Surveillance as Information and Constraint: Evidence from Exchange Price Bands

Anand Bharadwaj* Vinodh Madhavan†

October 2026

Abstract

Exchange surveillance conveys information and restricts trade, and the two functions can conflict. In a sequential trading model, disclosing how informed order flow is, which reveals nothing about fundamentals, can push the prices required for informative trades outside a legal range and leave the market strictly less informed after subsequent trading, a disclosure paradox. Using the complete history of security-level price bands on India's National Stock Exchange, we find that regulatory decisions identify trading states with little information about value, and that tight bands delay the incorporation of common information without reducing it permanently. Optimal surveillance separates communication from restriction.

Keywords: market surveillance; price limits; price discovery; financial regulation.
JEL codes: D82, G14, G18.

*Department of Economics, University of California, Berkeley. Email: anand_bharadwaj@berkeley.edu.
†Amrut Mody School of Management, Ahmedabad University. Email: vinodh.madhavan@ahduni.edu.in.

1 Introduction

When an exchange places a security under surveillance, its action has two consequences that are usually discussed separately. The intervention communicates a judgment by the market operator, which observes far more of the order flow than any participant, that trading in the security looks unusual. At the same time, it alters the rules of trade, and so determines which orders can execute and which private information can reach prices. This paper argues that the two consequences can work against each other, documents their joint footprint in one of the world's largest equity markets, and studies how surveillance should be designed once the interaction is taken into account.

What makes the interaction possible is the content of surveillance information. An exchange that flags a security conveys little about its value; it conveys how the orders arriving in that security ought to be interpreted. Suppose the message indicates that order flow is unusually well informed. Competitive liquidity suppliers then require greater compensation for adverse selection, and the prices at which they are willing to absorb orders move away from the current quote. Under a trading rule that caps daily price movements, those prices may lie outside the legal range, in which case the informative orders never execute and the market forgoes the trades from which it would have learned most. We refer to this outcome as the *disclosure paradox*: by disclosing more about the market, a regulator can leave the market knowing less about the asset.

The security-level price bands of the National Stock Exchange of India (NSE) allow both roles of surveillance to be observed in a single instrument. Each security is assigned to a daily band permitting movements of 2, 5, 10, or 20 percent around the previous close, and the exchange revises these assignments as part of its surveillance process, so that a band is simultaneously a public regulatory signal and a binding limit on prices. We assemble the complete history of 14,380 band revisions for 1,218 firms between 2020 and 2025 and reconstruct the legal price limits in force on more than one million security-days. Checked against the exchange's official limit-hit reports, the reconstruction agrees on about 98 percent of observations.

Our first set of results concerns the information carried by the exchange's decisions. Adding the realized action to a recursive out-of-sample forecast more than doubles the twenty-session return R^2, and among securities that start from the same band on the same day, tightening is followed by returns about twelve log percentage points lower than loosening, together with a collapse in turnover and much quieter prices. If these decisions revealed what securities are worth, the repricing ought to be largest when tightening came as a surprise to the market. The data show the opposite pattern, with the decline concentrated among

tightenings that public conditions had predicted. The exchange's information therefore appears to concern the trading in a security, which is the kind of information on which the disclosure paradox turns.

A second set of results concerns the band as a constraint. For each security and day, we compare the movement in its industry peers with the price movement that its band still permits in the same direction. In our sharpest design, every security in an industry-date cell is assigned the same peer return, drawn from a separate group of firms, so that securities in a cell differ only in their legal room. Those with little room move 1.23 log percentage points less in the direction of the news on the day it arrives and 1.29 points more over the following ten sessions, leaving the cumulative response essentially unchanged. The delay is accompanied by thinner trading and by bunching at the boundary toward which the news points, and roughly half of it exceeds what mechanical truncation of prices would produce.

We formalize the disclosure paradox in a sequential trading model in the tradition of Glosten and Milgrom (1985), in which the exchange observes how likely incoming orders are to be informed. Because this state carries no information about the payoff, disclosing it leaves beliefs about fundamentals unchanged on impact. Our main theorem nonetheless shows that, for a fixed trading episode under a fixed band, the public record generated when the exchange withholds the information strictly Blackwell dominates the record generated when it discloses, so that every decision maker who cares about the payoff prefers the market without disclosure. Since private information is held fixed throughout, the loss arises entirely through execution. Disclosure has an offsetting benefit for traders without private information, whose spreads narrow in securities the regulator clears even as they are excluded from securities it flags; surveillance transparency thus trades liquidity against learning. The same mechanism implies that recorded prices become quieter as they become less informative, which should give pause to any evaluation of surveillance based on its effect on volatility. The result extends to markets in which unexecuted orders are imperfectly visible, and daily recentering of the band, combined with continuing public information, generates the delayed but complete adjustment observed in the data.

Three implications for regulatory design follow. Under a fixed band, the optimal disclosure policy is partial, pooling surveillance states just enough to keep informative trading feasible; this policy Blackwell dominates every alternative while preserving trade for every liquidity trader. When the band is the regulator's only means of communication, adding a category can raise welfare even though its limits never bind, since its value lies entirely in the information conveyed by its label, and an empirical test of whether a limit binds may therefore miss its effects. Finally, a separate disclosure instrument allows the exchange to communicate without relaxing its trading rule merely to send a signal. To give restrictions a benefit, we study a

financing environment in which manipulative trades drain internal funds from constrained firms. A tight band blocks such trades at the cost of blocking informative ones as well, and the value of intervening peaks before manipulation is most frequent.

These results contribute to three literatures. In the work of Grossman and Stiglitz (1980), Paul (1993), and Goldstein and Yang (2015, 2019), public information affects what markets learn by changing incentives to acquire or trade on private information. In our setting, incentives and private information are held fixed, and disclosure affects learning solely by changing which trades are legally feasible; this channel is enough to make the market strictly less informed. Moreover, because the information disclosed concerns the market and not the asset, the result does not depend on public news about fundamentals displacing private research. We also contribute to the study of price limits and price discovery (Kim and Rhee, 1997; Hou and Moskowitz, 2005), introducing a limit that is itself a regulatory signal, measuring the timing of adjustment with a common shock assigned identically within cells, and quantifying the share of the resulting delay that reflects pure truncation. Finally, the paper connects market microstructure to the design of regulatory information (Kamenica and Gentzkow, 2011; Bond and Goldstein, 2015; Goldstein and Leitner, 2018) in a setting where the regulatory category is both a message and a constraint, and its protection motive links the settlement-manipulation incentive of Kumar and Seppi (1992) to the financing friction of Holmström and Tirole (1997).

Section 2 describes the NSE's surveillance system and our data. Sections 3 and 4 present the evidence on the band as a message and as a constraint, Section 5 develops the model, and Section 6 studies regulatory design. Section 7 concludes.

2 Institutional Setting and Data

Each price band in the NSE's surveillance system is both a public statement of the exchange's assessment of trading in a security and a legal limit on the prices at which that security can trade. This section describes the institutional features that make the two roles observable and the data we assemble to measure them.

2.1 Price bands as surveillance

For the securities we study, the NSE permits daily price movements of 2, 5, 10, or 20 percent above or below a reference closing price, rounded to the applicable tick (National Stock Exchange of India, 2026a). These are security-level limits, distinct from the market-wide circuit breakers that halt trading across the exchange and from the dynamic price ranges

applied to some other securities. Each security is assigned to exactly one band at any time, and the assignment is public.

Revisions to these assignments form part of the exchange's surveillance process, and the procedures differ by direction in an informative way. Downward revisions are considered daily, whereas the exchange describes upward revisions as subject to "bi-monthly" review and objective eligibility criteria (National Stock Exchange of India, 2026b). The exchange can therefore narrow a band quickly when a security attracts scrutiny, while widening follows a more deliberate, rule-based process. A band history is thus a record of the exchange's evolving judgment about trading in each security.

Alongside the bands, the exchange operates three other surveillance instruments. The Additional Surveillance Measure (ASM) targets securities with abnormal price or trading patterns, imposing higher margins and, at more severe stages, narrower bands and delivery-based settlement (National Stock Exchange of India, 2026f). The Graded Surveillance Measure (GSM) targets securities whose prices appear disconnected from their fundamentals, applying escalating restrictions that culminate in limits on trading frequency (National Stock Exchange of India, 2026c). Trade-for-Trade settlement requires every transaction to be settled by delivery, eliminating intraday netting. A band revision may accompany any of these measures. We reconstruct all three programs by security and effective date, combining membership annexures, stage-change circulars, and daily regulatory-indicator reports; membership is known for 99.6 percent of eligible security-days. Section 3.3 measures their role directly.

The NSE is an order-driven market in which orders are matched by price and time priority (National Stock Exchange of India, 2026e). A transaction requires both a legal price and an order on the opposite side at that price. The market feed also disseminates outstanding buy and sell quantities and depth at the best price levels (National Stock Exchange of India, 2024), so unexecuted demand can itself convey information. The model in Section 5 incorporates both features, representing competitive liquidity supply through break-even quotes and treating the information in unexecuted orders explicitly in Section 5.5.

2.2 Band histories and revision timing

We link the exchange's archived band histories (National Stock Exchange of India, 2026d) to daily price histories through issuer identifiers and dated symbols. A revision is a change between distinct members of the fixed-band menu. The daily comparison panel includes every issuer-day with the required price and band records, whether or not the issuer has a revision in the event sample.

Measuring the market's response requires precise timing, which exchange circulars allow us to establish. In each case we can match directly, the date recorded in the band history is the date on which the exchange published the revision, and the new band takes effect in the next equity session (National Stock Exchange of India, 2023a, 2024a,b). Appendix Table 12 reports these cases. We therefore assign the old band to the publication session, apply the new band from the following session, and measure subsequent returns from the publication-date close. The reconstructed program histories use explicit effective dates where circulars state them and next-session effectiveness otherwise.

2.3 Sample and returns

Our daily price archive contains 1,262,554 issuer-days for 1,247 issuers from January 2019 through March 2026, with observations outside 2020–2025 supplying preceding histories and forward outcomes. Eligibility is determined solely by coverage in the linked price and band histories; we impose no screen on size or accounting characteristics. The price histories provide closing prices, traded quantities, and shares outstanding, from which we construct market capitalization and share turnover. We identify 215 splits, bonus issues, and consolidations from large offsetting changes in shares outstanding and prices and adjust prices accordingly (Appendix B). Returns exclude dividends.

The revision panel contains 14,380 issuer-date events for 1,218 firms between January 2020 and December 2025: 6,451 tightenings and 7,929 loosenings. We exclude windows around the special operational sessions of March and May 2024. Table 1 reports the transitions. The 5–10 percent pair is the most common, followed by 2–5 percent, and both directions occur frequently within every observed pair. The revision history therefore offers substantial variation in the direction of revision from a common starting band, which Section 3 exploits.

Table 1: Revisions between price-band regimes

Band pair (percent)	Tightening	Loosening	Total
2-5	1,452	2,043	3,495
2-10	0	0	0
2-20	0	0	0
5-10	2,641	3,044	5,685
5-20	853	1,081	1,934
10-20	1,505	1,761	3,266
Total	6,451	7,929	14,380

Notes. Distinct issuer-date revisions in the estimation panel, 2020–2025. A tightening reduces the permitted percentage movement. The panel contains 14,380 revisions for 1,218 firms; no direct 2–10 or 2–20 revisions occur.

Let p_{it}^{A} denote the adjusted log price of security i on equity session t. The market-relative price change between sessions $t+s$ and $t+h$ is

$$R_{i,t}(s,h) = p_{i,t+h}^{A} - p_{i,t+s}^{A} - \sum_{j=s+1}^{h} r_{t+j}^{M}, \tag{1}$$

where r_t^M is the lagged-capitalization-weighted mean adjusted log return across the available universe, cumulated on the equity calendar. Return coefficients are multiplied by 100 and reported in log percentage points.

2.4 Legal price limits

To study price adjustment we need the legal limits in force on each security-day. Daily open, high, low, close, previous close, and volume come from the NSE's Bhavcopy records, which cover all 1,492 equity sessions from 2020 through 2025. Dated symbols and historical ISINs match 1,073,942 of 1,078,071 eligible security-days, a coverage rate of 99.6 percent (Appendix Table 13; Appendix B.2).

We reconstruct each day's upper and lower limits from the prevailing band, the exchange's reference close, and the tick-size convention in force on that date, including the tick revisions of 2024 and 2025 (National Stock Exchange of India, 2024c, 2025). These limits determine how much price movement remains available to a security in each direction and identify closes at the economically relevant boundary. On dates for which the exchange publishes official limit-hit reports, the reconstruction agrees with those reports for about 98 percent of observations (Appendix Table 14), yielding a security-day record of how far the law allows each price to move.

3 Regulatory Decisions and Subsequent Trading

When the NSE revises a security's band, it acts on information that only the exchange fully observes: the complete order flow, client-level concentration, broker alerts, and its own surveillance systems. This section examines what the market can learn from those decisions. The exchange's actions are strongly related to observable conditions yet carry substantial predictive content beyond them; that content behaves unlike news about value; and trading changes sharply after each decision.

3.1 The predictive content of regulatory decisions

We begin by asking how much of the exchange's decision can be anticipated. Conditional on a recorded revision, we forecast whether the NSE tightens or loosens the band. The predictors are market-relative returns over the preceding 1, 5, 10, 20, 60, and 125 sessions, twenty-session volatility and turnover, log market capitalization, recent upper- and lower-boundary activity, and average traded value; the full specification adds the old band and calendar month. We estimate a regularized logit recursively, forecasting each test year using only earlier years.

Panel A of Table 2 shows that the exchange's decisions are highly predictable. Part of this reflects the menu itself, since a security at 2 percent cannot be tightened and one at 20 percent cannot be loosened. Predictability remains strong without those mechanical features: using market variables alone, annual out-of-sample areas under the ROC curve range from 0.867 to 0.930, and they remain between 0.853 and 0.916 for the interior 5 and 10 percent bands, where either direction is feasible. The exchange evidently responds in a systematic way to the conditions it observes.

Whether the realized decision tells investors anything beyond those conditions is a separate question. For horizons of 1, 5, 10, and 20 sessions, we compare a ridge forecast of subsequent market-relative returns built from the full measured public state with an otherwise identical forecast that also observes whether the exchange tightened. Training uses earlier years only and excludes observations whose outcomes would not yet be known at the start of the test year. Panel B shows that the regulatory decision improves prediction at every horizon, and the gains are large. At twenty sessions, out-of-sample R^2 more than doubles, from 0.022 to 0.056. A flexible Extra Trees model, which can capture nonlinearities and interactions in the public state, shows a similar gain in Panel C, from 0.063 to 0.093. The exchange's decision identifies return-relevant states that even a flexible model of observable conditions does not capture.

3.2 Tightening and loosening from a common starting point

To compare decisions that begin from the same regulatory state, we restrict attention to securities starting under a 5 or 10 percent band, where both tightening and loosening are feasible. Let T_e indicate a tightening at event e. We estimate

$$R_e(0,20) = \alpha_i + \delta_{d(e)} + \xi_{b_0(e)} + \beta T_e + g(X_e) + u_e, \tag{2}$$

where $d(e)$ is the exact revision date, $b_0(e)$ the initial band, and X_e the measured public state: preceding returns over 1, 5, 10, 20, 60, and 125 sessions, volatility, turnover, market

Table 2: Public information, revision direction, and return forecasts

Panel A: Out-of-sample prediction of revision direction				
Year	Full public state	Market variables	Interior bands	Brier, full
2021	0.986	0.930	0.916	0.048
2022	0.978	0.879	0.859	0.057
2023	0.841	0.867	0.853	0.187
2024	0.940	0.896	0.890	0.098
2025	0.966	0.896	0.888	0.068
Panel B: Out-of-sample prediction of subsequent market-relative returns				
Horizon	Public R^2	Public + action R^2	Increment	N
1	0.026	0.051	+0.025	12,583
5	0.042	0.082	+0.040	12,489
10	0.035	0.075	+0.040	12,457
20	0.022	0.056	+0.034	12,377
Panel C: Flexible twenty-session return forecast				
Model	Public R^2	Public + action R^2	Increment	N
Extra Trees	0.063	0.093	+0.030	12,377

Notes. Panel A reports recursive out-of-sample classification of tightening versus loosening, conditional on a recorded revision. The full public state includes preceding market variables, the old band, and calendar month. The market-variables specification excludes the latter two predictors. The interior-band specification also restricts the old band to 5 or 10 percent. Panel B reports ridge out-of-sample R^2 for subsequent market-relative returns using the full measured public state, including the old band and calendar month, with and without the realized action. Panel C reports the corresponding twenty-session comparison for an Extra Trees forecast with fixed hyperparameters. Return-forecast training observations are purged when their outcomes would not yet be observed at the start of the test year.

capitalization, recent boundary activity, and traded value, entered linearly after standardization. Firm effects absorb persistent differences across securities, and exact-date effects compare revisions made on the same day under identical aggregate conditions. Standard errors are clustered by firm and exact revision date.[1]

Table 3 reports the results. The pooled tightening-minus-loosening difference is −11.91 log percentage points over twenty sessions, significant at the 1 percent level. The difference is larger for securities starting at 5 percent (−16.13) than for those starting at 10 percent (−9.08), and largest of all, −17.72, when we compare only adjacent transitions from 5 percent: securities moved down to 2 percent against those moved up to 10 percent on the same day. The broader comparison across all revisions, using unordered band-pair effects, gives −10.46. In every specification, two securities that start from the same band on the same day, with similar recent histories, follow sharply different paths depending on which way the exchange moves.

