Actively Passive: The Rise of Market Volatility

Lars Lochstoer and Tyler Muir*

March 23, 2026

Abstract

U.S. stock market volatility has roughly doubled since the 1950s, yet macroeconomic fundamentals have grown calmer. We provide causal evidence that the rise of index trading – via futures, index funds, extended trading hours, etc. – is a primary cause. We exploit the introduction of E-mini S&P 500 futures and historical trading-hour adjustments as natural experiments to show that index-level trading amplifies market variance through trading volume, with overnight futures trading alone accounting for roughly half the total rise. A model in which investors shift from trading individual stocks to trading the market index rationalizes these findings: arbitrageurs absorb larger systematic demand shocks, raising volatility, reducing return autocorrelation, and increasing systematic risk. Stock market volatility increasingly reflects how we trade rather than economic fundamentals.

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

1. Introduction

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

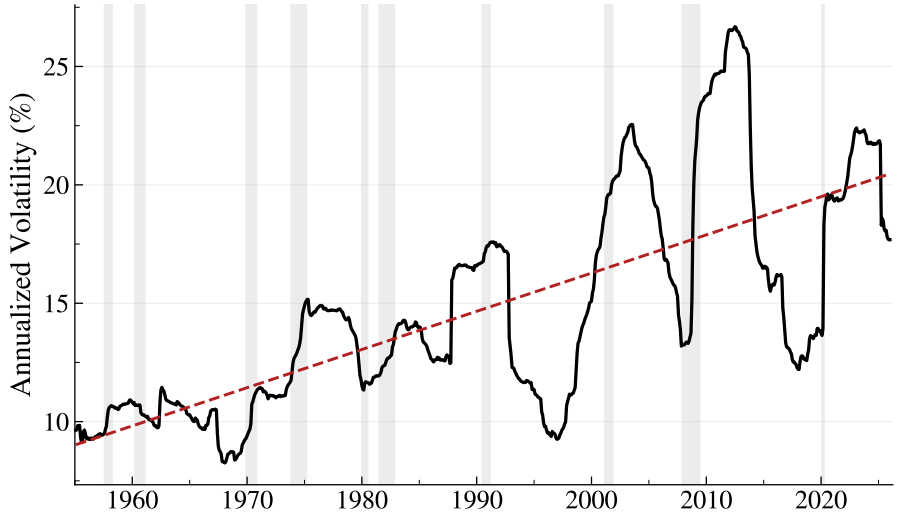


Figure 1: Stock Market Volatility. 5-year rolling standard deviation of daily CRSP value-weighted market return.

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

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

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

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

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

This correlation does not establish causation — volatility could drive trading rather than the reverse, or a third factor could drive both. We provide causal evidence from three 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 is proportional to trading time, not calendar time. We extend this logic to eight changes in NYSE trading hours between 1952 and 1985 — including the elimination of Saturday trading, extensions and reductions in daily hours, and the Wednesday closures. Using trading hours as an instrument for volume, we estimate that a 10% increase in trading volume raises market volatility by about 7%.

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

In the third experiment, we decompose overnight returns for individual stocks into systematic and idiosyncratic components. Because index futures allow investors to trade the market overnight but not individual stocks, our prediction is sharp: the systematic component of overnight variance should rise over time while the idiosyncratic component should not. Using a panel of the 500 largest U.S. stocks from 1993–2024, we estimate a triple-difference specification comparing systematic versus idiosyncratic overnight variance, on weekdays versus weekends, over time. The differential weekday trend is large and significant for the systematic component but negative for the idiosyncratic component. The implied economic magnitude is substantial: the overnight futures channel alone accounts for a roughly 58% increase in annual market variance over this 30-year period, or about half the total rise.

Not all index trading events are equally useful for identification: many major

products had negligible volume at launch and took decades to gain widespread adoption, making short-window event studies uninformative. Our experiments exploit changes where uptake was immediate, and we complement them with a triple-difference design suited to detecting gradual effects.

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

To interpret these patterns, we develop a simple demand-based 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. Arbitrageurs are thus required to absorb larger systematic demand shocks than before, which they only partially accommodate. This shift in trading structure raises market volatility, reduces return autocorrelation, and increases the

share of variance explained by market-wide factors. The predictions of the model are borne out in the data, which show both a decline in return autocorrelations and a rise in the systematic component of stock returns over time. Further, the model matches increased volume in market index trades and the large increase in market volatility.

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

Our findings imply that volatility is no longer solely a reflection of economic uncertainty, but also of how trading is organized. This matters for measuring macroeconomic uncertainty from daily stock returns (Bloom, 2009; Martin, 2017; Berger et al., 2020), for portfolio choice given the rise in systematic risk (Moreira and Muir, 2017; Lochstoer and Muir, 2022), and for understanding how market macrostructure shapes asset prices (Haddad and Muir, 2025; Harvey et al., 2025). It also raises policy questions: proposals to extend trading hours—for example, the NYSE's plan for 24-hour trading by the end of 2026—have the potential to increase volatility further.

Several strands of prior work inform our analysis. First, our paper contributes to a literature on excess volatility. Shiller (1981) first showed that stock prices move too much to be justified by subsequent dividends, and a large literature has sought expla-

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

nations. Our paper identifies a new source of excess volatility: the shift toward index-level trading, which forces arbitrageurs to absorb systematic rather than idiosyncratic demand shocks. Our model builds on the noise trader framework of De Long et al. (1990) and the limits-of-arbitrage logic of Shleifer and Vishny (1997), applied to the specific setting of index trading.

Second, we build on a literature studying how index inclusion and passive ownership affect prices. Shleifer (1986) shows that demand curves for stocks slope down using S&P 500 additions, and Barberis et al. (2005) show that stocks added to the index subsequently comove more with it—exactly the mechanism our model predicts. A key paper in this literature is Ben-David et al. (2018), who use Russell index reconstitutions to show that ETF ownership increases stock-level volatility, providing clean micro-level evidence that index products transmit non-fundamental demand shocks to underlying stocks. Our paper addresses a complementary question: whether the secular growth of index trading—across futures, ETFs, and extended trading hours—has raised *aggregate market* volatility over decades. Their cross-sectional design holds the total amount of index trading fixed; our time-series experiments exploit variation in that aggregate quantity (see also Da and Shive, 2018; Greenwood and Thesmar, 2011; Israeli et al., 2017; Basak and Pavlova, 2013; Sullivan and Xiong, 2012).

Third, our empirical strategy builds on the classic insight of French and Roll (1986) that trading itself can induce excess volatility. The positive relationship between volume and volatility is well established (Karpoff, 1987), but the canonical explanation attributes this to joint arrival of information (Admati and Pfleiderer, 1988). Bessembinder and Seguin (1992) study the contemporaneous relationship between daily futures activity and stock volatility, finding mixed results—deeper futures markets are associated with lower volatility, while unexpected surges in futures volume are associated with higher volatility. Overnight futures trading did not yet exist during their sample (1978–1989), so their analysis relies entirely on close-to-close returns and cannot isolate the overnight channel. We improve on this by decomposing variance into

overnight and intraday components and comparing systematic versus idiosyncratic variance within the overnight window, since individual stocks cannot trade overnight but the index can via futures. Hasbrouck (2003) shows that E-mini futures dominate price discovery among index instruments. Our overnight decomposition extends the framework of Lockwood and Linn (1990) and Barclay and Hendershott (2003), and connects to Lou et al. (2019), who show that overnight and intraday expected returns are driven by different forces.

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

Finally, Hong and Stein (2007) reviews models in which differences in beliefs generate volume and volatility. Our model falls within this tradition, as the demand shocks driving index trading arise from heterogeneous beliefs about market prospects.

Campbell et al. (1993) provides the demand-shock framework we build on, and Schwert (1989) establishes the benchmark findings on the relationship between stock market volatility and macroeconomic variables that our paper seeks to explain.

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

2. Market Volatility: Broad Facts

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

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

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

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

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

to increase in the future at this linear rate, or that the increase is necessarily linear from 1960-2025.

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

Column 2 uses weekly returns and finds an increase of about 8 percentage points, resulting in a total increase in volatility of about 77% over the sample. Column 3 constructs monthly data using compounded overlapping 4-week returns and finds an increase of about 5 percentage points, resulting in a total increase of around 38%. Thus, while volatility increases substantially across all return horizons, we can see that it is more heavily concentrated in short-horizon returns. Because compounded (log) returns over longer-horizons are equal to the sum of log returns at a daily horizon, the difference in the trend in volatility across horizons necessarily reflects changes in the autocorrelation of returns in the sample. For example, if return autocorrelation at all horizons is zero, then the expectation of squared weekly returns is equal to the expectation of the sum of squared daily returns within a week.

