

An Insider-Silence Precursor for Corporate Control Events and the Boundary of Its Transfer

Maksim Aleksandrovich Andreev, Joseph Howden
Quark Quantitative Research

Abstract

We report four measurements of insider behaviour around corporate control events, computed entirely from sources the Securities and Exchange Commission publishes at no charge, and the boundary that separates the one confirmed result from the three that failed. Under a design that holds the firm fixed and moves only the observation moment, an insider channel separates a deal moment from the same firm's deal-free moment at an area under the receiver operating characteristic of 0.6141 on 1,407 paired cases and 2,843 pairs, against a preregistered bar of 0.58 and a within-firm permutation 95th percentile of 0.5206. The direction inverts the premise that motivated the work. Every purchase-intensity feature scores below 0.5, and the only feature above it reads the *absence* of a non-routine open-market purchase in the trailing year. A follow-up unit asks whether the same feature conditions an outcome. On 2,853 announced reviews of strategic alternatives carrying 959 conversions, a base rate of 33.61% and complete feature coverage, the discrimination is 0.5176 against a frozen bar of 0.55. It exceeds a within-year permutation 95th percentile of 0.5148 at $p = 0.030$ in the preregistered direction and falls well short of the bar. The consequence bar fails alongside it: firms silent for a full year convert at 34.44% ($n = 1,925$, Wilson 95% interval [32.35%, 36.59%]) against 29.72% for firms whose officers purchased within the previous six months ($n = 461$, [25.73%, 34.04%]), a gap of 4.72 points against a 5.0-point bar with overlapping intervals. The intermediate stratum, 183 to 364 days of silence, converts at 34.05% ($n = 467$) and is not intermediate; the whole measured effect is one step at the recent-buyer boundary. A placebo arm scored three years before each announcement returns 0.5041 against a 0.53 ceiling, so what signal exists is announcement-linked rather than a permanent firm attribute. We attribute the failure of transfer to selection: the confirmed contrast is a firm against its own quiet moment, and once every subject has announced a review the variance that contrast reads has already been spent. Two earlier units are reported as they settled, a clustered purchase study that is null and powered with an inverted sign, and a case-control prevalence study sealed unread by a control-composition gate. Insider silence marks that a control process exists. It does not predict that the process will finish.

1. Introduction

Whether corporate insiders trade informatively before a takeover is among the oldest empirical questions in corporate finance. Keown and Pinkerton (1981) documented systematic pre-announcement price movement in targets; Seyhun (1986) and Lakonishok and Lee (2001) established that aggregate insider transactions carry modest predictive content for returns; Cohen, Malloy and Pomorski (2012) showed that the content is concentrated in a minority of opportunistic traders once mechanically recurring transactions are filtered out. Agrawal and Nasser (2012) reported that insiders of takeover targets trade

passively before an announcement, so that abnormal net purchase arises chiefly through withheld selling rather than through additional buying.

Two questions are routinely conflated in that literature and in practitioner use of it. The first asks whether insider behaviour marks that a control process is under way at a firm. The second asks whether it predicts that a process already visible will reach a definitive agreement. The two questions differ in their populations and in their nulls, and this paper reports that they differ in their answers. A signal can be real on the first and absent on the second, and the mechanism that produces that pattern is a property of the second population rather than a defect in the first measurement.

We report four units, all computed from the Commission’s quarterly ownership datasets and its full-text search and submissions endpoints, with no commercial transactions database at any point. One of the four confirms, and the other three fail in three different ways: a null that is powered and inverted in sign, a measurement sealed unread by an instrument gate, and a null that is small, directionally correct and far below the bar it had to clear.