[1]In the empirical tables, ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Table 3: Revision returns from a common initial regulatory state

Sample	Tightening − loosening	S.E.	N
All revisions, unordered-pair effects	-10.46***	(1.04)	13,995
Initial band 5 or 10 percent	-11.91***	(1.25)	9,749
Initial band 5 percent	-16.13***	(2.19)	5,454
Initial band 10 percent	-9.08***	(3.29)	4,295
Initial 5 percent, adjacent 2 vs. 10	-17.72***	(2.90)	4,393

Notes. Twenty-session market-relative return differences in log percentage points. The first row uses firm, exact revision-date, and unordered-pair effects. The remaining rows compare tightening and loosening from the same initial band; the pooled interior-band row also absorbs the initial-band category. All specifications control for the measured preceding public state: returns over 1, 5, 10, 20, 60, and 125 sessions, volatility, turnover, size, recent upper- and lower-boundary activity, and traded value. Standard errors are clustered by firm and exact revision date.

3.3 Band revisions and the broader surveillance response

Band revisions often form part of a wider surveillance response, raising the question of whether the revision simply proxies for the exchange's other measures. Immediately before a revision, ASM, GSM, and Trade-for-Trade status is known for 14,241 of our 14,380 events, and at least one program is active for 6,650. Within twenty sessions on either side, an entry, exit, or stage change occurs for 75.4 percent of tightenings and 55.7 percent of loosenings (Table 4).

Conditioning on current program membership and on program changes on the revision date leaves the return differences essentially unchanged. The pooled difference is −12.21 log percentage points, with −16.85 for securities starting at 5 percent, −8.93 for those starting at 10 percent, and −17.21 for the adjacent comparison. The predictive content of band revisions is not subsumed by the exchange's other surveillance programs.

3.4 The path of prices around regulatory decisions

Figure 1 traces the path of prices around revisions for a fixed cohort of 7,824 events with complete returns from twenty sessions before to sixty after, and it shows the conditions under which the exchange acts. Securities that are subsequently tightened have experienced a pronounced relative run-up in the preceding weeks, indicating that surveillance responds to rapid price appreciation. After the decision the path reverses, with most of the repricing concentrated in the first two weeks.

Over this horizon the revised band is persistent. Among windows fully covered by the band history, 93.5 percent of tightenings and 67.9 percent of loosenings see no further revision within twenty sessions (Appendix Table 11). The twenty-session comparison therefore

Table 4: Band revisions and other surveillance measures

Panel A. Surveillance around band revisions			
	Tightening	Loosening	All
Revision events	6,451	7,929	14,380
Known status before revision	6,393	7,848	14,241
A program active before revision	3,099	3,551	6,650
Program change on recorded date	530	33	563
Program change within ±20 sessions	4,867	4,413	9,280
Panel B. Twenty-session return differences with surveillance controls			
Initial band		Estimate	*N*
5 or 10 percent		-12.21*** (1.33)	9,690
5 percent		-16.85*** (2.33)	5,418
10 percent		-8.93*** (3.29)	4,272
$5 \to 2$ versus $5 \to 10$		-17.21*** (3.51)	4,367

Notes. Panel A uses all 14,380 revisions. Program changes comprise entries, exits, and recorded stage or series changes in ASM, GSM, or Trade-for-Trade. Panel B reports twenty-session tightening-minus-loosening coefficients in log percentage points, with firm and exact date effects, preceding public-state controls, and clustering by firm and date. The pooled row also includes initial-band effects. The specifications add current program membership and revision-date changes.

measures adjustment under a stable new regime. Over longer horizons further revisions become common, and part of the early difference unwinds as the regulatory and market state continues to evolve: on a common sample of complete sixty-session paths, the difference is −13.90 log points through session 20 and −7.34 through session 60.

3.5 Anticipated and unanticipated decisions

Were the return difference after a revision a response to private information about value revealed by the exchange, it would behave like a reaction to news. Public prices would already reflect whatever investors could anticipate, and the repricing at a decision would be proportional to its surprise, largest for tightenings that public conditions did not predict and smallest for those that were widely expected. We test this directly using the recursive out-of-sample probability of tightening, $\hat{p}_e$, from the market-variables logit in Panel A of Table 2. Because the logit trains only on earlier years, $\hat{p}_e$ is available for revisions from 2021 onward.

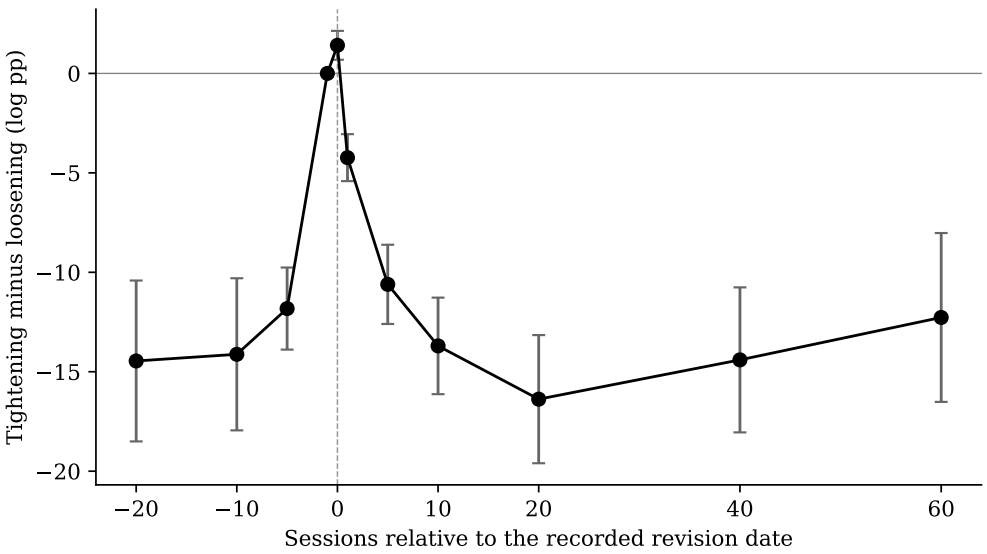


Figure 1: Price paths around regulatory decisions

Notes. Tightening-minus-loosening differences in cumulative market-relative log price changes relative to session −1, multiplied by 100. All plotted dates use the same 7,824 revisions from an initial 5 or 10 percent band with complete observed returns over the displayed window. Regressions absorb firm, exact revision-date, and initial-band effects. Predetermined return controls are omitted because the pre-event path is an outcome in this display. Intervals are pointwise 95 percent using firm/date clustering. The vertical line marks the recorded revision date.

The evidence in Table 5 runs against this interpretation. Panel A forms terciles of $\hat{p}_e$ in the pooled interior-band sample and reports returns separately by direction of revision. In the lowest tercile the tightening-minus-loosening gap is small, and it widens as tightening becomes more predictable, with loosening returns rising across terciles and tightening returns turning negative. Estimating the same relation continuously in the pooled specification of (2), Panel B finds that the gap widens by 5.8 log percentage points per unit of $\hat{p}_e$ over five sessions, significant at the 1 percent level and essentially unchanged when the ASM, GSM, and Trade-for-Trade controls are added. Under a news interpretation, unexpected tightenings would instead produce the largest declines.

On this evidence, the exchange's decisions identify securities in a distinct state and convey little about what they are worth. The states it flags, typically after sharp run-ups and heavy boundary activity, are those whose prices subsequently fall, and the fall is larger the more visibly a security is in that state. We accordingly read the predictive content in Table 2 as a summary of the exchange's assessment of trading in a security, which is the kind of information that the model in Section 5 assigns to the regulator.

Table 5: Repricing after anticipated and unanticipated decisions

A. Five-session returns by predicted probability of tightening

	Tightenings		Loosenings	
Tercile of $\hat{p}_e$	Mean $\hat{p}_e$	Return	Mean $\hat{p}_e$	Return
Low	0.10	0.89 (0.46)	0.11	1.24 (0.18)
Middle	0.36	-0.39 (0.24)	0.31	1.54 (0.22)
High	0.80	-0.71 (0.18)	0.70	1.89 (0.47)
Revisions		3,452		5,076

B. Pooled regressions on the predicted probability

	One session	Five sessions	Twenty sessions
$\hat{p}_e$ (loosening slope)	0.61 (0.79)	2.91* (1.60)	3.39 (2.90)
$T_e \times \hat{p}_e$	-1.44* (0.79)	-5.74*** (1.63)	-4.81* (2.80)
Implied tightening slope	-0.83 (0.72)	-2.83 (1.78)	-1.42 (3.17)
Observations	8,581	8,528	8,476

Notes. Revisions from an initial 5 or 10 percent band, 2021–2025. $\hat{p}_e$ is the recursive out-of-sample probability of tightening from the market-variables logit of Table 2, trained only on earlier years. Panel A reports mean five-session market-relative returns, in log percentage points, after forming terciles from the pooled distribution of $\hat{p}_e$ and then splitting by revision direction, after partialling out firm, exact-date, and initial-band effects. Panel B estimates (2) with $\hat{p}_e$ and its interaction with tightening added to the full public-state controls; coefficients are in log percentage points per unit of predicted probability. Standard errors in parentheses are standard errors of the mean in Panel A and are clustered by firm and exact revision date in Panel B.

3.6 Trading after regulatory decisions

Alongside the repricing, tightening changes how securities trade. For each revision, let $e_{i,t+j}$ be the capital-adjusted daily log return less the market benchmark, and measure subsequent variation by

$$\widehat{\sigma}_{e,H} = \left(\frac{252}{H} \sum_{j=1}^{H} e_{i,t+j}^2 \right)^{1/2}, \tag{3}$$

the daily root mean square on an annualized scale. Estimating (2) with this outcome, Table 6 shows that at twenty sessions volatility is 36.04 annualized percentage points lower after tightening than after loosening, equivalent to 2.27 percentage points of daily variation. Median annualized volatility is 42.7 percent after tightening and 57.1 percent after loosening, so the difference extends well beyond a few extreme cases. The contraction is largest immediately after the decision and fades over the following two months (Figure 2). It is robust to

alternative return definitions, tail treatments, comparisons restricted to same-day revisions, exchange-reported closes, and controls for other surveillance programs (Appendix Table 15). It is negative in every two-year subperiod and remains between −34.3 and −37.6 when any single calendar year is omitted.

Table 6: Returns and daily price variation after revisions

Outcome window	Return (log percentage points)	Volatility (annualized points)	Observations
Sessions 1–5	-10.49 (0.83)	-50.66 (2.56)	9,589
Sessions 1–10	-12.32 (1.09)	-43.22 (2.12)	9,364
Sessions 1–20	-12.79 (1.46)	-36.04 (1.66)	9,058
Sessions 1–60	-7.34 (2.06)	-17.64 (1.01)	8,409

Notes. Tightening-minus-loosening coefficients for revisions starting from 5 or 10 percent. Returns are cumulative market-relative adjusted log price changes multiplied by 100. Volatility is (3) multiplied by 100. Each row uses the same complete paths for both outcomes, with firm, exact date, and initial-band effects and the linear preceding public state in (2). Standard errors in parentheses are clustered by firm and exact date. The twenty-session row uses 9,058 events with complete return and volatility paths; Table 3 uses all 9,749 eligible twenty-session returns, while the −13.90 comparison in Section 3.4 uses the 8,409-event cohort observed through session 60. Annualization is a scale convention for daily variation.

Trading activity falls together with price variation. Table 7 shows that average daily share turnover falls by 1.75 log points after tightening relative to loosening, and traded value falls similarly. Sessions in which a security trades while its high, low, and close are identical become 19.0 percentage points more common, and the mean intraday range falls by 4.0 log percentage points. These effects strengthen when we compare tightenings and loosenings from the same band on the same day, and they are strongest for the adjacent comparison of moves from 5 to 2 percent against moves from 5 to 10 percent, where stale-price sessions become 33.3 percentage points more common.

These changes are concentrated where the new band should bind most tightly. For each security we measure how restrictive a prospective tightening would be by dividing its trailing twenty-session volatility by the width of the band it would move to, standardized within its initial band. A one-standard-deviation increase in this predetermined measure raises the post-tightening increase in stale-price sessions by 5.0 percentage points, deepens the contraction in intraday range by 0.53 log percentage points, and deepens the fall in volatility by 2.36 annualized points, each statistically significant. The strongest effects appear in the outcomes closest to execution.

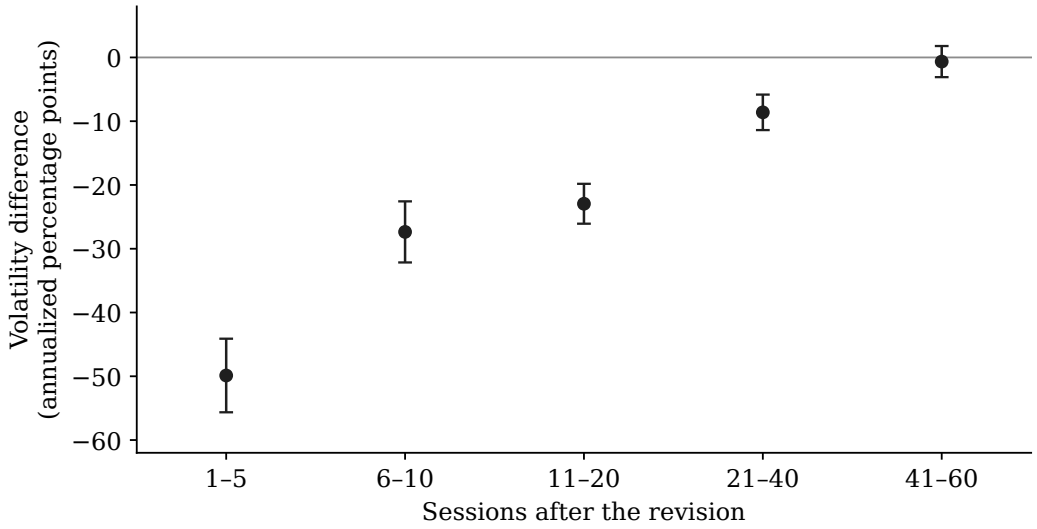


Figure 2: Volatility differences across subsequent sessions

Notes. Tightening-minus-loosening coefficients and 95 percent confidence intervals. Each outcome is the annualized root mean square of market-relative daily returns within the indicated disjoint window. All five estimates use the same 8,409 revisions starting from a 5 or 10 percent band, with complete returns through session 60. Controls and fixed effects match Table 6; standard errors are clustered by firm and exact date.

Table 7: Trading activity and stale prices after revisions

Sample	Log turnover	$H = L = C$ days (p.p.)	Intraday range (log p.p.)	Zero-return days (p.p.)
Common initial bands	-1.747 (0.087)	18.97 (1.63)	-4.02 (0.18)	2.07 (0.50)
Same initial band and date	-1.881 (0.124)	24.79 (2.45)	-4.73 (0.26)	2.93 (0.77)
5%→2% vs. 5%→10%	-2.149 (0.110)	33.30 (2.88)	-5.04 (0.21)	3.80 (1.14)

Notes. Twenty-session tightening-minus-loosening coefficients. Log turnover is the logarithm of mean daily traded quantity relative to shares outstanding. $H = L = C$ days are sessions with positive volume and no intraday price range. Intraday range is mean $\log(H/L)$, multiplied by 100; zero-return days are shares of sessions, also multiplied by 100. The first row uses revisions starting from 5 or 10 percent with firm, exact date, and initial-band effects and the linear preceding public state. The second absorbs initial-band-by-date effects and retains cells containing both directions. The third restricts the 5 percent starting state to adjacent revisions. Standard errors in parentheses are clustered by firm and exact date.

In sum, regulatory decisions identify securities in sharply different states, after which the market trades far less and records much quieter prices. Because a narrower band mechanically limits how far prices can move on any day, with a 2 percent band capping the annualized root mean square of raw daily returns near 32 percent, lower measured volatility says little by itself about what the market learns. To isolate how a legal range shapes the incorporation of information, we turn next to bands already in force, where the same common news reaches securities with different amounts of legal room.

4 Legal Price Limits and the Timing of Price Discovery

To isolate the effect of a legal range on the incorporation of information, we now turn from revisions to bands already in force. The approach is to observe how securities with different amounts of legal room respond to the same common news. Since the regulatory state is held fixed, the comparison does not depend on why the exchange assigned any security to its band.

4.1 Common news and available price capacity

Returns on industry peers provide a common information shock. For security i, let s_{it} be the equally weighted mean adjusted log return of at least five other securities in its industry on date t, excluding the focal security. Industry labels are the `smIndustry` classifications reported with the NSE's corporate announcements, held fixed throughout the sample. Retaining industries with at least six classified firms gives 209 firms in 19 industries.