Immediately, the pattern across horizons suggests differences in financial markets and trading behavior itself to explain the increase in volatility. For example, a story of time-varying volatility of fundamentals will not explain the change in volatility patterns across return horizons. One needs an increase in "discount rate" volatility – the volatility of a mean-reverting component of returns – to explain these differential

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

Table 1: Volatility Time Trends by Return Frequency

	Daily (1)	Weekly (2)	4-Week (3)
Constant	9.00	10.75	12.77
	(0.68)	(0.91)	(0.98)
Trend	11.44	8.24	4.85
	(2.03)	(1.94)	(1.99)
p-value	$<$0.001	$<$0.001	0.024
N	3,705	3,704	3,704
R^2	0.517	0.370	0.153
Vol increase	127%	77%	38%

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

patterns. Because these patterns are at relatively short horizons, this naturally suggests candidates based on trading behavior.[5] We will return to explanations based on fundamentals shortly.

2.1 Volatility, Volume, and Fundamentals

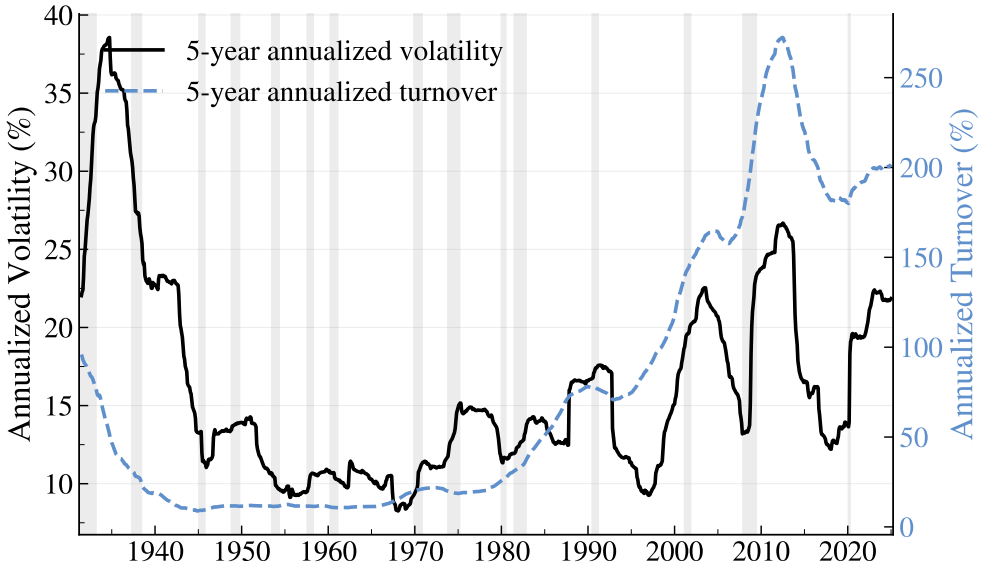


Figure 2: Volatility and Turnover

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

Figure 2 plots market volatility against total stock market volume as a share of total market capitalization over 5-year periods from 1931 to 2025. There is a strong correlation between trading volume and volatility since the 1960s, both in terms of low frequency trend and in the cyclical movements. A natural story is that higher trading activity leads to higher volatility, especially for shorter horizon returns. This

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

fits well with the return-horizon pattern we've shown so far. Of course, there could be other factors affecting both volatility and volume, and in addition high volatility could lead to more demand for trading activity, a reverse causality story. In Section 3.1 we document a causal relation between volume and volatility by using changes in trading hours as an instrument for volume as well as the introduction of the E-mini futures.

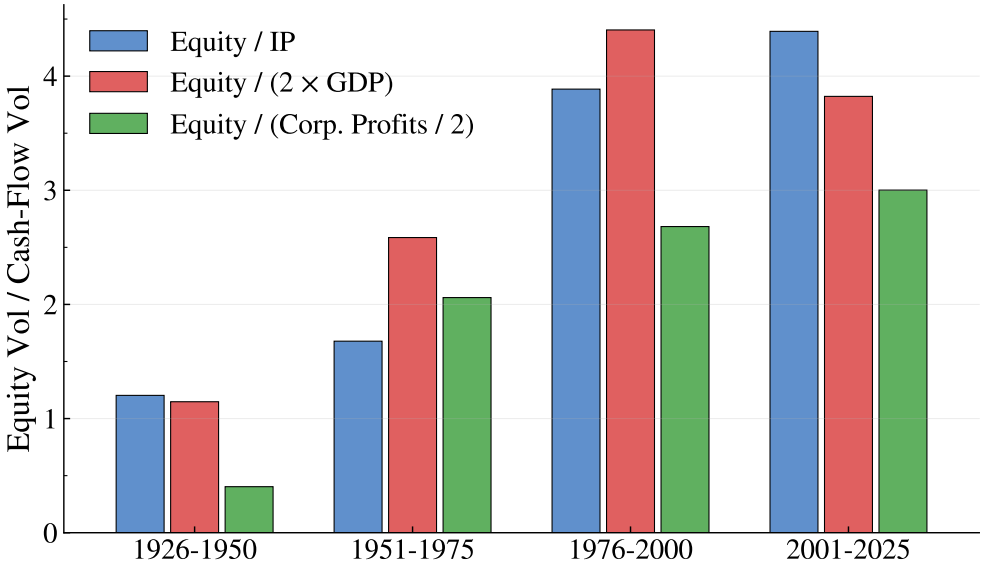


Figure 3: Equity Volatility Relative to Fundamental Volatility

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

Figure 3 documents that the increased volatility in the post-WW2 data is not driven by more volatile macroeconomic fundamentals. Specifically, the figure shows the ratio of annualized daily stock market volatility to annualized macro volatility over the four non-overlapping 25-year periods from 1926 to 2025. We consider three underlying series to measure macro volatility: industrial production growth, real per capita

GDP growth, and corporate profit growth. The first has monthly data going back to 1920 and the two other has quarterly post-WW2 data, which we splice with the annual data going back to 1929. Corporate profits are useful because they deal with changes in corporate leverage (they are a levered cash flow).

The figure shows that the ratio of stock market volatility to macroeconomic volatility was lowest in the early period, slightly increased in the 1951-1975 period, and then strongly increased in the 1976–2000 period and remained high thereafter. In other words, high volatility of fundamentals explains the high stock return volatility in the Great Depression and WWII. However, since then the volatility of macroeconomic fundamentals has declined substantially (through the "Great Moderation" period of the 1980s to mid 2000s). The macroeconomy has, if anything, looked *less* volatile over time – and hence is a poor candidate to explain the rise in volatility we've seen in the past 65 years. This echoes a long literature starting with Schwert (1989) which has difficulty tying variation in stock market volatility to fundamentals. Between these findings and the effects of return horizon, volatility of fundamentals is a poor explanation for return volatility, while trading activity appears a promising explanation.

2.2 Robustness of the rise in volatility

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

2.2.1 Outliers

It is well known that volatility is right skewed with some extreme events, such as the crash of 1987. Thus, our finding of increased volatility could potentially be driven by several extreme events. To address this, we compute volatility within a month using daily returns in that month, then compute the rolling 5-year mean of the 30 months with realized variances in the middle 50% of the distribution out of the 60 months in the period. We then annualize and take the square root to get the dependent variable

in the trend regression. Column (1) of Table 2 shows that the trend is almost identical at 11.3%. Thus, the rise in market volatility is not driven by outliers as events such as the 1987 crash or the Covid crisis have no effect on this alternative measure of volatility.

2.2.2 Trend robustness

Pension funds and other institutional investors were increasingly holding and trading market portfolios starting in the 1960s, with the first index fund (Vanguard) becoming available in 1976. In the early 1980s, index futures and options become available and increasingly traded. It is therefore natural to ask if the start date of the trend regressions are critical for documenting the rise in volatility. Column (2) of Table 2 shows the trend regression with a 1980 start date. The trend variable is the same, so the magnitude of the coefficient is still scaled as a 75-year change to be easily comparable to that in the full postwar sample regression. The rise in volatility with the 1980 starting date is slightly stronger than, but overall very similar to, the full sample trend at a little over 12%.

While we estimate a trend, we do not view it literally but as a convenient way to summarize the slow-moving rise in volatility over the postwar sample. This approach follows how (Campbell et al., 2001) characterize idiosyncratic volatility. In the Appendix, we instead estimate structural breaks in the mean of market volatility and find two breaks dated around 1980 and 2000. In each case, the mean of volatility increases substantially from the earlier mean and the overall increase in the mean is consistent with the trend result.

2.2.3 Concerns about stale prices

Volatility can appear low if there is no trade in a significant fraction of stocks on a given day. To address this, we weight stocks by total dollar volume on a given day, rather than weight by market capitalization. This ensures stocks that don't trade get

zero weight and more liquid, heavily traded stocks get higher weight. Column (3) of Table 2 shows that volume weighting produces a similar trend. Column (4) restricts to the top 100 stocks by market capitalization, which are the most liquid and least susceptible to stale-price issues, and finds a nearly identical increase.

