .29.$$

Thus sufficiently strong economies of scale in monitoring intensity rationalize $(\alpha_{\text{stk}}, \alpha_{\text{idx}}) = (.30, 1)$; the curvature is not separately identified. Under linear costs, re-solving the early-sample equilibrium gives $\alpha_{\text{idx}} \simeq .70$ (Section B.5). We hold both rates fixed in the main comparative statics to isolate changes in index participation and trading opportunities.

B.5 Variance decomposition and sensitivity to the index updating rate

Since fundamental news and transitory demand pressure are independent, daily market-return variance satisfies $\text{Var}(r_t) = \text{Var}(r_t^F) + \text{Var}(r_t^z)$. The annualized variance components in the early (E) and modern (M) benchmark economies are

	Early (1950–70)	Modern (2005–24)
Total variance	0.0132	0.0379
Fundamental component	0.0102	0.0146
Transitory z component	0.0030	0.0233
z share of total	23%	61%

The share of the modeled increase in variance attributable to transitory demand pressure is

$$\frac{\text{Var}_M(r^z) - \text{Var}_E(r^z)}{\text{Var}_M(r) - \text{Var}_E(r)} = 0.82. \quad (47)$$

Sensitivity to the index updating rate. Under linear monitoring costs, re-solving the early-sample equilibrium gives $(\alpha_{\text{stk}}, \alpha_{\text{idx}}) = (.30, .70)$ rather than the benchmark $(.30, 1)$. Holding the remaining primitives and the historical paths for ω and S fixed, slower index updating leaves more return continuation. Modern volatility falls from 19.5% to 18.1% and the 21-day log variance ratio from 0.38 to 0.25; the variance ratio also remains below the data after 1970.

Era	Benchmark (.30, 1)		Linear cost (.30, .70)		Data	
	Vol	log $VR_{1,21}$	Vol	log $VR_{1,21}$	Vol	log $VR_{1,21}$
1950–70	11.5%	−0.23	10.7%	−0.37	10.4%	−0.35 to −0.18
1971–90	13.6%	0.02	12.1%	−0.19	14.6%	−0.08 to 0.17
1991–2004	15.3%	0.14	14.0%	−0.02	16.3%	0.20
2005–24	19.5%	0.38	18.1%	0.25	19.5%	0.31 to 0.39

C. Empirical evidence on slow information diffusion

The early-sample variance ratios imply continuation in market returns. The model attributes this continuation to slow diffusion of common information across stocks. We test that mechanism directly, ask whether removing the estimated delay eliminates the continuation, examine whether indexation accelerated information diffusion, and then separate the secular rise in volatility from changes in the timing of permanent news.

C.1 Slow information diffusion

If some stocks incorporate common news before others, their returns should predict subsequent returns of slower stocks more strongly than the reverse. At the end of year $Y - 1$, we identify leaders within current S&P 500 constituents (the 500 largest stocks by market equity before 1958) with valid returns and market equity on at least 90% of trading days during the preceding five years. For each stock i , compute

$$\ell_i \equiv \sum_{k=1}^5 \left[\text{Cov}(M_t^{-i}, r_{i,t-k}) - \text{Cov}(r_{i,t}, M_{t-k}^{-i}) \right], \quad (48)$$

where M_t^{-i} is the universe return excluding stock i . Symmetric comovement cancels; a positive score indicates that stock i leads the rest of the market. We rank stocks by ℓ_i/β_i , where β_i is its market beta, excluding those with $|\beta_i| < 0.2$ from the ranking. The top 50 form a value-weighted Leader portfolio for year Y ; the remaining constituents form the Follower portfolio. The portfolio returns are out-of-sample relative to the selection of the Leader stocks.

Let L_t and F_t denote the daily Leader and Follower returns. Their directional covariance at lag j is

$$d_j \equiv \text{Cov}(F_t, L_{t-j}) - \text{Cov}(F_t, L_{t+j}), \quad j = 1, \dots, 5. \quad (49)$$

For each era we report $D \equiv \sum_{j=1}^5 d_j$ against the Follower portfolio and three 50-stock nonleader portfolios: matched on size, industry, and volatility; additionally matched on beta; and random. Positive D indicates that leaders predict the comparison portfolios more strongly than the reverse.