The sign of the peer return identifies which side of the focal security's legal range becomes relevant, and its magnitude determines how much pressure the news places on that range. Let P_{it}^{-} denote the exchange reference close and U_{it} and L_{it} the reconstructed upper and lower limits (Appendix B.3). The legal capacity in the direction of the peer movement and the corresponding pressure are

$$c_{it} = \begin{cases} U_{it}/P_{it}^{-} - 1, & s_{it} \geq 0, \\ 1 - L_{it}/P_{it}^{-}, & s_{it} < 0, \end{cases} \qquad q_{it} = \frac{|\exp(s_{it}) - 1|}{c_{it}}. \tag{4}$$

A pressure of one means that the industry has moved exactly as far as the focal security is legally permitted to move in the same direction. If automobile stocks rise 3 percent and a particular automobile stock can rise at most 2 percent, its pressure is 1.5. The same industry news therefore places high pressure on a security with little legal room and low pressure on one with ample room. The measure scales the news by the movement the law allows and imposes no one-for-one relation between focal and peer returns.

Requiring positive capacity, predetermined controls, and complete returns through ten sessions yields 248,832 security-days for 209 firms, including 3,659 observations with $q_{it} \geq 1$. For horizons $h \in \{0, 1, 2, 3, 5, 10\}$, let $R_{i,t:t+h} = \sum_{j=0}^{h} r_{i,t+j}$ and estimate

$$R_{i,t:t+h} = \alpha_i + \delta_t + \xi_{b_{it}} + \beta_h s_{it} + \gamma_h s_{it} D_{it} + \lambda_h D_{it} + g_h(s_{it}) + X'_{i,t-1}\theta_h + \varepsilon_{i,t,h}, \tag{5}$$

where $D_{it} = \mathbf{1}\{q_{it} \geq 1\}$, $\xi_{b_{it}}$ absorbs the band category, g_h contains quadratic and cubic terms in the peer return, and the predetermined controls are preceding 1-, 5-, and 20-session returns, trailing volatility and turnover, and lagged log capitalization. Standard errors are clustered by security and exact date (Appendix B.7).

The timing of the response is shown in Figure 3. On the day the news arrives, high-pressure securities load substantially less on their peers' return: the interaction γ_0 is -0.473, highly significant. The gap then closes steadily as subsequent returns accumulate, reaching -0.071 by session five and 0.074 by session ten. Over sessions one through ten alone, the interaction is $+0.548$. The same pattern holds separately for positive and negative peer returns (Table 8, Panel A), so a tight legal range delays a security's response to industry news without permanently reducing it.

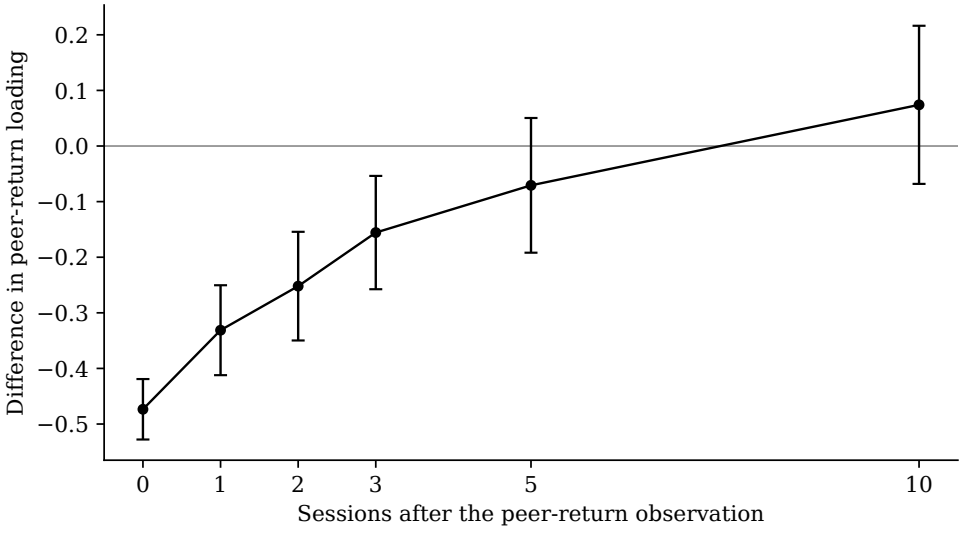


Figure 3: Timing of the response to industry news

Notes. The interaction γ_h in (5), with pointwise 95 percent confidence intervals based on security and exact-date clustering. All horizons use 248,832 observations, including 3,659 high-pressure observations. Returns are cumulative from the initial session through the indicated horizon. The coefficients measure differences in conditional loadings per unit of the peer log return.

4.2 Identical news, different legal room

Because each leave-one-out average excludes a different focal firm, the peer return varies slightly across firms in the same industry. To hold the news itself fixed, we divide each industry into two fixed groups by the ordering of security codes and assign every firm the mean return of the opposite group, requiring at least three source firms. Within an industry-date-group cell, every security receives precisely the same peer return, and none contributes to it. Writing $s_{g,-f,t}$ for the return of the group opposite focal group f in industry g, we estimate

$$\text{sign}(s_{g,-f,t})R_{i,t:t+h} = \alpha_i + \delta_{gft} + \xi_{b_{it}} + \tau_h D_{it} + X'_{i,t-1}\psi_h + u_{i,t,h}. \tag{6}$$

The fixed effect δ_{gft} absorbs the common news entirely, including its sign and magnitude. Multiplying returns by the sign of the news places upward and downward movements on a common scale. The only remaining difference among securities within a cell is how much legal room each has.

Panel B of Table 8 and Figure 4 report the results. Facing identical news, high-pressure securities move 1.23 log percentage points less in the direction of the news on the day it arrives and 1.29 points more over the following ten sessions. The two effects offset almost exactly, leaving a cumulative difference through session ten of 0.06 log percentage points with a standard error of 0.32, so that legal limits alter the timing of the response to common information while leaving its eventual size unchanged. The result holds when we require at least five source firms per group and when we control for prior ASM, GSM, and Trade-for-Trade membership.

4.3 Pressure, trading, and price location

As pressure rises, the delay grows steadily. Figure 5 estimates peer-return loadings within four fixed pressure intervals. Same-day loadings decline steadily as pressure rises and turn negative in the highest interval, while loadings through session ten remain positive and stable across intervals. The securities with the weakest initial response are those that catch up most afterward.

Pressure is also associated with changes in how securities trade. Table 9 relates continuous pressure to trading outcomes among 243,251 observations with positive volume and prices consistent with the reconstructed limits. A one-unit increase in pressure raises the probability of closing at the boundary in the direction of the news by 19.4 percentage points, more than tripling its sample frequency, and lowers the probability of closing at the opposite boundary

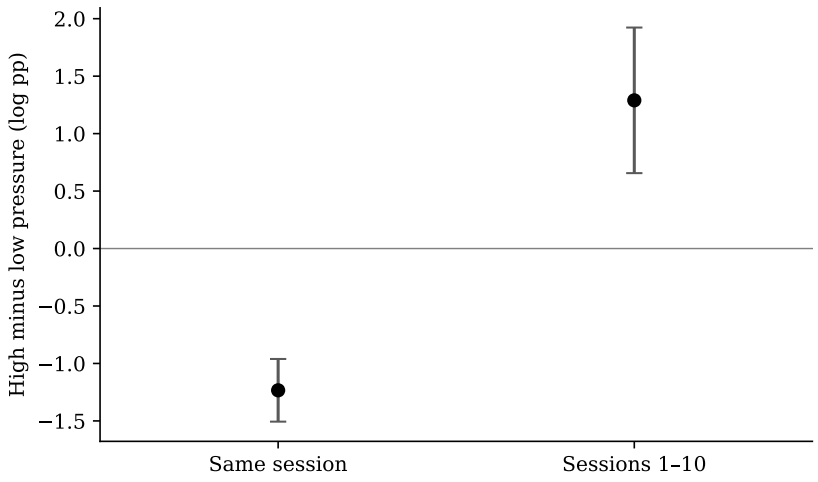


Figure 4: Timing of adjustment to identical news

Notes. High-pressure minus low-pressure signed return differences from (6), in log percentage points, with pointwise 95 percent confidence intervals. Every focal security within an industry-date-group cell is assigned the same return from the opposite fixed half of the industry. The same-session estimate is −1.23 (s.e. 0.14); the sessions 1–10 estimate is 1.29 (0.32). The cumulative difference through session ten is 0.06 (0.32).

Table 8: Immediate and subsequent responses to industry news

Comparison	Same session	Sessions 1–10	Sessions 0–10	N
A. Difference in peer-return loading				
All observations	-0.473	0.548	0.074	248,832
	(0.028)	(0.075)	(0.072)	
Positive peer returns	-0.493	0.821	0.328	130,928
	(0.076)	(0.248)	(0.262)	
Negative peer returns	-0.591	0.752	0.161	117,904
	(0.093)	(0.251)	(0.227)	
B. Common peer return: signed return difference (log percentage points)				
Common return from the other group	-1.23	1.29	0.06	253,658
	(0.14)	(0.32)	(0.32)	
At least five firms in the peer group	-1.33	1.12	-0.21	156,191
	(0.16)	(0.39)	(0.41)	
Controls for prior surveillance	-1.23	1.29	0.05	253,480
	(0.14)	(0.32)	(0.32)	

Notes. Panel A reports γ_h from (5); the middle column uses returns over sessions 1–10 excluding the initial session. Panel B reports $100\tau_h$ from (6), using a peer return drawn from the opposite fixed half of the industry, with a fixed effect for each combination of industry, date, and focal group. Its first row uses 253,658 observations and 5,481 high-pressure observations. The second requires at least five source peers, and the third adds known prior ASM/GSM/Trade-for-Trade membership. All entries within a row use identical observations. Standard errors in parentheses are clustered by security and exact date.

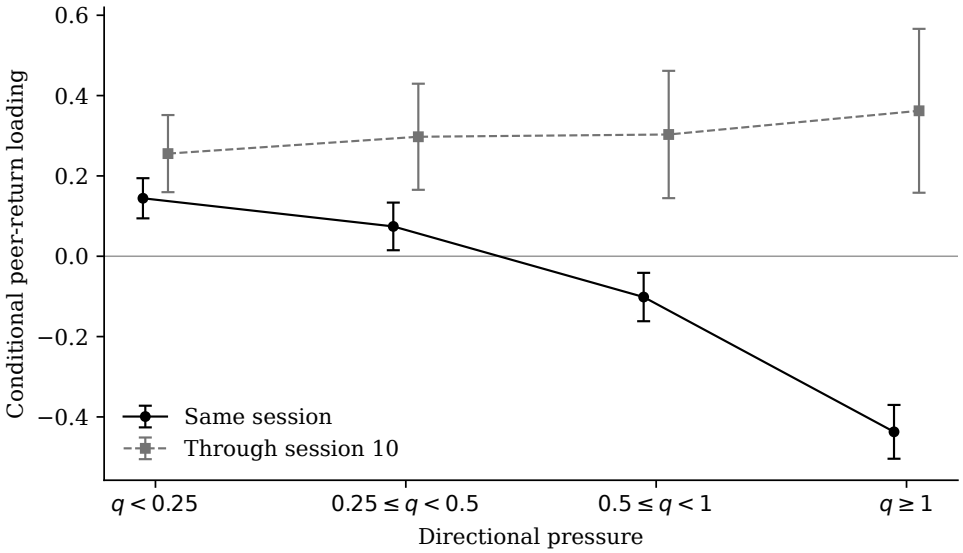


Figure 5: Peer-return loadings across directional pressure

Notes. Conditional slopes on the peer return in the fixed intervals $q < 0.25$, $0.25 \leq q < 0.5$, $0.5 \leq q < 1$, and $q \geq 1$. The regression interacts the peer return with the intervals and includes interval indicators, the baseline controls and fixed effects, and quadratic and cubic shock terms. Both horizons use the main 248,832-observation sample. Error bars are pointwise 95 percent intervals using security/date clustering.

by 3.4 points. Turnover falls, sessions with trading but no price range become more common, and the daily high–low range contracts.

Table 9: Trading and price location under directional pressure

Outcome	Sample mean	Change per unit of pressure
Close at the relevant boundary (percent)	8.40	19.39 (1.50)
Close at the opposite boundary (percent)	4.41	-3.39 (0.69)
Log daily share turnover	-6.89	-0.29 (0.03)
Sessions with volume and $H = L = C$ (percent)	4.25	5.20 (0.89)
Daily log high–low range (percent)	5.32	-2.21 (0.15)

Notes. Regressions of each outcome on continuous pressure, the signed peer return and its quadratic and cubic terms, predetermined controls, and security, date, and band effects. The common sample contains 243,251 observations for 209 firms. Means are unadjusted sample means. Shares and log ranges are multiplied by 100, so their coefficients are percentage-point or log-percentage-point differences per unit of pressure; log turnover remains in log units. Standard errors are clustered by security and exact date.

The bunching follows the direction of the news, with positive industry returns pushing securities toward their upper limit and away from their lower one and negative returns doing

the reverse. Explanations based on the general illiquidity or speculative character of securities in narrow bands have no reason to produce bunching on the particular side of the range that the day's news makes relevant. The pattern reflects the interaction of information with legal capacity.

Across a wide range of alternative constructions, the timing result is robust (Appendix Table 16). It survives flexible piecewise controls for the magnitude of the news, alternative pressure thresholds from 0.75 to 1.5, a continuous pressure interaction, value-weighted and median peer returns, larger peer sets, restriction to securities whose band is unchanged for twenty sessions, and restriction to prices fully consistent with the reconstructed limits. It holds within the narrow local window $0.75 \leq q \leq 1.25$, and it replicates with an independent industry classification from CMIE Prowess covering a different set of 305 firms in 24 industries, where the same-session interaction is -0.228 (standard error 0.038) and the interaction over sessions 1–10 is 0.324 (0.135).

4.4 Beyond mechanical truncation

Because a legal cap mechanically limits how far a price can move on any day, some attenuation of the immediate response would arise even if the price process were otherwise unchanged and simply clipped at the limits. We quantify this component directly. Let r^*_{it} denote a reference return before the legal range is applied, and define the clipping operator

$$\mathcal{C}_{it}(r) = \min\{\log(U_{it}/P^-_{it}), \max[\log(L_{it}/P^-_{it}), r]\}. \tag{7}$$

For a conditional reference distribution, write $m_{it} = \mathbb{E}[r^*_{it} \mid W_{it}]$ and $m^C_{it} = \mathbb{E}[\mathcal{C}_{it}(r^*_{it}) \mid W_{it}]$. Then

$$r_{it} - m_{it} = (m^C_{it} - m_{it}) + (r_{it} - m^C_{it}). \tag{8}$$

The first term is the attenuation implied by imposing the cap on the reference distribution; the second is everything else. Estimating (5) on each term in the same sample gives an exact decomposition of the headline interaction.

Table 10 reports three benchmarks. The first gives each security its peer return before clipping. The second estimates a conditional mean return recursively from earlier years, using only wide-band, low-pressure observations. The third, our preferred benchmark, applies the legal limits to the full estimated distribution of returns rather than to its mean, since idiosyncratic movements can reach the boundary even when the expected return does not. Under the distributional benchmark, mechanical clipping accounts for -0.233 of the -0.455 immediate attenuation. The remaining -0.222, comparable in magnitude to the mechanical

component, is economically large and highly significant. The unit-response benchmark gives a nearly identical division, and the conditional-mean benchmark attributes still more of the attenuation to forces beyond the cap.

Table 10: Mechanical clipping and remaining contemporaneous attenuation

Benchmark	Total deviation	Clipping component	Remaining deviation	N
Unit peer-return response	-0.466 (0.030)	-0.237 (0.012)	-0.229 (0.024)	242,127
Predicted conditional mean	-0.455 (0.031)	-0.107 (0.008)	-0.348 (0.028)	200,439
Predicted return distribution	-0.455 (0.031)	-0.233 (0.007)	-0.222 (0.028)	200,439

Notes. Coefficients on $s_{it}D_{it}$ from (5), using the three outcomes in (8); units are return loadings. All components in a row use identical observations. Recursive rules estimate the reference relation only from earlier-calendar-year, 20 percent band observations with $q < 0.5$, purged for outcome and capital-action availability. The distribution rule averages clipped values after adding earlier training residual quantiles. Samples require prices consistent with reconstructed bounds, exact exchange-return reconciliation, and no inferred capital action over sessions 0–10. Standard errors are clustered by security and date and condition on the estimated reference rule. Appendix B.10 provides its construction.

Since the benchmark transports the return process estimated under wide bands to securities under narrow ones, the residual measures attenuation beyond what a cap on that process would produce, whether through execution, liquidity supply, or attention. Combined with the directional bunching, the fall in trading, and the subsequent catch-up, the decomposition indicates that a binding range alters the trading process through which information reaches prices, over and above truncating them.

If a legal range changes which trades occur and when information enters prices, then what the regulator communicates about trading in a security may determine whether informative trades remain feasible at all. The model in the next section takes up this possibility.

5 A Model of Surveillance, Trading, and Learning

Section 3.5 showed that the exchange's decisions identify securities in distinct trading states without behaving like announcements of news about value, since the repricing after tightening is largest when tightening was most predictable. What the exchange appears to know, and what its actions convey, is how the order flow in a security should be read: how unusual, concentrated, or well informed current trading looks. Because the same actions also determine which trades are legal, it is natural to ask what happens when information of this kind meets a binding price range. The model developed here answers that question and delivers the

disclosure paradox, in which disclosing more about the market leaves the market knowing less about the asset.