Table 2: Volatility Trend Robustness

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

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

2.2.4 Additional Results on Horizon Effects and Return Autocorrelations

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

consistent with short-run mean reversion or reversals in returns.

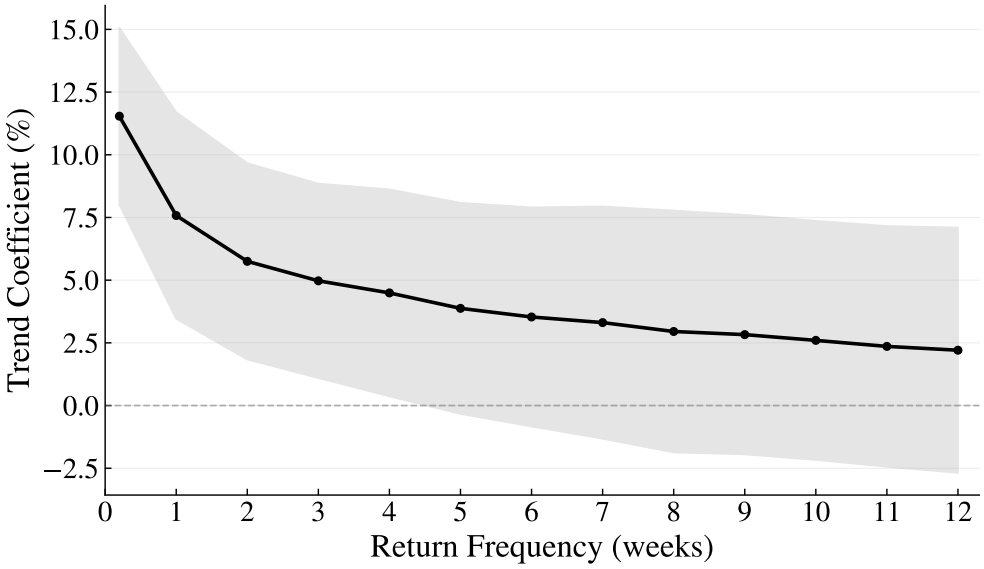


Figure 4: Time Trend in Realized Volatility by Return Frequency

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

Figure 5 shows the rolling ratio of 5-year return variance based on daily returns to 5-year return variance based on monthly returns. This difference is driven by the autocorrelation pattern. If daily returns-based variance is higher than monthly returns-based variance, there is mean reversion in the price of the market. In other words, short-term price pressure that mean-reverts quickly will show up in daily return variance, but not (as much) in monthly return variance. The figure shows the variance ratio over the sample plotted against the daily return volatility from Figure Figure 1. The two track each other quite closely, and the quantitative changes in the variance ratio are large, going from around 0.6 to over 2. Thus, the trading effects this pattern

suggests are quantitatively large.

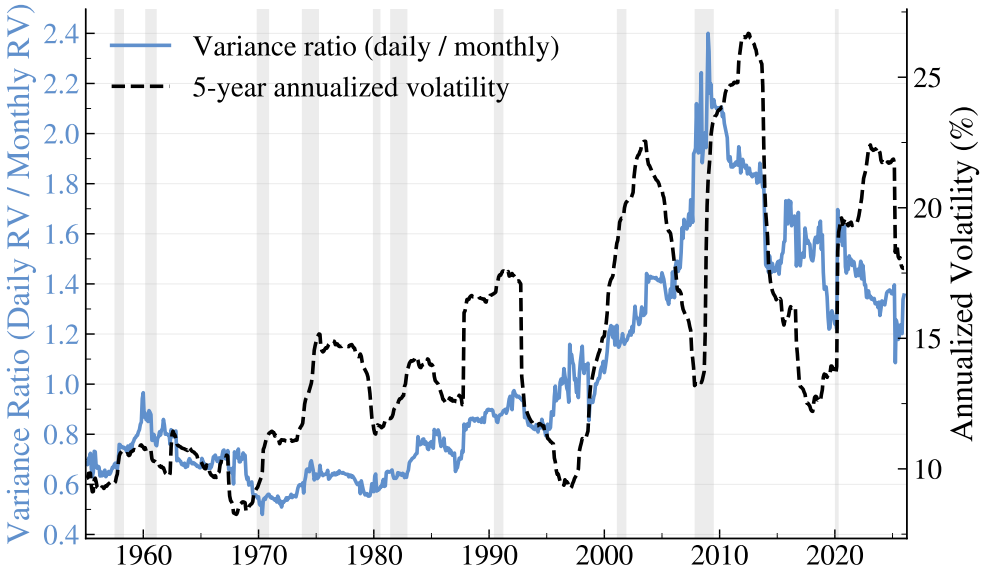


Figure 5: Variance Ratio and Volatility

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

3. Volume and Volatility: Stylized Facts and Causal Evidence

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

This concentration of trading volume in the overall stock market means that in-

vestors facilitating trade – e.g., arbitrageurs and liquidity providers – will have to absorb market risk, rather than idiosyncratic risk of individual stocks. Their compensation is given by the degree of short-term reversal, with higher reversal leading to more volatility of high frequency returns.

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

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

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

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

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

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

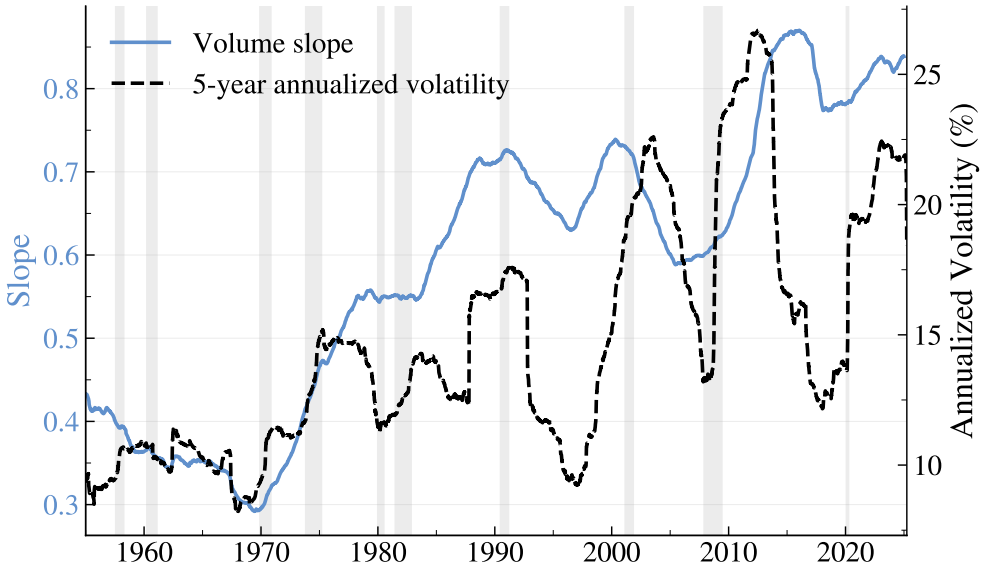


Figure 6: Market Trading vs Market Volatility

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

3.1 Causal Evidence: Volume and Volatility

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

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

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

We take these results and push them further. Specifically, we exploit 8 changes in trading hours on NYSE, 1952-1985. These include: the 1952 cut from 6 to 5 trading days per week (trading used to occur on Saturdays), the 1968 Wednesday closure, and the subsequent extensions of the trading day from 4 hours of trading per day in 1969 to eventually 6.5 hours today.

We present two versions of this analysis. The first, in Panel A of Table 3, aggregates to the weekly level using total trading hours per week as an instrument for weekly volume. This is the simpler specification. The second, in Panel B, works at the daily level and exploits the specific daily predictions of each schedule change, providing substantially more statistical power.

Panel A: Weekly IV. We compute total trading hours per week as our instrument. The first stage runs log total stock market volume on the log of trading hours. The

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

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

$$\ln \sigma_t = a_w + a_e + b \times \ln VM_t + \varepsilon_t \quad (3)$$

$$\ln VM_t = a_w + a_e + \gamma \times \ln \text{trading hours}_t + \nu_t. \quad (4)$$

We use trading hours *only* in 6 month windows before and after a change in trading hours as the instrument for weekly volume (VM) to see how trading activity affects volatility (σ). Restricting to shorter windows around changes in trading hours is important because trading hours have generally trended up over time along with volatility and volume. Further, the levels of volume and of volatility are different around each of these changes. To deal with this, we include event dummies, a_e, (equal to 1 in the 6 months before and after a change in trading hours). We thus estimate how the relative change in trading hours affects relative volume and volatility. In addition, we include week fixed effects, a_w, to control for seasonal patterns in volume and/or volatility.