Contribution. A confirmed within-firm insider precursor for corporate control events, with its direction inverted relative to the accumulation premise and its measured object corrected from purchase cadence to purchase absence. A preregistered transfer test of that precursor onto an outcome, on a population constructed so that the control-composition defect which sealed a predecessor cannot arise, and the null it returns. The observation that the stratum ordering carries no gradient, so that the residual effect is a single step at one boundary rather than a dose-response, together with an honest statement of what the available sample can and cannot exclude about that shape. A selection account of why the confirmed signal does not transfer, which bounds the confirmed result rather than contradicting it.

2. Sources, Instruments and the Routine Filter

Ownership transactions are read from the Commission’s quarterly structured datasets of Forms 3, 4 and 5, whose coverage across the runs reported here spans 3 January 2006 to 31 March 2026. Deal populations and outcomes come from the full-text search index and the per-filer submissions record, which are the same two endpoints that carry the companion probability ladder for corporate control events (Andreev and Howden, 2026f). Outcomes are typed by form code throughout and are never read from prose.

The transaction filter follows Cohen, Malloy and Pomorski (2012) and is applied at the level of an individual filing rather than by classifying a person once and for all. A purchase is treated as routine, and therefore excluded, when the same owner filed a comparable transaction in the same calendar month in prior years. One consequence of that construction

binds every result below and is stated wherever the results appear: the owner-history corpus available to the filter begins in January 2006 against a three-year lookback, so purchases cut before roughly 2009 are mechanically over-classified as non-routine because the history that would have marked them routine does not exist.

Every unit runs a point-in-time gate that refuses any record dated after the observation cut. In the conversion unit of section 5 the gate discarded 180,031 records across 10,328 calls, so it is an active constraint rather than a formality.

The feature. Days elapsed at the observation cut since the last non-routine officer or director open-market purchase, with a lookback window of 365 days. When no such purchase appears in the window the feature takes the value 365. Section 4.1 sets out what that construction can and cannot measure, and the answer is narrower than the feature’s name suggests.

3. Two Measurements That Shaped the Design

3.1 Clustered purchases and forward returns

The first unit tested the most direct reading of the insider literature, that clustered open-market purchases by several officers at one issuer within a short window predict positive abnormal returns over the following quarter. The primary horizon was 63 trading days and the design required at least 150 clustered events before a verdict could be written. It found 785, so the power bar cleared and the result is a settled null rather than a measurement limit.

The mean abnormal return at 63 trading days is -1.952% with a Newey-West t of -1.34 against a bar of 3.0. A 200-replication random-entry control into the same tickers has a 95th percentile of -0.955% , so the clustered basket performed worse than random entry rather than better. The effect is negative in all three time-thirds, which removes an era artefact as an explanation. The 21-day mean is $+0.203\%$ and is not the frozen primary. The free clustered-purchase channel therefore closes, and the register records that

any revival requires the opportunistic classification of Cohen, Malloy and Pomorski rather than raw clustering.

3.2 A case-control prevalence study, sealed unread

The second unit compared the prevalence of non-routine purchases in the window before a deal announcement against seeded controls, on 400 sampled deal episodes and 400 controls, with a two-proportion z of at least 3.0 as its primary bar. A prevalence gate frozen with the design required the control arm to fall between 8% and 60%. The observed control prevalence is 6.75%, below the floor, so the unit settled as measurement-limited and its statistics were sealed. A case-arm prevalence was computed and is recorded in the results file as non-citable under the frozen kill language, so it is not reproduced here; we note only that the seal was applied to a populated statistic rather than to an empty one. The control prevalence is quoted because it is the gate reading itself, which the frozen rule requires to be reported.

The defect the gate caught. A control pool seeded from all Commission registrants contains funds, shells and inactive issuers. Only 31% of control windows carried any ownership filing at all against 79% of case windows, so a contrast between the arms confounds anticipation of a deal with being an operating company. Every filing-derived feature inherits that confound, which is why the two successor designs either match controls on prior filing activity or eliminate the control pool entirely.

