Anticipated Orders and the Measurement of Stock Demand Elasticity: Evidence from Scheduled Index Rebalances

Daniel Nathan*

This version: August 2026

Abstract

The elasticity of stock demand is usually computed as a price response divided by an institutional trade. That division implicitly assumes that the institutional trade is what the price responded to. We follow a $92 billion ETF whose required trades are fixed five trading days ahead by a mechanical and public index rule, and we observe both sides of the auction that prices them. The fund's trade is known in advance. By the closing auction, roughly ten times as much buying and selling interest has gathered around it, 9.29 basis points per basis point of scheduled demand, against 0.44 in stocks that trade in the same auctions but carry no rebalance. Those two sides nearly cancel, and the target explains only 2.1 percent of the variation in the imbalance they leave. The closing price moves with that imbalance, not mechanically with the fund's original target. Dividing the price response by different candidate measures of demand on the same observations gives values from -3.27 to 2.38.

JEL Classification: G12, G14, G23

Keywords: demand elasticity, institutional trading, index rebalancing, anticipated demand, price impact

*Hong Kong Polytechnic University. Email: daniel.nathan@polyu.edu.hk. I thank Filipp Dokienko for excellent research assistance. All errors are my own.

1 Introduction

A large body of work measures how far a price moves when an institution has to trade, and divides the move by the trade to recover the elasticity of stock demand. Index additions are the canonical case: an abnormal return is divided by the percentage of the stock bought upon inclusion, and the ratio inverted. The division assumes the measured order is the demand the price responded to. The literature cannot check it at the moment the price is formed, because the trade and the market it lands in are almost never observable together.

We study a fully transparent institutional trade by a large index rebalancer. The setting is the quarterly reset of the S&P 500 Equal Weight index. The literature has mostly used additions and deletions, which are rare, arrive in ones and twos, and carry news about the firm alongside the trade. Capitalization-weighted tracking generates no rebalancing trade at all, because weights move with prices and stay correct on their own. Equal weighting is the opposite. Weights drift the moment prices move, so every quarter the reset forces a trade in a median of 437 stocks at once, none of it discretionary. And the trade is not news about the firm. It follows mechanically from price movements that are already public, and embodies no judgment about the company of the kind an index addition does.

The fund we follow is the Invesco S&P 500 Equal Weight ETF, with $92 billion in equity assets by the end of our sample, by far the largest of the equal-weight trackers, and a typical recent reset moves $2.9 billion in each direction. Its required trade follows from the published rule and the fund's own holdings, so anyone can compute it. Rebuilt from the inputs public twenty trading days ahead, it already correlates 0.86 with the trade the fund ends up having to do, and 0.97 five days ahead. It becomes mechanically exact at the reference close, the close whose prices the rule uses, typically five trading days before the fund trades. The fund's later holdings confirm the trade is carried out. We follow 62 rebalances from 2010 through mid-2026. For 32 of them we also have the exchanges' auction messages, which report the buying interest, the selling interest, and the imbalance between them second by second through the final ten minutes. For 30 of those our membership record reaches back far enough to build the scheduled target as well, and they carry 3.7 million of those messages.

We first show that the fund's planned trade is only one part of a much larger market. By the first auction update at 15:50 Eastern, stocks with larger scheduled trades also attract much more buying and selling interest: a one-basis-point increase in the scheduled trade is associated with 9.69 basis points of displayed auction interest. The relationship is positive in all 30 rebalances. The interest divides almost evenly, 5.10 basis points on the side the fund must trade and 4.59 against it. Every stock here is an S&P 500 member with a busy close of its own, so interest would gather around these names on any day. We use two controls to separate that baseline from what the rebalance adds. Holding the date and changing the stock, Russell 1000 members outside the index whose prices have moved the same way relative to their peers since the previous reset, and which trade in the same auctions, attract 0.44 basis points of buying and selling interest per basis point of that relative move, against 9.29 for index members, both estimated on the 25 rebalances where the two groups can be compared. Holding the stock and changing the date, the same names carrying the same target as a placebo on a matched control day attract 23.8 times less extra interest per basis point of that target. Take a stock with the average scheduled trade, 3.34 basis points of its shares outstanding. It sits in an auction displaying about 86.0 basis points of buying and selling interest, and about 54 of those, 63 percent, do not depend on how large its trade is. The trade the researcher divides by therefore governs only the part of the auction that scales with it, and at a typical event that part is the minority.

The closing price responds to the net imbalance between buyers and sellers, not to the fund's scheduled trade itself. The scheduled trade turns out to be a poor guide to that imbalance. Across stocks within the same auction, it explains only 2.1% of the variation in the displayed imbalance and points in the correct direction in only 53.2% of stock-events. On average, stocks the fund buys do lean toward buy imbalances, but the relationship is weak and varies substantially across rebalances. This matters because elasticity estimates divide the price response by the scheduled trade. If that trade is only loosely related to the imbalance the price actually clears, the resulting elasticity can be far too large or far too small, even when the underlying demand curve is unchanged.

Our setting lets us apply that inversion to four candidate quantities over the same price window, on the same observations, with the same estimator. Naming the scheduled target as the demand gives -2.91. Naming the displayed imbalance gives 2.38. Splitting that imbalance into the part our conditioning predicts and the part it leaves over gives -3.27 and 2.28. Nothing changes across those four numbers except which object is called the demand shock, and the sign of the answer changes with the naming. Nor is there a single price of the imbalance to invert. If the close paid one rate for the imbalance, the fitted component and the conditional residual of that same imbalance would have to carry the same coefficient, and they differ by -0.61 with an event-clustered standard error of 0.14 across 30 rebalances. The elasticity a researcher reports is therefore not read off the data alone. It is chosen, in the step where one quantity among several is called the demand, and our estimates put a range on how much that choice is worth.

One objection to our setting, and to what one can take from it, is that these trades are small. An index addition has trackers buying seven or eight percent of a stock's shares; the median trade we study is 1.35 basis points. At the scale of an addition the trade would dominate the auction it lands in, because the part of the auction unrelated to the trade is roughly fixed and a trade five hundred times larger swamps it. We do not claim our finding transfers one-to-one to additions. But additions are not what most of this literature estimates. Elasticities are recovered far more often from ordinary institutional flows such as mutual fund flows, benchmarking-driven trading, and instrumented order flow, and those are basis points of shares outstanding, not percent.[1] A

[1] Flow-induced trading by mutual funds (Coval and Stafford, 2007; Lou, 2012; Wardlaw, 2020); trading forced by benchmarking mandates (Pavlova and Sikorskaya, 2023); and demand systems estimated from institutional holdings (Koijen and Yogo, 2019; Gabaix and Koijen, 2021; van der Beck, 2026). In each the measured quantity is a small

flow can be entirely unrelated to information about the stock and still be a small part of the orders that set its price. The problem is worst where the literature works most, and our estimates say why. The smaller the trade relative to the market that meets it, the less of that market has anything to do with the trade.

These measurements bear on how demand elasticity is estimated. The object of interest is a price response divided by a demand quantity, and the usual experiment supplies a clean quantity measured at the moment the order is formed. Our estimates say that a clean initiating quantity is not the net of the orders the price clears. Nothing goes untraded. The question is what market forms around the known order before the benchmark price is set. Recovering a demand elasticity requires answering two distinct questions. First, is the institutional order a good proxy for the net order flow that the market clears? Second, if so, does the price response identify the slope of the demand curve? We study the first question. Exogeneity of the institutional order does not imply that it is the correct denominator for the price response.

Existing studies estimate demand elasticity by dividing a price response by a measured institutional order. That procedure implicitly assumes the institutional order is a good proxy for the net order flow that the market ultimately clears. If it is not, then the estimated elasticity mixes two things, how prices respond to net order flow and how well the observed institutional order predicts that net order flow. Our data let us observe part of that transformation directly. The scheduled order explains only 2.1% of the cross-sectional variation in the displayed imbalance that forms before the close, and points in the correct direction only 53.2% of the time.

We do not claim the numerical mapping we estimate applies to every institutional trade. Equal-weight rebalances are unusually predictable, so they are a setting in which anticipation is especially important. The broader point is general. Whenever other investors can react to a known institutional order before the price is set, the relationship between the initiating order and the quantity that ultimately reaches the market is an empirical question rather than an accounting identity. Our setting is valuable because the closing auction lets us observe part of that transformation directly. One piece of it travels outside our setting. In Russell reconstitutions, a different index family on a different calendar with a separately constructed predetermined quantity, the displayed imbalance is again far more dispersed on event dates than on matched control Fridays (Section 6.5).

Our work connects three literatures. First, a large literature estimates the elasticity of stock demand from the price response to measured institutional demand, using index changes, institutional holdings, or instrumented order flow. The designs differ, but each must connect a price response to a measured quantity (Koijen and Yogo, 2019; Gabaix and Koijen, 2021; van der Beck, 2026; Greenwood and Sammon, 2025; Shleifer, 1986; Harris and Gurel, 1986; Wurgler and Zhuravskaya, 2002; Petajisto, 2009; Chang, Hong, and Liskovich, 2015; Coles, Heath, and Ringgenberg, 2022; Chinco and Sammon, 2024). We examine the normally unobserved step

fraction of shares outstanding rather than the several percent an index addition delivers.

between that quantity and the market state present before the benchmark price is set. Greenwood and Sammon (2025) document that the price effect of S&P 500 additions has fallen by roughly a factor of twenty over three decades, and argue that institutions increasingly position to absorb index demand, and Li and Lin (2026) show that the effect is larger for less diversifiable order flow. Our sample does not span that historical transition. We observe instead the market that currently forms around a predictable institutional order shortly before its benchmark price is determined.

Second, recent work asks whether an exogenous demand shock can identify a structural elasticity once anticipation, expectations and market dynamics are taken into account. Bastianello (2026) shows that equilibrium multipliers need not identify elasticities when the shock moves expectations of future payoffs, He, Kondor, and Li (2025) show that absorbing a shock changes the covariance matrix investors face, Davis, Kargar, Li, and Silva (2026) argue that no single context-free elasticity exists in dynamic markets, and Fuchs, Fukuda, and Neuhann (2026) show a static failure through cross-asset price geometry. van Binsbergen, David, and Opp (2025) show that anticipation can prevent an exogenous demand shock from identifying a structural elasticity, even when the relevant quantity is taken as given. We ask a prior measurement question, whether the institutional order observed by the researcher is a good proxy for the market quantity that forms before the benchmark price is set. Their concern is how to interpret the price response conditional on a quantity. Ours is whether the observed order is the right quantity in the first place. What we add is an intermediate market state that these frameworks leave latent, the displayed buying and selling interest present as the closing auction develops.

Third, related work corrects the measured demand quantity, studies trading around predictable orders, and analyses closing auctions. Pavlova and Sikorskaya (2023) show that the passive holdings change understates the shock because benchmarked active funds trade mechanically alongside index funds, Sammon and Shim (2026) measure the share of index demand absorbed by firms through issuance and buybacks, Escobar, Pandolfi, Pedraza, and Williams (2026) use transaction-level data to trace who generates and who absorbs rebalancing demand, and Wardlaw (2020) shows that a widely used flow-induced-trading measure is mechanically related to the returns it is used to explain. Models of predictable trading show that announced orders attract counterparties (Admati and Pfleiderer, 1991; Lou, Yan, and Zhang, 2013), with consequences for market quality (Bessembinder, Carrion, Tuttle, and Venkataraman, 2016), price adjustment migrating into the pre-implementation window (Beneish and Whaley, 1996), passive funds executing predictable trades (Li, 2021), anticipatory speculators supplying liquidity at reconstitution (Pegoraro, Sammon, and Shim, 2026), and predictable closing demand (Zhao, 2026). Auction studies examine how displayed imbalances evolve and how they are priced (Jegadeesh and Wu, 2022), who trades at the close (Bogousslavsky and Muravyev, 2023; Hu and Murphy, 2024), and price impact across trading mechanisms (Goyal, Jegadeesh, and Wu, 2026). These papers model the transformation or infer it from prices and volume, and they define order flow relative to the auction's own contemporaneous imbalance. We measure both displayed sides directly, oriented by an institutional target fixed days before any auction message exists, which is what allows the scale of each side to be regressed on the size of the initiating order. The problem is not specific to a closing mechanism: recurring institutional cash needs generate predictable trading in ordinary continuous markets as well (Etula, Rinne, Suominen, and Vaittinen, 2020; Nathan, Suominen, and Tasa, 2026), Duffie (2010) shows that capital reaches a shock slowly enough for its price effect to depend on who is already present, and Behmaram and Davis (2026) find the measured elasticity varies with how heavily a stock is indexed.

2 The Rebalance and the Data

The experiment has four stages, set out in Figure 1. Public information predicts the target, the index rule fixes the final stock-level target, the exchanges reveal the evolving closing-auction interest, and the official close sets the benchmark price. The *reference close* is the close whose prices the index provider uses to fix the new weights. The *official close*, on the rebalance day five trading days later, is the close against which the fund's execution is judged. Two price objects follow. The *order–price multiplier* is the price response per basis point of scheduled demand, measured from the reference close to the official close. The *pricing of the displayed imbalance* is the price move from a given auction time to the close, per basis point of the imbalance displayed at that time.

2.1 The Predetermined Fund Target

The S&P 500 Equal Weight Index holds the same companies as the S&P 500 but resets the weights of the companies that stay in the index to approximately equal weights at scheduled quarterly rebalances. Weights drift between resets as relative prices move, so the next reset requires any fund tracking the index to sell stocks that have become overweight and buy stocks that have become underweight. A cap-weighted fund cannot supply the same experiment, however large it is. Cap weighting rebalances itself, because a stock whose price rises gains weight in the fund and in the index by the same amount, so the portfolio stays on target without anyone trading. Under equal weighting that same price move pushes the stock above the share it is entitled to while that share stays where it was, and closing the gap takes a trade. The much larger cap-weighted trackers do trade, on share-count and float updates, on index additions and deletions, and on flows, none of which produces a rule-fixed quantity that varies widely across stocks on the same day. The fund buys about as many dollars of stock as it sells, and both the buys and the sells are large-cap S&P 500 members. The rebalance therefore moves demand across stocks rather than adding demand to the market as a whole.

For the largest tracker, the Invesco S&P 500 Equal Weight ETF (RSP), the *scheduled target* in stock i at event e is

$$Q_{ie} = \frac{s_{ie}\, A_e/N_e - V_{ie}^{ref}}{ME_{ie}^{ref}}, \tag{1}$$

where A_e is the fund's equity assets at the reference close, which usually falls five trading days before the rebalance day, N_e is the number of index *companies*, V_{ie}^{ref} is the value of the fund's position in stock i at that close, and ME_{ie}^{ref} is the stock's market capitalization at the same close. The index gives each company one equal-weight slot. Almost every company in the index has a single listed share class and takes the whole slot, so s_{ie} equals one for those stocks. A few companies have more than one listed class, and those classes have to share the single slot. For them, s_{ie} is the class's own market value of freely tradable shares as a fraction of the company's, so the classes divide one slot between them and no class receives a whole slot. Section A.2.2 shows why this is the right allocation and what it costs to get it wrong. Everything on the right-hand side is public at that point, so the target follows mechanically from information anyone can assemble. It fixes the rebalance component at the reference close. Creations, redemptions or other holdings changes between that close and the trade can still move the shares the fund actually has to buy or sell, which is one reason the target and the executed order need not coincide. The numerator is the dollar amount the fund must trade in the stock and the denominator is the stock's market value, both measured at the same reference close. Dividing one by the other therefore gives the number of shares the fund must trade as a fraction of the shares outstanding. Unless stated otherwise, all order and auction quantities are expressed in basis points of shares outstanding, so one basis point of Q is 0.01% of its shares outstanding. For a company with a hundred million shares that is ten thousand shares. Every use of Q below refers to this final scheduled target.

A rebalance can be on the index calendar and still not move the fund's portfolio, so we check every scheduled date against what the fund actually did, comparing the change in its reported holdings with the change the equal-weight reset implies. 62 rebalances from 2010 through mid-2026 pass that check. 32 of them, spanning June 2018 through June 2026, also have the second-by-second auction messages the analysis needs. The last two of those rebalances postdate our historical membership record, so the scheduled target cannot be built for them. Analyses that depend on the target therefore use the 30 rebalances among these. Analyses of how the displayed imbalance is priced do not depend on the target and use all 32. Each design takes the largest set of events meeting its own mechanical support conditions, and when a design was settled is never what defines its sample. One comparison below, the within-event growth design of Section 3.3, uses the 25 quadruple-witching rebalances whose stock panel is balanced across the four clocks it compares, 15:50, 15:55, 15:59 and the last observation before the close. Realized holdings changes track the quantity the public rule implies closely (Table 3), with a full-sample slope of about 0.82, a median within-event R^2 of about 0.81, and a median correlation around 0.90. The scheduled target is therefore strongly reflected in later holdings, although the holdings data do not establish one-for-one execution of every required share. Table 2 reports the sample and the auction quantities, and Table 3 the validation, and Table A2 in the appendix compares this design with the Russell reconstitution design of Section 6. Later sections work with smaller subsets of these events, and each subset is defined by which auction messages exist rather than by how the results come out. Appendix A.1 gives the derivation.

