CORNERSTONE
INVESTMENT PARTNERS, LLC

"Should" and "Could" versus "Will" and "Do"

October 2026

Executive Summary

- Prices Move More Than Fundamentals: Our firm-wide philosophy is that stock prices are more volatile than the company fundamentals which drive value. That gap is the source of the mispricings we are paid to find, and we control risk where it originates, at the level of the individual stock.
- Two Tickers, One Bet: Stocks can appear independently mispriced while being mispriced for some of the same reasons. Two holdings that depend on the same credit conditions, the same trade flows or the same thematic narrative are one bet wearing two tickers, and no sector table will tell you when that has happened.
- Static Factor Models Misprice Risk: Models that forecast risk from past covariances assume the relationship between factors and risk is stationary. They report risk as cheap through calm stretches such as 2003 to 2007 and as expensive through dislocations such as 2009 and 2021, which is backward. That precision is artificial. False precision creates overconfidence, and acting on overconfidence can destroy capital.
- Market Structures Force Stocks Together: Thematic baskets, the labeling of companies as theme winners and losers, active exchange traded funds, and socially networked retail flows can push stocks into lockstep for reasons unrelated to the underlying business fundamentals.
- Market Modalities Measure It: Investors usually start with a hypothesis and then confirm with data. In this approach, we start with the data, and then ask why. We can identify relationships in the market based upon how stocks have actually traded, rather than how they "should" trade based upon heuristics.
- Our Approach Permits "Nowhere": Most clustering methods force every stock into a group, denigrating the homogeneity of the group and decreasing the value of the insight we most want, namely which stocks are trading on their own fundamentals. Our approach leaves those names unassigned, and its hierarchy cuts each group at its own density rather than at one threshold for the whole market.
- A Diagnostic, Not a Trigger: The modality map assists portfolio management decisions, by providing a view on how a stock actually trades relative to other stocks in the market. It does not size a position and it never trades for us. Our approach to market modalities raises questions we must answer.

Introduction

The Cornerstone Investment Team is a team of stock analysts focused on identifying idiosyncratic mispricings in the market, and we believe risk is best controlled at the stock level. Our firm-wide philosophy is that stock prices are more volatile than the company fundamentals which drive value, and that volatility is our opportunity rather than our enemy. But some of that movement is not about the companies at all. Index flows, thematic baskets and hedged option positions move stocks in groups, and two holdings bought for unrelated reasons will sometimes trade as one bet wearing two tickers. This paper describes the instrument we built to find those groups, why we built it this way rather than another, and what it shows about the market as it trades today.

Volatility Is the Opportunity, and Some of It Is Not About the Company

Cornerstone operates under a single firm-wide investment philosophy: stock prices are more volatile than the company fundamentals which drive value. Cognitive and behavioral errors of market participants create these mispricings that we are paid to find. These mispricings exist sometimes within near term volatility driven by emotions and other times within intermediate term indifference. We define risk as the permanent loss of capital, not as short-term price movement, and we control that risk where it originates, at the level of the individual stock.

That discipline leaves one gap, and we have written about it before. In a concentrated portfolio, small risks are amplified. Two stocks can appear independently mispriced may actually be mispriced for a similar reason, such as a macroeconomic, thematic, or clientele factor. Much of our portfolio construction effort goes into finding genuinely different idiosyncratic reasons across the portfolio, so that performance comes from the reconciliation of many independent mispricings rather than from a single concentrated exposure we did not know we had.

Our fundamental work seeks to understand how the market's misperception of a company is driving its mispricing. However, a stock is only worth what the market recognizes its value to be. Our thesis on a company can be correct, while its stock spends months trading on some other consideration entirely. We needed an instrument that interprets how the market is actually treating specific stocks, without diminishing the value of our fundamental analysis.

Why a Static Factor Model Will Not Close the Gap

The conventional answer is an *ex ante* risk model: explain returns through a rubric of style and industry factors, estimate the covariances among those factors from history, and forecast tomorrow's risk from them. We use these models and find them useful to identify our portfolio exposures versus those factors. But we do not trust the model's estimation of risk, and we do not believe that blind correlations divine persistent relationships in how the market prices individual stocks or the volatility of those stocks.