Table 20: Out-of-sample leader validation. The Leader 50 is selected annually using the previous five years’ beta-normalized scores in equation (48) within the S&P 500 (largest 500 stocks before 1958). Entries are $10^5 D$ against followers and 50-stock nonleader comparisons matched on size, industry, and volatility; additionally matched on beta; or selected at random. Brackets report 95% year-block bootstrap intervals.

Comparison portfolio	1955–70	1971–90	1991–2004	2005–24
$10^5 D$, Leader 50 against each comparison portfolio; 95% intervals in brackets				
Followers	1.23 [0.76, 1.81]	2.62 [1.64, 3.90]	1.77 [0.47, 3.26]	0.61 [-0.00, 1.30]
Size/industry matched	1.07 [0.45, 1.93]	2.36 [1.12, 3.95]	0.90 [0.41, 1.56]	0.29 [-0.11, 0.82]
Beta matched	0.93 [0.24, 1.85]	2.18 [1.12, 3.43]	0.84 [0.34, 1.54]	0.35 [-0.15, 0.93]
Random 50	1.34 [0.71, 2.00]	2.30 [1.33, 3.44]	1.71 [0.56, 2.88]	0.64 [-0.21, 1.58]

Table 20 shows positive directional propagation through 2004 across comparison portfolios, with no statistically detectable asymmetry in 2005–24. Figure 25 shows little continuation among leaders in 1955–70; they move into reversal roughly fifteen years before the broad market. The variance-ratio gap largely closes by 2005–24.

C.2 Slow diffusion explains the early continuation

We assess the contribution of delayed common news to early continuation using two corrections on the value-weighted S&P 500 universe.

Projection synchronization. Regress daily market returns on contemporaneous, lagged, and future Leader 50 returns,

$$r_t = a + \beta_0 L_t + \sum_{j=1}^K (b_j L_{t-j} + c_j L_{t+j}) + e_t, \quad K = 5. \quad (50)$$

With $\lambda_j \equiv b_j - c_j$, synchronize the delayed response using

$$r_t^\dagger = r_t + \sum_{j=1}^K \lambda_j (L_t - L_{t-j}). \quad (51)$$

In-sample this removes the fitted asymmetry by construction. To avoid using the corrected year in estimation, we instead apply coefficients averaged across the two preceding and two following years. Averaging the five-year rolling estimates, the log variance ratio rises from -0.307 to -0.089 in 1955–70 and from -0.202 to 0.069 in 1971–90 (Figure 26).




Figure 25: Five-year rolling log variance ratios for the Leader 50, the S&P 500 universe (largest 500 stocks before 1958), and the full market. The band is the 5th–95th percentile of 100 random 50-stock portfolios drawn annually from the same universe.

Covariance correction. The second method moves delayed permanent news to impact. Normalize Leader and Follower returns by market betas, $\tilde{L}_t = L_t/\beta_L$ and $\tilde{F}_t = F_t/\beta_F$, which proxy for cumulative permanent-news loadings. Assume serially uncorrelated news ε_t enters $\tilde{L}_t$ immediately with unit loading and $\tilde{F}_t$ with weights $a_0, \dots, a_K$ summing to one. Other return components are first differences of stationary price components, orthogonal to this news at all leads and lags, and generate symmetric Leader–Follower cross-covariances. Applying (49) to normalized returns gives $d_j = a_j \sigma_\varepsilon^2$ for $j \geq 1$, so

$$\begin{aligned} \Omega_{FL} &\equiv \sum_{k=-\infty}^{\infty} \text{Cov}(\tilde{F}_t, \tilde{L}_{t-k}) = \sigma_\varepsilon^2, \\ a_j &= \frac{d_j}{\Omega_{FL}}, \quad j = 1, \dots, K, \quad a_0 = 1 - \sum_{j=1}^K a_j. \end{aligned} \tag{52}$$

Let w be the Leader 50’s average market weight. Approximate the market by $w\beta_L \tilde{L}_t + (1-w)\beta_F \tilde{F}_t$, with permanent-news timing weights

$$m_0 = w\beta_L + (1-w)\beta_F a_0, \quad m_j = (1-w)\beta_F a_j \quad (j = 1, \dots, K),$$

and total loading $m_\Sigma \equiv \sum_{j=0}^K m_j = w\beta_L + (1-w)\beta_F \simeq 1$. Moving the response to

Table 21: Delayed information and corrected return dynamics of the S&P 500 universe. Starred moments move the estimated permanent-news response to impact under the normalization in the text. The final two columns report the sum of the first five daily autocorrelations.