Our framework is the sequential market of Glosten and Milgrom (1985), in which each order may convey private information and competitive liquidity suppliers set quotes accordingly. Since the NSE is an electronic limit-order market, the dealers of the model should be read as a stand-in for competitive liquidity supply, whose role is to fix the prices at which an incoming order can be absorbed. After establishing the paradox, we characterize the benefits that disclosure provides in return and then introduce public information from unexecuted orders and the daily recentering of the band.

5.1 Sequential trading under a price band

The asset pays

$$V = v_0 + \Delta X, \qquad \mathbb{P}(X = 1) = \mathbb{P}(X = -1) = \tfrac{1}{2}, \quad v_0 > \Delta > 0. \tag{9}$$

We express prices in normalized units $(P - v_0)/\Delta$. A band of width $b \in (0, 1)$ permits normalized prices in $[-b, b]$, so b measures the permitted price movement relative to the payoff spread.

A surveillance state $\theta \in \{l, h\}$, with $0 < l < h < 1$, determines the probability that an arriving customer is informed. Write

$$\pi = \mathbb{P}(\theta = h), \qquad \mu = (1 - \pi)l + \pi h. \tag{10}$$

Because the surveillance state is independent of the payoff, learning that order flow is highly informed leaves investors' beliefs about the asset's value unchanged while altering how they interpret the orders that follow. This independence isolates the execution channel rather than asserting that securities in different surveillance states have identical expected subsequent returns in the data. The revision evidence shows that the exchange acts in trading states with sharply different return paths, tied to the observable state the exchange flags, such as preceding run-ups. Allowing expected returns to vary with that observable state would not change the mechanism below: quote feasibility depends on how the report changes beliefs about the informativeness of order flow. We impose $\theta \perp X$ so that the disclosure result cannot be attributed mechanically to the regulator revealing fundamental news.

Conditional on (X, θ), customer types and liquidity directions are independent across arrival opportunities. With probability θ an arriving customer is informed, knows X, and buys when $X = 1$ and sells when $X = -1$. Otherwise the customer has an exogenous liquidity

need to buy or sell with equal probability and obtains benefit $v \geq b$ from execution, so she accepts any legal price in her intended direction. Customers trade one unit and live for one arrival opportunity. An arrival opportunity represents a round of economically meaningful order flow rather than an individual order, so N indexes such rounds within a fixed-band trading episode. Risk-neutral dealers compete in quotes and receive only the transaction payment. The public record contains quote availability, executed directions and prices, and inactivity, but not unsubmitted order intentions. The exchange observes θ, commits to its communication rule before observing it, and the rule is common knowledge.

If $Y \in \{-1, 1\}$ denotes intended order direction, then

$$\mathbb{P}(Y = y \mid X = x, \theta) = \frac{1 + \theta x y}{2}. \tag{11}$$

Suppose the exchange sends a public report M about θ, and let $m = \mathbb{E}[\theta \mid M]$. Because the report is uninformative about X, the public prior over the payoff remains symmetric, and the competitive ask and bid are

$$A(M) = m, \qquad B(M) = -m. \tag{12}$$

The half-spread thus equals the market's estimate of the informed share, and both sides are available precisely when $m \leq b$.

More generally, if $\theta = a$ is public and $z = \mathbb{E}[X \mid \mathcal{F}]$ is the current payoff posterior, the competitive quotes are

$$A_a(z) = \frac{z + a}{1 + az}, \qquad B_a(z) = \frac{z - a}{1 - az}. \tag{13}$$

A purchase can execute when $A_a(z) \leq b$, at ask $\max\{A_a(z), -b\}$, and a sale when $B_a(z) \geq -b$, at bid $\min\{B_a(z), b\}$. A legal break-even quote is thus offered whenever one exists. A boundary on the favorable side can leave the dealer a positive margin, while a boundary on the adverse side makes the quote unavailable. Customer trading at a break-even endpoint is resolved in favor of execution.

Lemma 1. *If $a > b$, neither direction can execute precisely when*

$$|z| < d(a, b), \qquad d(a, b) = \frac{a - b}{1 - ab}. \tag{14}$$

If order intentions that never reach the market are unobserved, this region is absorbing under a fixed band in the absence of other information.

In the region of two-sided withdrawal no transaction occurs whatever the customer knows, and when order intentions are not public the market learns nothing from that arrival opportunity. The legal range can affect information production even without a transaction at the boundary. If a purchase would require $A_a(z) > b$ while a sale remains feasible, the purchase cannot execute, and the last transaction may remain strictly inside the band although the feasible range prevents the next informative order from trading. Observed contact with the boundary is therefore unnecessary for the constraint to separate market beliefs from recorded prices, consistent with the finding that pressure predicts delayed adjustment well before closes reach the limit.

5.2 The disclosure paradox

Consider two regimes under the same fixed band, in one of which the exchange withholds its surveillance information and in the other publishes θ before trading. Although publication says nothing about the payoff, it tells dealers when an order is especially likely to be informed, and this can make the quote required to absorb the order illegal. We rank information experiments with respect to the payoff X in the sense of Blackwell (1953); the experiment includes the regulatory report and the complete public trading record.

Proposition 1. *Fix a finite number $N \geq 1$ of arrival opportunities. A customer arrives at each opportunity with probability $\rho \in (0, 1]$, independently of all states. Let*

$$c_N = \tanh(N \operatorname{arctanh} l)\,, \tag{15}$$

and suppose

$$0 < l < \mu < b < h < 1, \qquad c_N < \rho\mu. \tag{16}$$

Publishing θ conveys no information about X. Nevertheless, after the N trading opportunities the public experiment generated when the report is withheld strictly Blackwell dominates the experiment generated after publication. These inequalities hold on a nonempty open parameter set for every fixed finite N and every $\rho > 0$.

When the report is withheld, dealers price orders at the pooled informed probability $\mu < b$, every customer can trade, and each trade moves the posterior to $\pm\mu$. Publication separates the two states. In the high state the compensation for adverse selection that a dealer requires exceeds the legal range on both sides, so the orders that would have been most revealing never reach the record, while in the low state orders execute but each reveals relatively little. The market thus forgoes trading-based information in exactly the state where trading would have been most informative.

The proof constructs a state-independent garbling of the first record without disclosure that reproduces the entire final experiment after disclosure. For example, $l = 0.2$, $h = 0.9$, $\pi = 5/7$, $b = 0.8$, and $\rho = 1$ give $\mu = 0.7$ and satisfy the conditions for one arrival. Both regimes begin with a payoff mean squared error of one; after trading, the error is 0.51 under withholding and about 0.99 under publication, so that disclosure leaves the market almost as uninformed as it was before trading began. The Blackwell ranking is stronger than this quadratic comparison, since it implies that every decision about the payoff is better served by the record generated without disclosure.

With private information held fixed, the loss of learning arises entirely through competitive execution, and no change in information acquisition or possession is involved. This margin complements the responses of acquisition and trading studied by Grossman and Stiglitz (1980), Paul (1993), and Goldstein and Yang (2019), and it is distinct from the familiar concern that public news about fundamentals can crowd out private research, because the regulator here discloses nothing about fundamentals at all.

5.3 Liquidity and learning

Disclosure would be hard to justify if its only effect were to reduce what the market learns. Surveillance information is also valuable to those who trade without private information, however: revealing that order flow is benign lowers the spread that liquidity customers pay, even though revealing that it is dangerous can exclude them from trading altogether.

Proposition 2. *Consider one certain arrival ($N = \rho = 1$) under the conditions of Proposition 1.*

(i) Under withholding, every customer trades at half-spread μ.

(ii) Under publication, customers in the low state trade at half-spread $l < \mu$, and no customer trades in the high state.

(iii) Expected surplus of liquidity customers is higher under publication than under withholding if and only if

$$\upsilon < \upsilon^* \equiv \frac{(1-\mu)\mu - (1-\pi)(1-l)l}{\pi(1-h)}. \tag{17}$$

(iv) Final public information about X is strictly lower under publication.

Surveillance transparency therefore trades liquidity against learning. For moderate benefits of execution, the narrower spreads that disclosure brings to cleared securities outweigh the exclusion of liquidity traders from flagged ones, and uninformed traders gain even as the public record becomes poorer. When execution is very valuable, the loss of access in flagged states dominates, and withholding is preferable on both counts. In the numerical example

$v^* = 2.3$ in normalized units, so disclosure benefits liquidity customers for every execution benefit from the band width up to more than the full spread between the two payoff values. The empirical pattern has the same shape, with loosened securities trading more actively and tightened securities experiencing a collapse in turnover and a proliferation of stale prices.

5.4 Price variation and market learning

The disclosure paradox has a direct implication for the volatility evidence. Let p_k denote the normalized last transaction after arrival opportunity k, carried forward whenever no trade occurs, with $p_0 = 0$. Define recorded price variation by

$$Q_N = \sum_{k=1}^{N} (p_k - p_{k-1})^2. \tag{18}$$

Publication changes both the information entering each transaction and whether a transaction occurs. In the high state both sides withdraw, leaving the last price unchanged. In the low state, transactions remain feasible but reveal relatively little. Without publication, the first customer can trade at the more responsive pooled quotes.

Corollary 1. *Under the conditions of Proposition 1, let D denote publication and O withholding the report. Then*

$$\mathbb{E}Q_N^D \leq (1-\pi)c_N^2 < \rho\mu^2 \leq \mathbb{E}Q_N^O. \tag{19}$$

The same strict ordering holds for $\mathbb{E}\sqrt{Q_N/N}$. Publication therefore lowers recorded price variation, in either squared or root mean square units, while raising the expected posterior variance of the payoff.

Recorded prices become quieter because execution ceases in the states where a trade would have been most informative, so that, with the distribution of the payoff unchanged, lower recorded variation coexists with greater uncertainty after trading. In the spirit of French and Roll (1986), the result links price variation to the process by which information is incorporated. Quieter prices after a surveillance action are thus compatible with slower resolution of uncertainty whenever the trades that would have moved prices most no longer take place.

5.5 Information in unexecuted orders

Whether withdrawal destroys learning depends on the information available outside transactions. The NSE disseminates outstanding buy and sell quantities and several levels of quoted

depth (National Stock Exchange of India, 2024), so a missing transaction need not imply a missing public signal. Suppose that after each arriving customer's execution opportunity, the public observes $Q \in \{-1, 1\}$ with

$$\mathbb{P}(Q = q \mid Y = y, X, \theta) = \frac{1 + \eta q y}{2}, \qquad 0 \leq \eta \leq 1. \tag{20}$$

Errors are independent across customers conditional on their intentions. The signal is available whether or not the order executes; a transaction reveals Y exactly, making Q redundant for that customer. The signal is a reduced-form description of the information conveyed by unexecuted demand, and we do not model order submission and cancellation explicitly.

Proposition 3. *In the environment of Proposition 1, add the public signal* (20)*, and let* $c_N^Q = \tanh(N \operatorname{arctanh}(\eta h))$*. If*

$$c_N < \rho\mu, \qquad c_N^Q < \min\{d(h, b), \rho\mu\}, \tag{21}$$

withholding the report still strictly Blackwell dominates publication at the end of the N *opportunities. Publication also lowers expected recorded squared and root mean square variation and raises expected payoff prediction error. The conditions admit strictly positive* η*.*

In the high state the visible order record supplies some information, though too little over this horizon to restore either quote, and in the low state orders execute so that the order signal adds nothing further. In the numerical example, $\eta = 0.2$ satisfies the conditions, and publication raises final payoff mean squared error from 0.51 to about 0.97 even though unexecuted demand is partly informative.

The restriction on signal quality has content. If $\eta = 1$, every intention is observed, so the report and the sequence of intentions under publication can reproduce the entire withholding record, and publication weakly Blackwell dominates withholding. Nor does a small probability of perfectly observing an otherwise hidden intention satisfy (21): rare exact revelation creates high-state posteriors outside the support of the withholding experiment, and the experiments can become incomparable even while publication raises quadratic prediction error. Because the bound in Proposition 3 controls the information in every possible record of unexecuted orders, the average transparency of an order book is not a sufficient statistic for whether surveillance disclosure impairs learning.

5.6 Daily recentering and subsequent adjustment

The comparison under a fixed band concerns information at the end of a trading episode, whereas the data show adjustment that is delayed but eventually complete, and explaining the latter requires a continuation. In price units, let the next session's permitted range be

$$[L_t, U_t] = [(1-\beta)P_{t-1}, (1+\beta)P_{t-1}], \qquad 0 < \beta < 1, \tag{22}$$

where P_{t-1} is the previous close. This is the daily percentage-band rule, abstracting from ticks (National Stock Exchange of India, 2026b). A complete freeze does not reopen capacity on its own: if both quotes are unavailable, intentions remain unobserved, and no other information arrives, the close is unchanged and so is the next band. The absorbing state in Lemma 1 can persist despite daily recentering.

With the public order signal, learning can resume without any assumed path for prices. Repeat the one-opportunity market over successive sessions, keep X and θ fixed, and let a customer arrive each session. Apply (22), retain the latest transaction as the close, and observe (20) after each execution opportunity. Competitive break-even quotes are calculated from the full public history, with the favorable-endpoint rule of Section 5.1.

Proposition 4. *Suppose $0 < l < h < 1$, $v_0 > \Delta > 0$, and $\eta > 0$. Starting from $P_0 = v_0$, under either publication or withholding,*

$$\mathbb{E}[V \mid \mathcal{F}_t] \longrightarrow V, \qquad P_t \longrightarrow V \quad \textit{almost surely.} \tag{23}$$

Consequently, cumulative log-price changes under the two policies converge to the same terminal change when they share the same realized payoff. If the initial session satisfies Proposition 3 with $N = \rho = 1$, then

$$G_1 > 0, \qquad G_t \longrightarrow 0, \qquad G_t = \mathbb{E}\left[X \log\left(\frac{P_t^O}{P_t^D}\right)\right].$$

Publication therefore gives less payoff-aligned initial adjustment, followed by eventual recovery of the expected response gap.

Conditional on the payoff and the surveillance state, successive public order signals have mean $\eta\theta X$, so their signs eventually reveal the payoff even if transactions stop at first, and the break-even quotes approach the payoff as uncertainty resolves. Since $\theta < 1$, liquidity customers continue to arrive on both sides, and one who trades on the profitable side of an endpoint generates a new close that moves the following session's band toward the payoff. Catch-up therefore requires public information and continued liquidity trading in addition to

recentering, which cannot produce it alone. The result imposes neither monotone adjustment at every date nor convergence within a particular number of sessions.

Taken together, these results account for the pattern documented in Section 4, in which a constrained security responds less on the day information arrives, trades less, records quieter prices, and catches up later. They also clarify what the timing evidence measures. The peer-return results describe when recorded prices adjust, while the Blackwell comparison describes what is publicly known at a given horizon; because public knowledge can improve without a transaction and a later transaction can incorporate information already known, a market may close an initial gap in price responses that understated the full loss of information. The model is symmetric in the payoff by construction and leaves aside the unwinding of run-ups after tightening, which is associated with the state the exchange flags.

6 Regulatory Design

Once communication affects execution, disclosure policy and the trading rule have to be chosen jointly. Proposition 2 shows that full transparency trades liquidity against learning, lowering trading costs where order flow is benign and freezing trade where it is dangerous. With the band held fixed, the exchange can improve on both extremes by pooling surveillance states just enough to keep informative trades feasible. When the band is the exchange's only public message, the regulatory category has to carry information and impose the trading constraint at the same time. We solve both problems and then provide a financing foundation for the value of protection.

To obtain closed-form policies, we return to the one-opportunity benchmark in which unexecuted intentions are private. Informative order displays and later trading would reduce the costs of withdrawal, as Sections 5.5 and 5.6 show, without removing the tradeoff at the horizon over which the regulator values information.

6.1 Disclosure that preserves market learning

Suppose first that the band is fixed and the exchange decides how much of its surveillance information to disclose. Observing θ alone, it commits ex ante to a randomized reporting rule. There is one certain customer arrival and a fixed band b_T satisfying

$$0 < l < \mu < b_T < h < 1. \tag{24}$$

A report induces a distribution of posterior means $m \in [l, h]$ with mean μ, the Bayes-plausibility restriction in Kamenica and Gentzkow (2011). If $m \leq b_T$, the customer trades

and the final payoff posterior is $+m$ or $-m$ with equal probability. If $m > b_T$, both sides withdraw and the report contains no direct information about X, so the final payoff posterior remains zero.

Proposition 5. *Under* (24)*, consider the report inducing posterior means*

$$m = l \quad \textit{with probability } x = \frac{b_T - \mu}{b_T - l}, \qquad m = b_T \quad \textit{with probability } 1 - x. \tag{25}$$

The final payoff experiment induced by this report Blackwell dominates every alternative attainable through a report about θ under the fixed band. Every customer can trade under this disclosure. Its reduction in expected payoff variance is

$$Q_{b_T}(\mu) = \mu(l + b_T) - l b_T = \mu^2 + (\mu - l)(b_T - \mu). \tag{26}$$

The same report maximizes any nonnegative combination of this information value and completed liquidity-trading surplus, and, relative to withholding, it weakly raises the expected surplus of liquidity customers net of the half-spread they pay.