The target becomes exact only gradually, and what pins it down is mostly the movement of prices. That is the same force that generates the rebalance in the first place, since relative price drift pushes weights away from equal. By five trading days before the reference close, the fund has already published the holdings report that the target will be built from, and this is true in every event. What is still uncertain at that point is almost entirely prices, and Appendix A.1 attributes the forecast error to the rule's inputs one group at a time.

The trade is also highly predictable long before it is final. We rebuild each event's trade from the rule using only the holdings, fund assets, and market capitalizations that were public on a chosen date ahead of the reference close. Figure 2 traces that reconstruction across horizons. It has a correlation with the final target of 0.86 when it is built twenty trading days ahead, 0.93 at ten trading days, and 0.97 at five, and the regression slope relating the two is one at every horizon.[2] The trade is therefore highly forecastable well before it is final, and it becomes mechanically fixed only at the reference close. The size of the scheduled target has grown with the fund. Counting only one side of the trade, the median rebalance moved \$0.6 billion of stock in 2018, against \$2.9 billion from 2023+, with a maximum of \$6.1 billion.[3]

That forecastability is also this design's principal exposure, and we state it here rather than leave it implicit. We measure the order–price multiplier from the reference close forward, because that is the window in which the target is mechanically fixed. But the trade is substantially computable well before that close, and any price an outside investor extracted by positioning during the computable window lies outside our measurement window and is invisible to us. A reader is entitled to the reading that the price impact of these trades moved earlier, into the days or weeks over which the target became forecastable, and that the orderly auction we document is a symptom of that migration rather than evidence that the impact is small. We cannot rule this out. The natural test would compare pre-reference price paths against a counterparty-matched control, and the placebo constructions available to us do not support one: the same rule that makes the target computable makes it computable for every index member at once, so there is no contemporaneous set of comparable securities for which the trade was unforecastable. What we measure is therefore the terminal segment of a longer process, and every quantity in this paper should be read as conditional on that segment. The measurements stand on their own terms, and we do not claim they bound total impact.

RSP is not the only fund that runs this rule, and that matters for how our estimates should be read. RSP ended the sample period with \$92.2 billion in equity assets, and the other funds tracking the same equal-weight benchmark hold a fraction of that between them among the funds

[2] Sources are the forecastability rows of Table 3. The reconstruction uses public rule inputs, so it describes what any outside investor could have computed.

[3] We also ask how much of the final dollar trade a forecast gets right in both direction and size. For each stock we check whether the reconstruction made a month ahead points the same way as the final trade, and where it does we take the smaller of the two dollar amounts. Forecast dollars are valued at the date the forecast is made and final dollars at the reference close. We sum across stocks and take one half of that sum, so the figure is measured on the same one-sided basis as the turnover numbers above. At ten and five trading days before the reference close, the median share rises to 76 and 86 percent.

we can confirm, and mandates we cannot observe, internal or synthetic, plausibly add more. We use RSP because it publishes holdings, which is what makes the stock-by-stock target observable at all. Q is therefore one observed piece of the demand the rule creates in each stock, and the demand the rule creates in total is larger.

2.2 Following the Market to the Close

The rebalance is benchmarked to the official close, meaning the fund is judged on how close its execution comes to that day's closing prices. The auction messages do not identify RSP's own orders. What we see instead is the buying and selling interest the exchanges display as the auction fills up around a target that was fixed days earlier and will be measured against that close. The NYSE and Nasdaq closing auctions collect orders that ask to be filled at whatever the closing price turns out to be, orders that ask to be filled at the close only if the price is no worse than a stated limit, and a few related order types. In the final minutes before the auction runs, each exchange broadcasts what has arrived so far. At any moment the feed reports how much buying interest and how much selling interest are present, together with the imbalance left between them, the excess of one side over the other carried with a sign showing which side is larger. The auction feed we assemble contains 22,258,339 messages over the sample period, of which 18,510,938 fall on the control dates described below, ordinary busy closes that carry no equal-weight rebalance. The 30 rebalances themselves carry 3,747,401 messages, stock by stock and second by second, which is what lets us watch the closing state form.

What the feed shows is displayed interest, and its limits need stating. It reports quantities without the identities or motives behind them. It reports only what has been submitted in a form the exchange displays, so we do not see interest being worked in the continuous market, interest submitted after the last message we observe, or interest carried in order types the venue does not publish. One of the published quantities also carries no information of its own. The exchanges publish a paired quantity, the amount of buying interest that can be filled against selling interest, and the paired quantity is simply whichever of the two displayed sides is smaller. It is determined by the two displayed sides, so the quantities with distinct economic content are the two sides and the difference between them.

The messages are the exchanges' own auction broadcasts, Nasdaq's Net Order Imbalance Indicator and the NYSE's Closing Auction Imbalance Information. We obtain them as raw historical messages from Databento, which delivers the two venues in one common set of fields. Our main observation time is 15:50, the first moment at which both venues publish the state, and we follow the state at later times as the auction approaches. Every clock time in the paper is Eastern, the time zone in which both auctions run. We observe all eleven of these times on all 30 rebalances, so no gap in the data forces the choice of 15:50. We choose it because the scheduled target is fixed days earlier, while anything measured after 15:50 partly reflects investors reacting to an imbalance the exchange has already shown them. Appendix A.2 reproduces example messages, defines each field we use, and details how the two venues' auctions clear.

2.3 What the Auction Reveals

Stocks are indexed by i, rebalance events by e, and time within the closing auction by s, and we suppress i and e when unambiguous. The institution's scheduled target is Q. The part of that target represented in the displayed auction imbalance at auction time s is $q(s)$, and the net of all other displayed interest, carried with the same sign, is $C(s)$, whether it belongs to investors responding to the rebalance or to anyone else trading that day. The displayed imbalance is their sum,

$$R(s) = q(s) + C(s), \tag{2}$$

and $R(s)$ is the signed displayed auction imbalance at auction time s. We refer to it as the *displayed imbalance*. When a direction is needed we take the direction of the scheduled target as the reference, so that same-side interest is displayed interest pointing the same way as the target and opposite-side interest points against it. In plainer terms, same-side interest is displayed interest pushing the auction the same way the fund's trade would, and opposite-side interest pushes against it. The feed reports no identities or motives, so neither side is attributed to any particular investor, and the fund's own displayed order sits inside the same-side total.

One reading of R must be ruled out at the start. R is not the part of the scheduled target that fails to find a counterparty. Every share the fund buys is sold by someone at the clearing price, and the auction always clears. $R(s)$ is the net position of the displayed auction as it stands at time s, of which the fund's own order is one predetermined contribution. It can therefore be small when the fund is trading heavily, and it can point the opposite way from the fund's trade if everyone else together wants to sell more than they want to buy. When we say the target explains little of R, we mean it forecasts poorly the net the displayed orders combine to produce at the moment we observe them, not that some fraction of the target went untraded. That displayed net is an interim quantity: it excludes undisplayed interest and anything submitted after the observation time, so it is not the final quantity the cross clears.

We observe Q and R. We do not observe q and C separately, and that separation is the one ambiguity that runs through the paper. The holdings check establishes that the fund's holdings move to the new weights over the days on which the rebalance is implemented. Prior evidence indicates that index-tracking funds concentrate their rebalancing trades on the day the new weights take effect and execute most of them at closing prices (Li, 2021). Trading at the closing price is still not the same as trading in the closing auction: a broker can fill the fund away from the exchange at a price set equal to the official close, and the holdings report looks the same either way. The fund may not put its whole target into the displayed auction, either because it trades elsewhere or because it trades less than the full amount, and everyone else may add interest of their own. The data do not tell us how much of each is present.[4] Equation (2) decomposes the state into two unobserved components. It is an accounting statement, and it asserts nothing about whether C responds to Q.

Not being able to split the displayed imbalance into the fund's part and everyone else's leaves our measurement intact. Standard demand experiments observe or instrument Q and then estimate a price response, with the step in between invisible. We place our measurement inside that step, at the state that forms after other investors have responded and before the benchmark price is set. Throughout, λ_s is the loading of the displayed imbalance on the target at auction time s. We regress the displayed imbalance on the target within each event, $\lambda_s = \mathrm{Cov}(R(s), Q)/\,\mathrm{Var}(Q)$, and report two summaries of the resulting slopes: the pooled slope, which weights each rebalance by its within-event dispersion in the target, and the equal-event average, which weights each rebalance the same. Every slope quoted in the paper is tagged with which of the two it is.

In the usual convention, a demand multiplier M is the price response per unit of demand, and the demand elasticity ε is read off as its inverse, $|\varepsilon| = 1/|M|$, so an inelastic demand curve implies a large multiplier (Greenwood and Sammon, 2025; Gabaix and Koijen, 2021). For the scheduled target, denote the reduced-form multiplier by

$$M_Q = \frac{\mathrm{Cov}(\Delta p,\, Q)}{\mathrm{Var}(Q)}, \tag{3}$$

for the stated price window Δp. Everything in this paper turns on whether Q is the right quantity to place in that denominator.

2.4 Empirical Design

We ask two related questions about the displayed imbalance. How much of it does the scheduled target explain on its own? And how much more can be explained once we also use the opposite-side interest standing in the same auction message at the same moment? The first question is genuinely predictive, since the target is known days in advance. The second is not, because the opposite-side interest is measured at the same instant as the state it helps explain, so the answer describes how the pieces of the state hang together at a point in time rather than what anyone could have forecast. Let $L_{ie}(s)$ denote the displayed interest standing on the side opposite the scheduled target, read from the same auction message at auction time s as the state itself. Because L is defined relative to the direction of the trade while R carries the exchange's own sign, we orient the imbalance before estimating. Standing opposite-side interest pulls the exchange-signed imbalance down when the fund is buying and up when it is selling, so a single coefficient on L cannot represent both

[4] One explanation for the weak link between the target and the displayed imbalance documented below would be that the fund does not trade its full target. That does not appear to be what is happening. Replacing Q with the change in the fund's actual holdings, on the identical sample, moves the within-event correlation with the displayed imbalance from 0.13 to 0.14 and the out-of-event R^2 from 0.02 to 0.03. On the three events where the holdings change tracks the target almost exactly, the correlation moves by -0.01 to 0.02.

directions. Writing $\sigma_{ie} = \text{sign}(Q_{ie})$, for each event e and auction time s we regress the oriented imbalance $\sigma_{ie} R_{ie}(s)$ on $|Q_{ie}|$ and $L_{ie}(s)$ using all other events, so the coefficients used for event e come only from events $e' \neq e$. Let $(\widehat{a}_{es}, \widehat{b}_{es}, \widehat{c}_{es})$ denote the resulting coefficients. Rotating the fitted value back to the exchange's sign, the held-out fitted value and the conditional residual are

$$\widehat{R}_{ie}(s) = \sigma_{ie} \left(\widehat{a}_{es} + \widehat{b}_{es} \, |Q_{ie}| + \widehat{c}_{es} \, L_{ie}(s) \right), \tag{4}$$

$$u_{ie}(s) = R_{ie}(s) - \widehat{R}_{ie}(s). \tag{5}$$

We call $u_{ie}(s)$ the *conditional residual.* It is the part of the displayed imbalance left unexplained when the target and the opposite-side interest standing at the same moment are combined using coefficients estimated on other rebalances only. Holding out the whole event this way stops a rebalance from setting the standard it is then judged against. The rule is deliberately plain, an ordinary least squares fit in levels with no transformations, regularization, or variable selection. A more flexible rule might explain more of the displayed imbalance, and our evidence should be read as covering linear rules of this class. Appendix A.4 states the specification and the evaluation metric.

To know whether the conditional residual is unusually large on rebalance dates we need something to compare it with. We use monthly option-expiration Fridays, which already carry busy closes of their own and are therefore a demanding comparison. On each matched date we build the residual with the same coefficients and the same stocks as the rebalance the date is matched to. A control date has no scheduled order. We set $Q = 0$. There is no trade direction to orient by, so we leave the imbalance in the exchange's own sign, which is what $\sigma = +1$ means here. The conditioning variable becomes buy-side interest, because with no trade there is no direction against which to define an opposite side. Appendix A.4 states both, and the comparison should be read with them in view. The rebalance, and not the stock, is the experimental unit throughout. Thousands of stock observations identify a cross-sectional slope inside one date, but they all come from that one date, so they do not amount to thousands of independent rebalances. We therefore summarize event by event, give every event equal weight, and count rebalances rather than stock observations when we report how often a result comes out one way.

We cannot decompose u. It may contain other scheduled flows, strategic responses to the rebalance, ordinary institutional trading, information that arrives after the auction time we measure at, or part of the fund's own target that has not reached the displayed auction. The feed does not separate these. It does measure how much of the displayed imbalance remains unexplained once the target and the standing opposite side are accounted for, and how that quantity evolves as the close approaches.

Q_{ie} is the rule-implied scheduled target. $R_{ie}(s) = q_{ie}(s) + C_{ie}(s)$ is the displayed auction imbalance, $\widehat{R}_{ie}(s)$ its held-out fitted value, and $u_{ie}(s) = R_{ie}(s) - \widehat{R}_{ie}(s)$ the conditional residual. λ_s is the projection of the state on the target, and M_Q the reduced-form price multiplier on the target.

Inference throughout is clustered on the rebalance, which is the experimental unit, using a wild cluster bootstrap with the six-point weights of Webb (2023) wherever we test a coefficient on the scheduled target. We report bootstrap confidence intervals for magnitudes, t statistics only where the claim is that a difference exists, and specification tests as χ^2 or an exact p. Appendix A.3 gives the construction.

3 How the Market Forms Around Predetermined Demand

A multiplier regression compares the price response with the scheduled target. Before that comparison can mean anything, one has to know what else is present when the price is set. The buying and selling interest that assembles around the target is large and closely tied to it.

3.1 Buying and Selling Interest at First Publication

We begin by asking how much market has gathered by the time the auction state is first published. We regress the displayed quantity of stock i at rebalance e, measured in basis points of shares outstanding, on the size of the target the rule sets for that stock,

$$Y_{ie} \;=\; \alpha_e \;+\; \beta\,|Q_{ie}| \;+\; \varepsilon_{ie}, \tag{6}$$

where Q_{ie} is the scheduled target at the reference close and α_e is a rebalance fixed effect, so β is identified across stocks inside one auction. Y_{ie} is same-side interest, opposite-side interest, their sum, or the paired quantity, with the two sides oriented by the direction of the target.

Stocks with larger scheduled targets sit inside more displayed interest on both sides of the auction, and Table 4 reports the estimates. At 15:50, the first moment at which both venues publish the auction state, two stocks whose targets differ by one basis point differ by 9.69 basis points in *gross interest*, a marginal relationship on top of a large common level rather than the ratio of auction to order (Figure 3 and Table 4, Panel A), meaning the buying interest and the selling interest added together, with a bootstrap interval from 8.16 to 11.20. Each rebalance carries its own fixed effect, so the comparison runs across stocks inside the same auction and the common intraday path of a quadruple-witching close is absorbed. The relation is positive in all 30 rebalances, and it divides almost evenly, 5.10 basis points on the side of the target and 4.59 basis points against it. A single slope is the right summary here, and we verify the claim rather than assume it. A constant-elasticity alternative would imply a marginal response that falls steeply with the size of the scheduled target, so the two forms make sharply different predictions about the largest trades. Comparing them against binned conditional means of gross interest within each rebalance, both scored in levels across 300 cells, ten deciles of the target within each rebalance, the level form gives a root mean squared error of 18.2 against 39.6 for the constant-elasticity form, whose errors are positive in every decile and largest among the biggest trades. Adding a quadratic term in the target gives 0.0357 (0.0160), which is if anything mildly convex rather than the concavity a constant-elasticity form requires. The marginal response is therefore close to constant across the size distribution.

Because the equal-weight target is a common dollar amount, small members carry large targets in basis points of their own capitalization, and every quantity here is measured per share outstanding. The slope is therefore estimated on a regressor negatively related to firm size, with an outcome that is mechanically larger for smaller firms, so we report what happens when that gradient is controlled. Adding log market capitalization and log twenty-day ordinary turnover to the same regression on the 26 rebalances where both are available lowers the slope from 9.19 on those rebalances, against 9.69 on all 30, to 7.36, with an event-clustered standard error of 0.75, a reduction of about 20 percent. We report both and prefer neither: the uncontrolled slope is the response of the auction a fund actually meets, and the controlled slope is what remains once securities of comparable size and turnover are compared. The loading of the displayed imbalance on the signed target is unaffected by the same controls, moving from 0.405 to 0.413, so the result the paper rests on is not a size gradient.

What makes that consistent with a small proportional response is the intercept. At the mean, 53.79 basis points of gross interest, or 63 percent of it, does not move with the size of the scheduled target across stocks. The event fixed effect absorbs everything common to the rebalance, including any interest the rebalance itself draws in equally across the stocks it touches. It is therefore not a measure of what the auction would have held with no rebalance. The same-date comparison in Section 3.2 is built to answer it. The level slope measures what the rebalance adds at the margin, not what fraction of the auction it constitutes, and a large intercept is exactly the case in which those two readings differ without either being wrong. It also settles where the gap between the order and the auction is widest. Under the fitted relation, gross interest divided by the order is $53.79/|Q| + 9.69$, so the auction is a larger multiple of the order the smaller the order is. The arithmetic is straightforward; the economic content is the size of the intercept. At the median rebalance 63 percent of the auction does not scale with the trade, so the term that dominates for small orders is most of the auction rather than a rounding allowance. The intercept is an event-level average, and the fitted ratio $\alpha_e/|Q_{ie}| + \beta$ is event-specific, so it is not the intercept of any particular rebalance. The practical reading for other settings is to compare the shock to the auction it arrives in, not to shares outstanding.

