Reading VIX Against Itself

Austin Alhajj*
Massari Asset Management

October 2026
Massari Research Note No. 1 · Version 1.0

Abstract

Many tactical allocation strategies hold risk-on assets, such as equities, only while VIX looks calm. Calm can be judged against a fixed level or against VIX's own recent history, and this paper compares the two on daily VIX and S&P 500 closes from 1990 to 2026. Matched on how often they are on, a fixed level and a trailing z-score had been off before about the same number of the 28 VIX spikes in the sample. They part ways on the way out. After long spikes the trailing filter was back on within weeks and the fixed level within months or years, and the trailing filter switched about three times as often. A z-score also carries a mechanical limit: while a large spike is still in the window and VIX has settled back to a steady level, the reading can only get so calm, and the limit is set by how much of the window the spike fills. In practice, having the spike in the window rarely decided anything. In 19 short spikes it added about a month twice to the wait of the strictest filter, a z-score below −1, and in two others it actually brought the filter back sooner. The window length sets where a trailing filter falls between the two references. No strategy performance is reported.

Keywords: VIX; implied volatility; regime filters; rolling z-scores; robust scale estimation; volatility timing
JEL classification: C22, G11, G13, G17

Key findings

- Equal time on is the fair comparison. On that basis, fixed and trailing filters stepped aside before spikes about equally often, and the real difference is how fast each comes back.
- Coming back faster costs switching. A 20-day trailing filter changes state roughly 40 times a year, against about 13 for a fixed level. A band between two z-score thresholds cuts that to about 29, but a fixed level given its own band switches only about 5 times a year.
- A spike sitting in a z-score's window both helps and hurts re-entry. In 12 of 19 short spikes it made no difference to when the strictest filter, a z-score below −1, came back.
- Longer windows push a trailing filter toward a fixed one. At 250 days it comes close to a fixed level on re-entry and switching.

*Austin Alhajj is the founder and chief investment officer of Massari Asset Management, a Florida-registered investment adviser. The views expressed are the author's own. This paper is for informational and educational purposes only; see the disclosures at the end.

1 Introduction

In August 2024 VIX closed at 38.57, up from 23.39 the session before. Two weeks later it was back at 14.65, a little under its average for the month before the spike. A filter that holds equities while VIX is below its 20-day average was back on six trading days after the peak. A filter keyed to a 20-day z-score below −1 stayed off until late September, 35 trading days after the peak. Both filters read VIX against its own recent history. The difference between them comes from one design choice: dividing by recent dispersion. This paper works out what that choice does and how often it has mattered.

Tactical allocation strategies, discretionary and systematic, often use VIX this way. They hold risk-on assets such as equities when implied volatility is low relative to some reference and cut them otherwise. Published examples range from practitioner tests (Hanna 2008) and academic allocation studies (Cloutier, Djatej and Kiefer 2017) to index products that adjust equity exposure using realized volatility and the trend in implied volatility (S&P Dow Jones Indices 2024).

With VIX alone there are two candidate references. One is a fixed number, such as VIX above 20 or above 30. The other is VIX's own recent history, read as a deviation from a moving average or as a standardized score. A third family compares VIX with longer-dated implied volatility, usually the three-month index, and Johnson (2017) shows that the slope of that curve carries information about the price of variance risk. That approach needs a second data series and is left out here.

The main result is a comparison. Matched on how often they are on, a fixed level and a trailing z-score had been off before VIX spikes began in about the same number of episodes since 1990. So the choice between them shows up on the way back. After long spikes the trailing filter returned within weeks and the fixed level within months or years, and the trailing filter switched about three times as often to do it (section 6). Part of how a z-score comes back is mechanical, as the August 2024 case shows. Section 4 derives the limit that dispersion puts on a calm reading, and section 6 shows it mattered less often than the algebra suggests.

The paper draws on three areas of research. The first treats VIX as a gauge of risk and risk appetite and asks what its level and changes say about later returns (Copeland and Copeland 1999; Giot 2005; Bollerslev, Tauchen and Zhou 2009; Whaley 2009). The second scales exposure by recent volatility. Moreira and Muir (2017) find gains across many factors, Harvey et al. (2018) find them for equities and credit but not for other asset classes, and Cederburg et al. (2020) find them hard to capture out of sample. These studies take the volatility signal as given. The third, from robust statistics, documents how outliers inflate the scale estimates used to standardize them (Rousseeuw and Leroy 1987; Shiffler 1988) and how equally weighted estimates jump when an observation leaves the window (Alexander 1998, 2008). Here the third is applied to the first two: how do the choice of reference and the choice of scaling decide when a VIX signal reads calm?

The two references also differ in how steady they are from year to year. In this sample a filter keyed to VIX's own 20-day average was on for 42 to 72 percent of the days in every year. Nothing guarantees that. A VIX that fell steadily for a whole year would keep such a filter on every day. But an average that follows VIX tends to stay close to it, and in 36 years the filter never came near either extreme. Section 5 compares the references once exposure is matched, looking at how quickly each steps aside going into a spike as well as how quickly it returns, and section 6 checks both against every VIX spike since 1990. Sections 7 and 8 cover prolonged stress and what VIX cannot see.

Each day's signal uses that day's close and earlier ones only. Cboe sets the VIX close at about

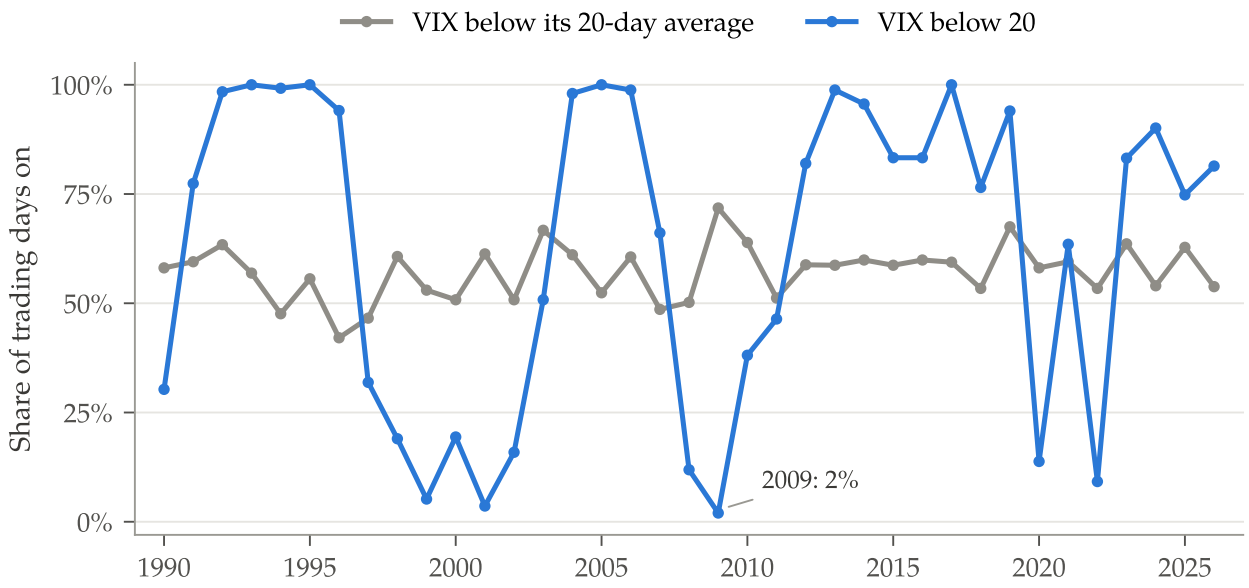


Figure 1. Share of trading days on, by year, for a fixed and a trailing VIX filter

Source: Cboe daily VIX closes, January 2, 1990 to October 16, 2026, NYSE trading days; trailing filter uses a 20-day simple average and starts in February 1990.

