Guidance That Never Closes: Measuring Abandoned Management Commitments in Earnings Calls

Author. Aravind R **Affiliation.** Inve (InveWise LLP), Independent Researcher **Correspondence.** dev@inve.money **ORCID.** 0009-0002-7413-248X

Date written: 4 September 2026 **Version:** 1.0 — Working paper. Not peer reviewed. **Data vintage:** All figures are a single snapshot of the corpus taken on 4 September 2026 and will change as the corpus grows. See §7.7.

Abstract

Management commitments made on earnings calls are usually studied at the firm level: a firm either issues guidance or it stops. This paper measures guidance at the level of the individual commitment and documents a form of disclosure attrition that firm-level designs cannot observe.

Using 8,220 earnings-call transcripts from 1,654 Indian listed companies, we extract 24,328 discrete forward-looking management commitments and track each one forward across every subsequent call by the same firm, assigning it a state from a closed twelve-state vocabulary.

We find that 2,271 commitments — 9.3% of all commitments, and 13.5% of the 16,842 that are old enough to have resolved — are never referred to again on any later call. They are not achieved, missed, withdrawn, or revised; they simply stop being mentioned. We term this *guidance abandonment*. It is neither rare nor concentrated: 1,000 of the 1,648 firms with a tracked commitment (60.7%) have at least one. Over the same corpus, the kept rate — commitments achieved, on track, achieved on diluted terms, or revised upward, as a share of resolvable commitments — is 57.8%.

We report the complete state distribution so that any rate can be recomputed under a reader's preferred definition, and we state two unbounded biases: short coverage depth and a vocabulary with no residual category both plausibly push the measured rate upward. We make no claim about intent: abandonment is a property of the earnings-call record, not a finding about managerial motive.

Keywords: earnings conference calls; management guidance; guidance abandonment; voluntary disclosure; forward-looking statements; disclosure attrition; managerial credibility; India; SEBI LODR; text analysis

JEL classification: G14, G34, M41, M48, D82, C81, G15

---

1. Introduction

A management team stands on an earnings call and says the company will reach a 4–5% margin in that business "in a few quarters." Two quarters later they repeat it. Then the target is never mentioned again —

not achieved, not missed, not withdrawn, not restated as something else. The next four calls proceed as though it had never been said.

The disclosure literature has a well-developed account of two adjacent phenomena. The first is **guidance cessation**: a firm publicly stops issuing quarterly earnings guidance, an announced, dateable, firm-level event whose causes and consequences have been studied at length (Houston, Lev, and Tucker 2010; Chen, Matsumoto, and Rajgopal 2011), including the mass withdrawals of early 2020 (Aaron, Kang, Ng, and Rusticus, SSRN No. 3794964). The second is the **non-answer**: an analyst asks a question during a call and management declines to answer it, a within-call event that occurs in approximately 11% of analyst questions and is systematically related to firm characteristics and managerial incentives (Gow, Larcker, and Zakolyukina 2021; Hollander, Pronk, and Roelofsen 2010).

Neither construct captures the episode above. The firm did not stop guiding — it kept issuing guidance on every subsequent call. No analyst was refused an answer at the moment of asking. What happened is that one specific, forward-looking, checkable commitment entered the record and never left it in any direction. It did not close.

This is invisible to a firm-level design, because the firm's guidance behaviour looks continuous. It is invisible to a call-level design, because within any single call nothing unusual occurs. It becomes visible only if the unit of observation is the *individual commitment*, carried forward across every subsequent call by the same firm until it resolves or the record ends.

That is what this paper does. We construct a commitment-level panel from 8,220 Indian earnings-call transcripts covering 1,654 listed companies, extract 24,328 discrete forward-looking commitments, and assign each one a state at the most recent call on which its status can be determined. Our central descriptive finding is that a substantial share of management guidance never closes at all.

Contribution. Three things.