Other funds track the same benchmark, so the demand the rule creates in each stock is larger than RSP's own target, and dividing by that target alone overstates the response per unit of rule-driven demand. Their holdings are near-exact scalar multiples of RSP's, so they add no variation of their own and cannot be separated from it as an independent component. What they change is the denominator. We rebuild the target in family units rebalance by rebalance, multiplying each rebalance's target by the ratio of confirmed same-rule assets to RSP's own, and we re-estimate equation (6) on the rescaled quantity rather than dividing the pooled coefficient by a median ratio, which would misweight rebalances whose ratio moves with their identifying variance. Buying and selling interest per basis point of rule-driven demand falls from 9.69 to 7.2 on the same-rule assets we can confirm, and to 6.1 under a deliberately generous assumption about the funds we cannot observe. The confirmed ratio is a floor, since internal mandates, direct indexing, and equal-weight derivatives are unobservable by construction.

RSP dominates the rule. Across the 30 rebalances it holds a median 74 percent of the same-rule assets we can confirm, ranging from 63 to 80 percent, and it is a median 4.8 times the size of the next largest fund running the rule. The other same-rule funds hold the same stocks in nearly the same proportions, so their required trades are close to scaled copies of RSP's own. Same-rule rebalancing therefore enlarges the quantity the response is divided by, which is the rescaling above, and it cannot account for the part of the displayed imbalance the target fails to explain. Demand tied to other index rules is a separate matter and is part of what we cannot separate from the fund's own.

For the multiplier, the point is that the known target predicts how much trading will gather far better than it predicts the imbalance that remains between the two sides, a contrast the estimates below quantify. Figure 4 shows the interest that gathers and the imbalance that remains together.

3.2 A Same-Date Counterfactual

This comparison sets stocks the rebalance touches against stocks trading in the same closing auctions, on the same dates, that it does not touch. We pool S&P 500 members with Russell 1000 members outside the index and regress displayed quantity on index membership, on the size of the stock's own predetermined demand, and on the interaction of the two,

$$Y_{ie} \;=\; \alpha_e \;+\; \gamma\,\mathrm{SP}_i \;+\; \beta_1\,|D_{ie}| \;+\; \beta_2\,(|D_{ie}| \times \mathrm{SP}_i) \;+\; \Gamma' X_{ie} \;+\; \varepsilon_{ie}, \tag{7}$$

where SP_i indicates S&P 500 membership and X_{ie} holds log market capitalization and log trailing turnover. For an index member, D_{ie} is the scheduled target Q_{ie}. For a non-member, it is the number the same rule produces from the same public inputs, computed the same way, except that no fund is trading it. We ask what the equal-weight rule would require if the stock had been given one company slot at the previous reset and had since drifted with its own relative price move. That gives $D_{ie} = (A_e/N_e)(1-g_{ie})/ME_{ie}^{ref}$. Here g_{ie} is the ratio of the stock's price at the reference close to its price at the previous reset, divided by the equal-weighted mean of that ratio across the non-member group, so g averages one and $1-g_{ie}$ is positive for a stock that has lagged. For an index member that expression is Q_{ie} itself, because the position carried into the rebalance is the slot value grown by g_{ie}. The two quantities are therefore in the same units, basis points of shares outstanding, built from the same company slot value, and differ only in whether a fund is scheduled to trade the number. The coefficient of interest is β_2, which asks whether displayed quantity responds more strongly to predetermined demand when a fund actually has to trade it.

The stocks with the largest targets could simply be the stocks whose closing auctions are busiest for reasons that have nothing to do with the rebalance, and that reading has to be ruled out. Almost every rebalance falls on a quadruple-witching date, when all closing auctions are busy. The rebalance fixed effect in equation (6) already removes whatever is common to an auction on such a date, so the estimate comes from differences across stocks inside one auction. It does not, by itself, separate exposure to the scheduled rebalance from anything else that varies across those stocks on that day.

We therefore run the same regression on stocks that trade in those same auctions but carry no rebalance. Russell 1000 members outside the S&P 500 carry a predetermined signal of their own, computed from the same public inputs by the same procedure. It measures the same drift in relative prices that fixes the target for index members, and no fund is scheduled to trade on it. Estimating equation (7) across the 25 rebalances on which the feed publishes usable auction data for both groups, the interest arriving on the side opposite the signal is 4.38 basis points per basis point among index members and 0.23 among non-members (Table 5). That is the comparison the question turns on: what distinguishes a rebalance is not merely more activity, but more interest standing ready to take the other side. Taking both sides together, buying and selling interest per basis point of the signal is 9.29 among members and 0.44 among non-members. The control signal carries roughly three times more cross-sectional variation, 16.13 basis points against 5.45, so the much smaller slope among non-members does not come from a signal that moves too little for a slope to be detected.

Holding market capitalization and turnover fixed, index members still gather 7.32 more basis points of buying and selling interest per basis point of the signal than non-members do. Letting market capitalization and turnover each carry their own slope against the signal gives 7.72, and comparing members only with non-members of similar market capitalization and turnover inside the same auction gives 7.34. Table 5 reports all three.

A second control moves the date instead of the stock, and answers the objection that these names simply have large closing auctions. We attach each rebalance stock's own target to its matched control date as a placebo, so the stock is held fixed and the dose is the same. The two dates still differ in more than the target, most obviously in quadruple-witching activity, and an event fixed effect does not remove differences that interact with $|Q_{ie}|$ across stocks. The same-date comparison in Section 3.2 holds the auction fixed and moves the stock instead. The levels confirm the objection's premise: at 15:50 mean gross interest is 86.0 basis points of shares outstanding on rebalance dates against 14.5 on control dates, so ordinary auctions in these names are substantial. The slope is what separates them. Across the 30 historical rebalances, gross interest rises 9.61 basis points per basis point of target at 15:50, with an event-clustered standard error of 0.71, and 11.40 at 15:55 and 12.11 at 15:59. The same regression on the 32 control dates, using the same stocks and the same placebo targets, gives 0.40, 0.63, and 0.71. The scheduled target therefore multiplies the slope by a factor of 23.8, computed on the unrounded slopes against the same stocks carrying the same nominal dose on a date when no such trade is scheduled. The remaining cell of that design, untreated stocks on control dates, is not reported. Coverage of never-member names on control dates is eighteen percent, and almost all of it consists of securities that entered or left the index at some point, which are the least suitable names for a cell whose purpose is to exclude index-linked activity. The two single differences above each control the confound that matters in their own direction, and neither depends on those rows. Mean absolute targets are close to identical across the two, 3.34 against 3.39, so the comparison is not driven by a different distribution of doses. Were large targets simply attached to large-auction names, the control slope would be near ten. It is under one. The extra interest scales with the scheduled order only when the order is actually scheduled, and this date-based control and the same-date member and non-member control reach that conclusion from different variation. The comparison therefore isolates what is specific to the stocks the rebalance touches, since the stocks it compares are trading in the same auctions on the same day. Other demand tied to that same rebalance cannot be separated from the fund's own.

3.3 How Displayed Interest Develops toward the Close

We next ask how much further market arrives once the state is public. We regress the change in displayed quantity since 15:50 on the size of the target, separately at each later auction time,

$$Y_{ie,s} - Y_{ie,15:50} \;=\; \alpha_{e,s} \;+\; \beta_s\,|Q_{ie}| \;+\; \varepsilon_{ie,s}, \tag{8}$$

with a fixed effect for each rebalance at each auction time s, so the common intraday path of a busy close is absorbed and β_s comes only from stocks with larger targets evolving differently inside the same auction. The coefficients in (8) are changes from 15:50 and are not comparable with the levels in (6).

Interest keeps arriving after the displayed auction imbalance becomes public. Between 15:50 and the last pre-close observation a further 3.7 basis points of gross interest accumulate per basis point of absolute demand, 2.0 against the target and 1.7 with it, and both sides grow in every one of the 25 rebalances whose stock panel is balanced across the four clocks the design compares.[5] The exchange feeds report these quantities without identifying who submitted them, so we cannot say who supplied this interest.

The scale is therefore settled and it is large: a market several times the size of the scheduled target is present before the price forms, most of it not scaling with the order's stock-level size, and it keeps growing until the close. What that market consists of, and whether it points where the order points, are separate questions, and they are the subjects of the next two sections.

[5]These are changes from 15:50 and are not comparable with the 15:50 levels, which are measured on 30 rebalances (Figure 4). Every auction time carries its own fixed effect within each rebalance.

4 Interest Grows Through Larger Revisions, Not More of Them

A market several times the size of the order forms around it. That fact on its own says nothing about what the market consists of, and the natural reading, that a larger scheduled target draws a larger crowd, turns out to be wrong. This section takes the interest apart along the only two margins it has, how many revisions to displayed interest arrive and how large they are, and reports what moves. The feed reports revisions, not orders: one revision can carry many orders, a cancellation, a replacement or a reprice.

The two margins are separable in the message feed. Treating each revision of the displayed interest between consecutive auction messages as an arrival, we regress each margin on the size of the market it sits in. Across stock-events, the log number of revisions on log gross interest gives 0.092 (0.023), and the log mean size of a revision on log gross interest gives 0.944 (0.047). The two sum to 1.04. Log revision volume is exactly log count plus log mean size, so a sum near one corroborates that revision volume tracks the displayed gross, but the two are not the same object and the sum is not an identity. Revision size carries roughly 91 percent of the two elasticities summed, with a small but measured count margin alongside it. Interest grows mainly through larger revisions rather than more of them. This split runs against market microstructure invariance (Kyle and Obizhaeva, 2016), which holds that the number of bets scales with the two-thirds power of trading activity and their size with the one-third power, so that variation in activity is carried mainly by count. Across securities in these auctions the split is the reverse, and close to the all-size corner. We flag two limits on the comparison rather than press it: invariance is a statement about bets over a trading horizon, whereas a revision is a change in displayed interest inside a five-minute window, and the mapping between the two is an assumption we cannot verify. Read narrowly, the contrast says that whatever assembles in a closing auction around a foreseeable trade does not obey the scaling that governs ordinary trading activity. This decomposition concerns the composition of gross interest itself, across securities at a given rebalance, and not the mapping from the scheduled target to that total. It is therefore unaffected by the functional form of that mapping, and nothing in the comparison of forms above bears on it.

Because a revision is a change in a published quantity, orders arriving within one snapshot interval net into one observation, which understates the count by more where arrivals are dense. We measure that bias rather than assume it away. Snapping the finest venue's message stream onto progressively coarser grids, which is what a slower feed would have published while the market itself is held fixed, moves the count elasticity from 0.092 at 4.2 seconds to 0.029 at 19.4 seconds and leaves the size elasticity unmoved throughout. Extrapolating that curve to instantaneous publication leaves the count margin close to its measured value, so netting attenuates it without generating it.

5 Which Way the Market Leans

Sections 3 and 4 established how much interest assembles around a scheduled target and what it is made of. Neither says which way it points. A market can be large and still be balanced, or large and lean against the order that drew it. What the target predicts about the state the price actually forms in is the paper's central measurement.

5.1 The Displayed Imbalance That Remains

We regress the displayed imbalance on the signed target, where the sign of the state is the side the exchange reports, with an excess of buying interest counting as positive and an excess of selling interest as negative,

$$R_{ie} = \alpha_e + \lambda Q_{ie} + \varepsilon_{ie}. \tag{9}$$

Equations (6) and (9) are estimated on the same stock-event observations at the same auction time, so the scale and the sign of the response are comparable by construction.

The displayed imbalance is a small residual of a large two-sided market, and that is the fact the loading of the state on the target summarizes. One basis point of scheduled demand is met by roughly ten basis points of gross interest, an amount specific to rebalance dates: across the historical rebalances, weighting each rebalance equally, same-side interest loads 5.31 basis points on the target and contra-side interest 5.03, for a gross 10.34. Table 4, Panel A reports the pooled counterparts of those side slopes alongside the equal-event imbalance loading. Because the two sides are each around five, a modest change in either converts near-exact cancellation into a positive displayed imbalance. The net directional wedge is a fraction 0.028 of gross interest over the historical sample and 0.080 late, so the rotation is small relative to the market that assembles around the target.

That share is unaffected by the one measurement limitation this design cannot remove. Other funds tracking the same rule submit orders we cannot separate from RSP's, so the scale of the initiating dose is uncertain. Rescaling the dose by the broader same-rule family lowers gross late-sample interest from 9.94 to 7.74 on confirmed assets and 6.58 under a generous ceiling, equal-event slopes throughout. The net wedge falls in exactly the same proportion, so the share is 0.080, 0.080, and 0.080 across the three conventions. The level of the two-sided market depends on how the dose is scaled. The balance between its two sides does not.

At the same moment, on the same observations, with the same estimator, the displayed imbalance, the excess of buying over selling or the reverse, moves with the target on average and tracks it poorly across stocks. One basis point of scheduled demand moves the displayed imbalance by 0.48 basis points (Table 4, Panel A), with a bootstrap interval from 0.10 to 0.84. That interval excludes one, which is the restriction a researcher dividing by the target has to assume: the displayed imbalance does not move one for one with the scheduled trade. Giving each rebalance equal weight gives 0.28. The pooled estimator weights each rebalance by how widely targets are spread across stocks within it, and that spread has grown over the sample, so the pooled figure is dominated by the later rebalances and the equal-weight figure is not. Predicting each rebalance from the others, the target accounts for 2.1% of the variation in the displayed imbalance across stocks, against 20.9% for buying and selling interest, and the sign of the state matches the sign of the target in 53.2% of stock-events. That loading is the product of two things, the correlation between the target and the displayed imbalance and the ratio of the state's cross-sectional standard deviation to the target's, so that $\lambda_s = \mathrm{Corr}(Q, R(s)) \times \mathrm{SD}(R(s))/\,\mathrm{SD}(Q)$. The dispersion of the state across stocks contracts steadily as the auction runs down to the close, so a coefficient measured at a later auction time is measured on a smaller state as well as on a better informed one. The difference between the two loadings just described, $\beta_s^{\mathrm{same}} - \beta_s^{\mathrm{contra}}$, is the corresponding coefficient for the target-oriented imbalance $\mathrm{sign}(Q)R$ on $|Q|$. That difference is the more useful way to read the loading. The two do not coincide exactly, because equation (9) uses signed Q and the exchange-signed imbalance. The loading is the net consequence of how the two sides balance, not the fraction of the target that survives into the auction, and we do not read it as such.

One feature of the later sample should be stated where the clock enters. On 2024 August 12, NYSE began including closing D-order interest in the published imbalance at 15:50 rather than 15:55, so what the earliest observation discloses is not identical throughout the late period, and NYSE supplies most of our observations. Six of the twelve late rebalances precede that change and six follow it, so the change cannot be the origin of the late-sample shift, which begins in 2023. But it does mean the late-sample coefficient at 15:50 is an average across two disclosure regimes, and that the drop between 15:54 and 15:55 in any intraday path partly reflects what becomes visible rather than clearing alone. The evidence we rely on is the broader decay across the auction and the post-disclosure path, not any single minute.

Pooled over the whole sample, then, the target supplies almost none of the variance of the displayed imbalance, and the mapping from the target into that state varies substantially across rebalances. A multiplier that divides a price response by Q divides by a quantity that, on average over this sample, barely describes the state the auction displays.

5.2 The Result across Rebalances

The pooled figures average over substantial movement, and that movement is confined to the displayed imbalance while the total amount of interest that gathers stays where it was. Splitting the sample at the end of 2022 (Table 4, Panel C), and reporting pooled slopes throughout this paragraph so that every figure comes from one estimator, gross interest per basis point of demand is 8.78 in the first 18 rebalances and 9.99 in the last 12. Over the same split the two sides rotate. Opposite-side interest moves from 4.55 to 4.60 and same-side interest from 4.24 to 5.39, and the net loading of the displayed imbalance on the target moves from -0.24 to 0.73. Under equal-event weighting the net rotation is the same in direction, but a different side carries it. Pooled, the rotation comes from same-side interest rising while the opposite side is flat. Equal-event, it comes from opposite-side interest falling while the same side is flat. One reading says more investors joined the fund's side, the other says the other side stepped back, and the net is the same either way (Table 4, Panel C). The movement in the net is present under both. Because the displayed imbalance is the difference between the two sides, the rotation of those sides is the movement in the transmission.

We chose the end of 2022 as a split point, which leaves open whether the data pick out a break at a particular date on their own. We do not read this as a dated regime change. Testing every split date that leaves enough rebalances on each side, the strongest evidence of a break in the transmission falls at the March 2023 rebalance, and once the procedure is charged for having searched over dates the absence of a break is not rejected. Weighting rebalances equally moves the selected split to a different date. Fitting a level shift together with a time trend leaves the trend negative and the shift positive, so the loading drifts down inside each period and steps up between them. Gross interest shows no break at any candidate date under the same procedure. Event by event the loading is high in 2018, falls through a trough in 2021 and 2022, and recovers afterwards. What the sample supports is that the rule initiating the trade is unchanged throughout while the balance of the market assembling around it moves.[6]

Section 6 reports how the displayed imbalance is priced across the same sample split. We do not treat the movement of the order–price regression across eras as evidence about either the transmission or the price of the displayed imbalance.