4:15 p.m. Eastern, a quarter of an hour after the S&P 500 closes (Cboe Global Markets), so a signal read from a day's close can be acted on no earlier than the next session. Comparisons across the whole sample, such as matching filters by their share of days on, are descriptive statistics computed after the fact. The episodes discussed in sections 4, 5 and 7 were chosen with hindsight, while section 6 takes every spike since 1990 by a fixed rule. The settings in the figures and tables are round illustrative values. The paper describes how signals behave and reports no strategy returns. The only market outcomes reported are how often the S&P 500 fell or rose 10 percent or more in the 20 trading days after a filter was on or off (section 6), and its realized volatility over those days. These describe risk, not what following a filter would have earned. Whether any filter is worth following is a separate question.

2 Level: the price of insurance

VIX measures the 30-day volatility implied by S&P 500 option prices, quoted as an annualized percentage (Whaley 2009). Reading it against a fixed number assumes the number means the same thing from one year to the next. VIX has averaged about 19 since 1990 and does revert toward that level, but the range it moves in shifts for years at a time.

A filter that holds risk-on assets while VIX is below 20 was on every trading day of 1993, 1995, 2005 and 2017, and on 2 percent of days in 2009. A filter that holds risk-on assets while VIX is below its own 20-day average was on between 42 and 72 percent of days in every year of the sample. Any fixed level behaves like the first filter, because the whole distribution of VIX shifts between regimes. Part of the trailing filter's steadiness comes from its reference moving with VIX.

The level is also ambiguous. Squared VIX is the sum of the variance the market expects and a premium for bearing variance risk (Bekaert and Hoerova 2014), which buyers of variance pay on

average (Carr and Wu 2009). That premium, measured as implied less realized variance rather than VIX itself, has predicted higher market returns at horizons of about a quarter (Bollerslev, Tauchen and Zhou 2009). Giot (2005) finds positive average returns after very high readings of implied volatility. So the usual case for holding fewer risk-on assets when VIX is high rests on risk, not on a forecast of losses. That is the logic of the volatility-managed portfolios of Moreira and Muir (2017), whose gains out of sample Cederburg et al. (2020) dispute.

A fixed level ignores recent history, so it stays off for as long as VIX stays high. A filter keyed to 20 was on for none of the 125 trading days from October 2008 to March 2009 (section 7).

3 Reading VIX against its own history

A trailing reference compares VIX with a statistic of its own recent past: a moving average, a percentage deviation from that average, or the deviation divided by recent dispersion. It measures how far VIX has moved from where it has lately been, and it needs only the one series.

These measures are long established. Copeland and Copeland (1999) already timed style and size rotations on changes in VIX, measured against its 75-day moving average. A moving average plus and minus a multiple of the standard deviation is a Bollinger Band, which Bollinger computes with the population standard deviation (Bollinger 2001). The textbook z-score of VIX against one window's mean and standard deviation is a linear rescaling of his %b. It is also the percentage deviation from the average divided by the window's coefficient of variation, $z = (V/\bar{V} - 1)/(s/\bar{V})$, so a z-score is a percentage deviation that is graded on how jumpy the window has been. On the trading side, the CVR3 rule, which Larry Connors developed with Dave Landry, buys the S&P at the close when VIX's intraday low stays above its 10-day moving average and VIX closes at least 10 percent above it. It exits when VIX trades back through the prior day's average or after two to four days (Connors 1998, 2002; Connors and Che 2001). Ding, Mazouz and Wang (2021) define a substantially high VIX as one at least 10 percent above its average over the prior 25 days.

The same measurement gets read in opposite directions. CVR3 treats a VIX well above its short average as fear that will unwind, and buys. A regime filter treats the same reading as a reason to hold fewer risk-on assets. Practitioner evidence at short windows points both ways. Hanna (2008), testing 1998 to 2008, found the S&P 500 did better on days when VIX was more than 10 percent below its 10-day average than on other days, though he noted that declines became more likely once VIX moved back toward its mean.

The window decides what kind of deviation a filter sees. Figure 2 measures how long deviations from trailing averages of different lengths last, and sets beside each the same calculation for a random walk. The half-lives come from first-order autocorrelation, as if each deviation were a simple autoregression. They are AR(1)-equivalent half-lives, a rough guide rather than a fitted model.

Deviations from a longer average last longer, as they must, since the average itself moves slowly. A random walk has no level to return to, so its half-lives come entirely from the moving average: subtracting a trailing average from any series makes the gap look as if it closes. VIX deviations die out much faster than the random walk's, and the difference widens with the window. Against a 250-day average the half-life is 18 trading days for VIX and about 115 for the random walk. That difference is consistent with VIX reverting toward a level, over and above what the moving average does by itself. Put as a probability: when VIX closed more than 10 percent above its 10-day average

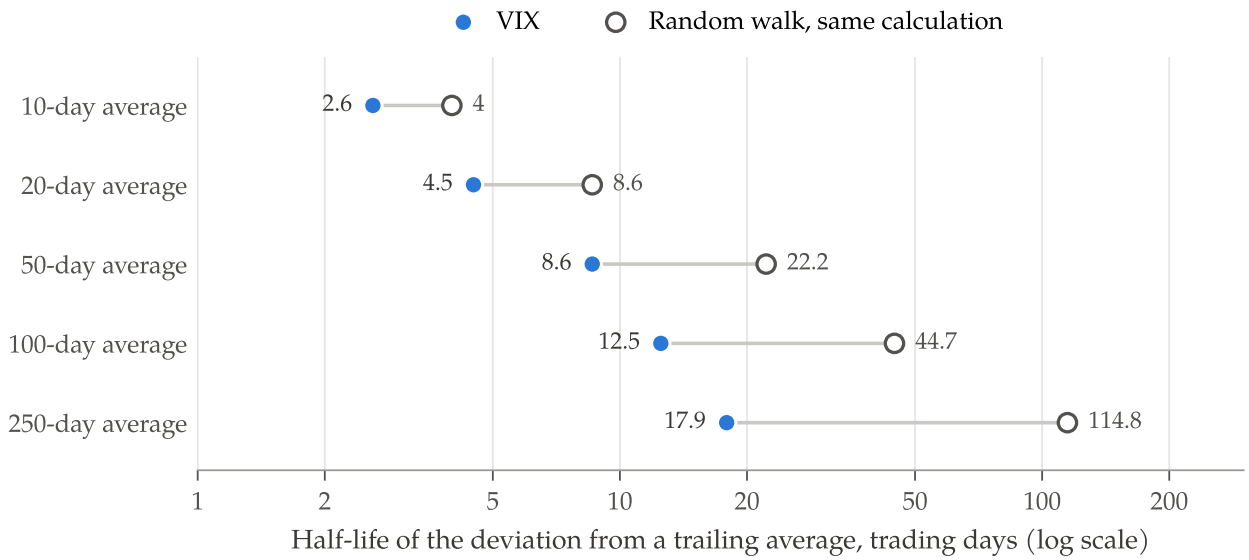


Figure 2. AR(1)-equivalent half-life of VIX deviations from trailing averages, against a random walk

Source: Cboe daily VIX closes, 1990 to October 16, 2026; half-life from the first-order autocorrelation of log VIX less the trailing simple average of log VIX; random walk measured against its own trailing simple average, simulated, 200,000 steps.

it was still that far above ten trading days later 12 percent of the time. Against its 100-day average the figure was 57 percent. A 10-day window mostly sees spikes that are gone within two weeks.

4 Adapts twice

A percentage deviation from a trailing average adapts once after a spike. The average rises while the spike sits in the window, so a VIX that has come back down looks calmer against it than against the pre-spike average. A z-score also divides by the window's standard deviation, and that rises with the spike too. The two effects pull in opposite directions. The higher mean brings the signal back on sooner, and the wider dispersion holds it back.