First, a construct and a measurement. We define guidance abandonment — a commitment that has survived at least one subsequent earnings call and that no subsequent call in the corpus resolves, revises, reiterates, or references — and we measure its incidence at 13.5% of resolvable commitments. It is a different object from the within-call non-answer: it is silence *across* calls rather than silence *within* one, and unlike a non-answer it leaves no trace at the moment it occurs. Nobody has to ask a question for a commitment to be abandoned. It is a disclosure event with no observable timestamp, which is precisely why it has not previously been counted. We note that our rate and the ≈11% non-answer rate of Gow et al. (2021) happen to be of similar order, but they are computed over different denominators — commitments in our case, analyst questions in theirs — and should not be read as comparable quantities.

Second, a base rate. Commitment-level disclosure attrition has, to our knowledge, no published incidence estimate. We supply one, along with the complete state distribution behind it, so the figure can be reproduced, contested, and recomputed under alternative definitions.

Third, a setting with an unusually complete record. India's disclosure regime since 2021 produces a near-census of earnings-call transcripts for firms that hold calls (§3.1), and Indian guidance is largely narrative

and operational rather than a single quarterly EPS figure — which means commitment-level tracking is not merely a refinement of firm-level measurement here, it is the only way to measure guidance at all.

What this paper is not. It is descriptive. We do not identify a causal effect, we do not estimate a market response, and we make no claim about managerial intent. A commitment recorded as abandoned may have been abandoned for entirely ordinary reasons: the metric was superseded, the business was reorganised, the person who made the commitment left. Our claim is narrower and, we think, more useful for being narrow: on the record of the earnings call itself — the forum management chose for the commitment, and the forum in which analysts and investors will look for its resolution — a large minority of guidance never closes.

Section 2 places abandonment against the existing literature. Section 3 describes the institutional setting and the corpus. Section 4 sets out the method. Section 5 reports results. Section 6 gives four verified illustrative cases. Section 7 sets out limitations, which are substantial and which we state in full. Section 8 concludes.

---

2. Related literature and the position of abandonment

Three strands of work bear directly on this paper.

Guidance cessation. Houston, Lev, and Tucker (2010) study firms that stop issuing quarterly earnings guidance and find that poor operating performance is the dominant cause; contrary to the argument that stoppers substitute richer long-horizon disclosure, they document that alternative forward-looking disclosure is *curtailed* after cessation, alongside deterioration in the information environment. Chen, Matsumoto, and Rajgopal (2011) reach broadly consistent conclusions. Aaron, Kang, Ng, and Rusticus (SSRN No. 3794964) study the pandemic-era withdrawals and find that exposure predicts withdrawal, and that subsequent performance improvement predicts reinstatement.

All three treat guidance as a firm-level policy with a binary state. That is the right unit for the question they ask. It is the wrong unit for ours: a firm can maintain an unbroken guidance policy while individual commitments inside that policy quietly fail to close.

Silence and non-answers within calls. Hollander, Pronk, and Roelofsen (2010) show that managers regularly leave call participants without an answer, and that non-response is predicted by firm size, CEO stock-price incentives, firm age, performance, and litigation risk. Gow, Larcker, and Zakolyukina (2021) measure the phenomenon directly and find a non-answer rate of approximately 11% of analyst questions, stable over time and similar across industries.

This literature is the closest antecedent for what we measure, and it establishes the key premise: managerial silence is a measurable, non-random disclosure choice. But its unit is the question–answer pair inside a single call. Abandonment happens between calls, and its defining property is that no question need be asked at all.

Guidance that lives only in the call. Lansford, Lee, and Tucker (SSRN No. 1270974) document that S&P 500 firms — including NYSE-listed firms, at a time when NYSE rules prohibited the disclosure of new material information in a call — disclose a non-trivial volume of management guidance *exclusively* in conference calls, never in the accompanying press release. This matters for our design in two ways. It establishes that the call is a genuine locus of guidance rather than a restatement of the release, which is why the call is the right corpus. And it establishes that a class of commitments exists with no press-release counterpart, which means for those commitments the call record is the only record.

Position. Abandonment sits in the gap between these strands. Cessation is the firm switching the policy off, announced. A non-answer is a refusal, observable at the instant of refusal. Abandonment is neither announced nor observable at any instant — it is an absence that only becomes measurable when a commitment is tracked forward and found never to have been closed. Table 3 sets the three side by side.