A factor model's load-bearing assumption is that the relationship between its factors and realized risk is stationary. It is not. In our view, the consequence is a model that prices risk exactly wrong at the moments the pricing matters most. Through a long calm stretch such as 2003 to 2007, measured covariances compress, forecast tracking error falls, and the model reports that risk is cheap. It reports this because the recent past was quiet, which is precisely when leverage and crowding are accumulating underneath. Run

the same machinery through 2009, or through 2021, and it does the opposite: covariances estimated across the dislocation are enormous, every position looks expensive in risk terms, and the model counsels caution at the exact moment dislocation has created the mispricings a valuation-driven process exists to harvest.

The damage is not the error itself. It is the false confidence. These models report to three decimal places, and the precision is artificial, because it rests on an assumption about stationarity that the model cannot check and the market does not honor. A manager who sizes positions against that number has outsourced judgment to an estimate of the past. We have watched that destroy value in both directions, through comfort in 2007 and through timidity in 2009. Further, relying upon a correlation without seeking to understand the underlying causation is a mental shortcut that puts tremendous value at risk. We believe this risk comes without alpha, so it can do nothing but decrease risk-adjusted returns.

There is a habit of language underneath this that we take seriously on our team. Culture is built from language and behaviors. We avoid saying that a stock ***should*** do something. "Should" is an academic claim about a world where prices track fundamentals on a reliable schedule. We are accountable for real capital belonging to real clients who care about real results, and the word that matches that accountability is ***will***. A risk model, by construction, describes how a portfolio should behave given how the past behaved. The question we have to answer is how our stocks will trade, over the horizon where we actually own them, in the market structure that actually exists now.

What Forces Stocks to Trade Together

Three mechanisms push stocks into lockstep for reasons that have nothing to do with the businesses underneath. None is new. All three have grown.

Thematic baskets and active exchange-traded funds. A stock enters a basket because someone writing a prospectus decided it belonged to a theme, and from that day forward a portion of its daily volume arrives as theme flow rather than company flow. For example, Cummins (CMI) makes diesel engines and other powertrain parts, primarily for Class 8 trucks. Its Power Systems segment was $7.5 billion of 2025 revenue, and the company's own analyst day disclosure puts data centers at 26% of that segment. That is roughly $2.0 billion against about $41 billion of segment revenue, or close to 5% of the company. Yet Cummins appears in artificial-intelligence thematic products, among them Roundhill's HALO exchange-traded fund, and on days when the AI trade moves, Cummins moves with it. It appears to us that five percent of a business is setting the trading behavior of the whole of it.

Labeling. Similarly, market participants often reduce companies to a single concept and then trade around that. A widely followed research letter calls life-science tools a loser from artificial intelligence, and IQVIA (IQV) and Thermo Fisher (TMO) underperform. The same letter later argues that AI may accelerate drug discovery, and the same two stocks reverse. Nothing in either company's fundamentals changed between those two days. The label changed, and that labeling is easy to post on a Reddit, Substack, or X group quickly. Labeling is a reductive cognitive error, and like other behavioral errors it creates the mispricings our process is built to exploit.

Hedging flow from the growth in option activity. This one is more often asserted than demonstrated, so it is worth being precise about what the evidence supports. The mechanism is real: a dealer who sells an option hedges it in the underlying, and the size of that hedge changes as the stock moves, so the dealer is a forced buyer or seller on a schedule set by the option book rather than by any view on the company. The

best available study, by Amaya, Garcia-Ares, Pearson and Vasquez using CBOE trade records reconstructed at one-minute frequency, finds that dealer gamma is typically positive, meaning that on an average day this hedging dampens volatility rather than amplifying it, and that the largest amplification they measure is about 3.3 percentage points of annualized daily volatility. We would not tell a client that retail option activity is making the market more volatile, because the research does not support it.