6 The Price of the Displayed Imbalance

Whether the displayed auction imbalance is priced can be measured directly, as the move from the auction reference price to the official close, where the reference price is the anchor the exchange publishes alongside each imbalance message at that auction time. Which rows count as valid follows from the structure of the feed and never from what the prices did.[7]

6.1 The Price–Imbalance Relationship

A displayed imbalance is associated with a substantial concession, meaning that the official close settles away from the reference price in the direction of the heavier side of the auction, and we

[6] Two NYSE closing-auction changes fall near the end of 2022, a rule that pulls the discretionary closing price toward the reference prices published with the imbalance messages, and new floor-broker connectivity for closing D orders, the discretionary orders floor brokers enter for the close. Both bear on quantities we measure. Neither is dated precisely enough to support an attribution, and the break itself is not established firmly enough to carry one. We make none. Figure 5 shows the series event by event with no smoothing.

[7] The outcome is $10^4[\log P^{close} - \log P^{ref}(s)]$, oriented by the sign of the state for presentation, with the official close taken from the security-master record and cross-checked against the final auction message, whose median absolute log difference from it is under two basis points. We exclude rows whose reference price is missing or not positive, rows carrying the placeholder values the feed publishes when it has no price, and rows with a zero state, and we tabulate every exclusion by reason.

measure the size of that move in basis points. On the pooled sample the fitted concession at 15:55 (Table 6, Panel A)[8] is 5.8 basis points at an imbalance of 5 basis points of shares outstanding, 10.1 at 10, and 15.0 at 20, and at 15:59 the corresponding values are 3.7, 5.5, and 8.9, all with event-clustered confidence intervals that exclude zero from an imbalance of 5 basis points upward. The relation between the size of the imbalance and the concession, which we call the pricing schedule, is concave on Nasdaq, where the nonlinearity test rejects linearity at both times with χ^2 of 49.7 at 15:55 and 39.3 at 15:59, while on NYSE linearity is not rejected over the observed range, with p of 0.3964 and 0.1903. Restricting to the historical rebalances that the transmission estimates use leaves both verdicts unchanged. The concavity is not a fixed property of the venue: split by era, the early Nasdaq schedule is close to linear, the marginal concession running 2.32, 2.50, and 2.54 basis points per basis point at imbalances of 5, 10, and 20, while the late schedule bends sharply, at 2.15, 1.79, and 1.14. The curvature is part of what changed rather than a constant we can condition on. The schedule is estimated on every rebalance meeting the mechanical support conditions for this design. We make no held-out claim for the target-dependent results, whose construction is not chronologically separable from these events. The price schedule is a different case. It is estimated on the displayed imbalance and the closing price and never uses the target at all, so the specification fixed on the earlier rebalances is applied unchanged to the later ones, and Panel C is a genuine out-of-sample check of it.

The auction is large relative to the fund's target. Final displayed paired volume, the quantity that already matches between the two sides, exceeds the full scheduled target in 100.0 percent of stock-events, at a median ratio of about 30, and the number of shares that actually trade in the close is at least that large.[9]

6.2 Stability across the Sample

The schedule is not a constant of the sample. At 15:55 the fitted concession at an imbalance of 10 basis points is 12.1 basis points before 2023, with a confidence interval of 7.8 to 16.5, and 6.8 afterward, with an interval of 5.7 to 7.9, and the test of era stability, clustered on the event, rejects the hypothesis that the two schedules are the same with a χ^2 of 40.9, and it rejects that hypothesis on each venue separately. On NYSE the late-era 15:55 schedule is statistically flat at moderate imbalances. At 15:59 stability is not rejected, and the design cannot detect a change smaller than a factor of about two at that time. A positive schedule exists in both eras and its level roughly halves at 15:55 across the sample split. Figure 6 plots the 15:55 comparison and Table 6 reports the fitted values behind it, at both times. Concavity is not uniform across venues and eras: it is pronounced for Nasdaq late, close to linear for Nasdaq early, and not rejected as linear on NYSE.

[8] The magnitude is in line with impact estimated on ordinary closing auctions. Applying the square-root impact function of Goyal, Jegadeesh, and Wu (2026) to our imbalances, scaled the way they scale theirs, gives a counterfactual concession of the same order as the schedule we fit. We do not use their function anywhere in our estimates and report the comparison only as an external check on magnitude.

[9] Source: the message panel's final-message paired quantity, described in Appendix A.2.

One comparison disciplines the level. Applying a published pricing function for ordinary closing auctions to our own displayed imbalances, row for row, places our schedule at roughly a quarter to a third of that benchmark's main reading. That benchmark's volatility term is ambiguous in its own source, described there as a daily standard deviation but winsorized in annualised terms. Read the first way the two schedules are close to equal, so the comparison brackets our estimate rather than pinning it.

The realized displayed imbalance is strongly associated with the closing-price concession, and so is the part of it the target and the standing opposite-side interest do not explain. Orienting the price move by the sign of the residual and regressing it on $|u|$ gives 0.58 basis points of price per basis point of the conditional residual, with an event-clustered standard error of 0.13. Appendix A.5.1 reports the corresponding specification on signed u. We nonetheless read $\widehat{g}(R)$, the fitted concession at a displayed imbalance of a given size, as a conditional equilibrium pricing schedule rather than the causal price effect of an exogenous shock to the state. R is the state the auction actually reached and u is the part of it left unexplained by the target and by the opposite-side interest standing at the same moment, and both are determined jointly with the price they are being related to, so neither coefficient is a causal price impact.

The displayed imbalance is priced, materially and nonlinearly, at magnitudes of a few to a dozen basis points, and its conditional price is itself an equilibrium object, one that moved across the same 2022 sample split as the loading of the displayed imbalance on the target.

6.3 What the Order–Price Multiplier Measures

The reduced-form regression of the closing price on the target is the object the literature inverts, and its coefficient, M_Q, is what we call the order–price multiplier, the closing-price move per unit of scheduled demand. We report it below, and we are cautious about reading it as an elasticity. Two properties of the estimate justify that caution. Its value depends on how rebalances are aggregated, and the same construction applied to matched dates carrying no scheduled target produces apparent multipliers of both signs. Table A1 and Appendix A.5 report the full set of aggregation rules we ran and the matched-date construction. Section 7.1 states what the coefficient does measure.

6.4 Why No Single Correction Factor Works

A stylized example shows what is at stake in the denominator, and that the error has no fixed direction. Let a fixed demand curve move the price by half a percent for every one percent of the company that is net demand where the price is set, an elasticity of two, and let a fund be required to buy one percent. If the net of the orders arriving when the price is set is a quarter of a percent, the curve moves the price by 0.125 percent, and dividing by the one-percent order implies an elasticity of eight, four times too elastic. If that quantity is instead two percent, the same curve implies an elasticity of one, too inelastic. One fixed curve, biases in either direction, and nothing involving expectations, dynamics, or a shifting schedule. The researcher normalized by the initiating target rather than by the net quantity in the market where the price formed.

The natural repair is to estimate the mapping and divide by it, and three facts together rule that repair out. The mapping is not a constant: across the two halves of our sample the scale of the market barely moves while its net loading on the target moves from -0.24 to 0.73 (Table 4, Panel C). It is not recoverable from the order: within a rebalance, the target accounts for 2.1% of the cross-sectional variation in the displayed imbalance. And it is not recoverable from history either. Predicting each rebalance's loading using nothing but the rebalances that precede it, scored strictly out of sample over 25 events, gives an R^2 of -0.14, and -0.07 over the 20 events left when the first predictions, which rest on only two or three training rebalances, are dropped. The same exercise on the gross scale of the market gives -0.25 and -0.08. Every figure is at or below zero and none is distinguishable from it: the predictor adds nothing, rather than doing harm. Extrapolating from a statistic's own history does no better than its historical mean, so the mean is the best ex-ante adjustment this evidence supports. That adjustment is a single constant, and the quantity it would correct does not sit still. The estimated event-level loadings run from -1.47 to 1.87. Event-level slopes are themselves estimated with error, so that spread overstates the spread of the true loadings, but only slightly: the between-event dispersion net of estimation error is 0.85, which is 96% of the observed variance. The repair fails not because there is nothing to divide by, but because the only divisor history supplies is a constant, and the quantity it would correct varies across rebalances by more than the constant itself. We test extrapolation from the statistic's own history, which is the adjustment a practitioner would attempt. Conditioning on observables such as event size, fund assets, target dispersion or market volatility is untested here, and we make no claim about it.

What can be said is how far apart the two reduced-form objects sit. A multiplier estimated on the scheduled target is related to the price sensitivity of the displayed imbalance through that event's own loading of the state on the target, but the two are not connected by that loading alone. Appendix A.5 splits the multiplier exactly into the part the fitted pricing schedule accounts for and a remainder, and the remainder is the larger piece, 5.79 of 6.44. The loading is therefore not a scalar that converts one multiplier into the other. Across our rebalances that loading has a median of 0.35 and an interquartile range of -0.31 to 0.93. Across the middle four-fifths it runs from -1.27 to 1.08, and it is negative in 30% of them. A researcher dividing a price response by scheduled target rather than by the state that formed is therefore working with a number that, in this sample, differs from the one built on the state by a factor spanning zero and both signs.

Nothing in this recovers a structural elasticity, and we do not offer it as an adjustment anyone should apply. The conditions under which an elasticity could be recovered are established elsewhere and we do not dispute them. What the range describes is the reduced-form object researchers actually estimate, and how far it can sit from the price sensitivity of the net of the orders that arrived. That quantity is hypothetical in this example. It is also not the displayed imbalance we observe, since the displayed imbalance is what the exchange publishes as the auction builds, not the net of the orders that finally arrive.[10]

The multiplier is harder to read as a single number than that. The price the auction charges for a displayed imbalance is nonlinear in that state, so no single coefficient converts quantities into prices across the observed range, and that schedule itself differs across eras. The reduced-form multiplier is therefore a compound equilibrium statistic whose components are separately observable and separately moving.

Refusing to divide the multiplier by a single correction factor still leaves something to quantify. The regression of prices on the gross institutional order is the object the conventional literature inverts, and the facts above say not to call its inverse an elasticity. The price response to the displayed imbalance is the economically relevant intermediate object we can measure. The structural demand elasticity is a different object, the amount by which investors would change the quantity they are willing to trade if the price were moved away from the level at which the market clears. Measuring it would require an exogenous shifter of the state that this experiment does not provide, because the displayed imbalance and the closing price are jointly determined and other traders can respond to the state. The second of the three objects, the price response to the displayed imbalance, still supports a calibration. The fitted concession is the move from the pre-auction reference price to the official close that accompanies a displayed imbalance of a given size. At a displayed imbalance of ten basis points of shares outstanding, the ratio of that imbalance to the fitted concession is about 0.8 before 2023 and 1.5 afterward, meaning about 0.8 basis points of imbalance for each basis point of price concession before 2023 and 1.5 afterward. Because the schedule is concave, this is an average ratio at that point on the schedule and not the local slope, which would be the reciprocal of the schedule's derivative there. The reciprocal is a unit conversion of the conditional pricing schedule, not a structural demand-elasticity estimate. Inverting the multiplier on the target can imply an elasticity bearing little relationship to the price sensitivity of the net of the orders that arrive, and in our data the multiplier on the target is unstable enough to imply radically different elasticities while the calibration sits on the order of one at moderate imbalances.

6.5 Corroboration from Russell Reconstitutions

Russell cannot reproduce the stock-level continuous-treatment design available for RSP, and it provides an independent setting in which the benchmark-driven quantity can be constructed before the event and the residual closing imbalance can be compared with matched non-reconstitution dates. Every June, FTSE Russell rebuilds its indexes from updated market capitalizations, and stocks that cross the size cutoff move between the Russell 1000 and the Russell 2000, which forces

[10]The example is stylized. In our data the state is not a fixed fraction of the order, and much of the price response measured over the order's price window happens before the late auction state forms, so the example understates the difficulty rather than describing the sample.

the funds tracking those indexes to trade them. These annual reconstitutions change the index rule and the affected stocks, along with the event frequency and the calendar, while keeping the same closing markets.

The migrating stock's required trade is predetermined in the same sense the fund's target is. We measure it before the event from the float held by the plain benchmark trackers of the index a stock is leaving and the one it is joining, and the tracker holdings reported afterwards confirm it. Across 389 genuine migrations in six verified reconstitutions from 2018 through 2023, the predetermined quantity agrees in sign with the tracker holdings change that follows in 99.0 percent of cases, with an equal-event correlation of 0.985 and an equal-event slope of 0.979 (Table 7).

What the auction leaves unresolved is again large. We measure the displayed imbalance for migrating stocks after taking out the part that the direction of the migration and the stock's float make predictable, so what remains is the part those two predictors do not explain. That residual imbalance is far more dispersed on reconstitution days than on matched non-reconstitution June Fridays, by a factor running from 20.5 to 241.3 across the six years and above one in 6 of 6 reconstitutions, an exact sign probability of 0.031 (Table 7). A rebalance the whole market can compute in advance therefore leaves a large conditional residual standing at the close in a second index family, on a different calendar, under a different rule.

We do not read anything further into Russell. The residual measured here is a different object from the conditional residual of Section 2.4, built on a different control benchmark, so the two dispersion ratios are not comparable. We estimate no Russell pricing schedule, and nothing here speaks to how Russell imbalances are priced.

7 Implications for Demand-Elasticity Estimates

7.1 Responding to Prices and Responding to Orders

An elasticity estimate needs the demand shock to move other investors along their supply schedule. Nothing in our argument says that a predetermined target Q is the wrong shifter for estimating the reduced-form equilibrium multiplier dp/dQ. Our question concerns the additional step of interpreting the inverse of that multiplier as a structural demand elasticity. Nor do we propose replacing Q with the realized auction imbalance in the price regression. That imbalance is endogenous and is an intermediate market quantity, not the structural demand shifter. Let $X(Q)$ be the quantity that the initiating target actually contributes to the net of the arriving orders once implementation is done, and let $S(p, Q)$ be the quantity other investors supply at price p, with their schedule allowed to depend directly on the known target Q as well as on the price. The auction price is whatever makes those two quantities equal, so $X(Q) = S(p, Q)$ holds at every value of the target. Push the target up by one unit and both sides have to move together, which pins down how far the price must travel,

$$\frac{dp}{dQ} = \frac{X_Q - S_Q}{S_p}. \tag{10}$$

A subscript here means a derivative. X_Q is how much an extra unit of target moves the net of the orders that arrive, S_Q is how much other investors shift their supply when the target changes with the price held fixed, and S_p is how much more they supply when the price rises, which is the slope the whole exercise is after. Reading the multiplier of the price on the target as one over that slope requires $X_Q - S_Q = 1$. The natural benchmark that delivers it is $X_Q = 1$ together with $S_Q = 0$, and these are distinct assumptions. $X_Q = 1$ says that one unit of the initiating scheduled target moves that net by one unit. $S_Q = 0$ says that other investors respond to the target only through the price, and never shift their supply because they know the order is coming. Selling because the price has risen is elasticity. Selling because a known buyer is coming is a shift in the supply schedule itself. We put a number on neither assumption separately. We do not observe X and we estimate neither X_Q nor S_Q structurally. What we measure is the wedge between the initiating order and the quantity the auction presents to the price, and what the data show is that this mapping is an empirical object rather than an identity, so the pair of assumptions inversion requires cannot be granted on the grounds that Q is mechanical. How the fund executes, and how other investors trade ahead of the order, can both affect the mapping from the order to the arriving net. Work on expectations and on dynamic demand, cited in the introduction, raises a further problem, that the schedules X and S need not stay put while the price is being set.

A canonical estimate makes the content of those two restrictions concrete. Greenwood and Sammon (2025) write price impact as the demand shock divided by the negative of the demand elasticity, where the shock is the percentage of a stock's capitalization bought upon index addition. Their order–price multiplier is therefore $M_Q = \Delta P/Q$ and the reported elasticity is $\hat{\varepsilon}_{\text{conv}} = -1/M_Q$. For index additions over the most recent decade of their sample they report $M_Q = 0.37$, or equivalently a demand elasticity of -2.72 on the conventional $-1/M$ conversion, which carries the opposite sign to the reciprocals in Table 1, against a multiplier of 6.75 in the late 1990s.[11] Their conversion is equation (10) with the numerator set to one. Carrying that numerator through instead, the same reduced form implies

$$\hat{\varepsilon}_{\text{struct}} \;=\; \left(X_Q - S_Q\right)\hat{\varepsilon}_{\text{conv}}, \tag{11}$$

so the conventional number obtains only at $X_Q = 1$ and $S_Q = 0$. Equation (11) looks as though it delivers a repaired number, and it does not. We do not estimate $X_Q - S_Q$, so the equation is neither a corrected elasticity nor a bound on one. Its purpose is to name the object that a published order–price multiplier sets to one when it is inverted, namely the difference between the order and the net of all the orders that arrive with it.

[11] Greenwood and Sammon, Section 4. The multiplier falls by roughly a factor of twenty across those two eras.