---

3. Institutional setting and data

3.1 Why India

Indian listed companies operate under a disclosure regime that makes an earnings-call corpus unusually complete. Regulation 46(2) of the SEBI (Listing Obligations and Disclosure Requirements) Regulations, as amended by the (Second Amendment) Regulations, 2021, notified 6 May 2021, requires a listed entity that holds post-earnings or quarterly analyst and investor calls to publish the audio recording on its website promptly — before the next trading day or within twenty-four hours of the call's conclusion, whichever is earlier — video recordings, where made, within forty-eight hours, and a written transcript within five working days, submitted simultaneously to the recognised stock exchanges. Transcripts are subject to a minimum five-year retention requirement under the entity's archival policy; audio and video recordings carry a shorter minimum.

Three consequences matter for measurement. Transcripts of calls that are held exist as a matter of regulation rather than firm discretion, removing a selection channel that operates in jurisdictions where transcript availability is commercial. They are published on a short and uniform clock, so the record is close to contemporaneous. And the transcript record persists for at least five years, so a commitment made in one period can be matched against the record of later periods without archive gaps.

One boundary should be stated at once, because it is the limit of the "near-census" claim: the regulation governs the *disclosure* of calls a company chooses to hold. It does not require any company to hold an earnings call. Selection into the corpus therefore operates at the level of whether a firm convenes calls at all, which is the restriction stated in §3.3, not at the level of whether a call it held is transcribed.

A fourth consideration is substantive rather than mechanical. Indian management guidance is predominantly narrative and operational — a margin band, a capacity addition, a target for a segment, a debt-reduction path, a commissioning date — rather than a single consensus-comparable quarterly EPS number. Firm-level "does this firm guide or not" measurement therefore carries little information in this

setting. Guidance here is a stream of heterogeneous, individually checkable commitments, which is exactly the object our design takes as its unit.

3.2 The corpus

The corpus comprises **8,220 earnings-call transcripts from 1,654 Indian listed companies**, from which **24,328 discrete forward-looking management commitments** are extracted, distributed across **1,648 companies**. Mean coverage depth is **4.18 distinct quarters per company**. The corpus is a snapshot as of 4 September 2026.[^1]

Coverage depth is the single most important descriptive fact about the corpus and it constrains everything that follows. A mean of 4.18 quarters means the median commitment has had only a handful of subsequent calls in which to resolve. Long-horizon guidance is therefore materially right-censored, a limitation we address in §4.3 by excluding not-yet-resolvable commitments from every rate we report, and revisit in §7.3.

[^1]: The corpus is the underlying dataset of Inve, a commercial earnings-call analysis platform (https://inve.money). The author is affiliated with the platform. See the competing-interests statement at the end of this paper.

3.3 Sample restrictions and what they exclude

The corpus is restricted to Indian listed companies that (a) hold earnings calls and (b) publish transcripts of them. This is not a random sample of listed India. Firms that hold calls are larger, more widely covered, and more institutionally held than the listed universe as a whole. Any rate reported here should be read as a rate among firms that participate in the earnings-call channel, not among Indian listed companies generally.

We do not condition on market capitalisation, sector, index membership, or analyst coverage, and we do not trim. Every Indian company with at least one extracted commitment is in.

One exclusion is applied and is worth stating precisely, because it corrects figures that appeared in an earlier draft of this work. The store that holds the Indian corpus also contained 51 commitments belonging to seven **US-listed** companies, carried under the Indian partition. These are excluded here. Their removal reduces the commitment count by 51 (24,379 to 24,328) and the company count by seven (1,655 to 1,648). None of the excluded rows carried an abandoned state — they were 42 `new`, 6 `on_track` and 3 `revised_up` — so the abandonment count and the headline abandonment and kept rates are unchanged. Three state counts move (`new` 7,528 to 7,486; `on_track` 5,209 to 5,203; `revised_up` 600 to 597), as do the revenue-scoped and persistent-deflection subsets in §5.4 and §5.5 and the per-company mean kept rate in §5.2. At one decimal place the `new` share moves from 30.9% to 30.8%; no other share does.

Two transcript figures also differ between drafts — 8,227 to 8,220 transcripts, and 1,661 to 1,654 companies with transcripts — and they move for a separate reason: the earlier count included seven transcripts held under the US partition. No commitment in this paper derives from them. We record it

rather than silently restating because the contaminated figures were circulated before this correction, and a reader comparing the two should know why they differ.