The balance between the two can be stated exactly. Write each close in the window as a baseline plus a spike, $x_i = b_i + \lambda e_i$, where the profile e_i is non-negative and is zero outside the spike and on today's close, and $\lambda > 0$ sets the spike's height. Let the window's mean and standard deviation use weights w_i that sum to one: equal weights set every w_i to $1/N$, and exponential weights let them decline with age. W is the total weight on the days the spike covers.

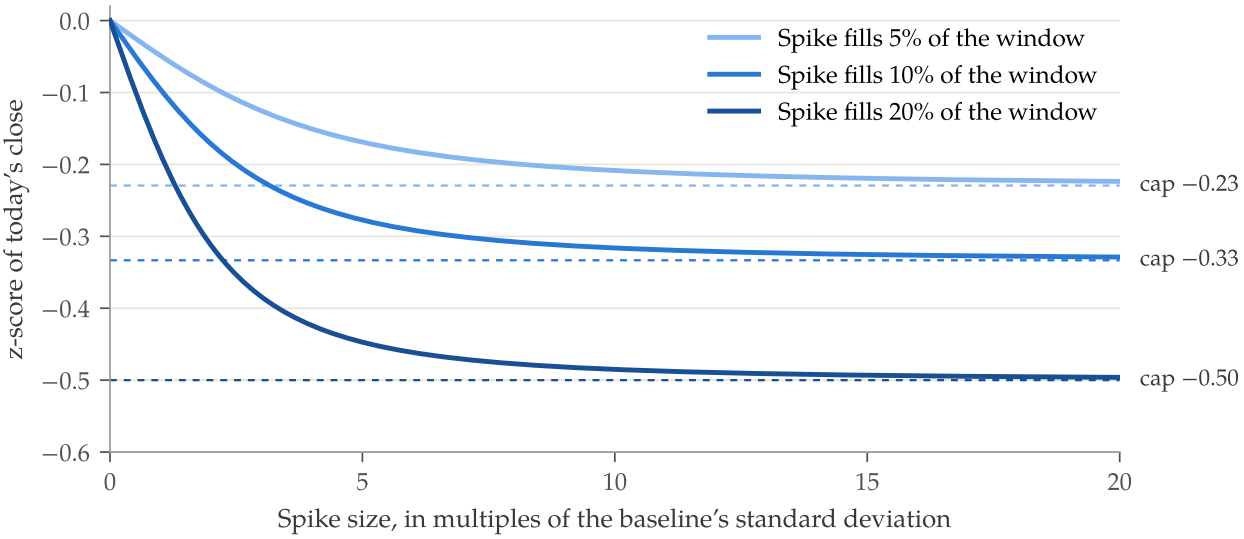


Figure 3. The calm reading of a z-score while a flat spike stays in the window

Notes: Computed from equation (1) for flat spikes filling 5, 10 and 20 percent of the window, with equal weights.

Proposition 1. *Suppose that on today's close the baseline sits at its own window mean, and that within the window the baseline is uncorrelated with the spike profile. Then the z-score of today's close is*

$$z(\lambda) = -\frac{\lambda\,\mu_e}{\sqrt{\sigma_b^2 + \lambda^2\sigma_e^2}} \tag{1}$$

where μ_e and σ_e are the weighted mean and standard deviation of the profile and σ_b is the weighted standard deviation of the baseline. As the spike grows, the reading moves steadily toward $-\mu_e/\sigma_e$ without reaching it if $\sigma_b > 0$, and

$$\frac{\mu_e}{\sigma_e} \le \sqrt{\frac{W}{1-W}} \tag{2}$$

with equality when the spike is flat. Without the two conditions on the baseline, $z(\lambda)$ still tends to $-\mu_e/\sigma_e$ as the spike grows.

The proof is in Appendix A.

The bound in equation (2) depends only on how much of the window the spike covers. The limit itself also depends on the spike's shape. Both are finite because today's close sits in its own window and is not part of the spike, so W is below one. Equation (2) is a benchmark for a spike on a steady base, not a bound on every observed reading, and August 2024, below, went past it. Height still matters for a stricter threshold: the taller the spike, the further VIX has to fall below its pre-spike level before the reading gets there. The sign is the conventional one, so calm readings are negative and the cap is a floor on how negative the reading can go.

With equal weights W is the share p of the window the spike covers. A flat spike filling 5 percent of the window caps the calm reading at −0.23, at 10 percent the cap is −0.33, and at 20 percent it is −0.50. A one-day spike in an N-day window caps it at $-1/\sqrt{N-1}$, which is −0.23 for 20 days. Turned around, a filter that waits for $z < -k$ cannot fire on a steady base until the spike covers more than $k^2/(1+k^2)$ of the window, which is half the window for $k = 1$. Any profile other than a

flat one has more dispersion relative to its mean and holds the reading closer to zero. The reading moves almost in step with the spike while $\lambda\sigma_e$ is small next to σ_b, and levels off toward the cap once $\lambda\sigma_e$ passes σ_b. For a flat spike that means a height above about $\sigma_b / \sqrt{p(1-p)}$. Simulated windows, described in Appendix B, match equation (1) to within sampling error.

None of the statistics here are new. Read from the other side, the same algebra gives the familiar bound on how large a single observation's z-score can be (Samuelson 1968; Shiffler 1988). The two papers state it as $\sqrt{N-1}$ and $(N-1)/\sqrt{N}$ because they divide by N and $N-1$ respectively. Outliers hiding behind the scale estimate they inflate is a staple of robust statistics (Rousseeuw and Leroy 1987). For a filter, the consequence is that a calm threshold beyond the cap cannot be reached while the spike remains in the window, unless VIX falls well below its pre-spike level or the elevated closes fill enough of the window to loosen the cap. Re-entry can end up waiting on the window length.

The cap describes a limit. On any given day the level VIX has to reach is exact. Let m' and s' be the mean and population standard deviation of the other $N-1$ closes in an equally weighted window. Then today's z-score is below $-k$, for $k^2 < N-1$, exactly when

$$x_N < m' - k\,s'\sqrt{\frac{N}{N-1-k^2}}\,. \tag{3}$$

The spike enters only through m' and s'. It raises m', which helps, and widens s', which hurts. The proof is in Appendix A.

August 2024 shows the cap at work, though not in the simple form of Figure 3. VIX closed at 38.57 on August 5, up from 23.39 the session before. By August 19 it was back at 14.65, about 26 percent below its 20-day average, and a filter keyed to a 10 percent discount to that average had turned back on seven trading days after the peak. The 20-day z-score that day read −0.91, the calmest it reached while the August 5 close stayed in the window. The other 19 closes that day had a mean of 20.19 and a standard deviation of 5.80, so by equation (3) a reading below −1 needed a close under 14.08.

That reading is beyond the flat-spike caps because the spike did not rise from a steady base. VIX had drifted up from about 12 in early July to the high teens by the end of the month, so on August 19 every other close in the window stood above that day's 14.65. A rise that fills most of the window loosens the cap. The single close at 38.57 still kept the dispersion wide enough to hold the reading short of −1. A filter keyed to a z-score below −1 stayed off until September 24, 35 trading days after the peak, once the spike had left the window and a second, smaller rise in early September had faded. The same filter returned 23 trading days after the February 2018 peak, three sessions after the peak close had left its window.

The cap binds hardest on a short spike. When stress lasts long enough to fill the window, the elevated closes become the baseline. On April 2, 2020, thirteen trading days after the March 16 peak, VIX closed at 50.91 and its 20-day z-score read −1.05, with every close in the window above 41. Compared with the prior month, a VIX of 51 looked calm. Section 7 returns to that side of the trade-off.

A second mechanical effect comes when the spike leaves. With equal weights an observation counts fully until the day it falls out of the window and then not at all, so the mean and standard deviation drop in one step. Alexander (1998, 2008) calls these the ghost features of equally weighted estimates. On September 3, 2024, twenty trading days after the peak, the August 5 close left the

window and the trailing standard deviation fell from 5.94 to 3.96. VIX also rose that day, to 20.72 from 15.00, and the 20-day z-score read 0.62. Over a 21-day window that still held August 5, the same close scores 0.26. Holding that close fixed, the drop-out alone added 0.36. Across the full sample the drop-out moves the 20-day z-score by more than 0.25 on about 1.6 percent of days.