Our identifying assumptions are as follows. First, increases in trading hours only affect volatility through increased trading activity. Second, the decision to increase trading hours does not depend on an expected increase in volume 6 months before / after the official change in hours. We do not assume that increased demand for trading doesn't affect trading hours in general, but that a decision to permanently increase trading hours does not depend on a predicted spike in volume during the event window. The second assumption is most appropriate when the event window is relatively narrow – our baseline uses 6 months because we face a tradeoff in statistical precision in using realized volatility. Our results are robust to a tighter 3-month window, at the expense of slightly larger standard errors.

Panel A of Table 3 presents the weekly results. Column 3 presents the first stage. Trading hours increase volume nearly one to one – when the stock market is open longer, people trade more. A coefficient of 1 means that if trading hours increase 10%, volume also increases by 10%. The first stage shows that trading hours are a strong instrument for volume, consistent with prior work. This increase in volume translates to increased market volatility, demonstrating a causal effect of trading activity on market volatility. Column 2 finds a positive causal effect of volume on volatility. A 10% increase in trading hours increases volume by 10% and leads to about a 14% increase in volatility. The coefficient of 1.4 in Column 2 is statistically significant at the 5% level. Column 1 presents OLS estimates of volatility on volume and finds a strong positive relationship as well. Column 4 regresses volatility directly on trading hours. Since the first stage coefficient was 1, the coefficient in column 4 is about the same as that in column 2. The relatively similar statistical significance in columns 4 vs 2 highlights that the first stage is strong, but the relationship between changes in trading hours and volatility is slightly noisier.

Panel B: Daily frequency. The weekly specification aggregates each schedule change into a single shift in weekly trading hours. We can gain statistical power by working at the daily level, where each event generates specific predictions about which days are affected. For example, the 1968 Wednesday closure means that Thursdays now span two calendar days instead of one, while Tuesdays are unaffected. The elimination of Saturday trading in 1952 means that Mondays now span the full weekend rather than returning from a half-day Saturday session.

Working at the daily level allows us to separate two distinct components of the time between market closes: *trading hours* (hours the market is open) and *calendar hours* (total hours elapsed from the previous close to the current close). These vary independently across event types. When the NYSE eliminates Saturday trading, Monday's calendar hours jump (the full weekend gap widens) while trading hours are

unchanged. When the NYSE extends closing time from 3:00 to 3:30 PM, trading hours increase but calendar hours remain at 24 for weekdays.

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

$$\ln r_t^2 = a_e + a_d + b \times \ln \text{turnover}_t + \beta_c \times \ln \text{calendar hours}_t + \varepsilon_t \quad (5)$$

$$\ln \text{turnover}_t = a_e + a_d + \gamma \times \ln \text{trading hours}_t + \gamma_c \times \ln \text{calendar hours}_t + \nu_t \quad (6)$$

where r_t^2 is the squared daily market return, turnover is daily dollar volume divided by market capitalization, and log trading hours instruments for log turnover. We include event fixed effects a_e, day-of-week fixed effects a_d, and individual holiday fixed effects (a separate dummy for each post-holiday day, which prevents holidays from contaminating the calendar hours coefficient). We use ± 22 trading day windows (roughly one month) around each schedule change, yielding 9 events and 403 daily observations. Newey-West standard errors with 10 lags account for serial correlation.

Panel B of Table 3 presents the results. The first stage (column 7) is very strong: a 10% increase in trading hours increases turnover by about 13%, with an F-statistic of 57 on the excluded instrument. The IV estimate in column 6 implies an elasticity of daily variance with respect to turnover of 2.3, significant at the 5% level. The OLS estimate in column 5 is similar at 2.6, suggesting minimal attenuation bias at the daily frequency. Column 8 presents the reduced form: log trading hours enters with a coefficient of 2.9, significant at the 5% level, while the coefficient on log calendar hours is 1.4 and statistically insignificant. Under the information-only null, the trading hours coefficient should be zero and the calendar hours coefficient should be one. The data strongly reject this null in favor of trading activity driving variance.

Table 3: IV: Instrumenting volume with trading hours

Notes. Panel A uses weekly data. The dependent variable is log realized volatility (annualized) of the weekly value-weighted market return. Volume is total weekly dollar volume. Trading hours is total trading hours per week. The sample includes 6-month windows around each of the 8 changes in NYSE trading hours, 1952–1985. Panel B uses daily data. The dependent variable is log squared daily market return. Turnover is daily dollar volume divided by market capitalization. Trading hours and calendar hours (hours from previous close to current close) are measured for each trading day. The sample uses ± 22 trading day windows around 9 NYSE schedule changes. Standard errors in parentheses: Newey-West in Panel A (with week and event FEs); Newey-West with 10 lags in Panel B (with event, day-of-week, and individual holiday FEs). ***, **, * denote significance at the 1%, 5%, and 10% level.

Panel A: Weekly

	(1) OLS	(2) IV	(3) First Stage	(4) Reduced Form
	σ on VM	σ on VM	VM on Hours	σ on Hours
Log Volume	0.72***	1.42**		
	(0.12)	(0.72)		
Log Trading Hours			1.02***	1.45**
			(0.31)	(0.67)
Observations	326	326	326	326
R-squared	0.54	0.49	0.90	0.49
FE: Week & Event	Y	Y	Y	Y

Panel B: Daily

	(5) OLS	(6) IV	(7) First Stage	(8) Reduced Form
	r^2 on Turnover	r^2 on Turnover	Turnover on Hours	r^2 on Hours
Log Turnover	2.59***	2.32**		
	(0.51)	(0.94)		
Log Trading Hours			1.27***	2.94**
			(0.17)	(1.21)
Log Calendar Hours	-0.10	0.16	0.55***	1.44
	(1.13)	(1.27)	(0.13)	(1.11)
Observations	403	403	403	403
R-squared	0.37		0.92	0.32
First-stage F		57.0		
FE: Event, DOW, Holiday	Y	Y	Y	Y

Overall, the evidence strongly supports a positive causal relationship whereby increased trading activity drives significantly higher volatility of the overall stock market. The weekly specification provides a simple and transparent test; the daily specification sharpens identification by exploiting the specific daily predictions of each schedule change and delivers a first-stage *F*-statistic well above conventional thresholds for instrument strength.

3.1.2 The E-mini Futures as a Natural Experiment

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

We use close to open and open to close returns as measures of overnight and intraday returns, respectively.[7] We compute the sum of squared returns within a rolling 22 day period (roughly 1 month). For individual stocks, we aggregate into a single measure of volatility by value weighting the resulting overnight and intraday variances at the stock level. We then compute $\left(var(r_{mkt}^{night})/var(r_{mkt}^{day})\right) / \left(var(r_{stock}^{night})/var(r_{stock}^{day})\right)$. This represents the ratio of the two variance ratios.

Figure 7 plots the results. Panel A shows the increase in this variance ratio. The first vertical line represents the introduction of the E-mini in September 1997. Since

[7] Thus, we also use weekend returns.

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

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

Table 4 tests the increase formally. Because the dependent variables are 22-day rolling averages, a simple pre/post dummy would misrepresent the timing of the shift — on September 10 only 1 of 22 observations in the rolling window reflects the post-launch regime. We therefore use a ramp variable that equals zero before the E-mini launch, increases linearly from 0 to 1 over the next 22 trading days, and equals one thereafter. The dependent variable in column (1) is the log ratio of overnight-to-intraday variance for the S&P 500 index, controlling for the corresponding log ratio for individual stocks. The ramp coefficient of 0.52 ($t = 8.7$) implies the index variance ratio increased by a factor of $e^{0.52} \approx 1.7$ after the E-mini launch, holding stock-level patterns fixed. Column (2) confirms that log futures volume increased by 0.60 over the same window.