3.4 A note on scope

Our corpus contains earnings calls and nothing else. It does not contain press releases, investor-day presentations, annual reports, regulatory filings, media interviews, or conference appearances. This boundary is a deliberate design choice — the call is the forum in which the commitment was made, and therefore the natural forum in which to look for its resolution — but it is also the paper's most consequential limitation, and we treat it as such in §7.2. Where we say a commitment was abandoned, we mean it was abandoned *on the earnings-call record*. We do not and cannot mean the company never addressed it anywhere.

3.5 Data availability

Aggregate statistics reported in this paper are fully reproducible from Tables 1 and 2: every rate in the text is an arithmetic function of counts printed in those tables, and we have deliberately reported the complete state distribution so that no reader must take a derived rate on trust.

Firm-level records are not distributed. The complete commitment-level register is maintained as a commercial product and is not released, and we did not prepare a public sub-sample: a capped extract was drafted and then withdrawn, because publishing an adverse state against a named issuer on a coding that no human had checked against the primary transcript is not a standard we were willing to meet permanently. We would rather say that plainly than ship a sample whose weakest rows would be the first thing a sceptical reader opened.

What this costs the reader is exact replication, and we do not minimise it. What it does not cost is the ability to check the paper's arithmetic, which is fully exposed in Tables 1 and 2, or the ability to test the construct independently: the definitions in §4.1 and §4.4 are complete and self-contained, and an independent estimate of the abandonment rate built on a separately constructed corpus would be a stronger test of this paper's central claim than any sample we could publish. §7.1 states the corresponding limitation plainly.

Requests for access to the coding of a specific case discussed in this paper may be sent to the corresponding author, and will be answered with the underlying transcript passage and the state history for that commitment.

---

4. Method

This section describes the design at the level a referee needs to judge validity. Implementation details — the extraction model, its configuration, the written extraction specification, and the internal output schema — are proprietary and are not disclosed. §7.1 states the replication cost this imposes.

4.1 Unit of observation

The unit is a **commitment**: a forward-looking statement made by a member of management on a scheduled earnings call that carries all three of

1. an identifiable **subject** — a metric, a deliverable, an asset, or a business line;
2. a **direction or level** — a number, a band, a threshold, or an unambiguous qualitative direction; and
3. a **horizon** — an explicit period, or a period recoverable from the surrounding text.

Statements failing any of the three are not commitments for our purposes. General optimism, descriptions of past performance, and answers that decline to commit are excluded at extraction. Commitments are attributed to the call on which they are first made; that call defines the commitment's origin period.

4.2 Tracking and state assignment

Each commitment is carried forward and evaluated against every subsequent earnings call by the same company. At each subsequent call the commitment is assigned a state drawn from a **closed, pre-specified vocabulary** (Appendix A). The states are mutually exclusive and exhaustive: every commitment holds exactly one state at any point. **Twelve of the vocabulary's values are populated in the Indian corpus**, and those twelve counts sum to the corpus total; two further values, `not_reported` and `unverifiable`, exist in the vocabulary but are not populated here (§4.3, §7.8). Where this paper refers to "the twelve states," it means the twelve populated in this corpus. Table 1 demonstrates the summing identity.

The variable analysed here is the **current state** — the state the commitment holds as of the most recent call in the corpus for that company. A commitment that was on track in one period and missed two periods later contributes once, as missed. This is a stock measure, not a flow measure, and it is the appropriate one for the question "what is the standing condition of guidance in this market."

4.3 The resolvable denominator

The platform's own logic excludes three states from any rate calculation on the grounds that they are not yet informative about outcome: `new`, `not_reported`, and `unverifiable`. In the Indian corpus only `new` is populated; the other two exclusions are inert here. The resolvable denominator is therefore

16,842 = 24,328 − 7,486 (total − `new`).

We adopt this denominator throughout and label it explicitly at every use. Rates on the full 24,328 are also reported, because the choice of denominator materially changes the headline and readers are entitled to both.