Exponential weighting removes that cliff. The spike's weight W then decays a little each day instead of dropping to zero on one, so its pull on the dispersion fades gradually. A robust dispersion estimate, such as the median absolute deviation, is pulled around less by a short spike. Neither removed the problem. Each was set to be on for the same share of days as a strict equal-weight filter. The equal-weight filter was still off 20 or more trading days after the peak, by which time the peak had left a 20-day window, after 7 of 19 short spikes. The exponential version was off that long after 6 and the median-based version after 4 (section 6). Most of the median-based version's improvement comes from its scale, not its center. Keeping the mean and dividing by the median absolute deviation also gives 4. Keeping the standard deviation and measuring from the median gives 8.

5 Reference versus threshold

Some of what separates a trailing filter from a fixed one comes from the reference itself. An average that follows VIX tends to stay near it, so a filter keyed to VIX's own 20-day average is rarely on for nearly every day of a year, or almost none, the way a fixed level often is. That is a tendency, not a rule, and it accounts for part of the steadiness in Figure 1. Going from a 20-day z-score below 0 to one below −1 cuts time on from 57 to 26 percent, and the two filters differ on exactly the 31 percent of days in between.

The comparison worth making is at equal exposure. The filters are built on daily closes since 1990, each on when conditions look calm. Five are on when VIX is below a fixed level of 15, 17.5, 20, 22.5 or 25. Five are on when VIX is below its 20-day average by 0, 2.5, 5, 7.5 or 10 percent, and five when the 20-day z-score is below 0, −0.25, −0.5, −0.75 or −1. A deviation and a z-score always share a sign, so the first filter in each trailing family is the same one, leaving fourteen in all. No threshold was tuned on returns. Each z-score filter is paired with the fixed level that VIX sits below on the same share of days, a match made after the fact, and Figure 4 counts the days on which each pair disagrees.

Matched pairs disagree on 35 to 43 percent of days, against 38 to 50 percent for independent filters with the same share of days on. Cohen's kappa, which measures agreement beyond what chance alone would produce (Cohen 1960), runs from 0.09 to 0.15. No 95 percent interval, from a bootstrap that resamples whole calendar years, extends beyond 0.05 to 0.21. Most of this is Figure 1 seen again. About four-fifths of the disagreement comes from years in which the two filters spend very different shares of days on, and timing within a year accounts for 6 to 9 percent of all days. For the filter keyed to VIX's own average, longer windows bring the references together. Kappa rises from 0.12 at 10 days to 0.35 at 250 and disagreement falls from 44 to 31 percent. Within the trailing family, matched z-score and percentage-deviation filters differ on at most 11 percent of days.

Table 1 gives each filter's base rate and switching and how it behaved around three spikes. VIX had nearly doubled in the month before the 2018 and 2024 peaks. Every trailing filter had been off for at least 15 trading days by then, while fixed filters at 17.5 and above turned off on the day of the peak or at most two sessions before it. In 2020 the build-up was long enough that every filter had

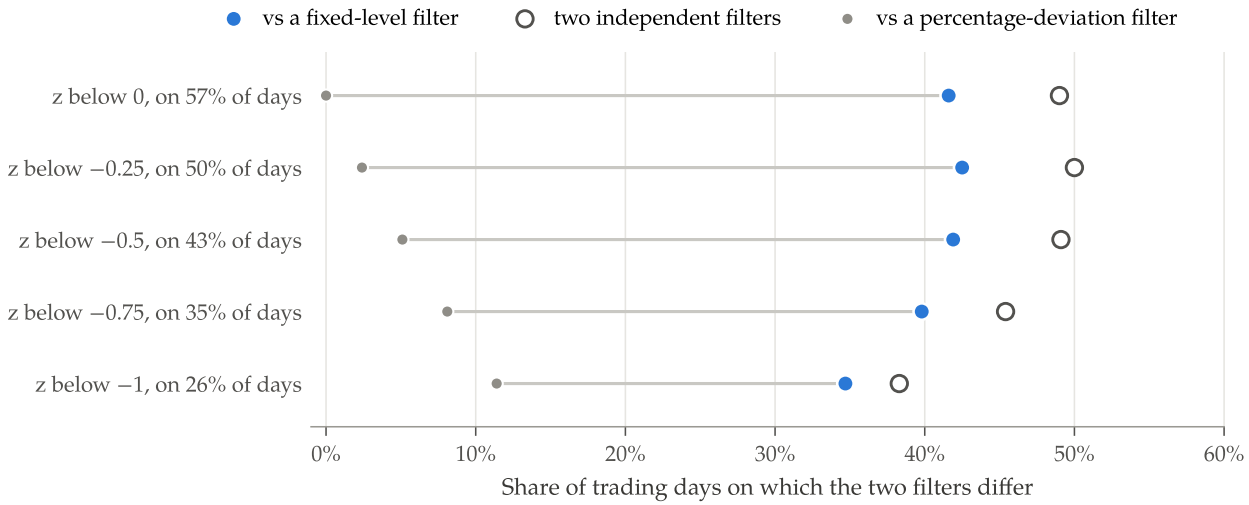


Figure 4. How often filters matched on share of days on disagree

Source: Cboe daily VIX closes, February 1990 to October 16, 2026, NYSE trading days; independence benchmark = 2q(1−q) for a share q of days on.

Table 1. Fourteen VIX filters: base rate, switching, and behavior around three spikes

			Feb 5, 2018		Mar 16, 2020		Aug 5, 2024	
Filter is on when	Days on	Switches per year	Off	Back	Off	Back	Off	Back
Panel A. Fixed level								
VIX below 15	32%	12.0	1	23	17	810	8	9
VIX below 17.5	50%	13.6	0	13	15	264	2	7
VIX below 20	63%	13.9	0	7	15	231	1	6
VIX below 22.5	74%	11.1	0	7	15	101	1	4
VIX below 25	83%	9.0	0	6	15	57	0	3
Panel B. Percentage deviation from the 20-day average								
VIX below its 20-day average	57%	40.4	16	8	17	10	15	6
2.5% below the average	47%	41.7	18	11	17	10	21	6
5% below	37%	38.1	19	12	17	11	34	6
7.5% below	27%	33.0	21	12	19	11	47	6
10% below	18%	25.6	36	12	60	11	50	7
Panel C. 20-day z-score								
20-day z-score below −0.25	50%	41.2	18	12	17	11	15	6
z-score below −0.5	43%	42.1	18	13	17	11	21	7
z-score below −0.75	35%	41.3	19	14	19	13	21	8
z-score below −1	26%	39.3	19	23	38	13	21	35

Notes: Days on = share of trading days the filter is on. Each episode gives two counts of trading days: how long the filter had been off at the close before the peak (0 means it was on at that close and turned off at the peak), then how long after the peak close it was first on again. A z-score below 0 is the same filter as VIX below its 20-day average.

been off for weeks. Coming out, the trailing filters were back on within 6 to 14 trading days in all three episodes, apart from the z-score at −1 in 2018 and 2024 (section 4). In 2020 they returned with VIX between 51 and 57. Fixed filters came back within 3 to 23 trading days in 2018 and 2024, when VIX fell quickly, and took 57 to 810 in 2020.

A trailing filter responds to change, so it steps aside during a build-up and returns during a decline whatever the level. It switches 26 to 42 times a year against 9 to 14 for the fixed levels in the table. All three spikes were preceded by weeks of rising VIX, so they say nothing about spikes that arrive from a steady VIX (section 8). The longer counts for the stricter filters also reflect, in part, how rarely those filters are on at all. Section 6 repeats the comparison across every spike since 1990.