3.1.3 The Rise of Futures and Overnight Variance

More generally, the volume of the E-mini futures has exploded since it was introduced, leading to a huge amount of trading volume in the market index. Because

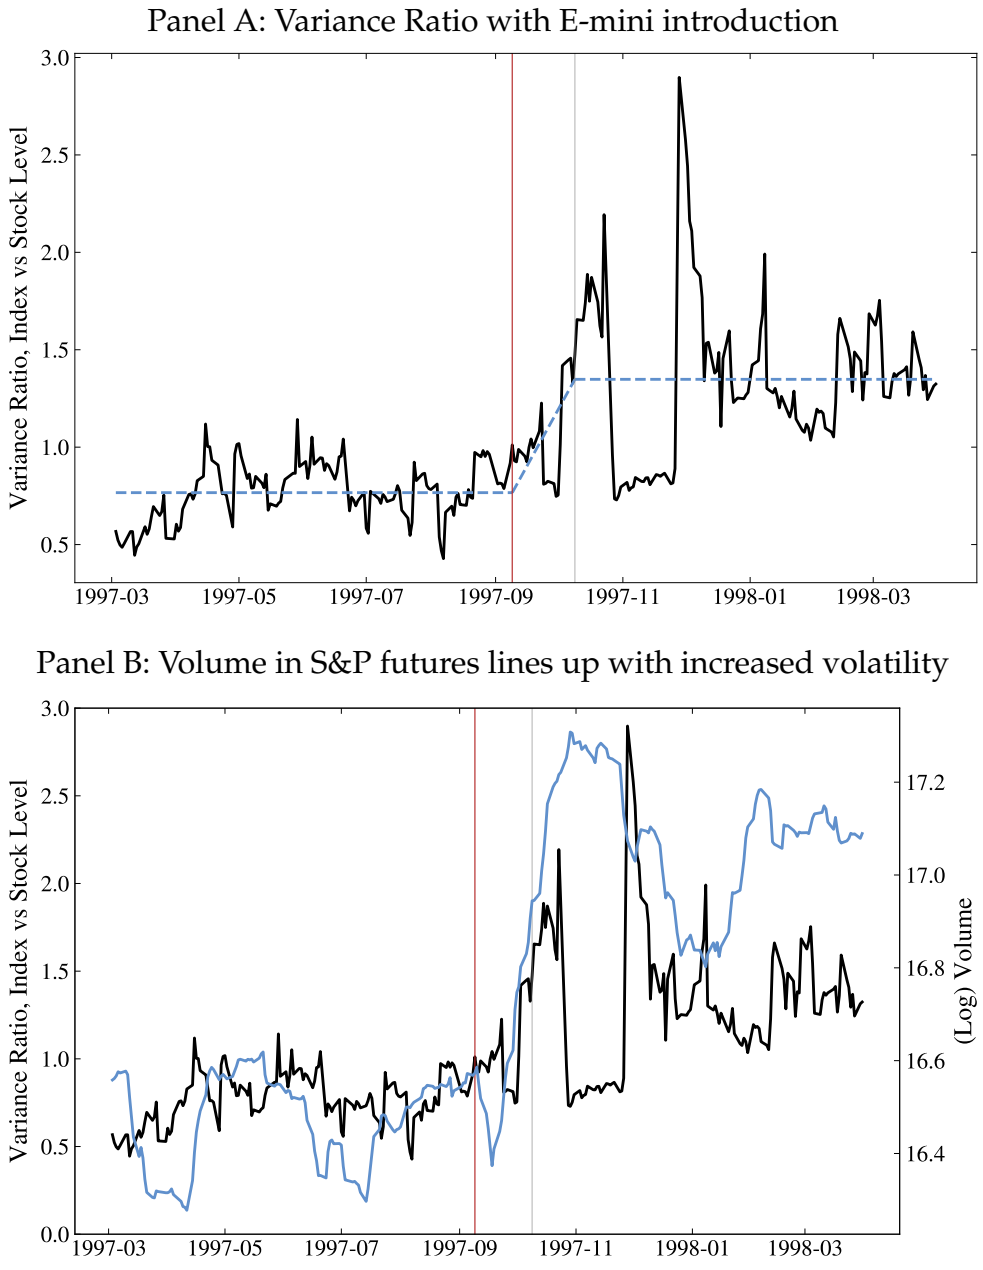


Figure 7: The Introduction of the E-mini Futures and the Effects on Volatility.

Panel A plots 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). Panel B adds total volume in S&P500 futures contracts (22-day rolling average, log scale).

Table 4: E-mini Futures Introduction and Overnight Volatility

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

	(1) Log Index Var Ratio	(2) Log Volume
E-mini Ramp	0.520***	0.599***
	(0.060)	(0.048)
Log Stock Var Ratio	1.486***	
	(0.340)	
Observations	273	273
R-squared	0.712	0.839

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

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

But is this rise in overnight volatility driven by the kind of trading our theory predicts? We now decompose overnight returns for individual stocks into a systematic component (driven by the market) and an idiosyncratic component (stock-specific).

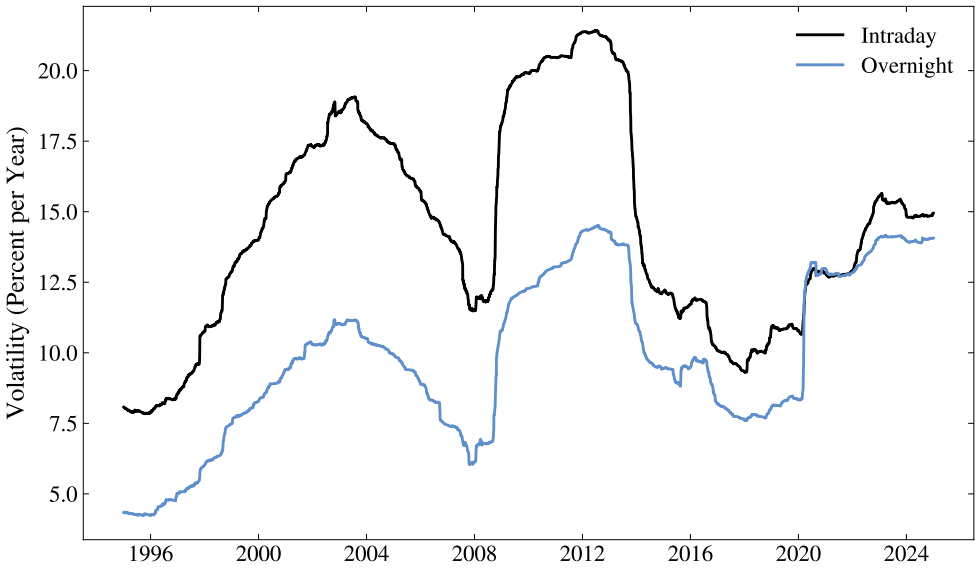


Figure 8: The Rise of Overnight Volatility on the Market.

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

Our prediction is sharp: because index futures allow investors to trade the market overnight but not individual stocks, the systematic component of overnight variance should rise over time while the idiosyncratic component should not.

Setup. We take the 500 largest U.S. common stocks by market capitalization, re-selected at the start of each year from 1993 to 2024, yielding a panel of 1,358 unique stocks and 3.5 million stock-day observations. For each stock i in year t, we estimate an overnight beta $\hat{\beta}_{i,t}$ by regressing the stock's overnight return (close-to-open) on the market overnight return (SPY close-to-open), using a rolling 3-year window of prior data to avoid look-ahead bias.

We then decompose each overnight return into systematic and idiosyncratic com-

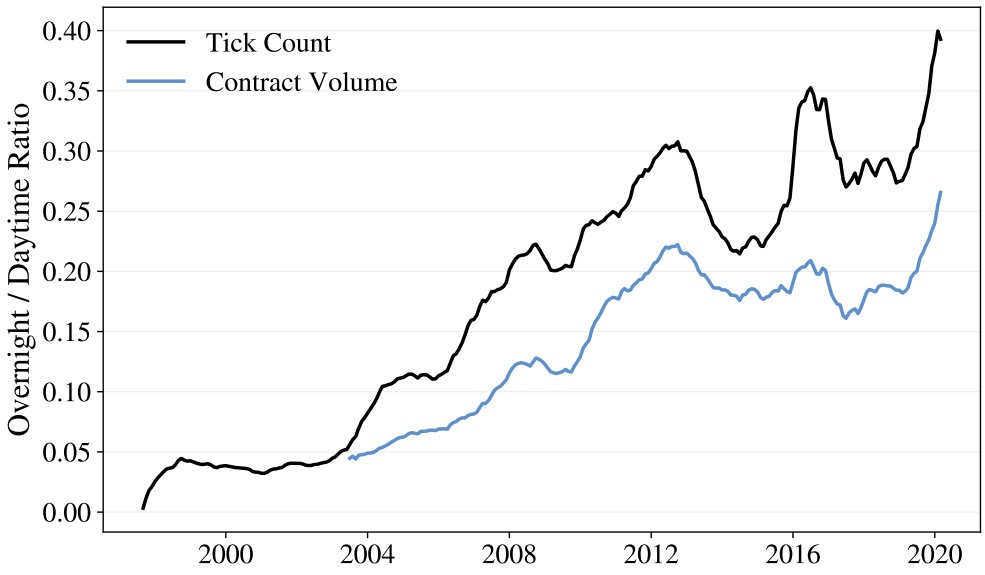


Figure 9: Overnight-to-Daytime Trading Activity in S&P 500 Futures

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

ponents:

$$r_{i,t}^{\text{ovn}} = \underbrace{\hat{\beta}_{i,t} \cdot r_{m,t}^{\text{ovn}}}_{\text{systematic}} + \underbrace{\hat{\varepsilon}_{i,t}}_{\text{idiosyncratic}} \tag{7}$$

Our dependent variable is the squared component normalized by the number of non-trading hours in the overnight gap: $(\hat{\beta}_i \cdot r_{m,t}^{\text{ovn}})^2 / h_t$ for systematic variance per hour and $\hat{\varepsilon}_{i,t}^2 / h_t$ for idiosyncratic variance per hour. We normalize by hours so that weekday overnights (17.5 hours), weekends (65.5 hours), and holidays (varying) are comparable.