Excluding `new` removes the commitments too young for the record to say anything about them. It does not, however, make the abandonment rate conservative, and we do not claim that it does. **The direction of the censoring bias in this corpus runs against us, and we state it here rather than in the limitations section because it bears on how the headline should be read.** A commitment is eligible to be coded abandoned once a single subsequent call has passed without mentioning it (§4.4, condition 2). With mean coverage of 4.18 quarters, many commitments have only one or two subsequent calls on file. A long-

horizon commitment — a plant to be commissioned three years out, a margin to be reached "over the medium term" — that simply has not yet come up again will be coded abandoned rather than pending. That inflates the measured rate. Pulling the other way, management routinely reiterates live commitments on the next call, which is what places a commitment in `on_track` rather than `ghosted`; and a commitment that is genuinely dropped early is correctly caught. We cannot bound the net effect with the data reported here. **The honest statement is that 13.5% is an estimate with an unquantified censoring bias whose most plausible direction is upward**, and the appropriate correction — a minimum-elapsed-calls threshold, or a survival specification with the horizon as a covariate — requires the period labels that §7.4 explains we do not trust.

4.4 Definition of abandonment

A commitment is coded **abandoned** (the platform's internal field value is `ghosted`; we use "abandoned" in prose and reserve the platform term for the field) when all of the following hold:

1. the commitment was made on an earnings call in the corpus;
2. at least one subsequent earnings call by the same company exists in the corpus;
3. no subsequent call resolves it (achieved, missed, achieved on diluted terms), revises it (up, down, delayed, diluted), flags it (at risk), or restates it as still standing (on track); and
4. no subsequent call references the commitment at all.

Condition (4) is what separates abandonment from every other state. A commitment that management repeats without progress is `on_track`, not abandoned. A commitment management explicitly withdraws is a revision, not abandonment. Abandonment requires the commitment to leave the conversation without leaving a trace of its departure.

4.5 Definition of the kept rate

The kept rate is

> *(`achieved` + `achieved_diluted` + `on_track` + `revised_up`) ÷ resolvable commitments*

This is the platform's operative definition, read from its implementation rather than assumed, and it is one defensible definition among several. It is generous in one respect — `on_track` counts as kept although the outcome is not yet realised — and strict in another, since the denominator includes abandoned, at-risk, delayed, revised-down and diluted commitments.

The kept rate is highly sensitive to the denominator. A narrower denominator restricted to unambiguously closed outcomes produces a materially higher figure; a wider one produces a lower figure. We do not adjudicate between definitions. We report the full state distribution in Table 1 precisely so that a reader who prefers a different numerator or denominator can compute it directly, and we would encourage anyone quoting a kept rate from this paper to state the definition alongside the number.

4.6 Validity

Four checks bear on whether the measurement can be trusted.

Internal consistency. The twelve state counts sum exactly to 24,328 (Table 1). The resolvable denominator, the kept numerator, and the three-way decomposition in Table 2 are all sums of the same printed counts. Every rate asserted in this paper is re-derivable by a reader with a calculator and Table 1 alone.

Vocabulary closure, and its cost. The state vocabulary is closed and pre-specified. Extraction cannot emit a state outside it, so there is no residual or "other" category quietly absorbing hard cases — a common failure mode in text-derived classification, and one that would be invisible to a reader if it existed here. The cost is the mirror image of the benefit, and we would rather state it than have it found: with no residual category populated in this corpus, a genuinely ambiguous case cannot be parked. It must be forced into a substantive state, and `ghosted` — the only state defined by an *absence* of subsequent evidence rather than the presence of it — is the state most available to absorb a case the extraction could not resolve. This is a second reason, independent of censoring, to treat 13.5% as an upper-leaning estimate. It is also the mechanism behind §7.8's refusal to compare markets: the parallel US corpus does populate `not_reported` and `unverifiable`, and those states absorb exactly the cases that must be forced here.

Case-level verification against primary sources. Every quotation in §6 was checked against the primary transcript in which it was spoken, and the two US cases were additionally checked end to end against independent external sources (company releases, filings, and contemporaneous press reporting). Coverage of this check is not uniform and we specify it rather than imply it: for the US cases it extends to the period attributions and the outcomes; for the Indian cases it extends to the quotations and the origin call, and the later steps rest on the corpus record alone, which is why they are presented without period labels (§7.4). This exercise is not decorative — it is how we found and discarded four candidate cases whose corpus coding did not survive external checking, including one where an apparent abandonment had been addressed by the company outside the earnings-call channel and one where our own period label was wrong. Those discards inform §7.2 and §7.4.