The harvest behind the sealed unit is sound and was retained: 5,595 ownership documents cached, with the parser exercised against planted purchase, sale and routine transactions, and six selftests passing before any scoring was attempted. The failure lies in the comparison population rather than in the instrument, and the harvest was carried forward into the successor designs.

4. The Within-Firm Design and the Confirmed Channel

The design that produced the confirmation removes the comparison population in a different way. Each deal episode is paired with moments at the same firm five, seven and nine years from its anchor, screened for admissibility by the control builder's own rules, with features cut 180 days before the anchor. Because the firm is held fixed and only the observation moment moves, the null that the score reflects a permanent firm attribute sits at 0.5 by construction. The identification argument is developed at length in a companion methods paper (Andreev and Howden, 2026g) and is not repeated here.

Of 2,853 case episodes, 1,976 contributed at least one admissible pairing, against 4,492 placebo moments, for 6,468 rows in total. Four channels were specified in advance: insider behaviour, ownership structure, fundamentals, and price and volume. A blocking coverage gate required 60% computability on both sides of the pairing before a channel's statistics could be read.

Table 1: Channel results under the within-firm design, bar 0.58

CHANNEL	PAIRED CASES	PAIRS	AUC	PERMUTATION P95	COVERAGE, CASE / PLACEBO	VERDICT
Insider	1,407	2,843	0.6141	0.5206	0.900 / 0.633	Passes
Ownership	1,976	4,492	0.5380	0.5122	1.000 / 1.000	Above permutation, below bar
Insider and ownership jointly	1,407	2,843	0.6086	0.5190	—	Passes, no better than insider alone
Fundamentals	—	—	sealed	—	0.751 / 0.373	Coverage gate refused
Price and volume	—	—	sealed	—	0.150 / 0.138	Coverage gate refused

The permutation null reassigns the pseudo-case uniformly within each pairing group and refits the out-of-fold model on every draw, 200 times. The insider channel clears both clauses of the primary bar. The ownership channel exceeds its own permutation ceiling and falls short of the bar, so it is weakly informative and does not certify. The joint model is no better than the insider channel alone. Two channels were never scored: the fundamentals arm reached 37.3% computability on the placebo side and the price

and volume arm reached 13.8%, both below the 60% floor, and the frozen rule seals a channel’s statistics rather than reporting them beneath a caveat.

The direction of the confirmed channel inverts the premise the unit inherited. Every feature that reads purchase intensity scores below 0.5, which is to say that a deal moment carries *less* insider buying than the same firm’s quiet moment. The single feature above 0.5 reads how long it has been since a purchase occurred.

Table 2: Insider features, case-weighted AUC against the within-firm null of 0.5

FEATURE	AUC	READING
Days since the last non-routine purchase	0.5751	Above the null
Net purchases minus sales	0.4740	Inverted
Non-routine purchase present	0.4330	Inverted
Distinct purchasing officers	0.4143	Inverted
Non-routine purchase count	0.4086	Inverted

The reading the register records is the legal blackout. A firm engaged in a control process restricts officer and director transactions in material nonpublic information, so the window closes once a deal is live and purchases stop. That is the opposite of accumulation ahead of an approach, and it aligns with the passive-trading result of Agrawal and Nasser (2012), in which target insiders withhold transactions rather than adding them. The inversion is the fifth such result the programme has recorded across unrelated units, and it is reported as measured rather than resolved by a sign flip.

Three verification passes attach to the confirmation. An independent agent recomputed every figure with its own code across three fold seeds, obtaining 0.6125, 0.6141 and 0.6163, reproduced all five per-feature values to four decimal places, confirmed that the pairing and the permutation both operate within firm, and found no leakage. Seventeen selftests passed on the panel builder and ten on the scorer, each paired with a control that must fail, and the panel builder was mutation-tested against its own suite, which found and killed one surviving mutant. The panel repair was load-bearing rather than cosmetic: re-running

admissibility with the first-filing rule omitted admits 148 truncated-window placebo moments, a panel share of 2.22% against a frozen void ceiling of 2.0%, so the run would have been void without it.