6 How often it mattered

Sections 4 and 5 rest on three episodes chosen with hindsight. To check them, this section takes every VIX spike since 1990. A spike day is a close at least 1.5 times the average of the 20 closes before it. Spike days within 20 trading days of each other count as one episode, and the peak is the highest close in it. That gives 28 episodes, listed in Appendix C. The peak is the highest close from the first spike day to five days after the last, so in five episodes (August 1990, August 1998, October 2008, October 2018 and April 2022) VIX closed higher shortly after the dated peak. In 1998 and 2008 it went higher still weeks later, to 45.7 and 80.9. Widening the window to 20 days moves those five peaks and turns August 1990 into a long episode. The medians barely change: 13.5 against 148 trading days after long spikes and 8.5 against 9.5 after short ones. Nineteen episodes count as short, meaning VIX was back within 25 percent of its pre-spike average within 20 trading days of the peak. The other 9 are long.

A filter keyed to a 20-day z-score below −1 stayed off until the peak had left its window, 20 or more trading days, after 7 of the 19 short spikes. How much of that came from the spike itself sitting in the window? To find out, the z-score was recomputed each day with the spike's closes left out. That removes the spike's pull on both the average and the dispersion, so it measures the effect of the spike being in the window rather than the cap alone. Of the seven, that explains two, January 2021 and August 2024, where the filter would have come back 19 and 27 trading days sooner. In the other five it would have stayed off anyway, because VIX hadn't come down far enough.

The effect runs both ways. Across all 19 short spikes, leaving the spike out changed the timing in seven. In five it would have brought the filter back sooner: by 19 and 27 trading days in those two cases, and by two days in each of three others. In the other two it would have brought the filter back later, by 3 trading days in March 2007 and by 24 in January 2022. There the spike's pull on the average outweighed its pull on the dispersion, the first of the two effects in section 4 winning out. Exponential and median-based versions, set to be on for the same share of days, stayed off past the window after 6 and 4 of the 19.

Table 2 sets each z-score filter beside the fixed level that VIX sits below on the same share of days. Measured from the day each spike began, the matched pairs had been off for 10 or more days in about the same number of episodes. The bracketed figures in Table 2 compare that with random dates. A trailing filter is in a 10-day off run on only 10 to 35 percent of all days, so its hits stand out more against that benchmark. But the comparison is descriptive. The spike rule and the trailing filter share the same 20-day average, and the runs mostly reflect VIX rising before the spike began. A shock that arrives from a steady VIX gives no warning (section 8). After long spikes the trailing

Table 2. Fixed and trailing filters matched on time on, across 28 VIX spikes since 1990

Filter is on when	Days on	Off 10+ days before spike began: of 28 (chance)	Median days back after spike: Short	Median days back after spike: Long	Switches per year	10% S&P fall in 20 days: on / off
20-day z-score below 0	57%	7 (2.7)	5	10	40.4	2.3% / 5.6%
VIX below 18.8	57%	7 (9.1)	3	114	13.0	1.1% / 7.1%
z-score below −0.5	43%	10 (5.4)	11	14	42.1	2.1% / 4.9%
VIX below 16.5	43%	12 (12.6)	12	149	14.9	1.0% / 5.8%
z-score below −1	26%	17 (9.8)	15	17	39.3	1.9% / 4.3%
VIX below 14.1	26%	17 (18.1)	71	563	11.0	0.8% / 4.7%

Notes: Each fixed level is on for the same share of days as the z-score filter above it. Off 10+ days counts episodes in which the filter had been off for at least 10 trading days at the close before the first spike day; the figure in brackets is the number expected if the 28 dates had been drawn at random. Days back is the median number of trading days from the peak to the first close at which the filter is on, over 19 short and 9 long episodes. The short-spike median for VIX below 18.8 is sensitive to the level: at exactly 18.80 it is 6. A 10 percent S&P fall counts days followed by an S&P 500 close at least 10 percent lower within the next 20 trading days. Off-minus-on gaps in that rate, in percentage points with 95 percent intervals from a calendar-year bootstrap, are 3.3 (1.0 to 6.1) and 6.0 (3.0 to 9.1) for the pair on 57 percent of days, 2.7 (0.8 to 5.1) and 4.8 (2.1 to 7.6) at 43 percent, and 2.4 (0.7 to 4.3) and 3.9 (1.2 to 6.4) at 26 percent. The fixed level's gap less the z-score filter's is 2.7 (0.04 to 5.1), 2.1 (−0.4 to 4.2) and 1.5 (−1.2 to 3.6). With 2001–02 or 2008–09 left out, the gaps are 1.8 to 2.5 points for the z-score filters and 2.9 to 4.8 for the fixed levels, and every interval still excludes zero.

filters were back on in a median of 10 to 17 trading days and the fixed levels in 114 to 563. Part of that gap is built in, since a long spike is one in which VIX stayed well above its pre-spike level, which by itself keeps a fixed level off. After short spikes the two were close, except at the strictest setting.

After a day a fixed level was off, the S&P 500 fell 10 percent or more within 20 trading days about six times as often as after a day it was on. For the matched z-score filter the ratio was a little over two. Each gap holds up in a bootstrap that resamples whole calendar years. Off days also came before most of the 10 percent rallies, at every setting. That was 2.4 to 3.3 percent of the z-score filters' off days against 1.1 to 1.2 percent of their on days, and 2.9 to 4.8 percent of the fixed levels' off days against 0.0 to 0.1 percent of their on days. A filter that is off is picking out large moves in either direction. These numbers rest on about 25 separate declines, and nearly half the days fall in 2001–02 and 2008–09. Leaving out each of those periods in turn shrinks every gap by about a fifth to a quarter, and each still holds up. The fixed level's gap was the larger at every setting, by 1.5 to 2.7 points. That difference is borderline for the pair on 57 percent of days, where its 95 percent interval only just excludes zero ($p \approx 0.05$), and it is not statistically clear at the other two settings or once either crisis is left out. Because the VIX close comes after the S&P close, the same rates were also measured from the next day's close. None moved by more than 0.3 percentage points.

Realized volatility, which uses no returns to a strategy, tells the same story. Over the 20 trading days after a day the fixed level at 16.5 was off, S&P 500 realized volatility had a median of 16.6 percent, against 9.7 percent after a day it was on. For the matched z-score filter the medians were 13.8 and 11.9 percent.

A further test asked whether either filter anticipated new VIX spikes (Table 6, Appendix C).

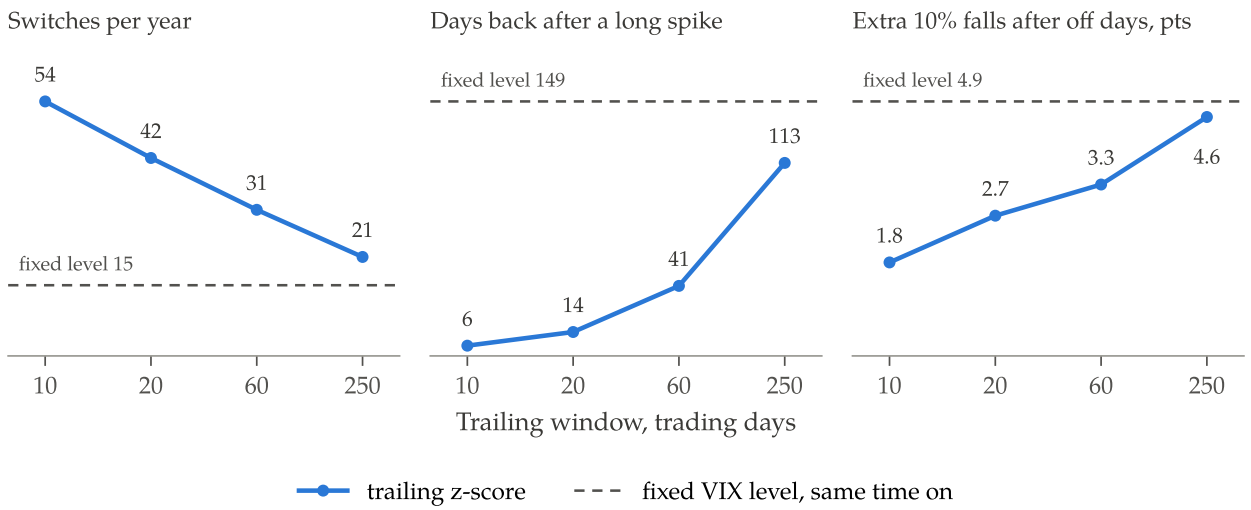


Figure 5. Trailing filters across window lengths, against a fixed level with the same time on

Source: Cboe daily VIX and S&P 500 closes, 1991 to October 16, 2026; trailing filter = the middle setting of Table 2 at each window, on for 43 to 44 percent of days, beside the fixed level on for the same share; days back = median over the nine long spike episodes.