Identification. We exploit a within-day-type comparison. On a regular weekday, the overnight gap is 17.5 hours. On a weekend, it stretches to 65.5 hours. In both cases, individual stocks cannot be traded overnight. But the *market* can be traded overnight via index futures, and futures trade more hours on weekdays than weekends. If futures trading drives overnight variance, the systematic component of overnight variance per hour should be higher on weekdays—where more futures trading occurs—than on weekends.

Moreover, this weekday-weekend gap in systematic variance should *grow over time* as overnight futures trading has expanded, from zero hours before 1982, to a few hours in the late 1980s, to nearly 24 hours today. The idiosyncratic component, in contrast, should show no such trend, because individual stocks remain untradeable overnight regardless of era.

Trend specification. The steady upward trends visible in Figures 8 and 9—rising overnight volatility alongside rising overnight futures activity—motivate a trend-based specification that directly connects to the central question of this paper: what is driv-

ing the long-run increase in market volatility? We estimate:

$$\frac{y_{i,t}^2}{h_t} = \alpha_i + \gamma_1 \cdot \text{Weekday}_t + \gamma_2 \cdot \text{Year}_t + \gamma_3 \cdot \text{Weekday}_t \times \text{Year}_t + \varepsilon_{i,t} \tag{8}$$

where $y_{i,t} \in \{\hat{\beta}_i \cdot r_{m,t}^{\text{ovn}},\ \hat{\varepsilon}_{i,t}\}$ denotes either the systematic or idiosyncratic component of the overnight return, with stock fixed effects and standard errors double-clustered by stock and date. Column (1) of Table 5 reports results for the systematic component and Column (2) for the idiosyncratic component.

The coefficient γ_3 captures the differential trend in weekday versus weekend overnight variance. For the systematic component, $\gamma_3 = 0.05$ ($t = 6.22$), indicating that systematic overnight variance per hour has been trending up significantly faster on weekdays than on weekends. For the idiosyncratic component, the trend is *negative* and smaller in magnitude ($\gamma_3 = -0.02$, $t = -3.51$), consistent with the idea that idiosyncratic risk has if anything declined while overnight trading has expanded.

Triple-difference test. Column (3) tests whether this difference is statistically significant. We stack the systematic and idiosyncratic observations, doubling the sample to 6.98 million, define an indicator $S_{i,t} = 1$ for the systematic component, and estimate:

$$\begin{aligned}\frac{y_{i,t}^2}{h_t} = \alpha_i^c + \gamma_1 \cdot \text{Weekday}_t + \gamma_2 \cdot \text{Year}_t + \gamma_3 \cdot \text{Weekday}_t \times \text{Year}_t \\ + \delta_1 \cdot S \times \text{Weekday}_t + \delta_2 \cdot S \times \text{Year}_t + \delta_3 \cdot S \times \text{Weekday}_t \times \text{Year}_t + \varepsilon_{i,t}\end{aligned} \tag{9}$$

where α_i^c denotes stock $\times$ component fixed effects. The base terms (γ coefficients) capture patterns in the idiosyncratic component; the δ coefficients capture how the systematic component *differs*. The triple interaction $\delta_3 = 0.07$ ($t = 8.91$) confirms that the differential weekday trend is significantly larger for systematic than idiosyncratic overnight variance.

Futures hours specification. Columns (4)–(6) replace the continuous year trend with a direct measure of overnight futures trading hours, which varies across years and across weekdays versus weekends based on the actual history of S&P 500 futures and Globex trading schedules. This specification includes year fixed effects, so identification comes purely from within-year variation in futures hours across day types. The coefficient on Weekday $\times$ FuturesHrs for the systematic component is 0.30 ($t = 9.50$), more than three times the idiosyncratic coefficient of 0.09 ($t = 3.57$). The triple-diff in column (6) confirms this difference ($t = 6.95$).

Economic magnitude. The overnight futures channel alone implies a roughly 58% increase in annual market variance over this 30-year period. To see this, note that the Weekday $\times$ Year coefficient of 0.05 (in units of $10^6 \times r^2$/hour/year) implies that each year, systematic overnight variance per hour grows by an additional 0.05 on weekdays relative to weekends. Over 30 years and 17.5 weekday overnight hours, this cumulates to an extra 0.27 basis points of systematic variance per weekday night. Aggregated over 260 weekday nights per year, this amounts to roughly 70 basis points of additional annual market variance. Relative to a baseline annual market variance of approximately 121 basis points (corresponding to 11% annualized volatility in the early 1990s), this represents a 58% increase—accounting for about half the total rise in market variance over this period.

4. A model of index investing and market returns

In this section, we present a simple model of the impact of increased index investment on market outcomes. The model borrows features from Merton (1987) and Campbell et al. (1993) (henceforth, CGM).

Following Merton, we assume that each investor only knows about a subset of the financial assets and therefore only trades these. This assumption is motivated by the

Table 5: Overnight Variance Decomposition: Systematic vs. Idiosyncratic

Notes: The dependent variable is the squared overnight return component per non-trading hour, multiplied by 10^6 (i.e., basis points $\times$ 100). Overnight returns are decomposed into systematic ($\hat{\beta}_i \cdot r_m^{\text{ovn}}$) and idiosyncratic ($\hat{\varepsilon}_{i,t}$) components using rolling 3-year betas. Columns (1)–(3) use a continuous year trend; columns (4)–(6) replace the trend with overnight futures hours and add year fixed effects. Columns (3) and (6) stack the systematic and idiosyncratic observations; the Sys $\times$ interaction tests whether the effect differs across components. Weekday = 1 for single-day overnight gaps, 0 for weekends and holidays. Sample: 500 largest U.S. stocks by market cap (re-selected annually), 1994–2024, yielding 1,358 unique stocks and 3.49 million stock-day observations. Standard errors double-clustered by stock and date in parentheses. */**/*** denote significance at 10/5/1%.

	Panel A: Trend			*Panel B: Futures Hours*		
	(1) Sys	(2) Idio	(3) Triple-Diff	(4) Sys	(5) Idio	(6) Triple-Diff
Weekday	1.34***	2.97***	2.97***	−3.27***	1.53***	1.72***
	(0.07)	(0.07)	(0.07)	(0.44)	(0.39)	(0.41)
Year	0.01	−0.02***	−0.02***			
	(0.00)	(0.00)	(0.00)			
Weekday × Year	0.05***	−0.02***	−0.02***			
	(0.01)	(0.01)	(0.01)			
Weekday × FutHrs				0.30***	0.09***	0.08***
				(0.03)	(0.03)	(0.03)
Sys × Weekday			−1.63***			−5.18***
			(0.08)			(0.50)
Sys × Year			0.03***			
			(0.01)			
Sys × Wkday × Year			0.07***			
			(0.01)			
Sys × Wkday × FutHrs						0.23***
						(0.03)
Stock FE	Y	Y	Y	Y	Y	Y
Year FE				Y	Y	Y
N	3.49M	3.49M	6.98M	3.49M	3.49M	6.98M
Stocks	1,358	1,358	1,358	1,358	1,358	1,358

fact that investors empirically only hold a few assets.[8] When an index fund is introduced, these investors shift their attention and investment to this more efficient asset. Following CGM, risk averse arbitrageurs accommodate buying and selling pressure arising from investors' demand shocks.

4.1 Market structure

There are N assets with per share dividends $D_{it} \sim N\left(\mu, \sigma_D^2\right)$. The correlation between any two assets' dividends is ρ_D. Stock i is in supply x_i. There are N investors and 1 arbitrageur. Every agent acts competitively as a price taker. The investors have extrapolative beliefs, where their beliefs about future dividends are extrapolated based on the dividends of the assets they hold. The arbitrageur is rational and can frictionlessly invest in all assets. The risk-free rate has gross return $R_f > 1$ and is supplied elastically. Each agent has CARA utility over next period wealth, which implies that they are one-period mean-variance optimizers. Investors have risk aversion a, while the arbitrageur has risk aversion $\tilde{a}$.

Let P_{it} refer to the share price of the claim with dividends D_{it}, and let D_t and P_t be the $N \times 1$ vectors of the corresponding dividends and prices for all individual stocks. Let X denote the $N \times 1$ vector of supply with row i equal to x_i, normalized such that $\sum_{i=1}^{N} x_i = 1$. The value of the market portfolio is the sum of all stock prices times the supply, $P_{mt} = X'P_t = \sum_{i=1}^{N} P_{it}x_i$, while market dividends are $D_{mt} = X'D_t = \sum_{i=1}^{N} D_{it}x_i$.