Absence of a corpus-scale human audit. We do not report an inter-rater reliability statistic, and no human audit of extraction accuracy at corpus scale has been conducted. This is a real gap and we do not minimise it; see §7.1.

4.7 No claim about intent

We make no claim, and the design supports no claim, about why any commitment was abandoned. Words such as "buried," "hidden," "dropped deliberately," or "misled" appear nowhere in this paper's characterisation of any company. Abandonment is defined mechanically as a property of what did and did not appear in a subsequent transcript. It is a measurement, not an accusation, and we ask that it be cited as one.

---

5. Results

5.1 The state distribution

Table 1. Current state of all tracked commitments, Indian corpus, as of 4 September 2026.

State	Count	Share of all commitments
`new` (not yet resolvable)	7,486	30.77%
`on_track`	5,203	21.39%
`achieved`	3,558	14.63%
`missed`	2,380	9.78%
`ghosted` (abandoned)	2,271	9.33%
`at_risk`	880	3.62%
`delayed`	800	3.29%
`revised_down`	638	2.62%
`revised_up`	597	2.45%
`achieved_diluted`	376	1.55%
`diluted`	137	0.56%
`achieved_delayed`	2	0.01%
Total	**24,328**	**100.00%**

Counts are corpus values. Percentages are computed from the counts in this table and are subject to rounding; the column sums to 100.00% after rounding. The two decimal places are arithmetic precision, not measurement precision: the counts are exact, the classification producing them is model-based and unaudited at corpus scale (§7.1), and no inference in this paper turns on the second decimal.

Two features of Table 1 are worth flagging before any rate is computed from it.

First, `new` is the largest single state at 30.77%. Nearly a third of the corpus consists of commitments too young to have resolved. This is a direct consequence of a 4.18-quarter mean coverage depth and it is the reason every rate in this paper is reported against the resolvable denominator.

Second, the abandoned state is the fifth largest, and it is larger than the explicit revisions taken together. The three states in which a commitment is openly walked back — `revised_down` (638), `delayed` (800) and `diluted` (137) — sum to 1,575, well below 2,271. Adding `revised_up` (597), so that the comparison covers *every* state in which management explicitly changed a standing commitment in either direction, gives 2,172 — still fewer than the 2,271 that were abandoned. In this corpus, a commitment is more likely to disappear from the record than to be explicitly revised on it.

5.2 The headline rates

Table 2. Three-way decomposition of resolvable commitments (n = 16,842).

Category	Component states	Count	Share of resolvable
Kept	`achieved` + `achieved_diluted` + `on_track` + `revised_up`	9,734	57.8%
Abandoned	`ghosted`	2,271	13.5%
Neither kept nor abandoned	`missed` + `at_risk` + `delayed` + `revised_down` + `diluted` + `achieved_delayed`	4,837	28.7%
Total		**16,842**	**100.0%**

Counts are sums of the state counts in Table 1. The kept rate of 57.8% and the abandonment rate of 13.5% are the platform's own reported figures under the definitions in §4.4–4.5; the residual category is the arithmetic remainder. Note that `achieved_delayed` *(n = 2) falls in the residual rather than the kept numerator, because the operative kept definition in §4.5 does not include it. We report the definition as implemented rather than as we might have specified it, and the state is too small to affect any figure.*

The principal results:

- **2,271 commitments are abandoned.** This is **9.3%** of all 24,328 tracked commitments and **13.5%** of the 16,842 resolvable ones.
- **The kept rate is 57.8%** pooled across all resolvable commitments. Computed instead as the unweighted mean of per-company kept rates — which weights a company with four commitments equally with one that has forty — it is **56.7%** across the 1,459 companies holding at least one resolvable commitment. (The remaining 189 of the 1,648 companies hold only `new` commitments and drop out of this calculation entirely.) The two figures are close, which is weak evidence that the pooled rate is not driven by a handful of high-volume companies. It is only weak evidence: two very different weight distributions can produce similar means, and we have not tested the sensitivity directly.