Leaving aside days inside an episode, a new spike began within 20 trading days after 7.1 to 7.9 percent of the trailing filters' off days and 4.7 to 5.4 percent of their on days. The gap of about 2.5 points has intervals that exclude zero. For the fixed levels the gap ran the other way, 5.6 to 6.1 percent after off days against 6.9 to 7.5 percent after on days, and its intervals include zero. Neither result is much of a forecast. The trailing filter's edge probably owes most to VIX momentum and to the spike rule sharing its 20-day average. The fixed level's reversal reflects that a 50 percent jump is easier to reach from a low base.

Figure 5 repeats the comparison with windows of 10, 20, 60 and 250 trading days. A longer window switches less, comes back more slowly after long spikes and lines up more with large S&P falls. At 250 days it comes close to a fixed level on re-entry and switching, although at that setting the two still disagree on about 34 percent of days. The comparison also holds under other spike rules (Table 5, Appendix C). With a threshold of 1.3 times the 20-day average there are 80 episodes, with 2 times only 7, and from September 2003, when the current VIX method began, 23. In each case the matched pair on 43 percent of days had been off before the spike began in about the same number of episodes. The trailing filter came back after long spikes in a median of 11 to 15 trading days, against 97 to 149 for the fixed level. The spike's presence in the window delayed a strict filter by five days or more in at most five short spikes, and sped it up by that much in at most two.

The window is not the only way to cut switching. A filter with a band turns off when the 20-day z-score rises above 0 and comes back on only when it falls below −0.5. It was on for 51 percent of days and switched 29 times a year, against 40 to 42 for the single thresholds on either side of it. It came back after long spikes in a median of 14 trading days and after short ones in 11, the same as the z-score below −0.5, and had been off 10 or more days before 8 of the 28 spikes. The fixed level on for the same share of days, about 17.6, switched 13.5 times a year and took a median of 130 days to come back after long spikes. Against that fixed level, the band narrows the switching gap from about 27 a year to about 16. That comparison flatters the band, though, because a fixed level

Table 3. Filters in two periods of persistently high VIX

Period	Median VIX	Days VIX above 20	Share of days each filter was on: VIX below 20	VIX below its 20-day average	20-day z-score below −0.5
Oct 2008 to Mar 2009 (125 trading days)	48.5	100%	0%	54%	40%
Calendar 2022 (251 trading days)	25.5	91%	9%	53%	40%

Notes: Episodes chosen with hindsight; see Appendix B.

can take a band too. Turned off above 18.82 and back on below 16.51, a fixed level was on for 50 percent of days and switched 5.2 times a year. With a band on each side, the gap is about 24 a year.

7 Trailing references in prolonged stress

A trailing reference measures change, and in a long stress period VIX can fall relative to its own recent past while staying high. From October 2008 through March 2009 VIX closed above 20 on every trading day, at a median of 48.5. A filter holding risk-on assets while VIX is below 20 was on for none of those 125 days. A filter holding risk-on assets while VIX is below its 20-day average was on for 54 percent of them, and one keyed to a 20-day z-score below −0.5 for 40 percent. In calendar 2022 VIX closed above 20 on 91 percent of days, and the fixed filter was on 9 percent of the time against 53 percent for the trailing filter.

By its own measure the trailing filter was right that conditions had calmed relative to the recent past. Given a long enough stress period, though, a trailing reference rises with VIX and comes to treat high levels as normal. The z-score filter that was back on at a VIX of 51 in April 2020 (section 4) is a short version of the same thing.

A fixed level has the opposite problem, shown in section 5: a filter keyed to 20 stayed off for 231 trading days after March 2020 while VIX worked its way down. Which mistake costs more depends on the strategy and the episode. Coming back at a VIX of 51 in April 2020 meant buying near the bottom. Coming back in the autumn of 2008 meant buying with months of decline still ahead.

8 What VIX can't see

A filter that reads only VIX inherits VIX's blind spots, and the most basic is timing. Implied volatility rises when demand for hedging picks up, often on the day of the shock itself. A filter that reads VIX at the close cannot step aside before the close at which VIX first moves, and, because the VIX close comes after the S&P close, it can act only the session after that. The three spikes in Table 1 were each preceded by weeks of rising VIX. A shock that arrives while VIX is steady gives a filter that reads VIX no warning at all.

VIX is also a price, set by the supply of index options as well as the demand for them. Todorov and Vilkov (2024) judge 0DTE options an unlikely main cause of the low VIX of 2023 and point instead to dealer hedging of yield-enhancing structured products. Supply-driven moves in VIX

move a filter too. And VIX prices 30-day S&P 500 variance only, so a shock confined to one sector or another asset class can leave it nearly unchanged.

9 Conclusion

The z-score's cap is real but rarely decisive. After a short spike, a strict filter can stay off until the spike leaves its window even once VIX has calmed down. In this sample, having the spike in the window added about a month twice in 19 short spikes. In two others the spike's pull on the average brought the filter back sooner, by 3 and 24 trading days.

The bigger choice is the reference. Held to the same time on, fixed and trailing filters still disagree on more than a third of days, mostly over which years to favor. They were off before spikes began about equally often. A fixed level means different things in different regimes and came back slowly after stress. In the episodes tested it never treated a long stress period as calm. The trailing reference came back within weeks, at about three times the switching, and spent 40 to 54 percent of late 2008, early 2009 and 2022 on while VIX sat mostly above 20. Off days for both were followed by more large S&P moves in either direction, somewhat more so for the fixed level, though that difference is borderline and rests on the 2001–02 and 2008–09 declines.

Whether coming back faster is worth having is a question this paper leaves open. The answer depends on the gross trade-off between missed recoveries and re-entry into continued stress, which is not measured here, and then on trading costs and, in a taxable account, realized gains. A longer window moves a trailing filter toward the fixed one on switching and re-entry. A band cuts switching too, though a fixed level can use one as well.

References

Alexander, C. (1998). Volatility and correlation: Measurement, models and applications. In C. Alexander (ed.), *Risk Management and Analysis, Vol. 1: Measuring and Modelling Financial Risk*, chapter 4. Wiley.

Alexander, C. (2008). Moving average models for volatility and correlation, and covariance matrices. In F. J. Fabozzi (ed.), *Handbook of Finance*, 711–724. Wiley.

Bekaert, G. & Hoerova, M. (2014). The VIX, the variance premium and stock market volatility. *Journal of Econometrics* 183(2), 181–192.

Bollerslev, T., Tauchen, G. & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies* 22(11), 4463–4492.

Bollinger, J. (2001). *Bollinger on Bollinger Bands*. McGraw-Hill. Calculation conventions from bollinger-bands.com.

Carr, P. & Wu, L. (2009). Variance risk premiums. *Review of Financial Studies* 22(3), 1311–1341.

Cboe Global Markets. VIX historical data and Cboe Volatility Index methodology.

Cederburg, S., O'Doherty, M. S., Wang, F. & Yan, X. S. (2020). On the performance of volatility-managed portfolios. *Journal of Financial Economics* 138(1), 95–117.