Two defects are disclosed with the confirmation. The frozen secondary bar, an incremental Brier score over the review-flag baseline, could not be computed from the delivered builds because the channel features cover only case and placebo moments while the baseline panel needs them on the control arm as well; it is reported as not computable rather than approximated. And the preregistration froze no lookback window for the insider features, so the builder declared 365 days before any outcome was read and the omission is recorded as a defect in the preregistration rather than as an author’s choice made afterwards.

4.1 What the feature measures, after correction

The confirmed feature was first described in the register as reading the firm’s own purchase cadence, and it cannot do so. The lookback and the censoring constant are both 365 days, so an absent purchase is assigned the value 365 and is coded identically to a

purchase made exactly one year earlier. The feature is right-censored at its own lookback and carries no information about the interval between a firm's purchases.

Recomputation on the delivered panel, holding the pairing fixed, separates the two things the feature conflates. The continuous feature scores 0.5704. A binary indicator for no non-routine purchase within 365 days scores 0.5624. The continuous feature restricted to uncensored pairs, where a purchase exists on both sides and the elapsed time is genuinely observed, scores 0.5444, which is below the 0.58 bar. Censoring runs at 57.62% overall, 63.46% on case moments and 53.96% on placebo moments, and 88.7% of the excess over 0.5 is carried by the binary indicator alone.

Correction of record. The measured object is the *absence* of a non-routine purchase in the trailing year, not the rhythm of purchases. The verdict stands, because the channel bar was cleared by the fitted channel score rather than by the single feature, and the direction is unaffected. The description was corrected in the register on 14 August 2026 and the follow-up unit had been designed around a quantity the data cannot express. Any revival must extend the lookback before the word cadence is used again.

This correction determines what section 6 can find, and the two results should be read together. A binary indicator has no dose-response to display, so the absence of a gradient reported below is the same fact observed on a different population.

5. The Transfer Test

The confirmed unit established that a signal exists. Its frozen consequence withheld any conditioning claim and recorded the insider channel as a candidate

conditioner requiring its own registration. That registration asks the question a desk actually needs answered: given that a review of strategic alternatives has been announced, does pre-announcement insider silence condition whether the review converts into a definitive process?

The design carries no control pool. Every subject has publicly announced a review, and the contrast is converters against non-converters inside that population, so the control-composition defect of section 3.2 cannot occur. The run measures that claim rather than asserting it: whether an issuer appears anywhere in the purchase corpus scores 0.4972 on conversion, so the operating-company distinction that sealed the predecessor is measurably absent here.

The population is the announcement-context eligible subset of the strategic review corpus, 2,853 episodes drawn from 6,273 harvested, reproducing exactly the count established by the sibling unit on announcement wording, with zero join failures and zero document-cache misses. There are 959 conversions, a base rate of 33.61%. Feature coverage is 100%, against a blocking floor of 80%, because the ownership corpus spans January 2006 to March 2026 and therefore covers every 365-day lookback in a panel running from 2010 to 2025.

Because the confirmed unit read a panel-restricted digest while this unit reads the universe purchase-only index, a reconciliation gate was placed ahead of every statistic. All 4,622 of 4,622 shared filer-and-cut pairs reproduce the earlier unit's four purchase-side feature values, with zero mismatches. The two sources are the same object where they overlap.