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

4.1.1 Asset market without an index fund

In this economy, each investor is assumed to only invest in one stock in addition to the risk-free asset. For notational simplicity, we assume investor i invests in stock i. For the asset they hold, investor i believes that:

$$E_t^i\left(D_{it+1}\right) = \mu\left(1-\phi\right) + \phi D_{it}, \tag{10}$$

where $1 > \phi > 0$. That is, they erroneously extrapolate from past dividends of the asset they hold when forming beliefs about future dividends. This time-variation in beliefs creates demand shocks that are not due to information. The rational arbitrageurs correctly believe expected dividends equal μ.

Proposition 1. *The equilibrium prices in the case where no index fund is traded are given by*

$$P_t = m + AX + BD_t, \tag{11}$$

where

$$A = \begin{bmatrix} \alpha & \delta & \cdots & \delta \\ \delta & \alpha & \cdots & \delta \\ \vdots & \vdots & \ddots & \vdots \\ \delta & \delta & \cdots & \alpha \end{bmatrix}, \quad B = \begin{bmatrix} \beta & \gamma & \cdots & \gamma \\ \gamma & \beta & \cdots & \gamma \\ \vdots & \vdots & \ddots & \vdots \\ \gamma & \gamma & \cdots & \beta \end{bmatrix}, \tag{12}$$

$m \propto \mu$, $\alpha < \delta < 0$, *and* $\beta > \gamma > 0$. *Further,* β *and* γ *are increasing in the degree of extrapolation,* ϕ. *See the Appendix for derivations.*

The equilibrium prices are lower when the supply of the asset is higher and, to a lesser extent, if the supply of other assets is higher. This is because higher asset supply means more risk has to be borne in equilibrium. Further, the extrapolative beliefs of the investors imply assets temporarily get over- or undervalued depending on the level of current dividends. The arbitrageur does lean against investors' demand shocks, but only partially accommodate them due to risk aversion. The fact that $\gamma > 0$

means asset i's price increases if asset j experiences positive demand shocks. The high price of asset i in this case comes from arbitrageur demand. The high investor demand for asset j means the arbitrageur will decrease their position in asset j. However, it is then, all else equal, optimal to increase their positions in the other risky assets to maintain an optimal amount of risk taking.

As in CGM, these demand-shocks lead to trading volume and autocorrelation in asset returns. While the extrapolation is based on firm-specific dividends, this bias also affects the market portfolio as dividends are correlated. The excess returns on the market claim are

$$R^e_{mt+1} = P_{mt+1} + D_{mt+1} - R_f P_{mt}. \tag{13}$$

Let C_D be the dividend correlation matrix, with ρ_D in the off-diagonal entries. The appendix shows that the market return variance and autocorrelation are given by

$$Var\left(R^e_{mt+1}\right) = X'\left(B + I_N\right)\sigma^2_D C_D\left(B + I_N\right)X + R^2_f X' B\sigma^2_D C_D B X, \tag{14}$$

and

$$Corr\left(R^e_{mt}, R^e_{mt+1}\right) = -R_f \frac{X' B\sigma^2_D C_D\left(B + I_N\right)X}{X'\left(B + I_N\right)\sigma^2_D C_D\left(B + I_N\right)X + R^2_f X' B\sigma^2_D C_D B X}.$$

The variance is increasing and the autocorrelation is decreasing in the degree of extrapolation, as this makes B "bigger." This is intuitive, and consistent with the analysis in CGM, as the demand shocks in this case temporarily push prices further away from fundamental value.

Finally, we consider the average fraction of market risk in individual asset returns:

$$\bar{R}^2 = \frac{\frac{1}{N}\sum_{i=1}^{N} Cov\left(R^e_{mt}, R^e_{it}\right)^2}{Var\left(R^e_{mt}\right) Var\left(R^e_{it}\right)}, \tag{15}$$

and trading volume:

$$dv_t = \left|\left(X^A_t - X^A_{t-1}\right) \odot P_t\right|, \tag{16}$$

where dv denotes the dollar volume vector across assets and where $\odot$ is the Hadamard product (element-by-element multiplication).[9] In the empirical section, we show that, relative to an asset market with more index trading, the level of market risk in asset returns is relatively low. Further, we also show that the cross-sectional regression

$$dv_{it} = a_t + b_t \frac{x_i P_{it}}{P_{mt}} \sum_{i=1}^{N} dv_{it} + \varepsilon_{it}, \tag{17}$$

yields a relatively low slope coefficient, b_t. In fact, the appendix shows that b_t is lower when, in the cross-section, larger-weight names have lower per-dollar-of-market weight share of aggregate volume. That is, when

$$q_{it} = \frac{dv_{it}}{\sum_{i=1}^{N} dv_{it}} \frac{\sum_{j=1}^{N} x_j P_{jt}}{x_i P_{it}} \tag{18}$$

is lower. If all investors trade the index, the share of total volume equals the market share of the firm and so in this special case $q_{it} = 1$. This economy, where investors trade individual assets with corresponding strong idiosyncratic components to their demand shocks, is a case where typically $q_{it} << 1$ for a large stock i and thus b_t is low.

4.1.2 Asset market with an index fund

We now introduce an index fund that buys the market index. Aggregate market dividends are:

$$D_{mt} = \sum_{i=1}^{N} x_i D_{it}. \tag{19}$$

The mean market dividends is thus μ and the variance is $\sigma^2_{Dm} = X' C_D X \sigma^2_D$. The arbitrageurs offer the index fund to investors competitively with no fees. It is ex ante

[9] As is well-known, prices can go negative in the CARA-Normal setup, which is awkward for the interpretation of the volume calculations. A sufficiently high mean level of dividends, μ, makes the probability of this negligible.

optimal for the investors, who only can trade one asset, to now trade the fund instead of an individual asset.

The investors again have an extrapolation-bias based on the dividends of the asset they hold. In this case, that is the index fund, so their expectation about future market dividends are given by

$$E_t^i\left(D_{mt+1}\right) = \mu\left(1-\phi\right) + \phi D_{mt}, \tag{20}$$

for all investors i.

Proposition 2. *The price of the market in the economy with an index fund is*

$$P_{mt} = a_m + b_m D_{mt}, \tag{21}$$

where $a_m = \frac{(b_m+1)\mu\left(1-\frac{N\phi}{N+a/\tilde{a}}\right)-\frac{a\sigma_{Dm}^2}{N+a/\tilde{a}}}{R_f-1}$ *and* $b_m = \frac{\phi}{R_f(N+a/\tilde{a})N^{-1}-\phi}$.

It is immediate that b_m is increasing in the level of extrapolation, ϕ, and arbitrageur risk aversion, $\tilde{a}$. This economy differs from the case without an index fund in two critical ways. First, all trade is along the market dimension, which immediately implies that $b_t = 1$ in the cross-sectional volume regressions. Further, investors trade more aggressively in the index fund than they did in individual assets as the Sharpe ratio of the index fund is (on average) higher than that of individual stocks due to better diversification. This forces the arbitrageur to take on more risk from the correlated component of the demand shocks, which they again only partially accommodate due to risk aversion. In sum, when an index fund is introduced the market portfolio return becomes more sensitive to the extrapolation bias.

The variance of these excess market returns is:

$$\sigma_m^2 = \left(\left(b_m+1\right)^2 + R_f^2 b_m^2\right)\sigma_{Dm}^2. \tag{22}$$

Since b_m is relatively high, market returns are more volatile than in the no-index-fund-

Table 6: Model Calibration

The table gives the main parameters of the model calibration. Here, μ is the mean of dividends, ρ_D is the dividend correlation, N is the number of shares and investors, $\tilde{a}$ is the arbitrageur risk aversion, R_f is the gross risk-free rate, a is the investor risk aversion, σ_D is the dividend volatility, and ϕ is the extrapolation bias parameter.

Model Calibration			
μ	0.05	R_f	1.02
ρ_D	0.20	a	1
N	503	$\tilde{a}$	0.05
σ_D	0.18	ϕ	0.60

case. The autocorrelation is

$$Corr\left(R^e_{mt+1}, R^e_{mt}\right) = -R_f \frac{b_m\left(b_m+1\right)}{\left(b_m+1\right)^2 + R_f^2 b_m^2}, \tag{23}$$

which for the same reason is more negative than in the no-index-fund-case.

From the equilibrium expressions for the individual market prices, we infer that average systematic risk, $\bar{R}^2$, increases when the index fund is introduced as asset returns are more highly correlated with the market. See the appendix for derivations.

4.2 Calibration and results

To get a sense of quantities, we calibrate the model and report key moments across the two economies. We let there be 503 firms, consistent with the current number of firms in the S&P500, with supply set equal to the market weight of each stock so that we mimic the size distribution. Table 6 gives the parameters of the model. The key parameter is ϕ, which is set to 0.6.

Figure 10 shows the standard deviation of market returns across the two economies. In the case without an index fund, the standard deviation is 13%, whereas in the case

with an index fund the standard deviation is 22%. This is a substantial increase and on par with what we find in the data.