The gap between the two denominators is worth stating plainly, because it is where a careless citation would go wrong. *9.3% of all tracked commitments* and *13.5% of commitments old enough to resolve* are both correct and they answer different questions. The second is the more informative figure, and it is the one we would use, because the first is diluted by 7,486 commitments that have not yet had a fair chance to close.

5.3 How widely abandonment is distributed

Abandonment is not a property of a few unusual companies.

1,000 of the 1,648 companies with at least one tracked commitment — 60.7% — carry at least one abandoned commitment.

This is the result that most changes how the headline rate should be read. A 13.5% abandonment rate concentrated in a hundred serial offenders would be a story about a disreputable tail. A 13.5% rate spread across three-fifths of the companies in the corpus is a story about a normal feature of the earnings-call channel. It is the latter.

It follows that abandonment cannot be used, on its own, as a firm-level red flag. A single abandoned commitment places a company in the majority. Whatever discriminating power the construct has must come from intensity, from the materiality of what was abandoned, or from the pattern across time — none of which this paper measures.

5.4 The revenue and growth subset

Commitments differ in how checkable they are. A capacity-commissioning date resolves cleanly. A stated ambition to "improve the customer experience" does not. To test whether abandonment is an artefact of vague commitments, we isolate the subset carrying an explicit revenue or growth scope — the most numerically checkable class in the corpus, and the class analysts are most likely to follow up.

Within this subset there are **4,468 commitments across 1,293 companies**, of which **261 of 2,920 resolvable commitments are abandoned: 8.9%**.

The subset is defined as commitments whose scope classifier assigned an explicit revenue scope — consolidated, segment, or new-segment. An earlier draft of this paper reported 8.9% on a wider filter that also admitted commitments the classifier had explicitly labelled *not* revenue-scoped; that filter was wrong and the figure it produced should be disregarded.

(Medium confidence. Scope classification is model-assigned and has not been externally validated; the subset boundary is softer than the state vocabulary. Treat 8.9% as indicative rather than precise.)

Two things follow, and the first cuts against the headline. Restricting attention to the most checkable commitments moves the abandonment rate from 13.5% to 8.9% — a decline of roughly 4.5 percentage points, or about a third in relative terms. Checkability matters. Some of what the corpus-wide rate measures is vagueness: a commitment stated loosely is materially more likely to go unmentioned than one stated as a number against a period, whether because loose commitments are genuinely easier to let lapse or because our extraction is worse at matching their restatements. We cannot separate those two explanations here, and a reader should treat 13.5% accordingly.

The second is that the effect does not go away. Among commitments carrying an explicit revenue scope — numeric, period-bound, and the class analysts are most likely to chase — **one in eleven is still never mentioned again**. That is the more defensible version of this paper's claim, and it is the one we would ask a reader to carry: abandonment is concentrated in less checkable guidance, and it persists at a non-trivial rate even where it should be hardest to sustain.

5.5 A secondary signal: persistent deflection

The corpus also records **deflected topics** — subjects the extraction flagged as having drawn a non-substantive answer, each carrying a recorded deflection streak. There are **1,849 such topics across 1,315**

companies, but the label is weaker than its name suggests: 1,240 of them carry a streak of one or zero, meaning the deflection was recorded on a single call. Only **609 topics across 452 companies** show a streak of two or more, which is the subset for which the word *persistent* is warranted. We report the smaller figure as the meaningful one.

(Medium confidence, and we flag it prominently. These are model-extracted with no external validation of any kind, and unlike the state vocabulary the construct has no closed definition anchored to a subsequent observable event.)

We report it for one narrow reason: it is the natural cross-call analogue of a phenomenon the literature has established within calls (Gow et al. 2021; Hollander et al. 2010), it appears at a broadly similar order of magnitude to abandonment in the same corpus, and its existence is worth recording so that someone with a validated instrument can measure it properly. We draw no comparison between this figure and the measured within-call non-answer rate — the constructs, the units and the denominators all differ, and our figure has no external validation to place against theirs. No result in this paper depends on this subsection, and it could be deleted without affecting anything in §5.1 to §5.4.

5.6 Summary against the literature

Table 3. Three forms of disclosure attrition.