Cloutier, R., Djatej, A. & Kiefer, D. (2017). A tactical asset allocation strategy that exploits variations in VIX. *Investment Management and Financial Innovations* 14(1), 27–34.

Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* 20(1), 37–46.

Connors, L. A. (1998). *Connors on Advanced Trading Strategies*. M. Gordon Publishing Group.

Connors, L. (2002, January 11). Market timing using the VIX. TradingMarkets.

Connors, L. A. & Che, G. (2001). *Trading Connors VIX Reversals*. M. Gordon Publishing Group.

Copeland, M. M. & Copeland, T. E. (1999). Market timing: Style and size rotation using the VIX. *Financial Analysts Journal* 55(2), 73–81.

Ding, W., Mazouz, K. & Wang, Q. (2021). Volatility timing, sentiment, and the short-term profitability of VIX-based cross-sectional trading strategies. *Journal of Empirical Finance* 63, 42–56.

Giot, P. (2005). Relationships between implied volatility indexes and stock index returns. *Journal of Portfolio Management* 31(3), 92–100.

Hanna, R. (2008, May 15). Is a low VIX a short trigger? Quantifiable Edges.

Harvey, C. R., Hoyle, E., Korgaonkar, R., Rattray, S., Sargaison, M. & Van Hemert, O. (2018). The impact of volatility targeting. *Journal of Portfolio Management* 45(1), 14–33.

Johnson, T. L. (2017). Risk premia and the VIX term structure. *Journal of Financial and Quantitative Analysis* 52(6), 2461–2490.

Moreira, A. & Muir, T. (2017). Volatility-managed portfolios. *Journal of Finance* 72(4), 1611–1644.

Rousseeuw, P. J. & Leroy, A. M. (1987). *Robust Regression and Outlier Detection*. Wiley.

S&P Dow Jones Indices (2024). *S&P 500 Dynamic VEQTOR Index Methodology*. May 2024.

Samuelson, P. A. (1968). How deviant can you be? *Journal of the American Statistical Association* 63(324), 1522–1525.

Shiffler, R. E. (1988). Maximum Z scores and outliers. *The American Statistician* 42(1), 79–80.

Todorov, K. & Vilkov, G. (2024). What could explain the recent drop in VIX? *BIS Quarterly Review*, March 2024.

Whaley, R. E. (2009). Understanding the VIX. *Journal of Portfolio Management* 35(3), 98–105.

Appendix A. Proofs

Proposition 1. For any series y over the window, write its weighted mean and variance as $\bar{y} = \sum w_i y_i$ and $s_y^2 = \sum w_i (y_i - \bar{y})^2$, with the weights summing to one. With $x_i = b_i + \lambda e_i$, the window mean is $\bar{b} + \lambda \mu_e$ and the variance is $\sigma_b^2 + 2\lambda c + \lambda^2 \sigma_e^2$, where c is the weighted covariance of the baseline and the profile. Today's close is $x_N = b_N$, because $e_N = 0$, so

$$z(\lambda) = \frac{b_N - \bar{b} - \lambda\, \mu_e}{\sqrt{\sigma_b^2 + 2\lambda c + \lambda^2 \sigma_e^2}} \tag{A1}$$

Dividing the numerator and denominator by λ shows that $z(\lambda)$ tends to $-\mu_e/\sigma_e$ as λ grows, provided $\sigma_e > 0$. Under the two conditions of the proposition, $b_N = \bar{b}$ and $c = 0$, and (A1) reduces to equation (1). Written as $z(\lambda) = -\mu_e / \sqrt{\sigma_b^2/\lambda^2 + \sigma_e^2}$, it shows that the size of the reading rises with λ and stays below μ_e/σ_e for every finite spike.

For the bound, let S be the set of days on which $e_i \neq 0$, so that W is the total weight on S. W is less than one, because today's close carries positive weight and lies outside S. By the Cauchy–Schwarz inequality,

$$\mu_e^2 = \Big(\sum_{i \in S} w_i e_i \Big)^2 \leq \Big(\sum_{i \in S} w_i \Big) \Big(\sum_{i \in S} w_i e_i^2 \Big) = W\,(\sigma_e^2 + \mu_e^2) \tag{A2}$$

which rearranges to $\mu_e^2(1-W) \le W\sigma_e^2$, or $\mu_e/\sigma_e \le \sqrt{W/(1-W)}$. Equality in (A2) holds only when e_i takes the same value on every day in S, that is, when the spike is flat. With equal weights W is the share p of the window the spike covers, and a one-day spike in an N-day window gives $p = 1/N$ and a limit of $-1/\sqrt{N-1}$. Setting $\sqrt{W/(1-W)} = k$ gives the occupancy $W = k^2/(1+k^2)$ that a reading of $-k$ requires on a steady base.

Equation (3). With equal weights, let m' and s'^2 be the mean and population variance of the $N-1$ closes other than today's, and let $d = x_N - m'$. The window mean is $m' + d/N$, so x_N lies $d(N-1)/N$ above it, and the window's population variance is $\frac{N-1}{N}s'^2 + \frac{N-1}{N^2}d^2$. Setting the z-score equal to $-k$ with $d < 0$ and squaring gives $d^2(N-1)(N-1-k^2) = k^2N(N-1)s'^2$, so $d = -k\,s'\sqrt{N/(N-1-k^2)}$ when $k^2 < N-1$. The z-score falls as d falls, which gives equation (3). Applied to every day of the sample, the formula matches the computed z-score to rounding error. ■

Appendix B. Data and method

VIX. Cboe's published daily history, January 2, 1990 to October 16, 2026, using closing values. Values before September 22, 2003, when the current S&P 500-based method was launched, are back-calculated by Cboe; the original S&P 100-based index was published separately as VXO until Cboe discontinued it in 2021. Only NYSE trading days are kept; 34 entries dated on exchange holidays are dropped, all but one of them since 2022. Four NYSE trading days have no entry (March 1, 1991; January 31, 1997; November 26, 1997; December 31, 1999), and rolling windows count the observations present.

Timing. Cboe's VIX close is set at about 4:15 p.m. Eastern, after the 4:00 p.m. close of the S&P 500. Signals in this paper are read at the VIX close and could be acted on no earlier than the next session. Forward S&P outcomes in section 6 are measured from the signal day's close and, as a check, from the next day's close.

Statistics. Moving averages and standard deviations are simple and equally weighted over the stated window, with the standard deviation in its population form, except for the exponential and median-based versions in sections 4 and 6. The exponential version applies a smoothing weight of 2/(N + 1) to both the mean and the squared deviations from it. The median-based version divides the deviation from the window median by 1.4826 times the median absolute deviation. The two mixed versions in section 4 pair the mean with that scaled median absolute deviation and the median with the standard deviation. Each day's value uses that day's close and earlier ones only. Half-lives in Figure 2 come from the first-order autocorrelation φ of the daily deviation of log VIX from the trailing simple average of log VIX, as $\ln(0.5)/\ln(\varphi)$. The random-walk baseline applies the same calculation to a simulated Gaussian random walk of 200,000 steps. For a random walk, the deviation from its own N-day simple average has first-order autocorrelation $(2N-4)/(2N-1)$, which gives the plotted random-walk half-lives to within simulation error.

The cap. The simulation uses windows of 100 observations with Gaussian noise of unit dispersion, inserts flat spikes covering 2, 5, 10 and 20 percent of the window at sizes from 1 to 1,000 times the noise, and computes the z-score once the series has returned to its pre-spike level, averaging 400 paths per case. In these windows the two conditions of Proposition 1 hold only on average. The

simulated readings still matched equation (1) to within sampling error and converged on the stated limits.