Our measurements say why setting that difference to one is a substantive assumption and not a bookkeeping convention. At 15:50, the first moment at which both venues publish the auction state, one basis point of scheduled demand is associated with 9.69 basis points of displayed buying and selling interest, the displayed imbalance moves by 0.48, and the target accounts for 2.1% of the cross-stock variation in that imbalance. Across the two halves of the sample the scale of the market is essentially unchanged, 8.78 against 9.99, while the imbalance per basis point of target moves from -0.24 to 0.73, all pooled slopes under an initiating rule that never changes. Time variation in an order–price multiplier can therefore reflect time variation in what assembles around the order, even when the initiating rule and the gross scale of the market around it are unchanged. None of this identifies X_Q or S_Q, separately or jointly. Equation (10) is used here as accounting and not as attribution: it states what must hold for an order–price multiplier to equal a structural slope, and we make no claim about which of its two conditions fails. What the data establish is the composite mapping from the scheduled target to the market state in which the price is formed. That mapping is an empirical object that moves, and so cannot be set to unity on the grounds that the initiating demand shock is mechanical, exogenous, or precisely measured.

7.2 The Same Data, Four Denominators

The conventional step reads a demand elasticity off a multiplier by inverting it. Our setting lets us apply that step to four candidate denominators over the same price window, on the same observations, with the same estimator, and see what it returns. Basis points of market capitalization and of shares outstanding coincide at a common price, so the four are directly comparable.

Inverting the multiplier using four candidate measures of demand gives values from -3.27 to 2.38, including a change of sign, and nothing varies across those rows except which object is named the demand shock. The sign depends on which empirically meaningful quantity is called the demand.

The two components of the displayed imbalance make the point without any naming convention at all. If the closing price paid a single linear price for the displayed imbalance, then $\widehat{R}$ and u, being components of R, would carry the same coefficient. In a joint projection of the price move on both components, equality of the two coefficients is rejected: they differ by -0.61 with an event-clustered standard error of 0.14, built from the full covariance of the pair, a t statistic of -4.42 across 30 event clusters. That rejects a common linear coefficient under this specification. It does not rule out a single nonlinear schedule for R, which is the case the concavity in Section 6 in fact supports, and either way there is no one number to invert.

A multiplier that divides a price response by a measured trade carries an implicit first stage: the mapping from the trade a researcher observes to the quantity the price ultimately clears. That stage is invisible in most designs, whereas here the mapping into it can be watched directly, and it is weak. The target loads on the displayed imbalance at 0.48, with a bootstrap interval that excludes one. The four rows are separate projections rather than Wald ratios, so that loading diagnoses the conventional multiplier on the target; it does not explain the spread across all four.

The instability is not confined to the choice of denominator. Holding the denominator fixed at the scheduled target, the multiplier itself moves with the level at which it is aggregated: 2.13 to 6.44, a factor of 3.0, across weighting schemes on identical events, and from 11.40 to -1.01 between the two halves of the sample. Within a rebalance, across stocks, the target explains 2.1% of the displayed imbalance the price forms in. A parameter of a demand curve is the same number at every level of aggregation. A statistic that mixes an order with whatever assembles around it need not be, because the assembly is what varies across stocks, across events, and across eras.

We state two limits on that reading. Aggregation instability is also what heterogeneous true elasticities would produce, and we cannot rule that out. A second piece of evidence comes from matched dates carrying no scheduled target. The same construction there produces apparent multipliers of both signs, and heterogeneous elasticities do not generate a multiplier where no trade occurred. And none of this recovers a structural elasticity or offers a corrected one.

7.3 What an Adequate Design Must Observe

Three objects should be kept apart, and only one of them is measured here. The first is the conditional price sensitivity of the displayed imbalance, the closing-price move that goes with an imbalance of a given size. That is what we measure, and its reciprocal is on the order of one at moderate imbalances, with meaningful concavity and a shift in level across eras. The second is the structural demand elasticity, which this design does not identify, because the displayed imbalance and the closing price are jointly determined and other investors can respond as the state forms. The third is an elasticity inferred from the fund's own order, which the demand-elasticity literature actually reports and which is not identified either, for a different reason: the order is associated with a sizeable change in the average displayed imbalance but explains little of which state materializes.

That leaves two questions for a design that cannot observe the arriving net, which will usually be the case. What does it need before its multiplier identifies a slope, and what should it report when it cannot meet that standard?

A design needs exogenous variation in the net of the orders that arrive, not merely a view of it. What we observe is the displayed imbalance, and regressing price on it yields a conditional price sensitivity, which is what we report and what Section 5 shows to be jointly determined with the price. Equation (10) makes inversion of the order-based multiplier legitimate only when the measured order maps one for one into the arriving net and other investors do not respond directly to its arrival. Neither can be granted here. The scheduled target loads on the displayed imbalance at 0.48, with a bootstrap interval that excludes one, and it accounts for 2.1% of the cross-sectional variation in it, so the mapping is neither one for one nor tight. We cannot say which of the two conditions fails, only that the pair cannot be assumed, because we do not separate the fund's arriving contribution from other investors' response. A design that cannot see the second quantity should report the order–price multiplier as a multiplier: a statistic describing what assembled around one particular order.

Both sides of that quantity, rather than the net, are what make the mapping interpretable. That is a diagnostic requirement, not a condition for inversion. The displayed imbalance is the difference between two much larger and separately moving flows, and in our sample same-side and opposite-side interest change quite differently over time while their sum barely moves. A design reporting only net order flow cannot tell a change in demand from a change in who is willing to take the other side.

A control that holds fixed whatever else makes the moment busy is needed before any of it can be attributed to the order. Our two controls are a same-date comparison across stocks and a same-stock comparison across dates, and neither alone would be persuasive: the first leaves open that index members are structurally different, the second that rebalance dates are. Together they triangulate the claim. Where implementation cannot be observed at all, the honest report is the multiplier plus a description of how the trade was implemented, how far ahead it was forecastable, when it executed, and what interest surrounded it.

Two propositions follow, and the first is something a researcher can act on without an auction feed.

A diagnostic for aggregation instability. A parameter of a demand curve is the same number at every level of aggregation. A statistic that divides a price response by an order need not be, because whatever assembles around that order varies across assets, across events, and across time. So estimate the multiplier three ways on the same data: within events across assets, across events under different weightings, and as a ratio of period-level means. If the denominator is the demand, the elasticity is common across assets, and the schedule is linear, the three agree. Disagreement rejects the three jointly rather than the denominator alone, which is why the diagnostic has a second step. In our setting they do not, and the gaps are large: the across-event multiplier moves by a factor of 3.0 with weighting alone, and the period-level figure changes sign between the halves of our sample. We report the quantity a within-event multiplier would have to rest on, the first stage from the target to the displayed imbalance. It explains 2.1% of that imbalance across stocks inside a rebalance and loads at 0.48, not one. The second step separates the denominator from the alternatives. Rerun the same construction on matched dates carrying no scheduled rebalance. Heterogeneous elasticities and a curved schedule cannot produce a multiplier where no trade occurred, so an apparent multiplier on those dates points at the denominator. Ours produces coefficients of both signs there. Both steps run on data a researcher already has, without an auction feed. Together they are a test rather than a correction, and they return nothing anyone should divide by.

A selection claim about the literature. The measurements above give the claim of Section 1 its content. A market several times the size of the order assembles before the price forms, most of it not scaling with the order's stock-level size, and the order predicts little of the net. That is what foreseeability buys the other side of the trade, and foreseeability is not an accident of our setting but a consequence of the criteria by which such settings are chosen. One step is needed to make the claim general, because two different kinds of property are involved. Exogeneity and precise measurement are properties a researcher verifies, often after the fact. Being mechanical and fixed by a public rule are properties the market can act on in advance. Only the second pair implies anticipation directly, and a reader can fairly object that a mutual-fund flow shock is exogenous without being published five days ahead. What connects them is verifiability. Instruments in this literature are built from public inputs, index rules, reported holdings, disclosed flows, because that is what allows exogeneity to be checked by anyone. Whatever a researcher can compute from public inputs, a trading desk can compute as well, and usually sooner. The claim is therefore not a caveat about equal-weight rebalances. Mechanical designs whose stock-level realization is computable from public inputs before the trade are especially exposed: what a researcher can compute that way, a trading desk can compute as well, and usually sooner. We do not extend the claim to designs whose exogeneity is verified after the fact without the order being publicly computable beforehand.

None of the assumptions inversion requires follows from exogeneity, mechanical construction, or precise measurement. Two clean routes to a structural elasticity remain: exogenous variation in the net of the orders that arrive, or enough structure to recover the mapping from the order to that quantity.

8 Conclusion

This paper opens the interval between a precisely measured institutional order and the benchmark price against which that order is judged. A quarterly index rule fixes a fund's scheduled stock-by-stock target in advance, 62 rebalances show that target strongly reflected in later holdings, and 30 carry the exchanges' own second-by-second record of the auction state that forms as the benchmark price is set.

The order first organizes a market many times its own size. One basis point of scheduled demand is met by 9.69 basis points of buying and selling interest, split almost evenly between the two sides, and that interest is largely absent when the same kind of signal reaches stocks carrying no rebalance, or when the same stocks are observed on dates carrying no rebalance. The imbalance those two sides leave over moves with the target on average, tracks it poorly stock by stock, and maps into it in a way that does not sit still across the sample. That imbalance, and in particular the part of it our conditioning does not explain, is what the closing price moves with, through a material and concave conditional schedule. The reduced-form multiplier of prices on the target combines two things that need not move together, how the target maps into the displayed imbalance and how that imbalance is priced. For the measurement of stock-demand elasticity, the prior question is whether the measured order maps one for one into the net of the orders that arrive when the price is set. What we observe is the first quantity to stand between a known institutional target and its benchmark price, and it shows that this mapping is an empirical object, not an identity. We do not claim that the quantitative mapping documented here applies to every demand shock. The broader requirement is that the mapping from a measured order to the price-setting quantity be justified rather than assumed. Whether the multiplier of the correct quantity traces a fixed demand curve is a further question, and the halving between eras of the price schedule for the displayed imbalance shows that the schedule need not sit still. Recovering a structural elasticity requires either exogenous variation in the net of the orders that arrive or enough structure to recover the mapping from the order to that quantity. A price response divided by an anticipated order is an equilibrium statistic of the process that assembles around the order.

What would settle the question is variation in how far ahead a given order is anticipated, holding the order itself fixed. Index rules that differ in how much discretion a committee retains offer that variation, and separating foreseeability from the return drift that generates the trade is the natural next design. We leave it to future work.

The fund's target is known. The market that forms around it is not.

References

Admati, Anat R., and Paul Pfleiderer, 1991, Sunshine trading and financial market equilibrium, *Review of Financial Studies* 4, 443–481.

Bastianello, Francesca, 2026, Expectations and inelastic markets, working paper, University of Chicago Booth School of Business, July 9.

Behmaram, Pouya, and Carter Davis, 2026, Indexing and the elasticity of stock demand, working paper, April 23.

Beneish, Messod D., and Robert E. Whaley, 1996, An anatomy of the "S&P game": The effects of changing the rules, *Journal of Finance* 51, 1909–1930.

Bessembinder, Hendrik, Allen Carrion, Laura Tuttle, and Kumar Venkataraman, 2016, Liquidity, resiliency and market quality around predictable trades: Theory and evidence, *Journal of Financial Economics* 121, 142–166.

Bogousslavsky, Vincent, and Dmitriy Muravyev, 2023, Who trades at the close? Implications for price discovery and liquidity, *Journal of Financial Markets* 66, 100852.

Chinco, Alex, and Marco Sammon, 2024, The passive-ownership share is double what you think it is, *Journal of Financial Economics* 157, 103860.

Coles, Jeffrey L., Davidson Heath, and Matthew C. Ringgenberg, 2022, On index investing, *Journal of Financial Economics* 145, 665–683.

Coval, Joshua, and Erik Stafford, 2007, Asset fire sales (and purchases) in equity markets, *Journal of Financial Economics* 86, 479–512.

Davis, Carter, Mahyar Kargar, Jiacui Li, and Dejanir Silva, 2026, On the recovery of demand elasticities in dynamic settings, working paper, SSRN 6497439.

Chang, Yen-Cheng, Harrison Hong, and Inessa Liskovich, 2015, Regression discontinuity and the price effects of stock market indexing, *Review of Financial Studies* 28, 212–246.

Duffie, Darrell, 2010, Presidential address: Asset price dynamics with slow-moving capital, *Journal of Finance* 65, 1237–1267.

Escobar, Luis, Lorenzo Pandolfi, Alvaro Pedraza, and Tomas Williams, 2026, Who trades index rebalancings? Evidence on benchmarking and inelastic demand, CEPR Discussion Paper 21526.

Etula, Erkko, Kalle Rinne, Matti Suominen, and Lauri Vaittinen, 2020, Dash for cash: Monthly market impact of institutional liquidity needs, *Review of Financial Studies* 33, 75–111.

Fuchs, William, Satoshi Fukuda, and Daniel Neuhann, 2026, A trilemma for asset demand estimation, working paper, May 8.

Gabaix, Xavier, and Ralph S. J. Koijen, 2021, In search of the origins of financial fluctuations: The inelastic markets hypothesis, NBER Working Paper 28967.

Goyal, Amit, Narasimhan Jegadeesh, and Yanbin Wu, 2026, Price impact in closing auctions, opening auctions, and continuous markets: A benchmark for cost of trading on anomalies, *Journal of Financial and Quantitative Analysis*, forthcoming.

Greenwood, Robin, and Marco Sammon, 2025, The disappearing index effect, *Journal of Finance* 80, 657–698.

Harris, Lawrence, and Eitan Gurel, 1986, Price and volume effects associated with changes in the S&P 500 list, *Journal of Finance* 41, 815–829.

He, Zhiguo, Péter Kondor, and Jessica Shi Li, 2025, Demand elasticity in dynamic asset pricing, NBER Working Paper 34450.

Hu, Edwin, and Dermot Murphy, 2024, Vestigial tails? Floor brokers at the close in modern electronic markets, *Management Science*.

Jegadeesh, Narasimhan, and Yanbin Wu, 2022, Closing auctions: Nasdaq versus NYSE, *Journal of Financial Economics* 143, 1120–1139.

Lou, Dong, 2012, A flow-based explanation for return predictability, *Review of Financial Studies* 25, 3457–3489.

Lou, Dong, Hongjun Yan, and Jinfan Zhang, 2013, Anticipated and repeated shocks in liquid markets, *Review of Financial Studies* 26, 1891–1912.

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

Kyle, Albert S., and Anna A. Obizhaeva, 2016, Market microstructure invariance: Empirical hypotheses, *Econometrica* 84, 1345–1404.

Li, Jiacui, and Zihan Lin, 2026, Price multipliers are larger for less diversifiable order flows, *Journal of Finance*, forthcoming.

Li, Sida, 2021, Should passive investors actively manage their trades?, working paper, SSRN 3967799.

Nathan, Daniel, Matti Suominen, and Tuomas Tasa, 2026, Intramonth liquidity demand and the momentum premium, working paper.

Pavlova, Anna, and Taisiya Sikorskaya, 2023, Benchmarking intensity, *Review of Financial Studies* 36, 859–903.

Pegoraro, Stefano, Marco Sammon, and John J. Shim, 2026, Optimal index-linked rebalancing with anticipatory trading, working paper.

Petajisto, Antti, 2009, Why do demand curves for stocks slope down?, *Journal of Financial and Quantitative Analysis* 44, 1013–1044.

Sammon, Marco, and John J. Shim, 2026, Who clears the market when passive investors trade?, *Review of Financial Studies*, forthcoming.

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

van Binsbergen, Jules H., Benjamin David, and Christian Opp, 2025, How (not) to identify demand elasticities in dynamic asset markets, NBER Working Paper 34528.

van der Beck, Philippe, 2026, On the estimation of demand-based asset pricing models, *Journal of Financial Economics*, forthcoming.

Wardlaw, Malcolm, 2020, Measuring mutual fund flow pressure as shock to stock returns, *Journal of Finance* 75, 1645–1690.

Webb, Matthew D., 2023, Reworking wild bootstrap-based inference for clustered errors, *Canadian Journal of Economics* 56, 839–858.

Wurgler, Jeffrey, and Ekaterina Zhuravskaya, 2002, Does arbitrage flatten demand curves for stocks?, *Journal of Business* 75, 583–608.

Zhao, Ruiyang, 2026, Preying on leveraged ETFs, working paper, arXiv 2608.03703.