What it does support is more useful to us. Hedging flow is mechanical and indiscriminate: it trades the constituents of a basket or an index without reference to any individual company's fundamentals, which is to say it is a machine for manufacturing correlation. And it is extraordinarily concentrated. In the data the SEC assembled for its April roundtable, the top ten underlying securities now account for 31.7% of all options volume, up from 24.0% in 2012; contracts expiring the same day they are traded rose from 19.6% of volume at the start of 2022 to 28% by 2025; and on expiration days in December 2025, individual customers were 73.5% of contract volume excluding market makers. That flow concentrates in the index products and the handful of mega-caps that are themselves the largest index weights, which means hedging in a few names transmits into every basket that holds them.

Jesse Cohn of Elliott Management described the result in a post on 12 October 2026: "The stock market feels different right now because it is different. Winners keep winning, losers get pounded, earnings reactions are increasingly unpredictable and larger." He attributes it to the growth of passive investing alongside high-frequency, retail and short-duration capital dominating trading volumes, and concludes that this creates short-term risk and volatility but also opportunity for those who can identify value. We agree with the diagnosis and with where it leads. We believe more volatility creates more opportunity for a disciplined valuation-driven process, provided the process survives it. Which means the one thing we must not do is mistake a structural move for a fundamental one and buy a value trap on the way down.

What We Measure: Market Modalities

A market modality is a group of stocks that have been actually trading together over a chosen window. That is the whole definition, and every word of it is load-bearing.

Actually trading together means the grouping is discovered from daily returns, not assigned from a classification. No sector code, no industry table, no analyst's judgment about which companies belong with which enters the calculation. We take the returns, remove what the whole market did that day so that we are looking at relative movement rather than beta, and ask which names moved with which. Whatever groups fall out, fall out. We name them afterwards, by reading the membership, and the name is a convenience for conversation rather than an input to the result. This matters more than it sounds: a method that starts from sector labels can only ever rediscover sector labels.

Over a chosen window means a modality is a statement about a period, not a property of a company. A stock does not have a modality the way it has a balance sheet. It had one over the last eight months. Ask a different window and the answer may differ, which is a feature we rely on rather than noise we apologize for.

We call these modalities rather than clusters or factors deliberately. "Factor" implies a persistent, priced, compensated source of return, which is a claim we are not making. "Cluster" is accurate as a description of the arithmetic and suggests more permanence than the thing deserves. A modality is a mode of trading, a

way a set of stocks is being handled by the market right now. It describes behavior in a period, and when the period changes, so do those relationships.

Why This Method, and What We Rejected

We use HDBSCAN, hierarchical density-based spatial clustering of applications with noise. Two properties decided it, and the second is the one that does the real work.

It does not force every stock into a group. Hand a "k-means" algorithm a thousand stocks and ask for fifteen groups and it returns fifteen groups containing a thousand stocks, because that is what it does. So will a Gaussian mixture model. Neither has any way to express *this stock belongs nowhere*. The consequence is not a rounding error at the edges. Every genuinely idiosyncratic stock, every company whose price is being set by its own numbers, gets assigned to whichever group it resembles least badly, and those forced assignments dilute the groups they land in until membership stops carrying information. A stock that moves with its group and a stock that merely had nowhere else to go become indistinguishable, and the intra-group relationships that are supposed to be the output mean nothing.

For a stock picker that failure is disqualifying, because the names with no modality are not a residue to be tidied away. They are the ones trading on their own fundamentals, which is a unique condition, and a valuable source of potential idiosyncratic opportunities. A method that cannot find them is answering a different question than the one we asked.

It cuts each group at its own density, not at one threshold for the whole market. This is why HDBSCAN rather than DBSCAN, and it is the hierarchical half of the name. Plain DBSCAN takes a single global density threshold and applies it everywhere. HDBSCAN instead builds a minimum spanning tree over the distances and removes its edges one at a time in decreasing order of weight, which produces a full dendrogram in which every horizontal slice is exactly the DBSCAN answer at that threshold. It then condenses that tree, using a minimum group size to decide which splits are genuine branchings rather than a handful of points falling away from a group that persists, and selects the final set by excess of mass, which chooses each group at the height where it survived longest as the threshold tightened.