Table 3: The four frozen bars of the transfer test

BAR	FROZEN THRESHOLD	OBSERVED	RESULT
1. Primary discrimination	$AUC \geq 0.55$ and above permutation p95	0.5176, p95 0.5148, $p = 0.030$	Fail
2. Consequence gap	≥ 5.0 points, Wilson intervals disjoint	4.72 points, intervals overlap	Fail
3. Placebo in time	$AUC \leq 0.53$ at the cut less three years, margin ≥ 0.02	0.5041, margin 0.0135	Ceiling passes, margin fails
4. Coverage, blocking	$\geq 80\%$	100%, 2,853 of 2,853	Pass

The primary is directionally correct and statistically detectable. It exceeds the within-year permutation 95th percentile at $p = 0.030$ and its sign is the one named before the data were read, which is the inverted direction the confirmed unit measured. A discrimination of 0.5176 against a bar of 0.55 is

nonetheless inert, and the bar exists to refuse exactly that trade. The second bar misses by 0.28 points, which invites a reading as a near miss, and the overlapping Wilson intervals rule that reading out: the gap fails on both of its clauses rather than on a rounding.

Table 4: Disclosed secondaries, none promoted

SECONDARY	N	AUC
One episode per firm, earliest	2,285	0.5114
Conversions restricted to a lag of 30 days or more	2,853	0.5116
Excluding 2010, where the routine lookback is degraded	2,680	0.5207
Restricted to episodes from 2013, primary	2,359	0.5295
Restricted to episodes from 2013, placebo	2,359	0.5040
Issuer appears anywhere in the purchase corpus	2,853	0.4972

The 2013 restriction was frozen as an amendment before any outcome statistic and is reported for that reason. On the subset with a complete three-year routine lookback the placebo margin reaches 0.0255 and would clear the third bar, on a primary of 0.5295 that remains far from 0.55. Suppressing a favourable secondary is as much a distortion of the record as promoting one, and neither operation changes the verdict.

6. The Stratum Ordering Carries No Gradient

The consequence bar was frozen as a contrast between firms silent for a full year and firms whose officers purchased within the previous six months, with an intermediate stratum specified so that the shape of the response could be read. The three rates are given below.

Table 5: Conversion by length of pre-announcement insider silence

STRATUM	N	SHARE OF POPULATION	CONVERSION	WILSON 95%
Recent buyer, 182 days or fewer	461	16.2%	29.72%	[25.73%, 34.04%]
Middle, 183 to 364 days	467	16.4%	34.05%	[29.89%, 38.46%]
Silent, 365 days or more	1,925	67.5%	34.44%	[32.35%, 36.59%]

Ordered by increasing silence the rates rise, and the increments are 4.33 points and then 0.39 points. A graded signal would distribute the response across the ordering; this response is spent at the first boundary. The middle stratum is statistically indistinguishable

from the fully silent stratum and is not intermediate between the endpoints. Everything the measurement detects is that the 16.2% of firms whose officers made a non-routine purchase in the six months before announcing a review converted less often than the rest.




Figure 1. Conversion of an announced review of strategic alternatives into a definitive process, by length of pre-announcement insider silence. Filled points are observed rates with Wilson 95% intervals; stratum sizes are printed beneath the axis and sum to the full population of 2,853 episodes carrying 959 conversions. The amber dashed line joins the two end strata and the hollow diamond marks 32.08%, where a linear dose-response in days of silence would place the middle stratum; that value is the midpoint of the two quoted endpoints and is a construction rather than a measurement. The observed middle point sits at 34.05%, above the gradient line and level with the fully silent stratum, so the response saturates at the first boundary instead of grading across the ordering. The middle stratum’s interval nonetheless spans the gradient prediction and both endpoints, so the sample establishes the location of the point estimate and does not exclude a gradient. Data: the transfer test’s results file (August 2026).

The qualification the figure makes visible belongs in the text as well. The middle stratum holds 467 episodes and its Wilson interval runs from 29.89% to 38.46%, which contains the gradient prediction and both endpoints. The sample therefore establishes

where the point estimate sits and does not exclude a graded response. What survives without qualification is that the point estimate does not fall between the

endpoints, and that the difference which produces the primary statistic is concentrated entirely at the recent-buyer boundary.