Era	log VR	log VR*	AC(1)	AC(1)*	$\sum_{k=1}^5 \text{AC}(k)$	$\sum_{k=1}^5 \text{AC}(k)^*$
1931–45	−0.305	−0.192	+0.052	−0.018	+0.076	+0.013
1946–54	−0.171	0.103	+0.154	+0.004	+0.096	−0.051
1955–70	−0.322	−0.039	+0.173	+0.048	+0.252	+0.055
1971–90	−0.095	0.169	+0.125	+0.007	+0.070	−0.071
1991–2004	0.296	0.482	−0.011	−0.056	−0.099	−0.181
2005–24	0.478	0.537	−0.121	−0.144	−0.148	−0.173

impact gives

$$\begin{aligned}\Gamma_M^*(0) &= \Gamma_M(0) + \Omega_{FL} \left(m_\Sigma^2 - \sum_{j=0}^K m_j^2 \right), \\ \Gamma_M^*(k) &= \Gamma_M(k) - \Omega_{FL} \sum_{j=0}^{K-k} m_j m_{j+k}, \quad 1 \leq k \leq K,\end{aligned}\tag{53}$$

with autocovariances at longer lags unchanged.

Estimation and results. Estimate Ω_{FL} by smoothing the real cross-periodogram near zero with the Dew-Becker (2017) RQS kernel:

$$\hat{\Omega}_{FL} = \frac{\sum_{s=1}^{T-1} W_T(\lambda_s) I_{FL}(\lambda_s)}{\sum_{s=1}^{T-1} W_T(\lambda_s)}, \quad \lambda_s = \frac{2\pi s}{T}.\tag{54}$$

Here W_T is the periodic extension of $B^{-1}W(\lambda/B)$, where $B = 1/H$ and

$$W(u) = \frac{1}{2\pi} \left[\frac{9}{8} - \frac{15}{8} \left(\frac{u}{2\pi} \right)^2 \right] \mathbf{1}\{|u| < 2\pi\}.$$

For delay estimation use $H = 50$ trading days, by era and in trailing 20-year windows. The kernel removes leading spectral-curvature bias; weights sum to one over nonzero Fourier frequencies, and negative estimates are retained. Finite bandwidth may also capture slowly reversing discount-rate movements. Section C.4 uses $H = 250$ for long-run variance.

Using moments estimated over each full era, Table 21 shows that the covariance correction removes most early continuation. In 1955–70, the log variance ratio rises from −0.322 to −0.039 and the five-lag autocorrelation sum falls from 0.252 to 0.055; in 1971–90, the corresponding changes are −0.095 to 0.169 and 0.070 to −0.071. The results are similar at $H = 250$ (corrected log variance ratios −0.031 and 0.163) and using zero-frequency instead of contemporaneous betas (−0.021 to −0.019 and 0.157




Figure 26: Five-year rolling observed, projection-synchronized, and covariance-corrected return dynamics of the S&P 500 universe: calendar-month log variance ratios, first-order daily autocorrelation, and the sum of the first five daily autocorrelations. Covariance corrections use trailing 20-year delay estimates.

to 0.189 across bandwidths). Thus both corrections attribute most of the early continuation to delayed common-information propagation.

Both corrections use the Leader 50 as a benchmark for the timing of common news. If leaders themselves incorporate permanent news gradually, delays shared by leaders and followers are not identified, and some slow information diffusion may remain in the adjusted returns. The corrections therefore need not recover the full-information

return process and may understate the contribution of delayed information to early return continuation.

C.3 Index membership and information diffusion

We next ask whether S&P 500 membership was associated with faster incorporation of common news when index products were concentrated in that index. Each month, we match the smallest 100 S&P 500 members by market equity to similarly sized NYSE/AMEX nonmembers; both stocks in each pair trade on NYSE or AMEX. Controls cannot be index members within twelve months of the match. Let r_t^{in} and r_t^{out} be the member and matched-nonmember portfolio returns, with each pair weighted by the member's market equity.