The market is why that matters. Cohesion varies enormously across the map. Utilities come in at 0.772 mean loading, among the tightest groups we have measured. Industrial distribution is 0.475. Any single global threshold tight enough to hold utilities as a distinct group would never resolve industrial distribution at all, and any threshold loose enough to find industrial distribution would already have swallowed utilities into something larger. Each group has to be cut at its own height, and a hierarchy is the only structure that permits it. A market contains both ferocious themes and loose affinities at the same moment, and an instrument that can see only one of those is not describing the market.

What we rejected. We tried modularity-based community detection, in the Louvain formulation, which partitions a correlation network by finding densely connected neighborhoods. It resolved the entire market into roughly four communities, the largest of which amounted to energy and not-energy, with mean within-community correlation around 0.12 to 0.17. Groups that large and that internally incoherent describe nothing a portfolio manager can act on. Louvain also partitions exhaustively, so it carries the same defect as k-means in a different costume: no stock may stand alone, and the method compensates for that constraint by drawing boundaries so coarse that membership stops distinguishing anything.

We also denoise before clustering, which is not optional. With roughly a thousand stocks, a sample correlation matrix contains a great deal of structure that is pure estimation noise, and random matrix theory locates the boundary between that noise and genuine common movement precisely. We discard everything below it. Skipping the step produces within-group correlations around 0.5 on the processed data that fall apart when measured on the real thing, which is worth naming as a failure mode because it is both flattering and invisible.

The Map Moves, and the Movement Is the Information

We run the analysis over several windows rather than one, and we compare the results. A risk model would treat the differences between those runs as estimation error to be smoothed. We treat them as the measurement.

In the August 2026 map there was a distinct group of speculative high-beta names sitting alongside precious metals, separate from the artificial-intelligence complex. By the September map it had dissolved. Twenty-five of its thirty-one members had moved into the AI-infrastructure group, which by then had grown to about a quarter of the index. The speculative complex did not change its fundamentals over five weeks. It stopped being a separate thing the market traded and became the market's expression of high beta, which is to say the AI trade had absorbed the function that the speculative basket used to perform.

In the same period, a single group containing utilities and defense primes split in two. Through the summer, regulated utilities and defense contractors had traded as one block, which makes no fundamental sense and perfect behavioral sense: both were being held as long-duration regulated cash flows at a moment when the market was trading duration. By October they had separated into a utilities group, the tightest in the entire map, and a separate group of defense primes and government services firms. Something in the way the market was handling duration had changed, and the separation showed up before any thesis we would have written about it.

This is the labeling cycle rendered as data. When a narrative forms, stocks that fit the label begin moving together regardless of what they own and operate. When the narrative shifts, they come apart. A map that was stable across every window would be describing industrial structure, which we can already read from a sector table. A map that moves is describing what the market is doing, which we cannot read anywhere else.