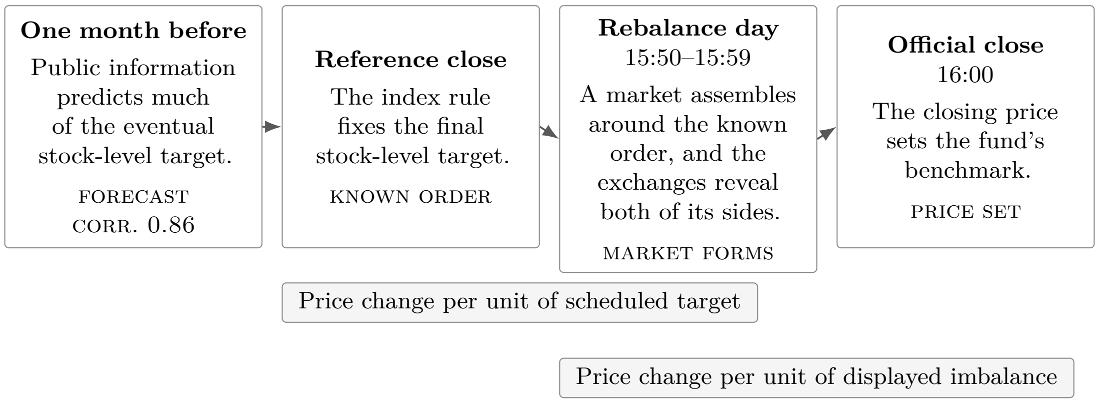

Figure 1: **Order of Events in the Rebalance.** The four panels give the order in which information arrives. The rule-implied target is the stock-level trade the published equal-weight rule requires. The forecast correlation is the within-rebalance correlation between the target reconstructed twenty trading days before the reference close and the target at the reference close, from Table 2, Panel C. Displayed buy and sell interest are the quantities the exchanges publish before the cross, observed from 15:50. The two bands give the windows over which the paper's two price objects are measured, and each is a price change divided by a quantity. The first is the price change from the reference close to the official close per unit of the scheduled target. The second is the price change from an auction snapshot to the official close per unit of the imbalance displayed at that snapshot. Neither is a causal effect of the quantity in the denominator.

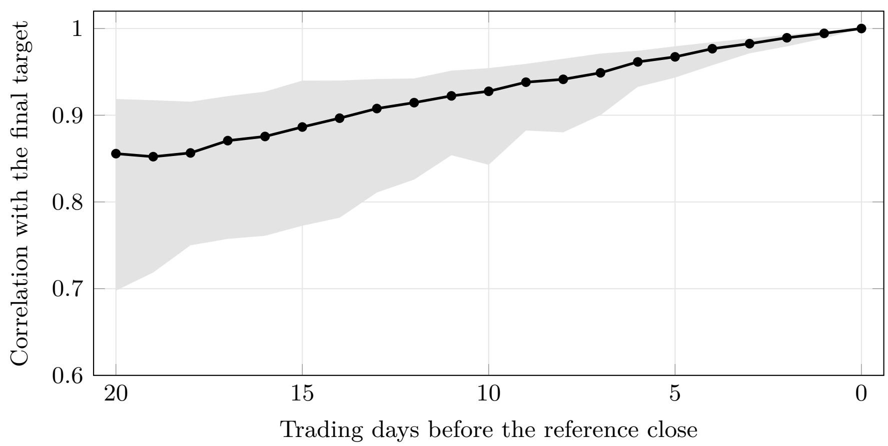

Figure 2: **How Early the Trade Can Be Computed.** Each point is the median across rebalances of the correlation between the target rebuilt from public inputs k trading days before the index reference close and the final rule-implied target. The shaded band spans the tenth to ninetieth percentile across rebalances. The reconstruction uses holdings, fund assets and market capitalizations as published on the day it is built. Horizons run from twenty trading days before the reference close to the close itself. The number of rebalances behind each point ranges from 52 to 62.

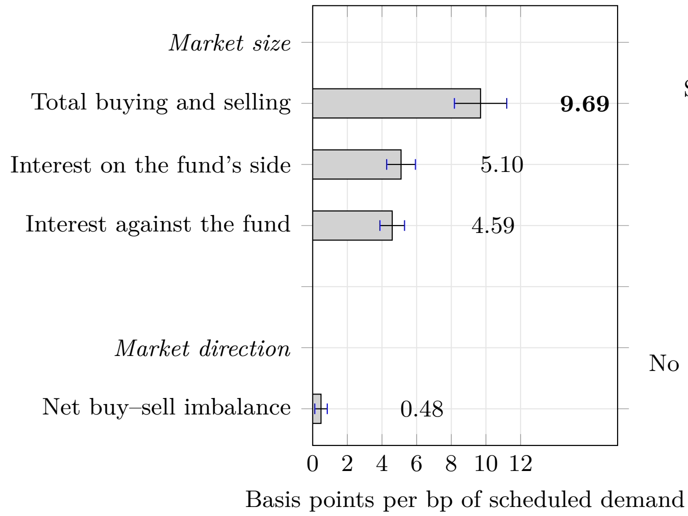

(a) Panel A: Index members, all 30 rebalances.

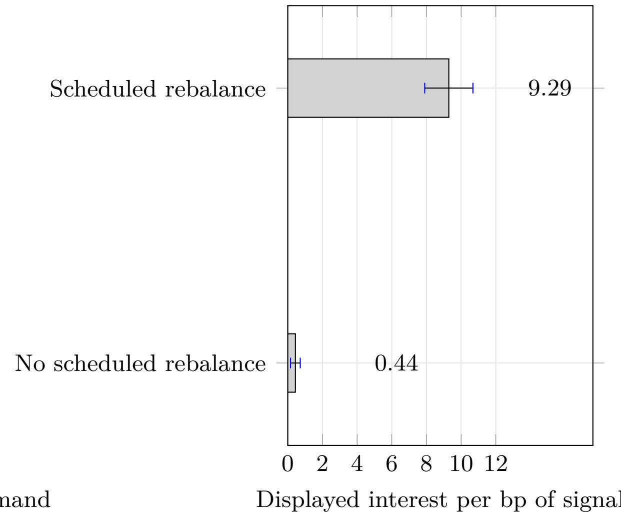

(b) Panel B: Same dates, 25 rebalances.

Figure 3: **The Scheduled Trade Predicts Market Size Far Better Than Market Direction.** Displayed interest at 15:50, the first moment both exchanges publish the auction state, in basis points of shares outstanding per basis point of scheduled demand. Same-side and opposite-side interest are the displayed buy and sell interest oriented by the sign of the fund's rule-implied target at the reference close; gross interest is their sum. The net-imbalance row comes from a different regression from the market-size rows, on the signed target rather than its absolute value, so 5.10 minus 4.59 is not expected to equal 0.48. The two groups are shown separately for that reason. In Panel B the signal is the fund's scheduled target for index members and the same quantity built by the same code path for non-members, who carry no fund trade. Panel A estimates equations (6) and (9) on 12,605 stock-event observations across 30 rebalances, with a rebalance fixed effect, so slopes are identified across stocks inside the same auction. Panel B estimates equation (6) separately for S&P 500 members and for Russell 1000 stocks outside the S&P 500 on the 25 rebalances where both groups have auction support; for nonmembers the regressor is the predetermined signal built by the same code path, and no fund trade is imputed to them. Bars are point estimates. Whiskers are 95% restricted wild cluster bootstrap intervals, Webb six-point weights, 9,999 replications, clustered on the rebalance.

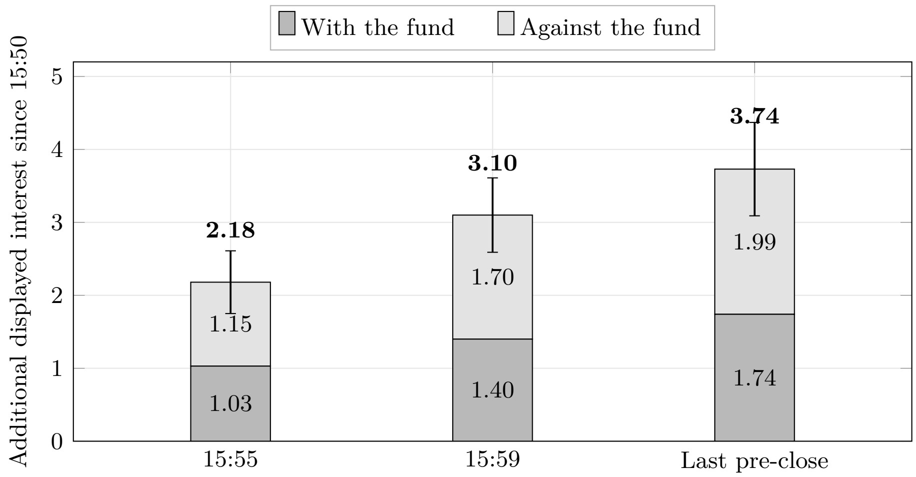

Figure 4: **More Interest Arrives after the Auction State Becomes Public.** After the auction state first becomes public, additional buying and selling interest continues to accumulate on both sides of the market. Each bar is the additional displayed interest accumulated since 15:50, in basis points of shares outstanding per basis point of the absolute scheduled target, split into the side the fund must trade and the side against it. Estimates are equal-event averages of within-rebalance slopes on the 25 quadruple-witching rebalances whose stock panel is balanced across the four clocks the design compares, from equation (8). Bold figures give the total. Whiskers are the 95% equal-event interval for the total additional interest, and the component estimates and their standard errors are reported in Table 4, Panel B.

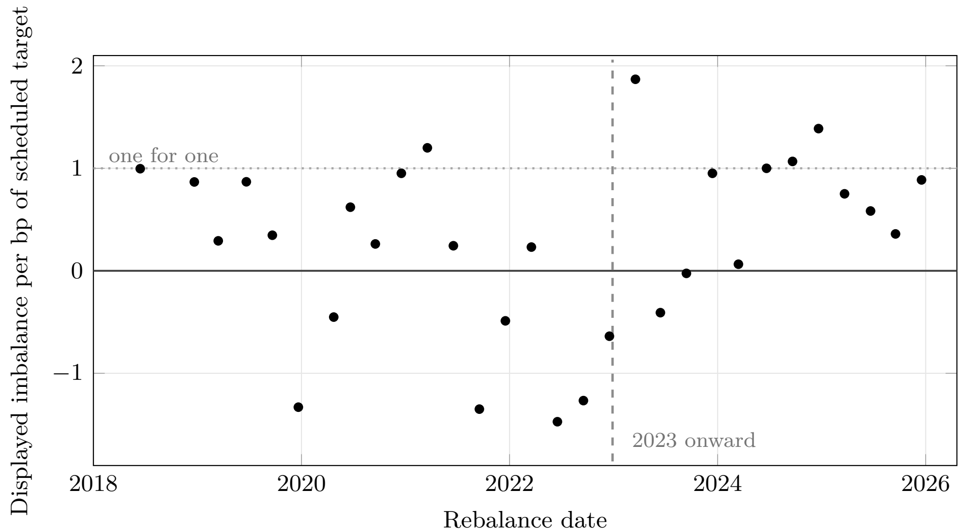

Figure 5: **The Target Predicts the Imbalance Very Differently across Rebalances.** One point is one rebalance's target-to-imbalance loading, with no smoothing and no line joining them, since each rebalance is a separate experiment. The vertical axis is the within-rebalance slope from equation (9), the signed displayed imbalance regressed on the signed rule-implied target, measured at 15:50, the first moment both exchanges publish the auction state. Quantities are basis points of shares outstanding, so the slope is basis points of imbalance per basis point of target. All 30 rebalances of the estimation sample are shown, using the same marker throughout. The solid line marks zero and the dotted line marks one, the value a researcher dividing by the target implicitly assumes.

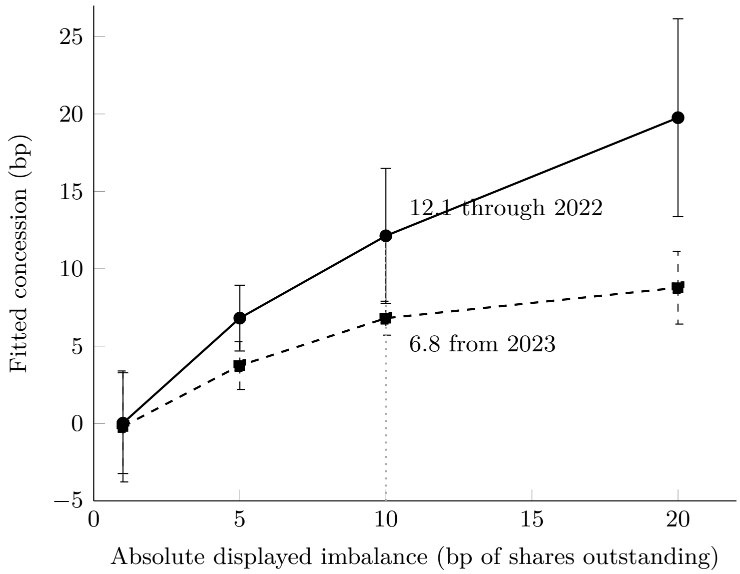

Figure 6: **The Price–Imbalance Schedule Is Lower after 2022.** The fitted concession against the displayed imbalance at 15:55, estimated separately on events through 2022 and from 2023, with event-clustered 95 percent intervals. At an imbalance of ten basis points the concession falls from 12.1 to 6.8 basis points, and the stability test across the split rejects with $\chi^2 = 40.9$. Estimated on the 32 rebalances with auction support. The plotted coordinates are generated from the same estimation output that produces Table 6.

Table 1: The Same Price Response, Four Candidate Denominators

Quantity treated as the demand shock	Multiplier M	Reciprocal $1/M$
Observed quantities		
Scheduled target Q	-0.34 (0.07)	-2.91
Displayed imbalance R	0.42 (0.08)	2.38
Decomposition of the displayed imbalance		
Fitted part $\widehat{R}$	-0.31 (0.15)	-3.27
Conditional residual u	0.44 (0.08)	2.28

Notes: Every row uses the same price move, the same observations and the same estimator. The only thing that changes is which quantity is treated as the demand shock. Each row estimates equation (3) with that quantity in place of Q, as a separate within-event regression on the same 30 historical rebalances, with event-clustered standard errors in parentheses. M is the price response per basis point of the quantity named in the row. The fitted part of the imbalance and the conditional residual are defined by equations (4) and (5). The conventional elasticity is $-1/M$. Reporting $1/M$ reverses every entry relative to that convention and leaves the disagreement across rows unchanged, which avoids an arbitrary sign flip. The price move is measured over the final five minutes of the auction, from the reference price at 15:55 to the official close, with each quantity taken from the 15:55 state, so these multipliers are short-window objects and are not comparable with multipliers measured from the reference close to the official close. Basis points of market capitalization and of shares outstanding coincide at a common price, so the four rows are directly comparable with one another.

Table 2: Sample and Economic Scale

Panel A: Fund and rebalance scale	
Absolute target across stocks, median (bp of shares outstanding)	1.35
75th percentile	3.68
90th percentile	8.10
standard deviation	6.20
One-way scheduled turnover per rebalance, median ($bn)	1.36
25th percentile	0.78
75th percentile	2.48
maximum	6.12
One-way turnover as % of fund assets, median over all rebalances	4.4%
Stocks traded per rebalance, median	437
minimum	258
maximum	481
Fund equity assets, start of sample ($bn)	14.9
Fund equity assets, end of sample ($bn)	92.2
Panel B: Auction data	
Verified rebalances, 2010–mid-2026	62
with message-level auction data	32
usable for target-dependent designs	30
Sample period	June 2018–June 2026
Stock-event observations at 15:50	12,605
Auction messages on the 30 rebalances	3,747,401
median messages per stock-event	297

Notes: Q_{ie} is the scheduled target for stock i at rebalance e, defined by equation (1) and computed from the published equal-weight rule at the reference close and expressed in basis points of shares outstanding. Panel A reports the cross-rebalance distribution of scheduled turnover. The $|Q|$ percentiles are computed over the pooled stock-event observations at 15:50 rather than as within-rebalance percentiles averaged across rebalances. One-way turnover is the dollar value of scheduled purchases at the reference close. Panel B counts rebalances passing the holdings verification and the subset carrying message-level auction data, and reports the auction messages falling on those rebalances. The raw feed additionally contains 18,510,938 messages on matched non-rebalance dates used as controls elsewhere in the paper. Table 3 reports the validation of the target against realized holdings and its predictability before the reference close.

Table 3: The Scheduled Target Is Real and Knowable in Advance

Does the fund trade the target?	
Realized-on-predicted holdings slope	0.82
Median within-event R^2	0.81
Median predicted-realized correlation	0.90
How early can the target be computed?	
Correlation with the final target, 20 trading days before	0.86
10 trading days before	0.93
5 trading days before	0.97

Notes: The upper block compares realized holdings changes with the scheduled target of equation (1) on the 62 verified rebalances, with slope standard errors of 0.01 to 0.02. The lower block reports the within-event correlation between the target reconstructed from information public k trading days before the reference close and the target at the reference close. Figure 2 plots that correlation at every horizon.