Under the optimal report, every high-θ state is pooled with enough low-θ states to bring the pool exactly to the legal quote threshold. Complete disclosure would identify the states in which trading stops, while complete silence preserves trade at the cost of charging every liquidity customer the pooled spread and leaving dealers needlessly uncertain about how to read an order. Partial opacity improves on silence on both counts, giving a share of cleared securities the narrow spreads of the benign state while preserving just enough uncertainty about the composition of order flow in the remainder to keep that market open and informative. Relative to full disclosure, it gives up part of the reduction in spreads for cleared securities in exchange for keeping flagged securities tradable.

6.2 Band design when the band is the message

Suppose next that the exchange has no separate means of communication and commits ex ante to a policy mapping θ into one of two publicly observed bands. The tight band T has width b_T, and the wider band W is wide enough to accommodate even the high-θ order flow. Since investors observe only the band, changing the restriction is also the exchange's only way of changing the market's prior about the quality of order flow.

Let Z denote the final payoff posterior. A downstream decision maker choosing e to maximize $2eX - e^2$ selects $e = Z$ and obtains ex-ante value $\mathbb{E}[Z^2]$. Completed liquidity trades

contribute normalized surplus $\gamma \geq 0$. Using W entails supervisory cost $\kappa \geq 0$, representing the direct protection forgone by relaxing the restriction. Thus

$$J = \mathbb{E}[Z^2] + \gamma \mathbb{P}(\text{liquidity customer executes}) - \kappa \mathbb{P}(W). \tag{27}$$

Section 6.4 provides a financing-based reason for protection to have value.

Define

$$\pi = \frac{\mu - l}{h - l}, \qquad x = \frac{b_T - \mu}{b_T - l}, \qquad J_P = \mu^2 + \gamma(1 - \mu). \tag{28}$$

Full separation assigns T to l and W to h. Selective relaxation uses W only in some low-θ states, leaving the restrictive category with posterior mean exactly b_T. Their values are

$$J_S = J_P + (\mu - l)(h - \mu) - \kappa\pi, \tag{29}$$

$$J_I = J_P + (\mu - l)(b_T - \mu) - \kappa x. \tag{30}$$

Proposition 6. *Under* (24)*, the global optimum over all randomized mappings from θ into $\{T, W\}$ is*

$$\mathcal{V}_B(\kappa) = \max\{J_P, J_S, J_I\}. \tag{31}$$

The three terms are attained by pooling all states under T, full separation, and selective relaxation. Under selective relaxation,

$$\mathbb{P}(W \mid \theta = h) = 0, \qquad \mathbb{P}(W \mid \theta = l) = \frac{x}{1 - \pi}, \tag{32}$$

and the posterior means are $m_T = b_T$ and $m_W = l$. Every order executes under each candidate policy.

The envelope includes all stochastic assignments, including those that shut down trade in one category. For primitives admitting an intermediate region of selective relaxation, the meaning of widening changes with its supervisory cost. When relaxing the restriction is inexpensive, the exchange uses the wide band for the state in which orders contain the most private information, allowing the larger competitive price response to occur. As κ rises, the optimum can reverse. The exchange keeps every high-θ state under the tight band and widens only in some low-θ states so that the posterior informed share in the remaining tight category rises exactly to the execution threshold. The wide category then carries smaller competitive price movements than the tight category. Nominal band width and equilibrium price impact need not have the same ordering.

A richer menu can separate communication from mechanical relaxation even when the exchange has no independent disclosure instrument. Introduce a third category C with width $b_C \in (l, b_T)$ and the same direct supervisory cost as T.

Corollary 2. *If $\kappa \geq (h-l)(h-b_T)$, the global optimum with $\{C, T, W\}$ assigns posterior l to C, posterior b_T to T, and never uses W. Every order executes and C is strictly nonbinding. Throughout the intermediate-cost region in which selective relaxation is optimal with only $\{T, W\}$, the additional category raises welfare by replacing mechanically unnecessary widening with a distinct message.*

Although the boundary of C never prevents a trade in the states assigned to it, the category is valuable because of the information carried by its label. A regulatory category can therefore improve price discovery without ever being observed to bind.

6.3 Separate disclosure and regulatory categories

With an independent report, the exchange can communicate a distinction without changing the trading range, and the resulting gain can be characterized exactly. Let

$$J_D = Q_{b_T}(\mu) + \gamma(1-\mu), \qquad \kappa_D = (h-l)(h-b_T). \tag{33}$$

The first expression is the value of optimal disclosure under the tight band; the second is the cost at which full separation and that disclosure have equal value.

Corollary 3. *With an independently chosen public report about θ as well as band choice, the optimal value is*

$$\mathcal{V}_D(\kappa) = \max\{J_S, J_D\}. \tag{34}$$

For $\kappa \geq \kappa_D$, optimal partial disclosure under T attains this value. Relative to a band-only optimum that uses selective relaxation, separate disclosure preserves the same final information and execution while eliminating widening, increasing welfare by κx. For $\kappa > \kappa_D$, no band-only assignment attains J_D.

The additional category in Corollary 2 achieves the same separation through the menu itself, giving the low-information states a distinct label while allowing their orders to execute under a narrow range. Separate disclosure thus delivers the same information and execution with less relaxation of the protective rule.

6.4 A protection motive from productive investment

In the design problem above, the value of protection enters only through κ, and a compact financing environment gives that term a concrete welfare interpretation. A trader with an outstanding cash-settled position can benefit from moving the settlement price, as in Kumar and Seppi (1992). The resulting payment reduces the counterparty's internal funds and, when external financing is limited by pledgeable income, its investment capacity. Following the financing friction in Holmström and Tirole (1997), a trading restriction can preserve productive investment even when the entrepreneur fully understands the public record, complementing the managerial-learning channel in Goldstein and Guembel (2008).

Consider one trading opportunity, conditional on a public probability ω of a potential manipulator. With probability $1-\omega$, a customer knows X and buys if $X=1$ or sells if $X=-1$. Otherwise an uninformed customer holds a pre-existing cash-settled position of $d>1$ units and may buy or remain inactive. A purchase incurs a brokerage fee c. The side position and its counterparty are fixed before the opportunity and arise only in this customer state. Position acquisition is sunk, positions cannot be renegotiated during trading, and unsubmitted intentions remain private.

Under a wide band admitting both payoff values, a manipulative purchase with probability σ gives the competitive buy price

$$p(\omega,\sigma)=\frac{1-\omega}{1-\omega+2\omega\sigma}. \tag{35}$$

The manipulator's gain is $(d-1)p-c$. Write $\bar{p}=c/(d-1)\in(0,1)$ and $\omega^*=(1-\bar{p})/(1+\bar{p})$. She always buys up to ω^* and thereafter mixes to keep the price at $\bar{p}$. The unconditional frequency $M(\omega)=\omega\sigma(\omega)$ first rises and then falls. A tight band with upper endpoint $u<\bar{p}$ deters these purchases. Its lower endpoint is assumed to admit the informed sale at -1. Without manipulative purchases the only buyer is informed, requiring an ask of one, so no legal buy quote is offered.

The restriction merges purchases with inactivity in the public record and can reduce the information it conveys. Its benefit comes from the settlement counterparty's financing constraint. This entrepreneur has liquid wealth $w>d$ and a project returning $R>1$ per unit invested. At most $\zeta\in(0,1)$ per unit invested can be pledged to outside lenders. After paying s, the maximum investment is $k(s)=(w-s)/(1-\zeta)$. Recipients of the settlement and trading payments have access to an outside technology returning one. Consolidating

all transfers, including the brokerage fee, leaves a loss of productive surplus following a manipulative purchase of

$$(R-1)\{k(0)-k(dp)\}=\Lambda p, \qquad \Lambda=\frac{(R-1)d}{1-\zeta}. \tag{36}$$

The payment redistributes wealth, and its social cost comes from the resulting change in investment: funds displaced from the entrepreneur's project earn the lower outside return. Equation (36) counts that loss after consolidating all payments. Appendix A.8 derives the financing limit and trading equilibrium.

Let $I_a=\mathbb{E}[(\mathbb{E}[X\mid\mathcal{F}_a])^2]$ measure the allocation value of final information under regime a. As in the preceding design problem, an operating decision producing $\nu(2eX-e^2)$ gives expected value νI_a. If administering a tight band costs $F>0$, its net gain is

$$G(\omega)=\Lambda M(\omega)p(\omega)-\nu(I_W-I_T)-F. \tag{37}$$

The first term is the expected surplus from investment preserved by deterring manipulation. The second measures how lost market information affects the operating decision, and the third is the cost of administering the restriction. The comparison holds public suspicion and the existing exposure fixed while trading, inference, and financing respond. The state in which manipulation is most frequent need not be the state in which this net gain is largest.

Proposition 7. *For every $\nu\geq 0$, the net gain in* (37) *is uniquely maximized at*

$$\widehat{\omega}=\min\{\omega^*,\sqrt{2}-1\}. \tag{38}$$

A tight band is strictly optimal for some public states if and only if

$$F<\Lambda\,\frac{\widehat{\omega}(1-\widehat{\omega})}{1+\widehat{\omega}}. \tag{39}$$

When this holds, there are unique $0<\omega_L<\widehat{\omega}<\omega_H<1$ such that the tight band is strictly optimal exactly on (ω_L,ω_H). If $\bar{p}<\sqrt{2}-1$, the welfare peak strictly precedes the peak of manipulation frequency. Holding $\bar{p}$ fixed, increasing Λ weakly expands the intervention region, while increasing ν weakly contracts it without changing the maximizing state.

The distinction between M and Mp drives this result. As suspicion rises from a low level, more potential manipulators generate more manipulative orders, but they also lead dealers to move the price less after each one. Productive losses can therefore begin to fall while manipulation is still becoming more frequent. With $\bar{p}=0.2$, for example, manipulation

frequency peaks when suspicion reaches two thirds, whereas the social value of protection peaks near 0.414. Figure 6 illustrates this separation. The optimal restriction consequently responds to the interaction of manipulative participation, price impact, and financial exposure.

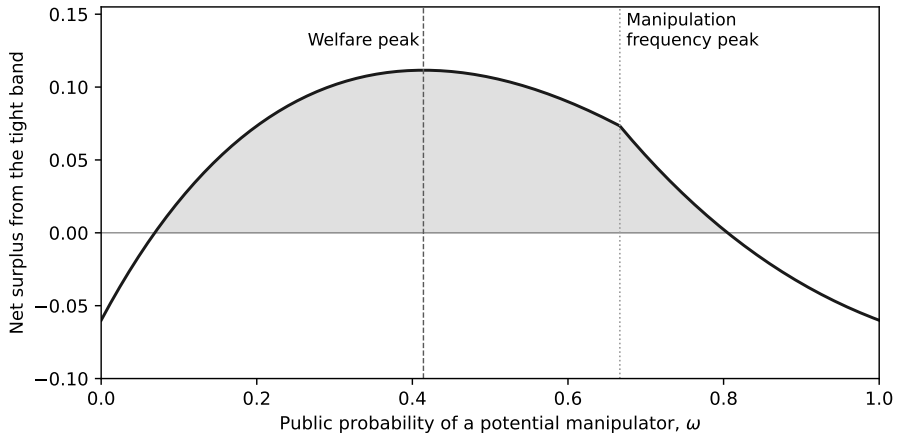


Figure 6: Manipulation frequency and the value of protection

Notes. Net surplus from the tight band in (37), with $R = 1.2$, $\zeta = 0.6$, $d = 2$, $c = 0.2$, $\nu = 2$, and $F = 0.06$. These primitives imply $\Lambda = 1$ and $\bar{p} = 0.2$. The shaded interval is the set of public suspicion states in which intervention raises welfare. The welfare peak precedes the peak of manipulation frequency.

In the financing environment, then, the protection term has a real allocative counterpart. Informative trading improves decisions, whereas manipulative trading can move a settlement price, drain internal funds, and reduce investment, so a restriction can protect productive activity by suppressing transactions even though it also removes information that the market would otherwise have produced.

7 Conclusion

Exchange surveillance communicates and constrains at the same time, and the two functions can conflict. A surveillance message tells liquidity suppliers how to interpret the order flow in a security, while the accompanying trading rule determines whether the prices they then require are legal. Where those prices fall outside the permitted range, the trades that would have been most informative do not occur, and a regulator that discloses more about the market can leave the market knowing less about the asset.

The NSE's price bands allow both functions to be observed. The exchange's decisions identify securities in sharply different states: the realized action more than doubles the out-of-sample predictability of subsequent returns, and tightening is followed by lower returns, a collapse in turnover, and much quieter prices. Since the repricing is concentrated among

tightenings that public conditions predicted, these decisions appear to convey information about trading in a security and little about its value. Under a band already in force, securities with little legal room respond less on the day industry news arrives and make up the difference over the following two weeks, while trading thins and closes bunch at the boundary toward which the news points. About half of the initial delay exceeds what mechanical truncation would produce.

In the model, over a fixed trading episode, disclosure that is uninformative about fundamentals can nonetheless leave the market strictly less informed about them, in the Blackwell sense, while making recorded prices quieter. Such disclosure still benefits traders without private information in securities the regulator clears, at the cost of excluding them from securities it flags. Imperfectly visible order books preserve the result and fully visible ones eliminate it, and a frozen market recovers once continuing public information and liquidity trading move the close, which recentering of the band cannot accomplish on its own.

For regulators, these forces turn surveillance into a problem of joint design. Partial disclosure keeps informative trading feasible while keeping every liquidity trader in the market; a regulatory category can raise welfare without ever binding, because of the information carried by its label; and a separate disclosure instrument lets the exchange communicate without relaxing its trading rule merely to send a signal. Restrictions retain their value where manipulative trades drain the internal funds of constrained firms, and the case for intervening is strongest before manipulation reaches its peak frequency. Disclosure policy and trading rules are, in this sense, two parts of a single instrument, and they are best chosen together.

References

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24(2): 265–272. doi:10.1214/aoms/1177729032.

Bond, Philip, and Itay Goldstein. 2015. "Government Intervention and Information Aggregation by Prices." *Journal of Finance* 70(6): 2777–2812. doi:10.1111/jofi.12303.

French, Kenneth R., and Richard Roll. 1986. "Stock Return Variances: The Arrival of Information and the Reaction of Traders." *Journal of Financial Economics* 17(1): 5–26. doi:10.1016/0304-405X(86)90004-8.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics* 14(1): 71–100. doi:10.1016/0304-405X(85)90044-3.

Goldstein, Itay, and Alexander Guembel. 2008. "Manipulation and the Allocational Role of Prices." *Review of Economic Studies* 75(1): 133–164. doi:10.1111/j.1467-937X.2007.00467.x.

Goldstein, Itay, and Yaron Leitner. 2018. "Stress Tests and Information Disclosure." *Journal of Economic Theory* 177: 34–69. doi:10.1016/j.jet.2018.05.013.

Goldstein, Itay, and Liyan Yang. 2015. "Information Diversity and Complementarities in Trading and Information Acquisition." *Journal of Finance* 70(4): 1723–1765. doi:10.1111/jofi.12226.

Goldstein, Itay, and Liyan Yang. 2019. "Good Disclosure, Bad Disclosure." *Journal of Financial Economics* 131(1): 118–138. doi:10.1016/j.jfineco.2018.08.004.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70(3): 393–408. JSTOR 1805228.

Holmström, Bengt, and Jean Tirole. 1997. "Financial Intermediation, Loanable Funds, and the Real Sector." *Quarterly Journal of Economics* 112(3): 663–691. JSTOR 2951252.

Hou, Kewei, and Tobias J. Moskowitz. 2005. "Market Frictions, Price Delay, and the Cross-Section of Expected Returns." *Review of Financial Studies* 18(3): 981–1020. doi:10.1093/rfs/hhi023.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101(6): 2590–2615. doi:10.1257/aer.101.6.2590.

Kim, Kenneth A., and S. Ghon Rhee. 1997. "Price Limit Performance: Evidence from the Tokyo Stock Exchange." *Journal of Finance* 52(2): 885–901. doi:10.1111/j.1540-6261.1997.tb04827.x.

Kumar, Praveen, and Duane J. Seppi. 1992. "Futures Manipulation with 'Cash Settlement'." *Journal of Finance* 47(4): 1485–1502. doi:10.1111/j.1540-6261.1992.tb04666.x.

National Stock Exchange of India. 2023a. "List of Securities Moving to Stage I of Graded Surveillance Measure (GSM)." Circular NSE/SURV/59662, December 8, 2023. Official NSE circular.

National Stock Exchange of India. 2024a. "List of Securities Moving to Stage I of Graded Surveillance Measure (GSM)." Circular NSE/SURV/60132, January 5, 2024. Official NSE circular.

National Stock Exchange of India. 2024b. "List of Securities Moving out of Graded Surveillance Measure (GSM) – Update." Circular NSE/SURV/61070, March 11, 2024. Official NSE circular.

National Stock Exchange of India. 2024c. "Revision in Tick Size of Securities in CM Segment." Circular NSE/CMTR/62174, May 24. Official NSE circular.

National Stock Exchange of India. 2025. "Revision in Tick Size of Securities in CM Segment – Update." Circular NSE/CMTR/67133, March 13. Official NSE circular.