The Russell 1000 as It Trades Today

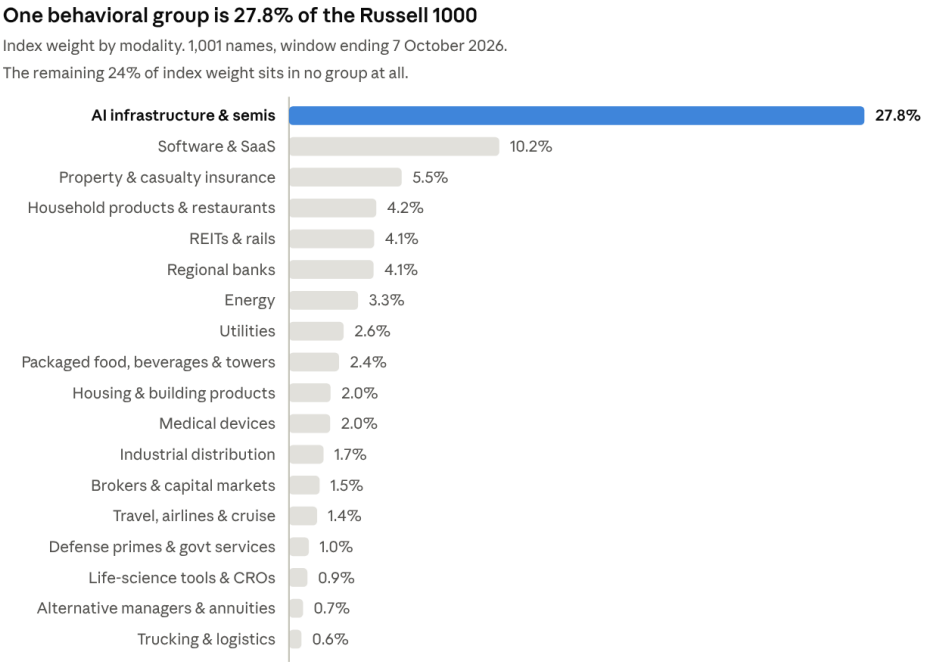


Exhibit 1. Cornerstone analysis of Russell 1000 constituents. 1,001 names, 2 February to 7 October 2026.

The shape is the finding. Eighteen groups, and the first is nearly three times the second. An investor holding the index is told they own eleven sectors, which sounds diversified. Over this window, better than a quarter of that index weight moved as one thing, and the eleven-sector description gave no warning of it. The long tail below Software is where the market's genuine behavioral variety lives, and it accounts for less index weight in total than the single block at the top.

One Bet Wearing Two Tickers

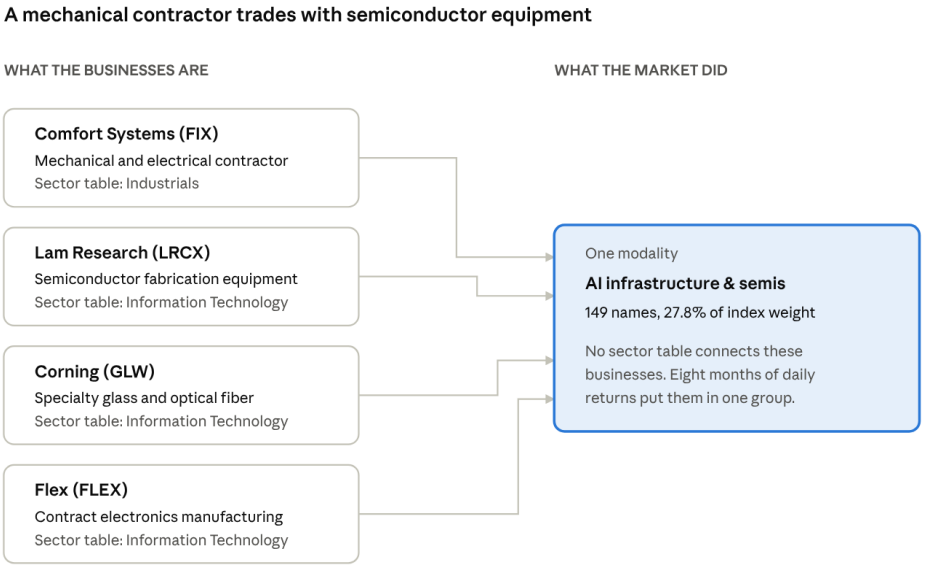


Exhibit 2. Four members of one Russell 1000 modality. Window ending 7 October 2026.

Comfort Systems (FIX) installs heating, ventilation and electrical systems in buildings. Lam Research (LRCX) builds the machines that etch silicon wafers. They share no customers, no inputs, no cost structure and no sector code. Over the eight months to October they moved as one group, because both are being bought as ways to own the data-center build, and because the capital expressing that view does not distinguish between the company that makes the equipment and the company that installs the cooling for it.