Section 4.1 supplies the reason to expect that shape. On the confirmed population the excess discrimination over the null is 88.7% attributable to a binary indicator for the absence of a purchase in the trailing year, with the continuous elapsed-time feature falling below its bar once the analysis is restricted to pairs where the elapsed time is genuinely observed. A binary indicator admits no dose-response. The step observed here on the conversion population and the censoring decomposition observed there on the precursor population are the same fact in two settings.

7. Ruling Out a Permanent Firm Attribute

A signal measured at an announcement can be an artefact of which firms announce rather than of when they announce. The programme's standing gate against that failure mode scores the identical feature at a cut three years earlier and requires the displaced score to stay beneath a ceiling. The gate was established after a sibling unit found that a placebo arm at deal-free moments five, seven and nine years from each anchor recovered 56% of a case-control gap, which established that most of what looked like a precursor was a permanent firm attribute (Andreev and Howden, 2026g).

Here the displaced score is 0.5041 against a ceiling of 0.53, comfortably inside it. The margin clause, which requires the primary to exceed the placebo by 0.02, fails at 0.0135. That failure is the primary's own weakness propagating into a relative test rather than evidence of an attribute confound, and it is recorded as such. The 2013 restriction of Table 4 makes the point arithmetically: on a subset with a complete routine lookback the margin reaches 0.0255 while the primary remains at 0.5295.

A second control is cleaner than the placebo because it is not relative. Whether an issuer appears anywhere in the purchase corpus, which is close to an indicator for being an operating company with active officers, scores 0.4972 on conversion. The confound that sealed

the case-control unit of section 3.2 is measurably absent from this design rather than argued away by construction.

8. Selection as the Mechanism

The two contrasts differ in what has already been conditioned on. The confirmed unit compares a firm at a deal moment against the same firm at a deal-free moment, so the comparison spans the full range of the underlying process, from no process to a live one. The transfer test compares one announced review against another. Announcing a review of strategic alternatives is itself public evidence that a board process exists; the register's own ladder measures that such an announcement raises the eighteen-month probability of entering a definitive process from a control rate of 3.44% to 31.71%. Every subject in the transfer population has therefore already supplied the evidence that the confirmed feature reads, and the variance the feature exploits is spent before the comparison begins.

This is selection on a variable that lies between the feature and the outcome, in the sense of Heckman (1979), and it predicts attenuation of exactly the kind observed: a signal that discriminates at 0.5751 within firm falls to 0.5176 once the population is restricted to firms that have announced. Attenuation of that magnitude bounds the confirmed result rather than contradicting it. Insider silence marks that a control process exists. It carries almost no additional information about whether a process already known to exist will reach a definitive agreement.

The transfer test is the second independent conditioner to fail on this population, and the two failures differ in kind. The sibling unit on announcement wording found conversion flat across four preregistered language classes, at 36.38%, 34.44% and 32.62% on the three occupied classes, a top-to-bottom gap of 3.76 points at $z = 1.57$ against bars of 3.0 and 5 points. Wording carried nothing measurable at all. Insider behaviour carries a little and not enough. Two conditioners of different families failing on the same population, one at zero and one just above it, is evidence about the population rather than about either instrument.

Consequence, as frozen. Nothing is wired. No multiplier is applied, no card text changes and no threshold moves. The conversion rate following an announced review of strategic alternatives is recorded as flat with respect to insider behaviour and should be displayed as flat, rather than as a quantity awaiting the right conditioner.

9. What Is Established and What Is Not

The transfer test resolves effects no smaller than a discrimination of about 0.53. At 2,853 episodes with a 33.61% positive rate the Hanley-McNeil standard error near 0.5 is approximately 0.0115, which places the 0.55 bar roughly 4.3 standard errors out and gives a minimum detectable effect of about 0.53 at 80% power. The null therefore reads as the absence of an effect of at least that size and not as the absence of an effect. That figure was written into the preregistration before the run, alongside a stated probability of 0.25 that both substantive bars would pass.