Table 4: How the Scheduled Target Maps into the Auction State

Panel A: First published state at 15:50. Cross-stock slope per 1 bp larger target, 30 rebalances					
Outcome	Pooled stock-events	95% CI	Equal rebalance	Positive in	Held-out R^2
Market size					
Total buying and selling	9.69	[8.16, 11.20]	10.34	30/30	20.9%
Interest on the fund's side	5.10	[4.25, 5.93]	5.31	30/30	—
Interest against the fund	4.59	[3.87, 5.30]	5.03	30/30	—
Market direction					
Net buy–sell imbalance	0.48	[0.10, 0.84]	0.28	—	2.1%

Panel B: Additional interest after 15:50. Changes, not levels, on 25 balanced rebalances, not addable to Panel A				
Outcome	Equal rebalance	Pooled stock-events	95% CI	Positive in
Total buying and selling	3.74	3.41	[2.76, 4.22]	25/25
With the fund	1.74	1.49	[0.92, 2.02]	25/25
Against the fund	1.99	1.92	[1.48, 2.42]	25/25

Panel C: First-state slopes by era, pooled over stock-events					
Outcome	Through 2022	95% CI	From 2023	95% CI	Change
Total buying and selling	8.78	[5.39, 13.85]	9.99	[8.43, 11.66]	1.21
Interest on the fund's side	4.24	[2.52, 6.74]	5.39	[4.47, 6.26]	1.15
Interest against the fund	4.55	[2.84, 7.16]	4.60	[3.83, 5.41]	0.05
Net buy–sell imbalance	-0.24	[-0.62, 0.27]	0.73	[0.30, 1.13]	0.97

Notes: Panels A and C report cross-stock slopes estimated on quantities observed at 15:50, not auction levels. Panel B reports cumulative changes after 15:50 on a different balanced sample and cannot be added to the Panel A estimates. The market-size rows regress each quantity on the absolute target while the net-imbalance row regresses the signed imbalance on the signed target. The difference between the two side slopes therefore need not equal the imbalance loading, and the small gap between 5.10 minus 4.59 and 0.48 is that difference in regressor rather than a rounding error. Panel A estimates equations (6) and (9) on 12,605 stock-event observations across 30 rebalances, with a rebalance fixed effect, so slopes are identified across stocks inside the same auction. Panel B estimates equation (8) on the 25 quadruple-witching rebalances whose stock panel is balanced across the four clocks the design compares, which is the same set of events used in Table 5. Panel C estimates the Panel A equations separately on each era, on 18 and 12 event clusters. The rotation in the net imbalance is 0.97 with a t of 3.75. Charging the procedure for having searched all 23 admissible split dates gives a permutation p of 0.14, so the era contrast describes the sample and is not a dated regime change, as Section 5 states. Quantities are basis points of shares outstanding. Q is the fund's scheduled target at the reference close. Interest on the fund's side and against it are the displayed buy and sell interest oriented by $\text{sign}(Q_{ie})$, and neither is defined relative to the auction's own contemporaneous imbalance. Total buying and selling is their sum. Pooled columns are within-event demeaned OLS, which weights each rebalance by its within-event dispersion in the regressor. Intervals are 95% restricted wild cluster bootstrap intervals, Webb six-point weights, 9,999 replications, clustered on the rebalance. The equal-rebalance columns weight every rebalance the same. The positive column counts rebalances with a positive within-event slope. The held-out R^2 is leave-one-rebalance-out, so each rebalance is predicted from the others.

Table 5: Predetermined Demand and Displayed Interest, Index Members and Same-Date Nonmembers

Panel A: Slope by group, per basis point of the predetermined signal

	S&P 500 members		Nonmembers	
Outcome	Slope	95% CI	Slope	95% CI
Opposite-side interest	4.38	[3.75, 5.12]	0.23	[0.09, 0.38]
Same-side interest	4.92	[4.16, 5.68]	0.21	[0.07, 0.37]
Buying and selling interest	9.29	[8.13, 10.68]	0.44	[0.16, 0.72]
Rebalances with a positive slope	25/25		19/25	
Cross-sectional SD of the signal (bp)	5.45		16.13	

Panel B: Difference between members and nonmembers

Specification	Difference	Std. error	95% CI (primary)
Control main effects	7.32	(0.66)	[5.90, 8.84]
Controls interacted with the signal	7.72	(0.73)	[6.14, 9.43]
Matched within auction on size and turnover	7.34	(0.68)	[5.86, 8.91]

Notes: The sample is 25 quadruple-witching rebalances on which both groups have auction support, 17,220 stock-event observations at 15:50. Quantities are basis points of shares outstanding. For S&P 500 members the signal is the fund's scheduled target at the reference close. For Russell 1000 stocks outside the S&P 500 it is the same drift signal, built by the same code path from the same public inputs. These stocks carry no scheduled rebalance exposure and no counterfactual fund trade is imputed to them. Same-side and opposite-side interest are the displayed buy and sell interest oriented by the sign of the signal. Buying and selling interest is their sum. Panel A estimates equation (6) separately by group, with a rebalance fixed effect, so slopes are identified across stocks inside the same auction. Panel B estimates equation (7) on the stacked sample. The reported coefficient is the interaction of the absolute signal with S&P 500 membership. Controls are log market capitalization and log twenty-day turnover, built for both groups from the same CRSP source at the same reference date. The matched row keeps only within-rebalance market-capitalization by turnover quintile cells containing both groups, dropping 2,458 observations. Intervals are 95% restricted wild cluster bootstrap intervals, Webb six-point weights, 9,999 replications, clustered on the rebalance. Standard errors in parentheses are event-clustered.

Table 6: The Price Schedule Rises with the Displayed Imbalance and Is Lower after 2022

All entries are the fitted remaining move to the close, in basis points		
Panel A. The overall schedule		
Displayed imbalance	At 15:55	At 15:59
5 bp of shares outstanding	5.8 [4.4, 7.3]	3.7 [2.2, 5.2]
10 bp	10.1 [7.1, 13.1]	5.5 [3.4, 7.6]
20 bp	15.0 [10.6, 19.4]	8.9 [5.5, 12.2]
Panel B. How the schedule changed, at an imbalance of 10 bp		
	Through 2022	From 2023
15:55	12.1 [7.8, 16.5]	6.8 [5.7, 7.9]
15:59	6.8 [3.8, 9.8]	3.0 [1.5, 4.5]
Panel C. Out-of-sample confirmation, six 2025 rebalances at 10 bp		
	Estimate	
15:55	6.5 (1.1)	
15:59	2.5 (0.6)	

Notes: The remaining move to the close is the concession of equation (2), the distance the official close settles from the reference price in the direction of the heavier side of the auction. It is fitted by restricted cubic splines on the displayed imbalance observed at the stated time, within events, with event-clustered inference. Figure 6 plots the 15:55 era comparison of Panel B and this table reports the fitted values behind it. Panel A pools the 32 auction-supported rebalances, with 95 percent event-clustered intervals in brackets, which exclude zero from an imbalance of five basis points upward. Panel B reports fits on each side of the sample split with 95 percent confidence intervals in brackets. The 15:55 stability test across the split rejects with $\chi^2 = 40.9$, while 15:59 stability is not rejected at a minimum detectable ratio near two. Panel C is a validation rather than six additional observations. The schedule never uses the fund's target, so every choice behind Panels A and B was fixed on the earlier rebalances before the 2025 events were acquired, and the same construction was then applied to them without adjustment. Standard errors in parentheses. The venue-level and conditional-residual specifications are in the appendix.

Table 7: Russell Reconstitutions: A Predetermined Quantity and Large Residual Imbalances

Panel A: Validation of the predetermined migration quantity			
Sign agreement with later tracker changes			99.0%
Equal-event correlation			0.985
Equal-event slope			0.979
Panel B: Dispersion of the residual displayed imbalance at 15:55			
	Dispersion (bp of shares outstanding)		
Reconstitution	Event day	Matched control	Ratio
2018	25.7	0.17	154.7
2019	27.0	0.11	241.3
2020	13.6	0.66	20.5
2021	23.6	0.22	109.1
2022	27.5	0.66	41.4
2023	18.7	0.57	32.5
Reconstitutions above their control		6 of 6, exact $p = 0.031$	

Notes: 389 genuine Russell 1000/2000 migrations across the six verified reconstitutions of 2018–2023. Panel A validates the pre-event tracker-based migration quantity against later tracker holdings changes. Panel B takes the displayed imbalance of migrating stocks, removes the part predictable from the direction of the migration and the stock's float, and reports the median absolute deviation of what remains, event day against a matched non-reconstitution June Friday. The last row counts reconstitutions whose event-day dispersion exceeds their own matched control. The event-day figure sits between 13 and 28 basis points in every year. The control side is what moves, because on an ordinary June Friday most of these stocks carry almost no displayed imbalance at 15:55, so the ratio is wide mainly through its denominator. No Russell pricing schedule is estimated.

A Appendix: Measurement and Robustness

A.1 The Rebalance Sample and the Scheduled Target

The RSP calendar is constructed from rule-implied quarterly reset dates and then verified against actual fund holdings. An event is retained only when the holdings interval bracketing implementation exhibits the equal-weight-reset signature and dominates placebo intervals. This procedure matters because scheduled calendars can contain displaced or postponed executions. The verification identifies 62 events from 2010 through mid-2026. The closing-auction message sample is the subset of 32 events with the required venue and clock support. Each design uses the maximal set of those events meeting its own mechanical support conditions, and acquisition vintage never determines inclusion: analyses that depend on the target and on observed historical membership run on the 30 historical events, while analyses of how the displayed imbalance is priced, which do not depend on the target, run on all 32. We make no held-out claim for the target-dependent analyses, because the target was rebuilt after the later events existed. The price schedule is the one exception, and only because it never uses the target: it is fitted on the displayed imbalance and the closing price alone, so a specification fixed on the earlier rebalances can be carried unchanged to the later ones. One criterion in the code defines the 25-event set. It keeps a rebalance when it falls on a quadruple-witching date and the same stocks are present at all four clocks the design compares, 15:50, 15:55, 15:59 and the last observation before the close. That is an inner join across the four snapshots, so it is a balanced-panel condition and not a claim that any clock went unpublished. Of the 30 target-supported rebalances it drops one off-cycle reset and the four most recent events. The growth design of equation (8) and the member and non-member comparison of equation (7) both run on that set, so their samples coincide by construction rather than by coincidence.

Two constructions of the early-vintage order exist in this project, and they agree. The forecastability path in Table 3 correlates the level of the reconstructed order at each horizon with the final rule-implied order. The second differences those same cumulative reconstructions into sequential revisions ($q_{10} = D_{10} - D_{20}$, and so on) and reports statistics on the revisions rather than on the levels, and its near-zero cross-correlations between early and late components are an empirical property of those revisions rather than a measure of level agreement. Differencing cumulative reconstructions is additive but does not by itself make the pieces orthogonal, so those correlations are a finding rather than a property of the decomposition. Reconciling the two pipelines stock by stock shows that they use the same holdings report dates, the same return paths, the same constituent universe, and an identical final target, equal to machine precision. Their level forecasts differ only through two minor vintage choices, the panel scaling to fund assets at the holdings report date rather than assets rolled to the information date and marking market capitalization at the information-date close rather than the prior close, which together move the horizon correlations by at most 0.02 in either direction. With the two vintage choices aligned the pipelines coincide to four decimal places. The reconstruction underlying Table 3 is therefore not sensitive to which pipeline computes it.

The holdings first stage is intentionally separate from the auction outcomes. The public-rule quantity is computed at the reference close. Later holdings then test whether the fund actually moved in the predicted stock-by-stock direction and size. No auction imbalance or price outcome enters the verification. The high within-event R^2 therefore validates the initiating quantity without using the displayed imbalance that is the paper's object of interest.

A forecast made a month ahead misses the final target because the inputs to the rule are still moving. We attribute that error to the inputs by updating them one group at a time, averaging over every order in which the groups can be updated. Relative price movement resolves almost all of it. Changes in the fund's assets and in its holdings roughly offset one another, and share counts and corporate actions account for a fraction of a percent. Prices are the largest contributor whatever order we update the groups in.

A.2 The Auction Message Feeds

A literature studies these same broadcasts on how closing imbalances evolve and how they are priced (Jegadeesh and Wu, 2022). The two venues begin publishing at different times, and where an early observation is missing at a venue we leave it missing rather than filling it in with a message from later in the auction. Section A.2.3 shows how the paired quantity follows from the two displayed sides, and Section A.2.1 sets out what each venue discloses and when.

Our auction data are the historical imbalance-message feeds of the two listing exchanges, Nasdaq's Net Order Imbalance Indicator on the TotalView-ITCH feed and the NYSE's Closing Auction Imbalance Information, obtained as raw historical messages from Databento, which delivers both venues in one common set of fields.

Interest reaches a closing auction through market-on-close and limit-on-close orders, which must be submitted before a cutoff, and through later offsetting order types the exchanges admit specifically to meet a published imbalance. NYSE additionally carries floor-broker discretionary interest, the D orders whose disclosure timing changed in 2024 and which Appendix A.2.1 discusses. The fields we use are few. The reference price anchors the auction, the published buy and sell interest give the two sides, and their difference is the displayed imbalance and the paired quantity is what already matches. We take the last message at or before each auction time, drop rows whose reference price is missing, non-positive, or carries one of the sentinel values the feed publishes when it has no price, and express every quantity in basis points of shares outstanding.

At 16:00 each exchange runs a single-price cross that maximises executable volume, resolving ties toward a smaller residual imbalance and then toward the prevailing continuous price. Nasdaq's cross is fully electronic. At the NYSE the designated market maker completes the auction and may supply liquidity against a residual gap. In our data the final published reference price before the cross is within two basis points of the official close at the median.

A.2.1 Venue Disclosure Regimes

NYSE and Nasdaq auction protocols change over the sample. The two features most important for interpretation are the NYSE D-interest disclosure boundary at 15:55 and the Nasdaq early-dissemination support. The NYSE 15:55 step can mechanically reveal additional interest and is therefore not treated as evidence that economic uncertainty arrives at 15:55. The clean NYSE anatomy is the post-disclosure decline. Nasdaq provides a cleaner early-to-late path in 2020–2023, when the feed disseminates from about 15:50 and becomes more frequent after 15:55.

These rule differences are why the paper emphasizes fixed-support absolute event levels. Near the close, matched-control residual MAD becomes small and exact zeros become common. Late event/control ratios are therefore partly denominator-driven even though the event residual itself remains economically meaningful. The early 15:50 level, by contrast, is far above the measurement floor.

A.2.2 Multiple Share Classes and the Company Slot

Equal weighting is assigned at the level of the company, so a company with more than one listed share class divides a single slot across its lines rather than receiving one slot per line. The exchange that lists futures on this index states the rule directly: for companies with multiple share class lines in the index, each line is assigned a weight proportional to its float-adjusted market capitalization, with weights set to 1/500 for each company at rebalancing.[12]

The fund's own filings show that structure without any need to interpret index language, on two independent dates. Its portfolio report for the period ending April 2026 lists 508 lines across 505 distinct issuers, the extra lines being exactly Alphabet, Fox and News Corp. Each of those companies' classes sums to a figure about the size of one ordinary company slot, 0.2386, 0.2099 and 0.2099 percent against a median single-line weight of 0.1945 percent, while each individual class line sits materially below one slot. The same pattern holds four years earlier: at July 2021, against a median line weight of 0.1989 percent, Alphabet's two classes sum to 0.2248 percent and Fox, News Corp, Under Armour and Discovery to between 0.1826 and 0.1890 percent, every company total inside the single-slot range and every class line far below it.

The distinction does not arise in the early part of the sample. No multi-class lines appear at any of the fifteen rebalances from September 2010 through March 2014. The first arrive in June 2014 and September 2014, and Discovery enters as two lines rather than three, matching the index's own composition.

[12]CME Group, *FAQ: E-mini S&P 500 Equal Weight Futures*, 25 February 2024. The index methodology uses company-level language throughout, stating that each company is equally weighted as of the rebalance reference date, while the float-adjusted weighting rule in the same document is evaluated line by line. No index document states a per-line rule. Structured-note prospectuses sometimes describe each constituent as equally weighted, which is a simplification contradicted by the filings below.

A.2.3 Feed Quantities and the Variance of the Displayed Imbalance

Paired quantity is not an independent feed quantity. It is algebraically

$$Paired(s) = \min\{Buy(s), Sell(s)\} = \frac{Buy(s) + Sell(s) - |I(s)|}{2}, \tag{12}$$

a transform of the two displayed sides, so when both sides grow paired quantity must grow. We therefore never treat it as an independent behavioral outcome, and the economically distinct facts are the sizes of the two sides and the residual difference between them.

The measurement claim in Section 2.4 can be stated formally. If the target Q_{ie} is fixed in the pre-event information set $\mathcal{I}_e^{pre}$, then

$$\operatorname{Var}(R_{ie}(s) \mid \mathcal{I}_e^{pre}) = \operatorname{Var}(q_{ie}(s) + C_{ie}(s) \mid \mathcal{I}_e^{pre}), \tag{13}$$

which need not be small. Knowing the initiating order removes uncertainty about the target, not about how much of it reaches the auction or about what everyone else brings to it.

A.3 Inference and Reporting Formats

Because several designs contain few rebalances, conventional cluster-robust standard errors are unreliable, so we use a wild cluster bootstrap instead. The bootstrap holds the regressors fixed and perturbs each rebalance's residuals by a random weight, drawn from the six-point distribution of Webb (2023). Those weights matter most in the smallest samples, where the more common two-point weights can produce only a handful of distinct draws. We compute the residuals with the null imposed, and we build a confidence interval by testing each candidate value of the coefficient and keeping the values the bootstrap does not reject. The pricing schedules of Section 6 and the appendix tables that accompany them report analytic event-clustered inference instead, which their notes state. Standard errors appear alone only in designs where the bootstrap is unavailable, and their notes say so.