National Stock Exchange of India. 2026a. "Price Bands." Official exchange documentation, accessed September 17, 2026.

National Stock Exchange of India. 2026b. "Daily Review of Price Bands." Official exchange documentation, accessed September 17, 2026.

National Stock Exchange of India. 2026c. "Graded Surveillance Measure (GSM)." Official exchange documentation, accessed September 17, 2026.

National Stock Exchange of India. 2026d. "Historical Surveillance Actions Reports: Price Band Changes." Official exchange archive, accessed September 17, 2026.

National Stock Exchange of India. 2026e. "Trading System." Official exchange documentation, accessed September 18, 2026.

National Stock Exchange of India. 2026f. "Additional Surveillance Measure (ASM)." Official exchange documentation, accessed September 19, 2026.

Paul, Jonathan M. 1993. "Crowding Out and the Informativeness of Security Prices." *Journal of Finance* 48(4): 1475–1496. doi:10.1111/j.1540-6261.1993.tb04763.x.

National Stock Exchange of India. 2024. "Market Feed: Capital Market (Level 1, Level 2 and Level 3) Specification." Version 1.25, June 14. Sections 4.4–4.5. Official market-feed documentation.

A Proofs and Additional Derivations

A.1 Quote feasibility

Proof of Lemma 1. Bayes' rule applied to (11) gives (13). A dealer selling below $A_a(z)$ or buying above $B_a(z)$ loses money conditional on the accepted side. The quotes specified in Section 5.1 minimize the dealer's margin subject to legality and nonnegative conditional profit. Solving $A_a(z) > b$ gives $z > (b-a)/(1-ab) = -d(a,b)$. Solving $B_a(z) < -b$ gives $z < (a-b)/(1-ab) = d(a,b)$. Their intersection is (14). When both directions are unavailable and unsubmitted intentions are not observed, every arrival opportunity has the same public outcome in both payoff states, so Bayes' rule leaves z unchanged. □

A.2 The disclosure paradox

Proof of Proposition 1. Because θ is independent of X, the report leaves the payoff posterior at $z = 0$; publication conveys no information about X. After publication, consider first state h. Since $h > b$, $d(h,b) > 0$, so $z = 0$ lies strictly inside the two-sided withdrawal region of Lemma 1. No order executes, nothing is learned, and the posterior remains zero at every opportunity. In state l, an observed direction updates the posterior by A_l or B_l, and

$$\operatorname{arctanh} A_l(z) = \operatorname{arctanh} z + \operatorname{arctanh} l, \qquad \operatorname{arctanh} B_l(z) = \operatorname{arctanh} z - \operatorname{arctanh} l.$$

After at most N directions, every disclosed-market posterior therefore has magnitude at most c_N in (15), and every quote required at an opportunity has magnitude at most c_N as well. Condition (16) implies $c_N < \rho\mu \le \mu < b$, so every low-state break-even quote remains legal.

Index the possible final outcomes after publication by j; each outcome includes the report. Let r_j be the unconditional probability of outcome j and $z_j = \mathbb{E}[X \mid j]$. Symmetry of the payoff implies

$$\mathbb{P}(j \mid X = x) = r_j(1 + x z_j), \qquad \sum_j r_j = 1, \qquad \sum_j r_j z_j = 0. \tag{40}$$

Without disclosure, the pooled quotes are $\pm\mu$, which are legal, and the first opportunity has outcomes $+$, $-$, and $\varnothing$, with probabilities conditional on $X = x$

$$\frac{\rho(1+\mu x)}{2}, \qquad \frac{\rho(1-\mu x)}{2}, \qquad 1-\rho.$$

Define a stochastic kernel from these three outcomes to j by

$$K(j \mid +) = r_j \left(1 + \frac{z_j}{\rho\mu}\right), \qquad K(j \mid -) = r_j \left(1 - \frac{z_j}{\rho\mu}\right), \qquad K(j \mid \varnothing) = r_j.$$

Since $|z_j| \le c_N < \rho\mu$, all entries are nonnegative, and by (40) each row sums to one. For $X = x$,

$$\frac{\rho(1+\mu x)}{2} K(j \mid +) + \frac{\rho(1-\mu x)}{2} K(j \mid -) + (1-\rho) K(j \mid \varnothing) = r_j \left(1 + x z_j\right),$$

so the kernel reproduces the conditional distribution of every disclosed-market outcome in both payoff states. The first undisclosed record therefore Blackwell dominates the entire disclosed record, and the final undisclosed history is at least as informative as that first record.

The dominance is strict. Because the posterior is a martingale and $|\cdot|$ is convex, the final undisclosed posterior satisfies $\mathbb{E}|z_O| \ge \mathbb{E}|z_{O,1}| = \rho\mu$, whereas every disclosed posterior satisfies $|z| \le c_N < \rho\mu$. A garbling cannot raise the expectation of a convex function of the posterior, so no garbling of the disclosed experiment reproduces the undisclosed one. Finally, for any fixed N and $\rho > 0$, choose $0 < \mu < b < h < 1$ and then $l > 0$ small enough that $l < \mu$ and $c_N < \rho\mu$, setting $\pi = (\mu - l)/(h - l) \in (0, 1)$. All inequalities are strict and continuous in the parameters, so they hold on a neighborhood. □

A.3 Liquidity and learning

Proof of Proposition 2. Under withholding, (12) with $m = \mu < b$ gives legal quotes $\pm\mu$, so every customer trades at half-spread μ. Under publication, the low state has quotes $\pm l$, which are legal, while in the high state $h > b$ makes both quotes infeasible by Lemma 1. This establishes (i) and (ii).

A liquidity customer who trades at half-spread s obtains $\upsilon - s$, since the payoff has conditional mean zero at the time of trade. A customer is a liquidity customer with probability $1 - \theta$. Expected liquidity surplus is therefore $(1-\mu)(\upsilon - \mu)$ under withholding and $(1-\pi)(1-l)(\upsilon - l)$ under publication. Write $a = (1-\pi)(1-l)$ and note from (10) that $1-\mu = a+\pi(1-h)$. Publication yields higher surplus exactly when $a\upsilon - al > (1-\mu)\upsilon - (1-\mu)\mu$, or $\pi(1-h)\upsilon < (1-\mu)\mu - al$, which is (17). The right side is positive because $1 - \mu > a$ and $\mu > l$. Part (iv) is Proposition 1 with $N = \rho = 1$. □

A.4 Recorded price variation

Proof of Corollary 1. After publication, both sides withdraw in state h, so $Q_N = 0$. In state l, every feasible history and every next-order posterior through opportunity N lies in $[-c_N, c_N] \subset (-b, b)$, so every arriving customer executes at a break-even quote and the recorded price after the first execution equals the public payoff posterior. Before the first execution the recorded price is zero, which is also the posterior. The recorded price is therefore a martingale under the public filtration, and independent nonarrival leaves it unchanged.

Conditional on $\theta = l$, expand

$$p_k^2 - p_{k-1}^2 = 2p_{k-1}(p_k - p_{k-1}) + (p_k - p_{k-1})^2.$$

The cross term has conditional expectation zero by the martingale property. Summing and taking expectations gives

$$\mathbb{E}[Q_N^D \mid \theta = l] = \mathbb{E}[(p_N^D)^2 \mid \theta = l] \le c_N^2.$$

Consequently $\mathbb{E}Q_N^D \le (1-\pi)c_N^2$. In the undisclosed market the first customer, if present, executes at $+\mu$ or $-\mu$. The first squared increment has expectation $\rho\mu^2$, and all later squared increments are nonnegative, giving $\mathbb{E}Q_N^O \ge \rho\mu^2$. Finally, $c_N < \rho\mu$, $\rho \le 1$, and $\pi > 0$ imply $(1-\pi)c_N^2 < \rho\mu^2$.

For root mean square variation, conditional Jensen's inequality and the bound above give $\mathbb{E}\sqrt{Q_N^D} = (1-\pi)\mathbb{E}[\sqrt{Q_N^D} \mid \theta = l] \le (1-\pi)\{\mathbb{E}[Q_N^D \mid \theta = l]\}^{1/2} \le (1-\pi)c_N$, whereas the first undisclosed opportunity gives $\mathbb{E}\sqrt{Q_N^O} \ge \rho\mu$. These bounds are strictly ordered by (16); division by $\sqrt{N}$ preserves the comparison.

For the uncertainty comparison, every final disclosed posterior has magnitude at most c_N, so expected posterior variance is at least $1 - c_N^2$. The first undisclosed record already reduces expected payoff variance to $1 - \rho\mu^2$, and further public information weakly reduces it. The inequality $c_N^2 < \rho\mu^2$ is strict. The stronger Blackwell comparison follows from Proposition 1. □

A.5 Public order information and recentering

Proof of Proposition 3. After publication, θ is known. In the high state, as long as both directions are unavailable, the only payoff information from an arriving customer is Q, whose conditional likelihood is

$$\mathbb{P}(Q = q \mid X = x, \theta = h) = \frac{1 + \eta h q x}{2}.$$

Starting from $z = 0$, every posterior after at most N such signals has absolute value at most c_N^Q. The bound $c_N^Q < d(h, b)$ proves by induction that neither quote is ever feasible. No-arrival observations carry no payoff information. In the low state, any executed direction changes posterior log odds by $2 \operatorname{arctanh} l$, and the bound $c_N < \rho\mu < b$ guarantees that both quotes remain inside the band at every potential execution. Since Q is conditionally independent of (X, θ) given the executed direction, it supplies no additional information there.

Every terminal disclosed posterior therefore belongs to $[-C, C]$, where $C = \max(c_N, c_N^Q) < \rho\mu$. Without publication, the first customer can trade and the first record has posteriors $-\mu, 0, \mu$, with unconditional probabilities $\rho/2, 1-\rho, \rho/2$; the additional Q is redundant when a transaction occurs. If terminal disclosed outcome j has unconditional probability w_j and posterior z_j, set

$$K(j \mid 0) = w_j, \qquad K(j \mid \pm) = w_j \left(1 \pm \frac{z_j}{\rho\mu}\right).$$

These probabilities are nonnegative and sum to one because $\sum_j w_j z_j = 0$. Applying this state-independent garbling to the first withholding record reproduces the entire disclosed experiment. Moreover, $\mathbb{E}_D|z| \leq C < \rho\mu \leq \mathbb{E}_O|z|$, ruling out a reverse garbling. This establishes strict Blackwell dominance.

For the quadratic comparison, $\mathbb{E}_D z^2 \leq C^2 < (\rho\mu)^2 \leq \rho\mu^2 \leq \mathbb{E}_O z^2$. The high state has no price changes after disclosure. In the low state the transaction-price martingale gives $\mathbb{E}Q_N^D \leq (1-\pi)c_N^2$, while the first withholding trade alone gives $\mathbb{E}Q_N^O \geq \rho\mu^2$. Jensen's inequality yields $\mathbb{E}\sqrt{Q_N^D/N} \leq (1-\pi)c_N/\sqrt{N}$, whereas $\mathbb{E}\sqrt{Q_N^O/N} \geq \rho\mu/\sqrt{N}$. Strictness follows from $c_N < \rho\mu$.

Finally, at $\eta = 0$ the conditions reduce to those of Proposition 1, since $c_N^Q = 0 < d(h, b)$. They are continuous in η and therefore hold on a nonempty interval of positive signal quality. □

Perfect and rare exact observation. If $\eta = 1$, publication reveals θ and the complete sequence of customer intentions and arrivals. Using those intentions, the observer can compute every quote and execution under the known withholding rule. Forgetting θ and retaining this simulated record gives a state-independent garbling into the withholding experiment. For rare exact observation, consider $N = \rho = 1$ and reveal Y with probability $a > 0$, otherwise reporting a blank. Under withholding the first trade reveals only Y, so all payoff posteriors have absolute value μ. Publication can produce the posterior $h > \mu$ when the high state and a revealed intention coincide. No garbling of the withholding experiment can produce a posterior outside $[-\mu, \mu]$. For sufficiently small $a > 0$, publication still has higher quadratic

prediction error by continuity from $a = 0$. Hence neither experiment dominates the other. This distinguishes uniformly noisy information from a low-frequency but exact record.

Proof of Proposition 4. Write the finite payoff support as $v_- = v_0 - \Delta > 0$ and $v_+ = v_0 + \Delta$. Conditional on (X, θ), the public signals are independent with mean $\eta\theta X$. Because $\eta > 0$ and $\theta \geq l > 0$, their limiting sample mean identifies the sign of X. The full public history contains these signals, so the posterior probability of the true payoff converges to one, with or without the initial report. Each intended order has conditional probability at least $(1-h)/2 > 0$ under every state. Conditioning on either order direction therefore preserves this concentration, and both break-even quotes converge to V.

Prices remain in $[v_-, v_+]$. If the break-even ask A_t is feasible, its execution price is $\max(A_t, L_t)$; if the break-even bid B_t is feasible, its price is $\min(B_t, U_t)$. The endpoint rule places each such price between the previous close and the corresponding break-even quote. Infeasible orders leave the price unchanged.

Fix $\delta > 0$. After a finite random time, both break-even quotes lie within δ of V. If the close is below $V - 2\delta$, a sell order executes at $\min(B_t, (1+\beta)P_{t-1})$ and raises it by at least $\min(\delta, \beta v_-) > 0$. Other orders either leave the close unchanged or bring it within δ of V. A sell order has conditional probability bounded below by $(1-h)/2$, so infinitely many such opportunities occur and only finitely many are required to reach the 2δ neighborhood. Above $V + 2\delta$, the analogous argument uses a liquidity purchase and the lower endpoint. Once the price enters this neighborhood, the endpoint rule and quotes within δ keep it there. Taking a sequence $\delta \downarrow 0$ proves $P_t \to V$. Since all prices are positive, continuity of the logarithm gives the claim about cumulative log-price changes. For the initial-session comparison, set $a = \Delta/v_0 \in (0,1)$ and let $p = (P_1 - v_0)/\Delta$. An executed price equals the payoff posterior, while no execution leaves $p = 0$. Therefore $\mathbb{E}[X \log(1+ap)] = \mathbb{E}[p \log(1+ap)]$. Symmetry of the price distribution changes the latter expectation to $\mathbb{E}[|p| \operatorname{arctanh}(a|p|)]$. Under withholding $|p| = \mu$, whereas under publication $|p| \leq c_1 < \mu$; the high state has $p = 0$. Since $x \operatorname{arctanh}(ax)$ is strictly increasing for $x > 0$, $G_1 > 0$. Both prices converge to the same positive payoff and remain within $[v_-, v_+]$, so dominated convergence gives $G_t \to 0$. No monotonicity of G_t is required. □

A.6 Optimal disclosure under a fixed band

Proof of Proposition 5. Every report about binary θ induces a posterior mean $m \in [l, h]$ with $\mathbb{E}[m] = \mu$, and every distribution of such means satisfying this Bayes-plausibility condition is implementable. For any convex function ϕ of the final payoff posterior, define

$$\psi(m) = \frac{\phi(m) + \phi(-m)}{2}.$$

The function ψ is convex and nondecreasing for $m \geq 0$. On $[l, b_T]$, it lies below the chord joining $(l, \psi(l))$ and $(b_T, \psi(b_T))$. For $m > b_T$, no order executes and the final payoff posterior is zero, so the conditional value is $\phi(0)$. The same chord, extended beyond b_T, is at least $\psi(l) \geq \phi(0)$. Hence every feasible disclosure has expected ϕ bounded by this chord evaluated at the prior mean μ. The two-point distribution in (25) attains that bound for every convex ϕ. For a binary payoff, dominance for all convex functions of the posterior mean is equivalent to Blackwell dominance.

Taking $\phi(z) = z^2$ gives (26). Both induced means are at or below b_T, so every order executes. The unconditional probability of a liquidity customer is $1 - \mu$ and cannot be increased by disclosure, so completed liquidity-trading surplus is maximal, which together with Blackwell dominance establishes the statement about nonnegative combinations. Finally, a liquidity customer facing posterior mean $m \leq b_T$ arrives with probability $1 - m$ and pays half-spread m, giving expected surplus $\ell(m) = (1 - m)(v - m)$. This function is convex, so the report, whose induced means lie in $[l, b_T]$ and average to μ, yields $\mathbb{E}\,\ell(m) \geq \ell(\mu)$, the surplus under withholding. □

A.7 Optimal band assignment and menu design

Proof of Proposition 6. Let x_l and x_h denote the joint probabilities of observing W together with states l and h. They lie in $[0, 1 - \pi] \times [0, \pi]$. Put

$$w = x_l + x_h, \qquad r = l x_l + h x_h,$$

and, when the relevant category is nonempty,

$$m_W = \frac{r}{w}, \qquad m_T = \frac{\mu - r}{1 - w}.$$

The wide category always trades. The tight category trades exactly when $m_T \leq b_T$, or

$$\mu - b_T + (b_T - l)x_l + (b_T - h)x_h \leq 0. \tag{41}$$

On this polygon, the policy objective is

$$\frac{r^2}{w} + \frac{(\mu - r)^2}{1 - w} + \gamma(1 - \mu) - \kappa w. \tag{42}$$

The two quadratic terms are perspective functions and therefore convex; the remaining terms are affine. A maximum over the polygon is attained at a vertex. Under $l < \mu < b_T < h$, the undominated vertices are pooling under T, full separation, and the selective boundary policy with $x_h = 0$ and $x_l = (b_T - \mu)/(b_T - l)$, yielding J_P, J_S, and J_I.