Censoring does not bind. The longest observed conversion lag is 547 days and the last episode falls on 23 January 2025 against outcomes evaluated as of 12 August 2026, so every episode carries a complete observation window. That was asserted in the run rather than assumed.

Four further limits are recorded rather than implied. Sale-side features cannot be computed from the universe index and are forced to a missing value so that no line can read a spurious zero; the corresponding feature scored 0.4740 in the confirmed unit, below the null, so nothing measured is lost. The routine filter's incomplete history degrades 2010 windows, and excluding them moves the primary to 0.5207, which does not change the verdict. Whether insider silence conditions conversion on the narrower rung, in which a target discloses that it has received an acquisition proposal, is not resolved here; that population holds 123 scored episodes in the companion ladder and this panel is too coarse to isolate it. And the fundamentals and price-and-volume channels of Table 1 were never scored at all, so no statement about them, in either direction, is available from this work.

No result in this paper has been validated out of sample. The confirmed channel is an in-sample measurement on historical filings, the transfer test is an in-sample measurement on a historical population, and neither has a forward ledger behind it.

10. Instrument Discipline and Reproducibility

Each unit separates harvesting, selftesting and scoring into modes run in a fixed order. A harvest queries the Commission's endpoints and caches every response to disk; a selftest exercises the scoring code against planted fixtures and must pass before scoring is permitted; scoring reads only cached material. No network request is made during scoring, so a replicator who has completed a harvest can re-derive every statistic without returning to the Commission.

Selftest counts are recorded with their failures. The confirmed unit ran 17 checks on its panel builder and 10 on its scorer, each paired with a control that must fail, and mutation-tested the panel builder against its own suite. The transfer test ran 8, covering detection, rejection, a planted inverse effect that must be reported as inverted rather than scored as a pass, leakage, the placebo gate, the permutation machinery, the reconciliation comparator and the eligibility gate. Its first synthetic fixture was rejected before the run because it tied silence and announcement year to the same loop index, which made silence a deterministic function of the stratum and rendered the within-year permutation null degenerate; the null reproduced the observed statistic exactly. The defect was in the fixture, and its discovery is also a demonstration that the stratified control binds.

Calibration record. Each unit froze a single probability that its primary bar would pass, scored at settlement as a Brier contribution. The within-firm precursor stood at 0.35 against a pass, for 0.4225. The transfer test stood at 0.25 against a failure, for 0.0625. The sealed case-control unit stood at 0.35 against a measurement limit scored as a failure, for 0.1225. The clustered-purchase unit was entered into the register by a back-fill and carries no frozen probability, so it contributes nothing to the curve. The programme's intent is that after roughly twenty such units the calibration curve replaces argument about confidence with arithmetic.

Every figure reported in this paper is transcribed either from the results file a unit wrote at settlement or from the register entry recording its verdict. Where the two are read together the results file governs, since the verdict is a prose summary of it. The single exception is the 32.08% of Figure 1, which is the midpoint of two quoted endpoints and is identified as a construction at both occurrences.

11. Discussion

The programme set out to find insiders accumulating ahead of an approach and found the opposite. Under a design in which the firm is held fixed, purchase intensity runs below the null on every measure and the informative quantity is the absence of purchases. That result is consistent with a blackout mechanism, it is consistent with the passive-trading finding of Agrawal and Nasser (2012), and it required a within-firm contrast to see, because a case-control contrast on the same data confounds the signal with what kind of company files ownership forms at all.

The transfer test then supplies the boundary. A signal that separates a deal moment from a quiet moment at 0.5751 separates a converting review from a non-converting one at 0.5176, which is detectable and inert. The population explains the drop: everyone in it has announced. A measurement that establishes existence does not license a claim about conditioning, and the two questions require separate registrations even when they use the identical feature computed by the identical code. That is the design lesson we would

carry to any other precursor, and it is why the confirmed unit's frozen consequence withheld the conditioning claim rather than assuming it.