For either portfolio p , estimate

$$r_t^p = a + \beta_0 r_t^M + \sum_{j=1}^2 (b_j r_{t-j}^M + c_j r_{t+j}^M) + e_t, \quad (55)$$

where r_t^M is the value-weighted market return. The lead-minus-lag statistic $\delta^p \equiv \sum_{j=1}^2 (b_j - c_j)$ measures delayed incorporation of market-wide news, net of symmetric comovement. We test $\Delta \equiv \delta^{\text{in}} - \delta^{\text{out}}$, estimated by applying the same regression to $r_t^{\text{in}} - r_t^{\text{out}}$ within each period, with heteroskedasticity- and autocorrelation-consistent standard errors. A negative Δ indicates faster incorporation by index members.

Figure 27 shows $\Delta = +0.004$ in 1957–72, -0.047 in 1973–94 ($t = -2.08$), and 0.003 in 1995–2024. The middle-period estimate differs from the early and late estimates by -0.052 ($t = -1.82$) and -0.051 ($t = -1.77$), respectively, in a pooled regression allowing coefficients to vary by period. The pattern is consistent with faster common-news incorporation among index members when index products were concentrated in the S&P 500.

C.4 Daily and long-run volatility

Slow information diffusion spreads permanent news across days, lowering measured daily volatility. To gauge its contribution to the secular rise, we estimate the regressions of Table 1 in the main text on the synchronized returns of equation (51). For the full market, the 1955–2025 increase in annualized volatility falls from 11.27 to 10.01 percentage points for daily returns, from 8.05 to 7.83 for weekly returns, and from 4.77 to 4.73 for 4-week returns. For the S&P 500 universe over 1955–2024, the corresponding changes are 10.70 to 9.60, 7.27 to 7.17, and 4.28 to 4.25. Most of the rise remains.

Long-run variance measures the return variation that survives aggregation over long horizons:

$$\Omega^{LR} = \sum_{k=-\infty}^{\infty} \Gamma(k) = \lim_{H \rightarrow \infty} \frac{1}{H} \text{Var} \left(\sum_{h=1}^H r_{t+h} \right), \quad (56)$$




Figure 27: Member–nonmember difference in delayed response to common news, $\Delta = \delta^{\text{in}} - \delta^{\text{out}}$, for the smallest 100 S&P 500 members and size-matched nonmembers, by period. Both stocks in each pair trade on NYSE or AMEX. Bars show 95% confidence intervals based on heteroskedasticity- and autocorrelation-consistent standard errors.

where $\Gamma(k)$ is the market-return autocovariance at lag k . First differences of stationary price components contribute zero long-run variance because their cumulative effects reverse. We estimate Ω^{LR} in trailing 20-year windows with the RQS estimator of Section C.2, using $H = 250$, and report annualized volatility $\sigma_t^{LR} = \sqrt{n_t \Omega_t^{LR}}$, where n_t is the number of trading days per year in the window. At finite bandwidth, slowly reversing discount-rate movements may also contribute.

Figure 28 shows that from 1970 to 2024 long-run volatility rises only from 15.9% to 17.3%, versus 10.3% to 19.0% for observed daily volatility and 11.6% to 19.8% for synchronized returns. Daily variance increases by 256 percentage-points squared, compared with 48 for long-run variance; the difference is 208 (s.e. 99, $t = 2.11$), and 209 ($t = 2.12$) for synchronized returns. The increase in long-run variance is thus about one fifth of the increase in daily variance. Both comparisons suggest that most of the secular rise in market volatility is transitory. Because long-run variance is estimated using the RQS estimator over finite 20-year windows, we regard this evidence as suggestive rather than definitive. Nevertheless, the point estimates indicate that long-run variance has increased only modestly relative to daily variance.




Figure 28: Observed daily, synchronized daily, and long-run volatility of the S&P 500 universe in trailing 20-year windows, plotted by end year. Synchronized returns follow equation (51); annualization uses each window's trading days per year.