On the complementary region where the tight category cannot trade, take the closure to bound its attainable value. At an execution boundary the actual policy additionally permits tight-category trading, so its value weakly exceeds the inactive limiting value. Within the inactive region the objective is

$$\frac{r^2}{w} + \gamma(w - r) - \kappa w,$$

which is again convex. Its relevant vertices are the selective boundary, pooling under W, and assignment of every low state to W while every high state remains in an inactive T category. Pooling under W is dominated by pooling under T, which delivers identical information and execution without the widening cost. The remaining inactive policy has value no greater than J_P because $\mu^2 \geq (1 - \pi)l^2$ and $1 - \mu \geq (1 - \pi)(1 - l)$. No other stochastic assignment can therefore exceed the envelope in (31). □

Proof of Corollary 2. Define the affine function

$$L(m) = m(l + b_T) - lb_T + \gamma(1 - m).$$

For any executing category with posterior mean $m \in [l, b_T]$, convexity of m^2 gives

$$m^2 + \gamma(1 - m) \leq L(m).$$

An inactive C or T category yields zero and also lies below L. For an executing wide category,

$$m^2 + \gamma(1 - m) - \kappa - L(m) = (m - l)(m - b_T) - \kappa \leq 0$$

whenever $\kappa \geq (h - l)(h - b_T)$. Thus every policy has value at most $L(\mu)$. Assigning posterior l to C and posterior b_T to T attains the bound, with both categories executing. Since $b_C > l$, C is strictly nonbinding in the states assigned to it. Relative to selective relaxation, the same posterior distribution and execution are obtained without paying the widening cost on the low-state message. □

Proof of Corollary 3. A report and band jointly induce a posterior mean $m \in [l, h]$. For an executing category the value is $m^2 + \gamma(1 - m)$, less κ if W is used. Since the wide band provides no additional execution when $m \le b_T$, it is weakly dominated by T on that interval. On each of $[l, b_T]$ and $[b_T, h]$, the executing payoff is convex in m, so its upper concave envelope uses only l, b_T, h. At b_T, the tight category dominates the wide one. The slope from l to b_T is $l + b_T - \gamma$; the slope from b_T to h, net of widening, is $h + b_T - \gamma - \kappa/(h - b_T)$. The former is at least the latter exactly when $\kappa \ge (h - l)(h - b_T)$. Because the prior satisfies $\mu < b_T$, the optimal experiment then uses l, b_T and has value J_D. For the reverse inequality the chord from l to h dominates and gives J_S. Inactive categories cannot improve either bound: for the former case the affine bound $m(l + b_T) - l b_T + \gamma(1 - m)$ is positive on $[l, h]$; in the latter case the full-separation chord is positive, since $\kappa < \kappa_D < h^2$. Their zero payoffs therefore lie below the relevant bound.

Selective relaxation has value $J_I = J_D - \kappa x$, establishing the stated gain. For $\kappa > \kappa_D$, $J_S < J_D$; also $J_P < J_D$ and $J_I < J_D$ because $l < \mu < b_T$ and $\kappa x > 0$. Proposition 6 then gives strict dominance over every band-only assignment. □

A.8 Manipulation and productive financing

The break-even price and participation threshold used in Section 6.4 are

$$\bar{p} = \frac{c}{d - 1}, \qquad \omega^* = \frac{1 - \bar{p}}{1 + \bar{p}}. \tag{43}$$

The trading equilibrium is

$$(\sigma(\omega), p(\omega)) = \begin{cases} \left(1, \dfrac{1 - \omega}{1 + \omega}\right), & 0 < \omega \le \omega^*, \\ \left(\dfrac{(1 - \omega)(1 - \bar{p})}{2\omega\bar{p}}, \bar{p}\right), & \omega^* < \omega < 1. \end{cases} \tag{44}$$

Informed customers accept a legal break-even quote when indifferent. The loss of information from the tight band is

$$I_W - I_T = \begin{cases} 0, & \omega \le \omega^*, \\ (1 - \omega)\left\{\dfrac{1 + \bar{p}}{2} - \dfrac{1}{1 + \omega}\right\}, & \omega > \omega^*. \end{cases} \tag{45}$$

With settlement payment s and external funding D, feasible project scale satisfies

$$k \le w - s + D, \qquad D \le \zeta k. \tag{46}$$

The proofs below establish these expressions and the policy characterization.

Trading equilibrium and public information. Conditional on a buy, the likelihood in the high-payoff state is $(1-\omega)+\omega\sigma$ and that in the low-payoff state is $\omega\sigma$. With equal prior probabilities, Bayes' rule gives (35). A manipulator buys when $(d-1)p-c\geq 0$. If she always buys, the quote is $(1-\omega)/(1+\omega)$, which is at least $\bar{p}$ exactly for $\omega\leq\omega^*$. Above ω^*, always buying would be unprofitable, whereas no manipulation would give a deviator the quote $1>\bar{p}$. The unique equilibrium therefore mixes at the probability solving $p=\bar{p}$, as in (44). At $\omega=\omega^*$ this coincides with the pure strategy. A sale reveals $X=-1$ and occurs at -1. The endpoint cases are understood by continuity; the proposition concerns $0<\omega<1$.

Under a cap $u<\bar{p}$, every permissible purchase is unprofitable for the manipulator. Any purchase accepted at such a price must therefore be from the informed high type and has conditional payoff $1>u$. A dealer who supplied a legal ask would lose money, so no such ask is offered. The sale remains feasible by the maintained lower-bound assumption. A suppressed intended buy and voluntary manipulator inactivity are both recorded as no transaction; no independent observation of the intended order is available.

Under the wide band the public outcomes are sale, buy, and inactivity. Their probabilities conditional on $(X=-1, X=1)$ are respectively

$$(1-\omega,0),\quad (\omega\sigma,1-\omega+\omega\sigma),\quad (\omega(1-\sigma),\omega(1-\sigma)).$$

Under the tight band the outcomes are sale and no transaction, with respective probabilities $(1-\omega,0)$ and $(\omega,1)$. Merging the wide buy and inactivity outcomes therefore reproduces the tight experiment. This is an exact state-independent garbling, so an agent making any decision solely from this public record cannot prefer the tight experiment on informational grounds.

For $\omega\leq\omega^*$, wide-regime inactivity has probability zero and the experiments are equivalent. Both yield $I=(1-\omega)/(1+\omega)$. Above ω^*, a wide-regime sale has probability $(1-\omega)/2$ and posterior -1; a buy has probability $(1-\omega)/(2\bar{p})$ and posterior $\bar{p}$; and inactivity has posterior zero. Hence $I_W=(1-\omega)(1+\bar{p})/2$. Under the tight band, a sale still reveals -1 while the no-transaction posterior is $(1-\omega)/(1+\omega)$, giving $I_T=(1-\omega)/(1+\omega)$. Subtraction yields (45). □

Financing and consolidated welfare. With high effort a project of scale k returns Rk. Diverting instead gives the entrepreneur private benefits $(R-\zeta)k$ and no verifiable output. A repayment D is incentive compatible only if $Rk-D\geq(R-\zeta)k$, or $D\leq\zeta k$. At equality the entrepreneur chooses productive effort. Competitive lenders advance funds at the unit

outside return, yielding (46). Since $R > 1$, the entrepreneur exhausts the financing limit and chooses $k(s) = (w - s)/(1 - \zeta)$. The condition $w > d$ ensures positive investment for every possible settlement payment $s = dp \leq d$.

Conditional on manipulation, the entrepreneur pays dp. In expectation the dealer earns p on the spot transaction, the manipulator earns $(d-1)p - c$, and the broker receives c. These receipts sum to dp, so the financial transfer has no direct aggregate welfare cost. Investment displaced from the productive entrepreneur is available to the outside technology, which returns one per unit. Consolidated productive surplus therefore falls by $(R-1)\{k(0) - k(dp)\} = \Lambda p$, not by $dp + \Lambda p$. Fees are transfers and do not generate an additional term proportional to cM.

For the operating decision, conditional expected surplus is $\nu(2eZ - e^2)$, where $Z = \mathbb{E}[X \mid \mathcal{F}_a]$. Maximization gives $e = Z$ and ex-ante surplus νI_a. Combining its loss with avoided settlement-induced investment losses and the implementation cost gives (37). □

Proof of Proposition 7. For $\omega \leq \omega^*$, manipulation is pure and the information loss is zero, so

$$G(\omega) = \Lambda f(\omega) - F, \qquad f(\omega) = \frac{\omega(1-\omega)}{1+\omega}. \tag{47}$$

The derivatives are

$$f'(\omega) = -1 + \frac{2}{(1+\omega)^2}, \qquad f''(\omega) = -\frac{4}{(1+\omega)^3} < 0.$$

Thus f increases up to $\sqrt{2} - 1$ and decreases thereafter. For $\omega > \omega^*$, put $A = \Lambda(1 - \bar{p})/2$. The gain is

$$G(\omega) = (1-\omega)\left[A - \nu\left\{\frac{1+\bar{p}}{2} - \frac{1}{1+\omega}\right\}\right] - F. \tag{48}$$

Here

$$G''(\omega) = \frac{4\nu}{(1+\omega)^3} \geq 0, \qquad G(1) = -F, \qquad G(\omega^*) = \Lambda\omega^*\bar{p} - F > -F.$$

Consequently no point in the mixed region exceeds $G(\omega^*)$, irrespective of ν. Combining this with strict concavity on the pure region yields the unique maximizer in (38). The maximal information loss is zero at that state, so its net value is positive exactly under (39).

Since $G(0) = G(1) = -F < 0$, positivity at $\widehat{\omega}$ implies a lower crossing in the increasing part of (47). It is unique. The decreasing part of that strictly concave branch has at most one further crossing. If it crosses before ω^*, convexity of (48) and its nonpositive endpoints exclude any later positive region. Otherwise $G(\omega^*) > 0$; convexity of the mixed branch and $G(1) < 0$ imply a unique crossing there, followed by negative values through one. The boundary case

$G(\omega^*) = 0$ is handled by continuity. This proves the exact interval characterization without an upper bound on ν.

The manipulation frequency peaks at ω^*. The welfare peak precedes it precisely when $\omega^* > \sqrt{2} - 1$, equivalent to $\bar{p} < \sqrt{2} - 1$. Finally, $\partial G/\partial \Lambda = Mp > 0$ for interior states, while $\partial G/\partial \nu = -(I_W - I_T) \leq 0$. Their effects on the intervention set follow pointwise; neither changes the maximizing state because the information term is zero throughout the pure-manipulation region. □

B Data Construction and Estimation

B.1 Band histories, calendar, and returns

The band history contains 17,659 records for 1,248 issuer codes during 2020–2025. When an issuer has several records on a date, the final recorded state determines its assignment. This gives 16,762 dated changes, of which 16,756 can be linked to prices. The price history contains 1,262,554 issuer-days for 1,247 issuers between January 2019 and March 2026. Of these issuers, 1,246 have prices in the 2020–2025 event period. Observations outside that period supply preceding histories and forward returns.

A recorded band change takes effect in the next equity session. Before the first change, the old band in that record supplies the preceding assignment wherever prices are available. The age of this initial assignment is measured from the first observed price; after a change, age is measured in equity sessions since the recorded date. Thus the stability restrictions use the length of the observed assignment history.

The equity calendar is the union of dates in the price data and contains 1,492 sessions during 2020–2025. Estimation excludes March 1–4 and May 17–21, 2024, around the special operational sessions. The market benchmark is the mean of valid adjusted log returns weighted by the previous session's market capitalization. Absolute daily log changes above 0.35 are excluded from the benchmark. Its daily values are cumulated on the equity calendar, and each stock outcome requires the specified session endpoints. Returns exclude dividends.

Capital adjustments are inferred from offsetting changes in shares and prices. A candidate requires a share change exceeding 10 percent and a new-to-old share ratio between 0.05 and 60. From two observations before the share change through one afterward, we identify the price movement closest in logs to the inverse share change. A factor is retained when that distance is at most 0.15 and the absolute log price change is at least 0.075. Adding the cumulative log factor to prices from the selected date onward gives the adjusted series. The rule identifies 215 factors. Excluding revisions within twenty sessions of a factor removes

94 fitted observations from the broader twenty-session return comparison, leaving 13,901 observations and a coefficient that rounds to −10.90 points.

Daily turnover is traded quantity divided by shares outstanding. Size is the logarithm of closing price times shares outstanding. The preceding return, volatility, and turnover controls use observations available before the revision date.

B.2 Daily prices and identifiers

Daily ranges are obtained from the EOD2 archive of NSE Bhavcopy files, supplemented by direct exchange files for February 1 and November 14, 2020, and May 11, 2022. The records use the capital-market format before the July 2024 change to CM-UDiFF and the latter format thereafter. Together they cover every scheduled date in the panel.

Contemporaneous symbols identify 1,044,710 security-days, and historical ISINs identify another 29,232. The resulting 1,073,942 matches cover 99.617 percent of the 1,078,071 eligible observations. The 4,129 remaining security-days, spread across fourteen issuers, are left missing. Agreement with the closing-price history selects among candidate identifier matches, so close agreement within the selected records follows partly from that matching rule. Table 13 gives annual coverage.

The Bhavcopy supplies open, high, low, close, previous close, and volume. The reported daily close is used as a closing-price observation; it does not identify the time or order-book conditions of the last individual trade.

B.3 Daily endpoints and boundary validation

With tick size τ_{it}, band width b_{it} in percent, and the exchange reference close P_{it}^{prev}, the reconstructed endpoints are

$$\widehat{U}_{it} = \tau_{it} \left\lfloor \frac{P_{it}^{\text{prev}}(1 + b_{it}/100)}{\tau_{it}} \right\rfloor, \qquad \widehat{L}_{it} = \tau_{it} \left\lceil \frac{P_{it}^{\text{prev}}(1 - b_{it}/100)}{\tau_{it}} \right\rceil. \tag{49}$$

Boundary comparisons allow one tenth of the assigned tick, with a minimum tolerance of 10^{-7}. The tick is Rs. 0.05 before June 10, 2024, and subsequently Rs. 0.01 below Rs. 250 and Rs. 0.05 above that threshold. From April 15, 2025, the higher-price tiers are Rs. 0.05 through Rs. 1,000, Rs. 0.10 through Rs. 5,000, Rs. 0.50 through Rs. 10,000, Rs. 1 through Rs. 20,000, and Rs. 5 thereafter. The last available prior-month close assigns the tier. Where that reference is missing, the closing-price observation for the current session is used.

The NSE circulars specify tick alignment, applicable series, and review dates (National Stock Exchange of India, 2024c, 2025). Equation (49) estimates the legal range from those

inputs. Series-specific assignments, changes in reference prices, and rounding conventions can produce differences from actual exchange endpoints. Regressions using boundary location therefore require positive prices and volume, valid OHLC ordering, $\widehat{L} < \widehat{U}$, and an observed range inside the reconstructed interval up to the tolerance. Observations failing these conditions have an unresolved boundary location and are excluded from those regressions. The return regressions in Section 4.1 do not impose the intraday-range restriction.

Official hit reports provide a comparison on the dates for which those reports are available. Agreement is 97.21 percent for upper hits and 96.94 percent for lower hits before the price-range consistency restriction, and 97.82 and 98.25 percent afterward. Table 14 reports sensitivity and precision as well as overall agreement. These comparisons assess reconstructed boundary locations against recorded hits; official endpoints are not independently observed for every security-day.

B.4 Other surveillance measures

ASM, GSM, and Trade-for-Trade histories combine membership lists, stage-change circulars, and daily regulatory-indicator reports. The circular collection has 6,258 metadata entries, and 2,856 ASM/GSM documents supply the relevant membership and transition information, including December 2019 observations establishing initial status. Machine-readable annexures provide long- and short-term ASM membership, GSM lists, and dated transitions. Scanned GSM circulars contribute an identifier only when it matches an ISIN exactly. Membership and stage are recorded separately. A membership criterion is not treated as a stage, and an unobserved stage remains unknown.

Transitions use explicit effective dates when available and the next eligible session otherwise. Daily regulatory-indicator reports take precedence over status inferred from preceding documents. The 482 report dates establish 481 following-session dates within the sample. Trade-for-Trade membership is identified from contemporaneous series BE, BZ, ST, and SZ. The history contains 20,389 transitions: 12,795 in ASM, 1,159 in GSM, and 6,435 in Trade-for-Trade, each linked to its publication and effective dates.

For each band revision, we record membership immediately before the revision, membership on that date, and changes within one, five, and twenty equity sessions on either side. Changes require adjacent sessions with known status and include recorded ASM/GSM stages and Trade-for-Trade series. The revision regressions use current program membership and revision-date transitions as controls; unknown membership or transition status remains missing rather than being recoded as inactivity.