Reporting the shape of the residual effect matters more than reporting its size. A 4.72-point gap with overlapping intervals invites the reading that the signal is real and the sample merely small, and the stratum ordering removes that reading: the middle stratum sits with the silent stratum rather than between the endpoints, and the discrimination decomposes on the parent population into a binary absence indicator that has no gradient to offer. The honest limit, stated in section 6 and drawn in Figure 1, is that the middle stratum's interval is too wide to exclude a gradient outright. What the two populations establish together is that the quantity being measured is closer to a switch than to a dose.

Gates written before the data were seen refused a measurement twice in this family of four units. The prevalence gate sealed an entire case-control comparison whose control arm was composed wrongly, and the coverage gate sealed two of the four channels in the unit that confirmed. A gate is worth what it costs only if it fires on the analyst's own work, and in this family the two that fired removed one unit and half of another from the citable record. Nothing in this paper licenses a change to a live decision rule, and the conversion rate following an announced review is recorded as flat with respect to everything tested against it so far.

References

- [1] A. Agrawal & T. Nasser. Insider trading in takeover targets. *Journal of Corporate Finance*, 18(3):598–625, 2012.
- [2] L. Cohen, C. Malloy & L. Pomorski. Decoding inside information. *Journal of Finance*, 67(3):1009–1043, 2012. Source of the routine versus opportunistic transaction filter applied throughout.
- [3] J.A. Hanley & B.J. McNeil. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1):29–36, 1982. Source of the

standard error used in the power block of §9.

[4] C.R. Harvey, Y. Liu & H. Zhu. ...and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016. The $t > 3.0$ convention adopted as the programme’s default significance bar.

[5] J.J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–161, 1979.

[6] A.J. Keown & J.M. Pinkerton. Merger announcements and insider trading activity: an empirical investigation. *Journal of Finance*, 36(4):855–869, 1981.

[7] J. Lakonishok & I. Lee. Are insider trades informative? *Review of Financial Studies*, 14(1):79–111, 2001.

[8] W.K. Newey & K.D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.

[9] H.N. Seyhun. Insiders’ profits, costs of trading, and market efficiency. *Journal of Financial Economics*, 16(2):189–212, 1986.

[10] E.B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212, 1927. Source of every binomial interval reported here.

[11] M.A. Andreev & J. Howden. A measured ladder of pre-announcement and in-process deal probabilities, reconstructed entirely from free EDGAR filings. *Quark Quantitative Research Working Paper*, 2026f. Companion paper; source of the strategic-review and disclosed-proposal rungs quoted in §§8–9.

[12] M.A. Andreev & J. Howden. Distinguishing a firm attribute from a moment in event-prediction studies: a placebo-in-time design. *Quark Quantitative Research Working Paper*, 2026g. Companion methods paper; develops the within-firm identification argument of §4 and the placebo gate of §7.

[13] M.A. Andreev & J. Howden. The reachable boundary of free filing data for event-driven research: a catalogue of measurement limits with named defect classes. *Quark Quantitative Research Working Paper*, 2026. Companion paper; catalogues the control-composition and coverage boundaries of §§3.2 and 4.

[14] Preregistrations for the four units reported here, each frozen before its unit read any outcome statistic, held in the experiments directory of the platform repository under the register identifier of the unit it governs.

[15] Register of hypotheses and verdicts, maintained in the platform’s research engine, carrying for each unit its frozen hypothesis text, its terminal status, its verdict and the paths of the programs and results files behind it.

© 2026 Maksim Aleksandrovich Andreev & Joseph Howden.
All rights reserved.

Quark Quantitative Research Platform — Working Paper
Generated 2026-08-17 21:46 UTC

Every quantity reported here is recoverable at no cost from the Securities and Exchange Commission’s ownership datasets, full-text search index and per-filer submissions endpoint. One unit confirms and three do not; all results are in-sample and none is forward-proven. Correspondence: M.A. Andreev, Quark Quantitative Research.