An investor holding both believes they hold a semiconductor position and a construction position. For as long as this modality persists, they hold one position twice. That is not an argument against owning either. It is an argument for knowing which one you are making, because the diversification that two sector codes imply is not there, and it will not come back until the narrative that created the group moves on.

How We Use It, and What It Does Not Do

Investors have two roles. First, they are analysts, identifying whether a stock is potentially attractive on its own merits. Second, they are portfolio managers, determining whether, when, and at what weight they should buy that attractive stock.

The modality map is a diagnostic. It surfaces a question for the team to investigate and debate. It does not size a position, it does not trigger a trade, and it does not overrule fundamental work. This is the same stance we take toward conventional risk models, and for the same reason: an instrument that reads market structure is describing the present, and the present is not an instruction.

When the map shows two holdings sitting in one group, the debate that follows is a fundamental one. Do we understand why the market is treating them as one thing? Do we believe the idiosyncratic characteristics of each will reassert themselves, and over what horizon? Are we being paid for the fact that they will trade together in the meantime? Sometimes the answer is that the shared behavior is temporary and the position is fine. Sometimes it is that we have accidentally doubled a bet, and we take one side down.

Four limits we hold ourselves to.

- It is descriptive, not predictive. The map says how a set of stocks traded over a window that has already happened. It makes no claim about tomorrow, and a group that has been tight for eight months can come apart in a week, as the ones above did.
- It is silent on valuation. A modality tells us nothing about whether a stock is attractively mispriced. It tells us what else will move when it moves. Those are different questions and we do not let the second one answer the first.
- It is a snapshot with a date on it. Every figure in this paper describes the market over a specific window ending in October 2026. Republished in six months, the groups will differ. That is the instrument working rather than failing, but it means no number here should be quoted as a standing fact.
- It inherits the limits of its data. The groups are built from price behavior, so anything that distorts prices distorts the map. We spend more effort on the integrity of the underlying price series than on the clustering itself, and we would encourage anyone building something similar to expect the same ratio.

Appendix: Numerical Methods of the Analysis

Key methodological details:

1. **Returns.** Daily total returns for every member of the universe, dividend-adjusted. Price-only series inject a false one-day move on every ex-dividend date, which matters because high-yielding stocks go ex-dividend on a shared calendar and would otherwise appear to move together for that reason alone.
2. **Winsorize.** Clip the most extreme half of one percent of daily moves at each tail. One takeover or one earnings gap will otherwise dominate a pair's correlation and manufacture a two-stock group out of a single news event.
3. **Market-normalize.** Subtract the universe's own return each day, so that what remains is relative movement. Without this step the first and largest common factor is simply the market, and the clustering collapses into one or two blocks.
4. **Denoise.** Discard the portion of the correlation matrix that is statistically indistinguishable from the noise a random matrix of the same dimensions would produce. On a thousand names and eight months of data, only about a dozen independent directions of genuine common movement survive this test.
5. **Cluster.** Density-based clustering on the denoised correlation distances, with a minimum group size of ten. Stocks in sparse regions are left unassigned.
6. **Tier.** Correlate every stock, including the unassigned ones, against each group's own return series, and sort them: core members assigned by the clustering, affiliated names correlating at 0.30 or better, weak tilts between 0.20 and 0.30, and idiosyncratic names below 0.20.
7. **Test stability.** Split the window in half, run both halves independently, and measure how much of the grouping survives. Bootstrap the same way. A map published without both numbers is an assertion rather than a result.

Parameters and Sensitivity

The settings are choices rather than discoveries, and we have tested the output against each.

Setting	Value	What moves if it changes
Winsorization	0.5% pooled tails	Looser produces spurious two-stock groups built from a single news event
Minimum history	Complete over the window	At 99%, surviving names still carry gaps, correlations go undefined and the pipeline returns zero groups with no error raised
Minimum group size	10	Raising it reduces group count roughly linearly without improving tightness
Cluster selection	Excess of mass	Leaf selection yields more, smaller, more homogeneous groups
Denoising	Marchenko-Pastur edge	Omitting it yields within-group correlations near 0.5 on processed data that collapse on the real thing

Why the Distance Metric Matters

We cluster on the Mantegna distance, the square root of twice one minus the correlation. It is a true metric, which means the triangle inequality holds and density-based clustering is well defined on it. Clustering directly on correlations, or on one minus correlation, is common and is not a metric, and the groups that result are not guaranteed to mean what they appear to mean.