B.5 Volatility and the shape of subsequent returns

Future variation uses consecutive daily returns beginning in the session after the recorded revision. Missing prices remain missing, and daily returns require adjacent equity-session closes. Each primary outcome uses a complete path. Among revisions beginning at 5 or 10 percent, 87.2 percent of tightenings and 93.3 percent of loosenings have complete twenty-session paths. Allowing at least eighteen observed one-session returns and dividing squared variation by their actual count gives an annualized coefficient of −36.28 points. Figure 2 uses the same complete sixty-session paths for all disjoint windows.

The root mean square in (3) includes the realized mean. The alternative annualized sample standard deviation is

$$\widehat{s}_{e,H} = \left[\frac{252}{H-1} \sum_{j=1}^{H} (e_{i,t+j} - \bar{e}_{i,H})^2 \right]^{1/2}.$$

The identity $\sum_j e_j^2 = \sum_j (e_j - \bar{e})^2 + H\bar{e}^2$ separates dispersion around the within-window mean from the squared cumulative movement. These measures describe daily variation along each path. They do not estimate the variance of the cumulative H-session return across events. Multiplying by $\sqrt{252}$ expresses daily RMS or SD in annualized units without assuming independent daily returns. The sampling frequency remains daily.

Directionality is $|\sum_j r_j| / \sum_j |r_j|$, using adjusted raw log returns and leaving zero-total-movement paths undefined. It lies in $[0, 1]$. Sign persistence is the share of adjacent session pairs with the same nonzero sign. Pairs containing a zero return are omitted without joining observations across them; two complete paths have no eligible pair. The negative-day share counts negative returns over all twenty sessions.

Tail cutoffs use the pooled interior-band sample and symmetric 1 or 5 percent trimming and winsorization. Returns calculated directly from the Bhavcopy use the reported close relative to the reported previous close. That comparison excludes inferred capital actions within the outcome window and two adjacent sessions and requires all twenty daily returns. The range comparison uses the RMS of log(High/Low) without converting it to a diffusion-based volatility measure. The common-date comparison requires both revision directions within an initial-band/date cell after imposing complete-path availability. The specification excluding subsequent band revisions conditions on later regulatory decisions. Table 15 reports the principal comparisons.

B.6 Trading activity and pressure

Daily share turnover is quantity divided by shares outstanding, and traded value is quantity times the closing price. Table 7 uses the logarithm of twenty-session mean turnover. A session with no intraday range has positive volume and $H = L = C$. The window share counts these recorded sessions out of all twenty dates; a missing OHLC record contributes no recorded such event. Requiring complete reconciled OHLC paths provides a separate check on this missingness. The intraday range is $\log(H/L)$. Activity regressions require complete twenty-session return paths.

The measure of a prospective tightening uses only conditions before the revision. The candidate tighter band is 2 percent for a stock starting at 5 percent and 5 percent for one starting at 10 percent. Trailing twenty-session volatility divided by the candidate width is standardized within the initial-band category. The regressions include its level, its interaction with tightening, and the public-state controls.

B.7 Industry classifications and peer returns

The primary industry labels are the `smIndustry` field accompanying the NSE's corporate financial-results announcements. Securities are linked to these labels by ISIN, dated symbol, and normalized company name. Requiring at least six classified securities per industry gives 209 firms in 19 industries. The labels are fixed throughout the estimation period and need not coincide with the classification available at each historical date.

The alternative uses CMIE Prowess industry classifications for small and medium-sized listed enterprises. Each security receives the most frequent `industry_group` label across its fiscal-year records, with alphabetical resolution of frequency ties. Of the securities in the daily-price universe, 500 have a classification and 332 belong to industries with at least six members. Daily peer availability and complete outcomes give 305 firms in 24 industries in the fitted comparison. This classification provides a distinct definition of industry peers and covers a different set of securities. The two classification systems are applied separately.

Daily adjusted log returns use the closing-price histories and inferred capital factors described above. They require consecutive equity sessions, with absolute log changes above $\log(1.35)$ treated as missing in the peer-return analysis. Peers are selected before imposing the focal stock's control or forward-return availability. The equally weighted proxy averages the adjusted log returns of at least five other securities in the same industry. The value-weighted alternative uses the previous session's capitalization and requires at least five other usable weighted observations. The median alternative takes the median after removing the focal return. The focal security is excluded from each proxy.

Pressure in (4) divides the absolute simple-return equivalent of the peer log return by the reconstructed capacity in its direction. Capacities no greater than 10^{-12} in return units are treated as zero and excluded. A comparison using nominal band width assesses the importance of tick adjustment. The 209-firm classification yields 276,918 daily records, 258,303 with a current stock return and usable peer return, and 254,878 with positive capacity. Complete controls and all forward horizons leave 248,832 observations. The common sample makes the cumulative and future-only coefficients comparable through the return-addition identity. Price-location regressions additionally require the observed daily range to lie within the reconstructed endpoints.

Predetermined controls are the previous daily return, sums over the previous 5 and 20 equity sessions requiring at least 3 and 10 returns, trailing twenty-session sample SD and mean turnover each requiring ten observations, and the previous session's log capitalization. Flexible peer-return controls add $\max\{|s| - k, 0\}$ and $\text{sign}(s)\max\{|s| - k, 0\}$ for $k \in \{0.0025, 0.005, 0.01, 0.02\}$. The pressure intervals, thresholds, and local window in Table 16 are fixed across outcomes. These comparisons assess the shape of the relationship; their thresholds do not determine regulatory assignment.

For the common-peer comparison, security identifiers are sorted within each industry and assigned alternately to two fixed groups. Each focal group uses only returns from the opposite group, and industry-date-focal-group effects hold that observed peer return fixed. Requiring at least three source firms gives 253,658 complete observations and 5,481 high-pressure observations. Of these, 17,147 lie in 3,250 cells containing both pressure states. The specifications reported in Table 8 also show the result with at least five source firms and with prior ASM, GSM, and Trade-for-Trade membership included as controls.

The main covariance estimator clusters by security and exact date, subtracting their intersection. Each cluster component uses $G/(G-1)$ and a residual-degree-of-freedom correction based on slope rank and fixed-effect levels. Confidence intervals use the smaller cluster count minus one degrees of freedom. Clustering by security and industry-date or by industry and date retains the signs and similar uncertainty for the principal interactions. Exact-date clustering permits dependence across industries within the same day; industry-date clustering imposes a narrower same-day grouping.

B.8 Revision regressions and forecasts

The common-initial-band return regression has firm, exact-date, and initial-band effects and linearly standardized preceding controls. The pre/post figure instead absorbs firm, exact-date, and initial-band effects without preceding-return controls, using firm/date clustering. Its

cohort must have a complete path from session −20 through 60. The return-persistence comparisons retain the linear preceding controls and cluster by firm/date. The headline return comparison clusters by firm and exact revision date. Program-status comparisons use the same return splines and remaining public-state controls with firm and exact-date effects and firm/date clustering. Their status controls are current membership and revision-date transition indicators. These regressions require known preceding membership and use it when current membership is missing. Missing transition indicators are excluded by the regression sample rule.

Forecasts train only on earlier calendar years. Missing continuous predictors receive training-sample medians, and observations whose forward return is unavailable at the start of the test year are excluded from return-forecast training. The public predictors are identical with and without revision direction. Ridge uses penalty 10 and training-sample standardization of all predictors, including numerical band and month codes. The classification model treats band and month as categorical. Extra Trees uses 400 trees, minimum leaf size eight, feature fraction 0.8, and the same fixed random seed for both forecasts.

B.9 Anticipated and unanticipated decisions

The predicted probability $\hat{p}_e$ comes from the recursive regularized logit for revision direction using the market-variables specification of Panel A of Table 2, which excludes the old band and calendar month so that predictions reflect market conditions rather than the mechanics of the menu. Each test year is predicted from a model trained only on earlier years, so predictions are available for 2021–2025. The sample comprises revisions from an initial 5 or 10 percent band with available predictions and outcomes. Panel A of Table 5 forms terciles from the pooled distribution of $\hat{p}_e$ across eligible revisions and then reports the two revision directions separately after partialling out firm, exact-date, and initial-band effects. Panel B adds $\hat{p}_e$ and $T_e \times \hat{p}_e$ to (2) with the full public-state controls; because $\hat{p}_e$ is a nonlinear function of predetermined variables, identification relies on that nonlinearity. Adding current ASM, GSM, and Trade-for-Trade membership and revision-date transitions leaves the five-session interaction at −6.01 (standard error 1.66).

B.10 Clipping benchmark

The benchmark sample requires valid reconstructed endpoints, an observed range within them, agreement within 10^{-6} between the adjusted daily log return and log(close/previous close), complete outcomes through session ten, no inferred capital action in that window, and no special operational session. This separates endpoint clipping of the economic return from

corporate-action adjustments. The unit-response rule has 242,127 observations. Recursive rules have 200,439 observations in 2021–2025 because 2020 supplies initial training data.

For each test calendar year, estimation uses only previous years under a 20 percent band with $q < 0.5$. The last ten sessions before the test year are purged so forward availability and the capital-action exclusion are known by the training cutoff. The fitted mean is linear in the peer return, preceding returns over 1, 5, and 20 sessions, trailing volatility, $\log(1 + 100\,\text{turnover})$, and lagged capitalization. It includes industry indicators and interactions of the peer return with volatility, transformed turnover, and size. Continuous regressors are standardized using training observations.

The distributional rule adds 199 evenly spaced interior empirical residual quantiles to the fitted mean and clips every resulting return at the focal endpoints before averaging. Residuals use the earlier training industry when it has at least one hundred observations, otherwise the pooled training sample; their discrete mean is removed. This benchmarks the distribution of idiosyncratic movements rather than clipping only their conditional mean. The standard errors in Table 10 condition on this estimated rule, so they do not incorporate estimation uncertainty from fitting the counterfactual distribution. Neither reference law identifies an unconstrained return without the transport assumption described in Section 4.4.

B.11 Band persistence and return reversal

Return comparisons at twenty and sixty sessions use a common complete sixty-session cohort. Additional returns over sessions 21–60 are constructed as the difference of the cumulative outcomes on that cohort. Persistence requires no further recorded band change between the event and the outcome horizon. These indicators are evaluated only when the entire horizon lies within the band-history coverage ending in December 2025; later price availability alone cannot establish later regulatory persistence. The persistent-band restriction conditions on later decisions.

C Supplementary Empirical Results

The tables in this appendix give the timing and measurement comparisons used in the main analysis and the sensitivity of the return and volatility results.

C.1 Revision persistence and the return horizon

Table 11 supports the distinction between the return horizon and the duration of the revised rule in Section 3.4.

Table 11: Return reversal and persistence of the revised band

A. Persistence of the revised band			
Horizon	Tightening: band unchanged	Loosening: band unchanged	Events
5 sessions	99.7%	87.7%	9,964
10 sessions	98.5%	79.5%	9,940
20 sessions	93.5%	67.9%	9,759
60 sessions	37.2%	27.2%	9,450
B. Conditional return differences on the same 8,409 complete sixty-session paths			
Sessions 1–20	-13.90 (1.50)		8,409
Sessions 21–60	6.56 (1.66)		8,409
Sessions 1–60	-7.34 (2.06)		8,409
C. Twenty sessions with no subsequent revision			
Cumulative return	-6.78 (1.47)		6,923
Annualized daily RMS	-22.89 (1.75)		6,923

Notes. Panel A reports the proportion without a further revision, requiring the horizon to lie within observed regulatory coverage. Panel B uses the same 8,409 complete sixty-session paths and the linear controls, security, date, and initial-band effects of the main revision comparison. Panel C conditions on no subsequent revision; it is a selected subset rather than a causal specification. Returns are log percentage points and volatility is annualized percentage points. All regression standard errors are clustered by firm and exact date.

C.2 Timing, coverage, and boundary validation

The empirical design requires the revision date to precede the new trading range and the reconstructed daily endpoints to match the exchange's observed boundary events. Tables 12–14 document these inputs.

Table 12: Circular matches for revision timing

Security	Revision	Recorded / circular date	Effective date	Context
SILLYMONKS	$10 \rightarrow 5$	Dec. 8, 2023	Dec. 11, 2023	GSM Stage I
HCL-INSYS	$10 \rightarrow 5$	Jan. 5, 2024	Jan. 8, 2024	GSM Stage I
HINDMOTORS	$5 \rightarrow 10$	Mar. 11, 2024	Mar. 12, 2024	Exit from GSM

Notes. The recorded date equals the circular date in these directly matched examples, with the revised band effective in the following trading session. Sources: National Stock Exchange of India (2023a, 2024a,b). The remaining revision dates come from the archived band history.

Table 13: Coverage of daily prices

Year	Sessions	Eligible days	Matched days	Coverage (%)
2020	252	126,972	124,893	98.363
2021	248	136,871	135,604	99.074
2022	248	152,692	152,082	99.601
2023	246	177,057	176,890	99.906
2024	249	225,201	225,195	99.997
2025	249	259,278	259,278	100.000
Total	1,492	1,078,071	1,073,942	99.617

Notes. Eligible security-days in the mapped universe, OHLC matches, and residual missing observations. All 1,492 expected dates are represented. Contemporaneous symbol or historical ISIN matching supplies the daily range; the closing-price histories described in Section 2.3 supply return outcomes and preceding controls.

Table 14: Comparison with official intraday-hit reports

	N	True positives	Agreement (%)	Sensitivity (%)	Precision (%)
All matched report rows					
Upper	91,319	11,991	97.21	93.42	87.54
Lower	91,319	8,223	96.94	93.59	78.64
Endpoint-consistent report rows					
Upper	88,095	11,199	97.82	93.22	90.98
Lower	88,095	6,514	98.25	92.32	86.74

Notes. Matched observations on the predetermined official report dates, shown before and after price-range consistency. True positives are hits recorded by both classifications. Sensitivity is the share of official hits detected by the reconstruction; precision is the share of reconstructed hits confirmed by the report. The second panel excludes unresolved records and does not reclassify them as no hits.

C.3 Price-variation robustness

Table 15 shows that the lower post-tightening daily variation reported in Table 6 is not specific to one return definition, tail treatment, clustering convention, or overlap with other recorded surveillance measures.

Table 15: Robustness of subsequent volatility differences

Specification	Volatility difference	Observations
Linear public-state controls	-36.04 (1.66)	9,058
Spline controls for preceding returns	-37.54 (1.62)	9,058
Standard deviation, demeaned	-37.78 (1.68)	9,058
Raw returns	-40.83 (1.77)	9,058
Trim both tails at 5 percent	-26.26 (1.33)	8,152
At most two zero-return days	-37.16 (1.59)	8,191
Same initial band and date, both directions	-34.45 (2.38)	3,470
No subsequent band revision	-22.92 (1.71)	7,085
Exchange close/reference-close returns	-41.01 (1.80)	8,845
Recorded surveillance controls	-35.67 (1.58)	9,037
Inactive prior status, no recorded nearby change	-12.34 (5.03)	2,122

Notes. Twenty-session tightening-minus-loosening coefficients in annualized volatility percentage points. The baseline uses complete paths, the linear public state, and firm, date, and initial-band effects. Each row changes only the stated feature, retaining firm/date clustering. The common-date row requires both directions within each initial-band/date cell and absorbs that interaction. The exchange-return row uses raw exchange close/reference-close returns on complete paths without inferred capital actions. Program controls comprise membership and revision-date changes in ASM, GSM, and Trade-for-Trade.

C.4 Validation of the peer-return comparison

Table 16 reports the alternative constructions discussed in Section 4.3.

Table 16: Sensitivity of the timing of peer-return responses

Specification	Same session	Sessions 1–10	Sessions 0–10	N
Flexible peer-return magnitude	-0.479 (0.029)	0.557 (0.074)	0.078 (0.071)	248,832
Value-weighted peers	-0.346 (0.027)	0.344 (0.054)	-0.002 (0.052)	248,832
Median peer return	-0.418 (0.032)	0.435 (0.082)	0.017 (0.077)	248,832
At least ten peers	-0.530 (0.041)	0.606 (0.110)	0.077 (0.121)	123,402
Band unchanged for twenty sessions	-0.420 (0.039)	0.464 (0.093)	0.044 (0.097)	163,872
Range-consistent prices	-0.465 (0.030)	0.555 (0.082)	0.089 (0.079)	243,251
Threshold $q \geq 0.75$	-0.383 (0.026)	0.434 (0.065)	0.051 (0.063)	248,832
Threshold $q \geq 1.25$	-0.549 (0.029)	0.632 (0.112)	0.083 (0.110)	248,832
Threshold $q \geq 1.5$	-0.583 (0.030)	0.625 (0.136)	0.042 (0.131)	248,832
Local window $0.75 \leq q \leq 1.25$	-0.135 (0.026)	0.295 (0.151)	0.160 (0.159)	6,431
Continuous pressure interaction	-0.426 (0.029)	0.510 (0.067)	0.084 (0.062)	248,832
Alternative industry classification	-0.228 (0.038)	0.324 (0.135)	0.097 (0.141)	62,344

Notes. Differences in peer-return loadings; the continuous row instead interacts the peer return with pressure. Outcomes within a row share the same sample. All regressions include security, date, and band effects and predetermined controls, with security/date clustered standard errors. Flexible and local specifications add piecewise-linear signed-magnitude controls. Value weights use lagged capitalization. The alternative classification uses CMIE Prowess industry labels for small and medium-sized listed enterprises and covers 305 firms. Appendix B.7 defines the specifications.