Filters. A filter is on when its condition holds at the close. In Table 1, days off count the consecutive closes before the peak close at which the filter was off, and days back count trading days from the peak close to the first close at which the filter is on. Each matched pair uses the fixed level below which VIX closed on the same full-sample share of days as the trailing filter, shown rounded to 0.1. The exact levels are 18.82, 16.51 and 14.10 for the three pairs in Table 2 and 17.63 for the band. Because VIX is quoted to two decimals, a level can match the target share only to within one quoted step, and the shares agree to within about 0.05 percentage points. Days back run to the first close after the peak at which the filter is on. That close does not mean the filter stayed on, and a filter can turn off again soon after. Agreement beyond chance is Cohen's kappa, with intervals from 1,000 bootstrap resamples of whole calendar years. The between-year share sums, year by year, the difference in the two filters' share of days on, which is the least disagreement possible given those shares. The band filter in section 6 turns off at the first close with a 20-day z-score above 0 and back on at the first close below −0.5, and otherwise keeps its previous state.

Episodes. The episodes in sections 4, 5 and 7 were chosen with hindsight: the closes of February 5, 2018, March 16, 2020 and August 5, 2024 for spikes, and October 2008 through March 2009 and calendar 2022 for persistent stress. All settings are illustrative.

Spike episodes. For section 6, a spike day is a VIX close at least 1.5 times the simple average of the 20 closes before it. Spike days no more than 20 trading days apart form one episode, whose peak is the highest close from its first spike day to five days after its last. An episode is short when VIX closes within 25 percent of its pre-spike average within 20 trading days of the peak. Days back are counted as in Table 1. A filter stayed off past the window when it was first back on 20 or more trading days after the peak. Early warning is measured at the close before each episode's first spike day, and the chance benchmark is the share of all days on which the filter had been off for at least 10 straight closes, times 28. To separate the effect of the spike sitting in the window, each z-score is also recomputed with the episode's spike days left out of the window, which then holds fewer closes until the spike days have aged out. Both directions of change are reported. Figure 5 starts in 1991, so that every window has a full year of history, which moves its 20-day values slightly from those in Table 2; all nine long episodes fall after 1991.

New spikes. For Table 6, a day counts when the first spike day of an episode falls within the next 20 trading days. Days from an episode's first spike day through 20 trading days after its peak are left out. Intervals resample whole calendar years, with 100,000 draws.

S&P 500. Cboe's published daily SPX closes, on the same dates as VIX. A day counts toward a 10 percent fall when the lowest close in the next 20 trading days is at least 10 percent below that day's close, and toward a 10 percent rise when the highest close is at least 10 percent above it. Realized volatility is the annualized standard deviation of daily log returns over the next 20 trading days. Intervals on the difference between days off and days on, and on the difference between two matched filters' gaps, resample whole calendar years, with 100,000 draws (1,000,000 for the difference at 57 percent, whose interval lies closest to zero). Both filters in a pair are resampled on the same years, and p-values are two-sided. The robustness check leaves out calendar 2001–02 or

2008–09 in turn.

Appendix C. Spike episodes and robustness

Table 4. The 28 VIX spike episodes since 1990

Peak	Peak VIX	Pre-spike average	Type	Days back: z-score below −0.5	Days back: VIX below 16.5
Aug 6, 1990	35.9	18.9	short	30	151
Apr 4, 1994	23.9	15.1	short	13	6
Oct 30, 1997	38.2	22.1	short	14	177
Aug 11, 1998	31.9	20.5	long	29	1311*
Sep 20, 2001	43.7	23.9	long	14	531
Mar 5, 2007	19.6	10.6	short	11	1
Oct 17, 2008	70.3	22.8	long	11	358
May 20, 2010	45.8	18.4	long	8	147
Aug 8, 2011	48.0	20.4	long	15	149
Feb 3, 2014	21.4	14.2	short	7	4
Oct 15, 2014	26.2	15.2	short	4	4
Dec 16, 2014	23.6	13.9	short	24	3
Aug 24, 2015	40.7	13.6	long	15	34
Jun 24, 2016	25.8	16.6	short	4	4
Aug 10, 2017	16.0	10.0	short	14	1
Feb 5, 2018	37.3	11.3	long	13	13
Oct 11, 2018	25.0	12.9	long	16	19
Dec 24, 2018	36.1	21.9	short	7	26
Aug 5, 2019	24.6	13.8	short	12	12
Mar 16, 2020	82.7	15.6	long	11	274*
Jan 27, 2021	37.2	23.2	short	6	55
May 12, 2021	27.6	18.1	short	8	16
Dec 1, 2021	31.1	17.1	short	7	346*
Jan 26, 2022	32.0	19.1	short	9	308*
Apr 26, 2022	33.5	21.9	short	15	246
Aug 5, 2024	38.6	15.0	short	7	7
Dec 18, 2024	27.6	14.4	short	17	4
Apr 8, 2025	52.3	21.9	short	11	55

Notes: Peak is the highest close in the episode; the pre-spike average is the mean of the 20 closes before the first spike day. Days back count trading days from the peak to the first close at which the filter is on. An asterisk marks a return that came only after the next episode had begun. In five episodes VIX closed higher shortly after the dated peak; see section 6.

Table 5. The matched comparison under other spike rules

Spike rule	Episodes (short / long)	Off 10+ days before spike: z / fixed	Days back, long spikes: z / fixed	Days back, short spikes: z / fixed	Strict filter, 5+ days: delayed by spike	Strict filter, 5+ days: sped up by spike
1.5 times the 20-day average	28 (19 / 9)	10 / 12	14 / 149	11 / 12	2	1
1.3 times	80 (66 / 14)	24 / 24	11 / 97	8 / 6.5	5	2
2 times	7 (2 / 5)	4 / 3	15 / 149	9 / 31	1	0
1.5 times, from Sep 22, 2003	23 (16 / 7)	9 / 11	13 / 149	8.5 / 10.5	2	1

Notes: The matched pair on 43 percent of days (a 20-day z-score below −0.5 and VIX below 16.5, or 16.0 from 2003) under other spike rules. Days back are medians in trading days. The last two columns count short spikes in which recomputing the z-score without the spike's closes would have brought a filter keyed to a z-score below −1 back at least five trading days sooner (delayed by the spike) or at least five trading days later (sped up by the spike).

Table 6. Did either filter anticipate new spikes?

Filter is on when	New spike within 20 days: after on days	New spike within 20 days: after off days	Off-minus-on gap, pts (95% interval)
20-day z-score below 0	5.4%	7.9%	2.5 (0.5 to 4.7)
VIX below 18.8	7.1%	5.6%	−1.5 (−5.5 to 2.2)
z-score below −0.5	5.0%	7.7%	2.7 (0.9 to 4.7)
VIX below 16.5	6.9%	6.1%	−0.8 (−4.9 to 3.0)
z-score below −1	4.7%	7.1%	2.4 (0.4 to 4.5)
VIX below 14.1	7.5%	6.1%	−1.4 (−6.6 to 2.9)

Notes: Share of days on which the first spike day of a new episode (section 6 rule) fell within the next 20 trading days, excluding days from each episode's first spike day through 20 trading days after its peak. Intervals from 100,000 bootstrap resamples of whole calendar years. The trailing filter and the spike rule share the same 20-day average.

Disclosures

This paper is for informational and educational purposes only. It is not investment advice, and it is not an offer or a solicitation to buy or sell any security or to adopt any investment strategy. It reports no performance, actual or hypothetical, of any strategy or account. The signals examined are illustrative constructions on public data, with settings chosen for exposition. Massari uses measures related to those studied here in some of its strategies. This paper does not describe how any strategy is implemented, and the results reported here are not the results of any Massari strategy or account. The past behavior of market indicators does not guarantee their future behavior. References to third-party research describe the findings of those authors.

Massari Asset Management is registered as an investment adviser with the Florida Office of Financial Regulation. Registration does not imply a certain level of skill or training. VIX® is a registered trademark of Cboe Exchange, Inc. S&P 500® is a registered trademark of Standard & Poor's Financial Services LLC.

Suggested citation: Alhajj, A. (2026). Reading VIX Against Itself. Massari Research Note No. 1, version 1.0. Massari Asset Management. *Contact:* contact@massari.am.