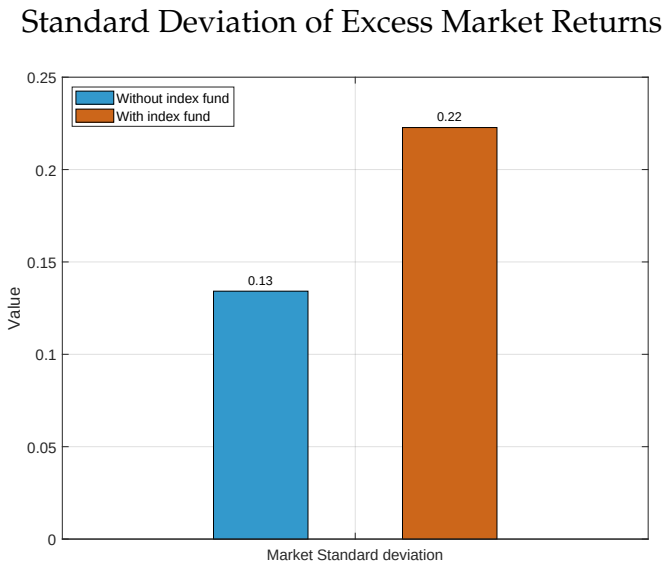


Figure 10: Model-implied Market Volatility.

The figure shows the standard deviation of excess market return for the economy without an index fund and the economy with an index fund.

Next, Figure 11 shows the change in the autocorrelation of excess market returns, the share of average systematic risk, and the slope coefficient on the volume regression. The changes in these are -0.12, 0.24, and 1.00, respectively. Again this is of similar magnitudes to those in the data, although the model is of course too simple to account for all empirical values. For instance, the volume regression in the index fund case yields $b_t = 1$ since all trade is along the market dimension, but this is not true in the data. That said, empirically, we do find this coefficient doubles from about 0.4 to 0.8 over the sample we consider.

In sum, a simple model based on well-accepted concepts from classic papers delivers both qualitative and quantitative predictions that line up with our main empirical arguments. Increased index trading increases market return variance, but decreases

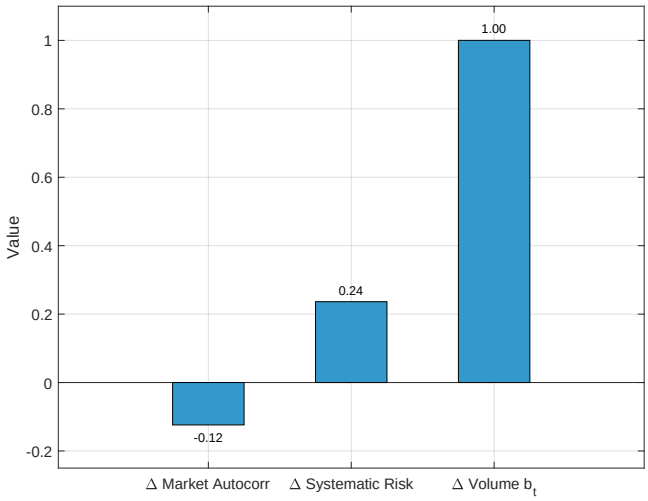


Figure 11: Change in Autocorrelation, Systematic Risk, and Volume Regression Slope.

The figure shows the change in key market outcomes as the economy goes from no index fund traded to having an index fund traded. The moments are the change in the autocorrelation of excess market returns, the average amount of systematic (market) risk in individual stock returns, and the slope coefficient in the volume regressions discussed in the main table.

market autocorrelation so the volatility effect is less pronounced at lower return frequencies. Index trading further increases the fraction of systematic risk in individual stock returns and lines volume up more strongly along the market dimension.

5. Conclusion

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

References

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

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

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

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

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

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

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

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

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

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

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

Campbell, John Y, Martin Lettau, Burton G Malkiel, and Yexiao Xu, 2022, Idiosyncratic equity risk two decades later .

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

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

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

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

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.

Gabaix, Xavier, and Ralph S.J. Koijen, 2024, In search of the origins of financial fluctuations: The inelastic markets hypothesis, *Econometrica* 92, 1347–1381.

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

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

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

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

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

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

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

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

Koijen, Ralph S.J., and Motohiro Yogo, 2019, A demand system approach to asset pricing, *Journal of Political Economy* 127, 1475–1515.

Koudijs, Peter, 2016, The boats that did not sail: Asset price volatility in a natural experiment, *The Journal of Finance* 71, 1185–1226.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

A. Appendix

A.1 Systematic vs. Idiosyncratic Volatility

The rise in stock market volatility is primarily at the market level rather than the individual stock level. That is, the majority of the rise we have documented is market risk and not idiosyncratic risk.

To show this, we run CAPM regressions for the universe of stocks in CRSP on the value-weighted excess market return using rolling 5-year windows. We compute the share of idiosyncratic risk as the ratio of the variance of the residual in this regression to the total variance of the individual stock return. The systematic share is one minus this ratio (or the R-squared of the regression). We aggregate this up to a single measure by value-weighting these shares for all stocks.

Figure 12 plots the results. The systematic share rises from 19% in 1960 to around 35% today, while the idiosyncratic share falls by the same amount. This roughly equates to a doubling of systematic risk over the sample. Thus, we conclude that the majority of the rise in volatility we document is coming from market risk but that this pattern is not equally mirrored in idiosyncratic risk.

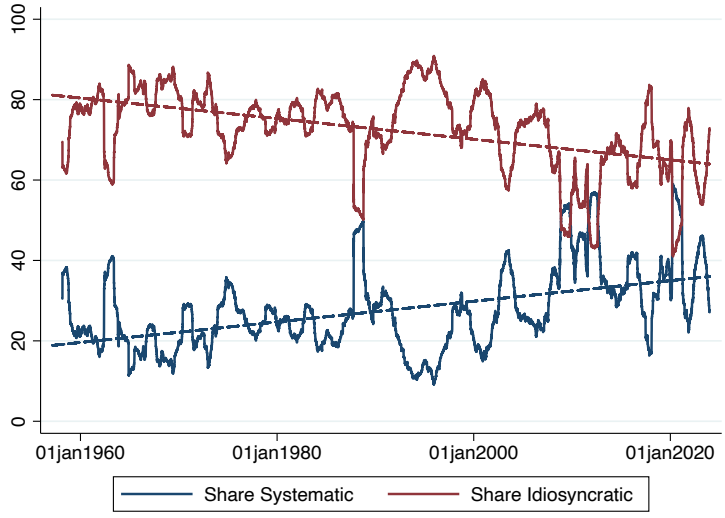


Figure 12: The share of systematic vs idiosyncratic risk.
Value-weighted average of R-squared from rolling 5-year CAPM regressions for individual stocks. The systematic share is the R-squared; the idiosyncratic share is one minus R-squared.

These results are consistent with idiosyncratic volatility having changed – they show that systematic volatility has changed by more (Campbell et al., 2001). They are consistent with updates to this work in Campbell et al. (2022) and Chiah et al. (2020)

that market volatility has increased by substantially more than idiosyncratic volatility in the last 50 years.

A.2 Robustness to the trend specification: Structural breaks

As an alternative specification of slow-moving market volatility dynamics, we here estimate a structural breaks model for the daily 5-year stock volatility series in Figure 1. Figure 13 shows that an estimation that allows for two structural breaks date these breaks to around 1975 and 2000. The increase in mean volatility over the sample is about the same as for the trend specification.

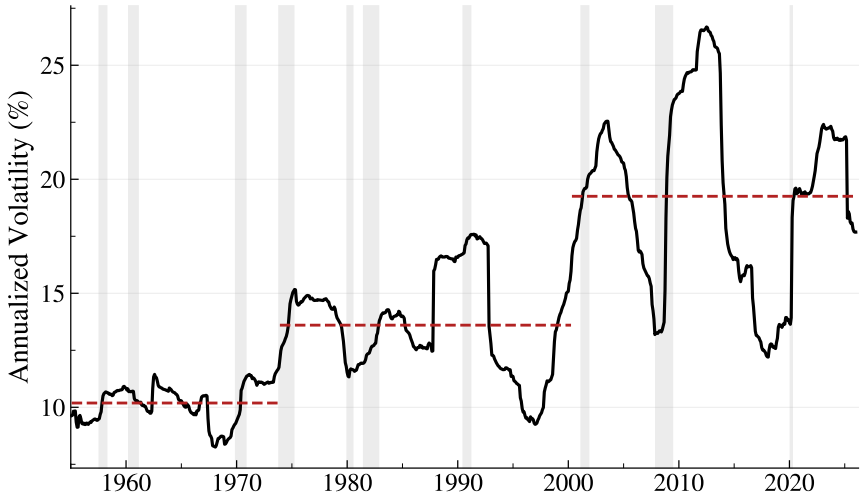


Figure 13: Structural Breaks in 5-Year Rolling Market Volatility

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