A.4 The Prediction Rule and the Unexplained Imbalance

The RSP prediction rule leaves one entire event out. For each held-out event, the relationship between the scheduled target, the standing opposite-side interest, and the 15:55 displayed imbalance is learned only from the other events. This prevents the within-event cross-section from fitting away the same event's idiosyncratic state. We fit it by ordinary least squares on an intercept and the conditioning variables in levels, scaled in basis points of shares outstanding, with no transformations, no regularization, and no variable selection, so each held-out event's fitted values come from a regression estimated on the pooled stock-event-venue observations of every other event. The out-of-event R^2 compares the stacked held-out predictions against the unconditional mean of the realized displayed imbalance over the same evaluation sample, so a negative value says the rule forecasts worse than that constant. Weighting observations equally and weighting events equally give the same values at the precision we report, and the evidence therefore speaks to linear forecasts of this class rather than to unforecastability in principle. Both reported specifications, Q alone and Q with the standing state, deliver out-of-event R^2 values that are negative and nearly identical under equal-observation and equal-event weighting, and the conditional residual used throughout the paper comes from the second.

A.4.1 Out-of-Event Prediction of the Displayed Imbalance at 15:55

The known target does not explain the displayed imbalance out of event. Explaining the 15:55 state with Q alone yields an out-of-event R^2 of -0.0581. Handing the regression the opposite-side interest from the same snapshot yields -0.1193, so even given one contemporaneous side of the auction, the rule still fails to predict a held-out rebalance. Both figures average event-level R^2 with events weighted equally, the experimental unit throughout. Scoring instead against a single pooled mean over all held-out observations gives 0.0246 and 0.0294, which sit near zero because that baseline absorbs cross-event differences in level. The two conventions differ in the benchmark they score against, not in how observations are weighted within an event, and Appendix A.4 states the metric. The conditional residual carries essentially the entire variance of the state, $\mathrm{Var}(u)/\mathrm{Var}(R) = 0.971$. This is not a separate estimate: pooled out-of-event R^2 is $1 - \mathrm{SSE}/\mathrm{SST}$, and the two coincide when the conditional residual has zero mean over the evaluation sample, which is the convention used here.

The unexplained part is large. At 15:55 the conditional residual u has a mean absolute deviation of 6.25 basis points of shares outstanding on rebalance dates against 1.55 on the 28 matched option-expiration Fridays, a ratio of 4.04 on the 26 events the control calendar references and 4.43 across all 32 events. A rebalance day therefore carries roughly four times the unexplained dispersion of an already busy control close.

The dispersion of u contracts as the cross approaches. On the panel of names observed at every auction time, it falls from 9.59 basis points at 15:50 to 6.79 at 15:55 and 4.72 at 15:59, a decline that appears in all 32 of 32 events with an exact two-sided sign probability below 0.001, so most of what stands early is resolved before the cross while what remains is still several times an ordinary day's. The contraction itself, however, is not special to rebalances. Applying the same construction to the matched control dates, mean $|u|$ falls to 0.70 of its 15:50 value by 15:55 and to 0.49 by 15:59 on rebalance dates, against 0.69 and 0.53 on control dates. The proportional contraction is therefore broadly comparable on the two kinds of date. The rebalance changes the state the auction begins with, not the rate at which it works a revealed imbalance down. These are point estimates, not a test of equality, and the control residual is built from buy-side interest rather than the target-oriented opposite side, so the two constructions are not identical.[13] Table A5 collects

[13]The event is the experimental unit throughout, with equal weighting across events and clustering on the event wherever a pooled regression appears. The stock panel defines the within-event objects.

these estimates.

Whatever the price response to Q turns out to be, the quantity the auction displays is not a small perturbation around the target. It is large in its own right, and knowing the target in advance does not shrink it.

A.4.2 Control-Date Residual Construction

Control dates have no scheduled rebalance order, so the counterpart to u has to be built explicitly. Each supported control date is matched to the next verified event in the same calendar year, and the observation set on that date is the matched event's stock universe, so the comparison holds the names fixed. We then set $Q = 0$, compute the displayed imbalance and the conditioning variable from the control-date feed, and apply the matched event's leave-one-event-out coefficients without re-fitting anything on control data. The control residual is the control-date state minus that prediction.

Two consequences of setting $Q = 0$ should be stated. The prediction collapses to the intercept plus the loading on the conditioning state, since the term in Q drops out. The coefficients are fitted in the target-oriented frame, so their intercept carries the oriented lean. Transporting it to a control date, where there is no direction to orient by, shifts every residual there by a constant. The control benchmark is therefore evaluated after demeaning, so that event and control statistics both measure dispersion about their own centres. And the conditioning state on control dates is buy-side interest rather than the side opposite a trade direction, because direction is undefined without a scheduled order. The control residual is therefore not the same object as u in every respect. It is the same machinery evaluated where the scheduled quantity is zero, which is the relevant benchmark for asking how much dispersion a rebalance adds. This asymmetry is retained exactly rather than repaired, since orienting by a direction that does not exist would be a choice made after the fact.

Date-level dispersion is the mean absolute deviation about zero. Tier means are equal-weighted averages across dates within a tier, so the ratio of 4.04 compares the average rebalance date with the average option-expiration date rather than pooling stock observations across dates.

A.4.3 Scale of the Residual

The magnitude of the conditional residual rises with the size of the scheduled target, and not weakly. Across the 30 historical rebalances, one basis point of absolute scheduled demand is associated with 0.80 basis points of additional absolute residual imbalance, with an event-clustered standard error of 0.11, and 0.62 (0.13) after controlling for size and turnover. We had fixed 0.25 in advance as the value below which we would treat the residual as scaling only weakly with the size of the trade. Both estimates are well above it, so that reading is rejected.

One caution attaches to the estimate. A coefficient of this size does not establish that the residual is proportional matching error, in the sense of a random fraction of the scheduled target failing to find a counterparty. It says that residual magnitude rises with trade size, which is a weaker and different statement, and one that has to be read alongside the very large endogenous two-sided market documented in Section 3.

The result sharpens rather than weakens the paper. The scheduled target predicts how much two-sided market assembles, and it predicts how large the remaining clearing surprise tends to be. What it continues to predict poorly is the sign and the stock-by-stock realization of that surprise.

A.5 The Order–Price Multiplier

The reduced-form multiplier is the within-event regression of the reference-close-to-implementation-close return on the scheduled target, estimated on 30 historical rebalances. Write $\widehat{g}$ for the fitted pricing function of the displayed imbalance, evaluated in levels at each stock-event's own displayed imbalance, and let $\widehat{e}_{ie} = \Delta p_{ie} - \widehat{g}(R_{ie})$ be what it leaves. The multiplier then splits as $M_Q = T_1 + T_2$ with $T_1 = \mathrm{Cov}(\widehat{g}(R), Q)/\,\mathrm{Var}(Q)$ and $T_2 = \mathrm{Cov}(\widehat{e}, Q)/\,\mathrm{Var}(Q)$. The split is exact by construction, since $\Delta p = \widehat{g}(R) + \widehat{e}$. It is an accounting check on the fit rather than a result, and T_2 is the part of the multiplier the fitted schedule does not account for. The identity holds event by event to machine precision. The largest absolute discrepancy across all events is 0.000000000000007.

Two things follow, and neither supports treating the decomposition as structural. First, the reduced form itself depends heavily on how events are aggregated: 6.44 weighting events equally, 3.04 weighting by precision, and 2.13 at the median. Second, and more telling, the component running through the observable displayed imbalance is small and stable while the remainder carries essentially all of the variation. Equal-event, T_1 is 0.65 against a remainder of 5.79; at the median the two are 0.98 and 1.52.

The era split is the clearest case. Across the 18 early rebalances the multiplier averages 11.40, with T_1 of 0.51 and a remainder of 10.89. Across the 12 later ones it averages -1.01, with T_1 of 0.85 and a remainder of -1.87. The reduced form moves by more than twelve units across the break while the part of it attributable to the observable clearing channel moves by a fraction of one. Restricting the price window to the final five minutes, where the auction is essentially all that can happen, the multiplier averaged equally across events is -0.28, against the pooled within-event -0.34 that Table 1 reports over the same window, with T_1 of 0.12 and a remainder of -0.40: the two components point in opposite directions even there.

The arithmetic is exact. The economic identification is not there, and we do not use the decomposition to claim recovery of its separate parts.

A.5.1 The Price of the Conditional Residual

The displayed imbalance is priced. Is its conditional residual priced too? The outcome is the log official closing-auction execution price relative to the listing-venue continuous-market midpoint in the five seconds ending at 15:55, on 10497 stock-event-venue observations across 26 events. That is the set of events for which this quote and statistics construction is available.

Orienting the price move by the sign of the residual and regressing it on $|u|$ gives 0.58 basis points of price per basis point of the conditional residual, with an event-clustered standard error of 0.13 and a t statistic of 4.56. Stating the same hypothesis the other way, regressing the unoriented price move on signed u while controlling for the scheduled target and the standing opposite-side interest, gives 0.71 with an event-clustered standard error of 0.16. The two forms agree. In that second specification the scheduled target itself carries -0.06. Because u is built from Q and L, that regression is a reparameterisation of one on the displayed imbalance and the same two controls, so this coefficient is not the effect of the target holding the displayed imbalance fixed. A larger conditional residual is associated with a larger concession in the direction of the shortage.

We do not read this as causal price impact. The residual u is an equilibrium object, jointly determined with the price it is being related to, and Appendix A.5 shows that a price regression of this family can produce large coefficients on dates carrying no scheduled target at all. The claim is the conditional association, and the role it plays here is to complete the chain: the order organizes a market far larger than itself, that market leaves a displayed imbalance whose direction is largely unpredictable, and that state is what the closing price responds to.

Table A4 reports both specifications.

A.6 Russell Migrations

The Russell sample begins from official membership histories. Genuine migrations are stocks that move between the Russell 1000 and Russell 2000. Direct additions from outside the Russell universe, direct deletions, non-switchers, and ambiguous transitions are excluded. Identifier mapping loses six migration observations across the six primary years. The final matched sample contains 389 migrations.

For institutional validation we use plain benchmark trackers only. Let $\kappa_{j,e}$ denote the fraction of float held by the plain trackers of index family j before event e. For a migration from source index s to destination index d, the pre-event share proxy is

$$Q_{ie}^{proxy} = (\kappa_{d,e} - \kappa_{s,e}) F_{ie}^{May}, \tag{14}$$

where F^{May} is May float-adjusted shares. This proxy uses no June price, implementation-close weight, or post-event holdings in its construction. Later tracker holdings are used only to validate the sign and scale. Because normalization by shares outstanding turns most within-direction variation into the May float factor, the proxy is not used as a stock-level regressor.

The proxy also provides a check on the direction-and-float residualization used in Section 6.5. After shares-outstanding normalization the proxy is the event-by-direction ownership-fraction gap times May float, so the float adjustment spans almost all of its within-cell predictable component. We therefore keep the simpler predetermined direction-plus-float residual. Applying the ownership-fraction gap to genuine migrations predicts later tracker changes with 99.0 percent sign agreement, an equal-event correlation of 0.985, and an equal-event slope of 0.979.

Table A1: Order–Price Multiplier: Aggregation and Price-Sensitivity Calibration

Panel A: Order–price multiplier and its decomposition, by aggregation, 30 rebalances			
Estimator	M_Q	T_1	T_2
Equal event weight	6.44 (2.57)	0.65	5.79
Precision weighted	3.04 (2.32)	0.62	2.41
Median	2.13	0.98	1.52
Panel B: Reciprocal price-sensitivity calibration			
	Concession at 10 bp	Imbalance per bp of concession	
Through 2022	12.1 bp	0.8	
From 2023	6.8 bp	1.5	

Notes: In the median row the components do not sum to M_Q, because each median is taken separately across events. Panel A aggregates the event-level slopes of equation (3) under different weightings, with between-event variance parameters of 8 to 66 across samples and standard errors in parentheses. The exact decomposition of M_Q, and the evidence that its components are not stable across specifications, are reported in Appendix A.5. Panel B converts the era schedules of Table 6 by the usual reciprocal convention. The calibration is a unit conversion of the conditional pricing schedule, evaluated at a ten-basis-point imbalance, and not an estimate of a structural elasticity, which would require an exogenous shifter of the net of the orders that arrive.

Table A2: Two Scheduled-Rebalance Designs

	S&P 500 Equal Weight (RSP)	Russell 1000/2000 reconstitution
Economic event	Quarterly equal-weight reset of continuing S&P 500 constituents	Annual migration between Russell 1000 and Russell 2000
Scheduled quantity	Rule-implied stock-level RSP target fixed at reference close	Migration direction and plain-tracker scale validated pre-event. Exact RSP-style target not used
Residual-analysis sample	32 verified auction-supported events. 14,314 stock-event-venue observations at 15:55	6 verified events, 2018–2023. 389 genuine migrations at 15:55
Primary object	Conditional residual u after conditioning on Q and the same-snapshot opposite-side interest	Date-specific direction- and May-float-adjusted residual $\widetilde{R}$
Control benchmark	Monthly option-expiration benchmark	Matched non-reconstitution June Fridays
Experimental unit	Rebalance event	Annual reconstitution

Notes: The table summarizes the two designs used in the paper. RSP denotes the Invesco S&P 500 Equal Weight ETF. The RSP scheduled quantity is the public-rule trade of that fund at the reference close, scaled by shares outstanding and verified in later holdings. The Russell sample contains only genuine migrations between the Russell 1000 and Russell 2000. Direct additions, deletions, non-switchers, and ambiguous transitions are excluded. u and $\widetilde{R}$ are different residual objects. Russell is an external test rather than a mechanical replication of the RSP projection. Inference treats the event, not the stock, as the experimental unit.

Table A3: How the Fund's Target Predicts the Displayed Imbalance over Time

Clock	Loading (equal-event)	Event SE	t
15:50	0.2789	0.1588	1.76
15:55	0.2141	0.1357	1.58
15:59	−0.0135	0.0999	−0.13

Notes: λ is the loading of equation (9), the linear projection of the displayed imbalance on the signed scheduled target, both in basis points of shares outstanding, computed within event and averaged equally across the 30 historical RSP rebalances. Standard errors treat the event as the experimental unit. Pooled within-event weighting gives 0.4833, 0.3195, and 0.1602 at 15:50, 15:55, and 15:59. Support falls from 7,173 observations at 15:50 to 5,432 from 15:55 onward, and both counts are below the Table 4 sample because this table requires support at every clock it reports.

Table A4: The Price of the Conditional Residual

Specification	Estimate	Event SE
Oriented price move on $\|u\|$	0.58	0.13
Price move on signed u, with Q and standing interest	0.71	0.16
Scheduled target Q in the same specification	-0.06	0.15

Notes: The conditional residual is defined by equations (4) and (5). These estimates differ from the residual row of Table 1 in three ways and are not directly comparable with it: the outcome here is the continuous-market midpoint rather than the auction reference price, the sample is the rebalances carrying that quote construction rather than all 30, and event-level slopes are averaged with equal weight rather than pooled within event. Estimates are basis points of price per basis point of the conditional residual, on 10497 stock-event-venue observations across 26 rebalances. The price move runs from the listing-venue continuous-market midpoint in the five seconds ending at 15:55 to the official closing-auction price. The sample is the 26 events for which that same Databento quote and statistics construction is available. u is the held-out 15:55 conditional residual used throughout. Each row is an equal-event mean of within-event regressions, with the event as the experimental unit. The first row orients the price move by the sign of u and regresses it on $|u|$; the second regresses the unoriented price move on signed u while controlling for the scheduled target and the standing opposite-side interest. The two give the same qualitative conclusion. A larger conditional residual is associated with a larger closing concession in the direction of the shortage. u is an equilibrium residual rather than an exogenous shock, so these are conditional associations and are not interpreted as causal price impact.

Table A5: The Scheduled Target and the Displayed Imbalance

Panel A: Out-of-event prediction of the displayed imbalance at 15:55			
Prediction using the scheduled target	-0.0581		
Prediction adding opposite-side interest	-0.1193		
Unexplained variance, relative to the imbalance	0.971		
Dispersion, rebalance dates (bp)	6.25		
Dispersion, matched control dates (bp)	1.55		
Event/control ratio, historical common sample	4.04		
Event/control ratio, all 32 events	4.43		
Panel B: Scheduled demand and the displayed imbalance, by clock			
	Loading		
15:50	0.2789 (0.1588)		
15:55	0.2141 (0.1357)		
15:59	-0.0135 (0.0999)		
Panel C: The conditional residual through the final minutes			
Balanced-panel MAD of u (bp)	15:50: 9.59	15:55: 6.79	15:59: 4.72
Events with decline, 15:50 to 15:59	32/32	exact sign test, $p < 0.001$	

Notes: Panel A reports the out-of-event fit of equation (4) and Panel B the loading of equation (9). Panel C reports the path of the conditional residual. Panels A and C use the 32 rebalances with 15:55 auction support, 14,314 stock-event-venue observations for the prediction rows and 13,671 for the fixed-composition path. The event/control dispersion ratio is reported both on all 32 and on the 26 events the frozen control calendar references, against 28 matched option-expiration control dates. Panel B is the loading of the state on the target, which depends on the target and therefore uses the 30 historical rebalances. The event is the experimental unit throughout and event-level standard errors are in parentheses. u is the conditional residual after conditioning on Q and the same-snapshot opposite-side interest, computed out of event, and out-of-event R^2 is computed on held-out events. λ is the loading of the displayed imbalance on the target, equal-event weighted. The full clock-by-clock path is in Table A3.