An Honest Note on Our Judgement for the Density Setting

We run HDBSCAN with the minimum-samples parameter set to one. That setting makes every core distance zero, which collapses the mutual reachability distance back to the raw metric. What we are running is therefore a single-linkage dendrogram over the denoised correlation distances, condensed at minimum size ten and cut by excess of mass. It remains hierarchical and the variable-height cut is intact, which is the property we need. But the density smoothing that normally guards single linkage against chaining, where two real groups fuse because a handful of intermediate stocks bridge them, is not engaged. The denoising step, the size floor and the persistence-based selection are carrying that protection instead.

The symptom to look for is a group that is large, internally loose and reads as two themes stapled together. Raising minimum samples to five or ten would test it directly, and that test is open rather than settled.

Stability

We report two numbers with every map, using the adjusted Rand index, which measures how much two groupings of the same names agree after correcting for the agreement chance alone would produce. The half-split figure splits the window, clusters each half independently and compares. The bootstrap figure resamples names. The October 2026 Russell 1000 map returned 0.773 and 0.893 respectively. We publish both because a density clustering reported without them is an assertion rather than a result, and because the half-split figure in particular is noisy. It has ranged from 0.52 to 0.80 across our runs, and we read it as a diagnostic on the estimate rather than as a signal about the market.

Sources

- U.S. Securities and Exchange Commission, *Roundtable on Options Market Structure: Supporting Data*. Figures on zero-days-to-expiry share of volume, individual customer share on expiration days, and concentration in the top ten underlying securities.
- Diego Amaya, Pedro A. Garcia-Ares, Neil D. Pearson and Aurelio Vasquez, *0DTE Index Options and Market Volatility: How Large is Their Impact?*, working paper, January 2025. Dealer gamma reconstruction and its measured effect on realized volatility.
- Cornerstone Investment Partners, *A Stock Picker's Approach to Portfolio Construction and Risk Management*, July 2026. The portfolio construction framework this paper extends.
- Cummins Inc., Analyst Day 2026 investor presentation. Power Systems segment revenue of $7.5 billion for 2025, with data centers at 26% of that segment.
- Jesse Cohn, Elliott Management, October 12, 2026
- Modality figures from Cornerstone's own analysis of Russell 1000 constituents, windows ending August, September and October 2026.

For more information, please contact:

www.cornerstone-ip.com | marketing@cornerstone-ip.com | (404) 751-3850

Cornerstone Investment Partners is a registered investment advisor with the SEC. SEC registration does not imply a certain level of skill or training. Certain information contained herein has been obtained from third party sources and such information has not been independently verified by Cornerstone Investment Partners. No representation, warranty, or undertaking, expressed or implied, is given to the accuracy or completeness of such information by Cornerstone Investment Partners or any other person. Except where otherwise indicated, the information contained in this presentation is based on matters as they exist as of the date of preparation of such material and not as of the date of distribution or any future date. Recipients should not rely on this material in making any future investment decisions. The opinions expressed are those of the investment professionals and do not constitute investment advice. Past performance is not indicative of future results. References to specific securities, indices, and market data are for illustrative purposes only and do not constitute a recommendation to buy or sell any security. Portfolio characteristics are subject to change. The Market Modalities analysis is a proprietary analytical tool. Its output is a hypothetical analysis of historical index constituent returns, is not investment performance, and is subject to the limitations described in this paper. Data are as of 7 October 2026 unless otherwise noted. Cornerstone may hold, or may have held, positions in securities mentioned in this paper on behalf of clients. This should not be deemed as an offer or solicitation to buy or sell any product or security.