	Unit	**Observable when it happens?**	**Announced?**	**Established incidence**
Guidance cessation	Firm-policy	Yes	Yes	Studied at firm level; Houston et al. (2010), Chen et al. (2011)
Non-answer	Question–answer pair, within one call	Yes	No	≈11% of analyst questions (Gow et al. 2021)
Abandonment	**Individual commitment, across calls**	**No**	**No**	**13.5% of resolvable commitments (this paper)**

The incidence column reports each literature's own rate over its own denominator. The rates in the last two rows are not comparable quantities — one is a share of analyst questions in a US sample, the other a share of commitments in an Indian one — and the column is a summary of what has been measured, not a ranking. Our figure additionally carries the upward-leaning bias set out in §4.3 and §4.6.

The row that matters is the third column. Cessation and non-answers both leave a mark at the moment they occur: a press release in the first case, a visible refusal in the transcript in the second. Abandonment leaves nothing. There is no call on which anything happens; there is only a growing set of calls on which something fails to happen. This is why it needs a longitudinal commitment-level panel to detect, and it is our best explanation for why a phenomenon of this magnitude has not previously been counted.

---

6. Illustrative cases

The four cases below are worked examples of the taxonomy, not evidence for any rate. Each was verified against the primary transcript; the two US cases were additionally verified against independent external sources. Quotations are verbatim.

The US cases are included because the taxonomy is not India-specific and the mechanics are clearer in firms most readers already know. **They are illustrations only. This paper reports no US rates, for reasons set out in §7.8.**

6.1 Abandonment — Eternal Ltd (formerly Zomato; NSE: ETERNAL), India

On the Q1 FY25 call, management set a target for the food-delivery adjusted EBITDA margin of 4–5%, with a horizon: *"in a few quarters from now on, we should get to that range … we're not very far from that now."* The target was reiterated on a later call. On every call after that in the corpus it is absent. It is not achieved, not missed, not withdrawn, and not revised. Corpus state: `ghosted`, and still `ghosted` as of the most recent parsed quarter.

A second commitment on the same business shows a different mechanism at work. Food-delivery gross order value growth was guided at *"20% plus … currently trending at 27%–28% YoY."* On a subsequent call, as reported growth fell to 16%, the target was recharacterised as *"a 4–5 year CAGR, not an annual target."* Reported growth fell further, to 13%, on the following call.

The pair is instructive. The GOV target was redefined in public — an observable event, on the record, which a reader of the transcript can see and evaluate. The margin target was not redefined. It simply stopped appearing. The second is harder to notice than the first, and it is the one our construct is built to catch.

(Verification and its limits. The two quotations and the origin call were verified line-by-line against the primary Q1 FY25 transcript published by the company. The subsequent steps are presented as an ordered sequence without quarter labels because those labels are derived and carry the known error described in §7.4 — we have not checked them against the later transcripts individually, and we do not assert them. Separately, and importantly: we did not search outside the earnings-call channel to establish whether the company addressed the margin target in a press release, investor presentation or filing. We make no claim that it did not. The claim here is exactly the mechanical one — the target entered the call record and did not leave it — and §7.2 explains why that is the only claim this design can make about any named firm.)

6.2 Counter-example — Navin Fluorine International (NSE: NAVINFLUOR), India

Not every upward revision is a ratchet, and a paper that only showed abandonment would be misleading about the corpus.

Navin Fluorine reset its EBITDA margin guidance to 25%, with management stating *"for now we will hold it to that number."* Each subsequent raise then **trailed a margin the company had already delivered**: guidance moved to *"north of 25%"* after the company printed 28.5%; to *"28% to 30%"* after it printed

32.5%; and then to *"30% ... plus/minus 200 basis points."* The fiscal year closed at a 32.6% margin, above the revised target. Corpus state: `achieved`.

The distinguishing feature is the ordering, and it is directly observable in the transcript record. Here the guidance follows the print rather than leading it: each raise ratifies performance already delivered rather than claiming performance yet to come. We draw no inference about motive from this — the ordering is the fact, and it is the fact a reader can check. It is also, on the evidence of Table 2, the more common shape in this corpus: 9,734 resolvable commitments are kept.

(We present the revision sequence in order rather than by quarter label. Our Indian period labels are derived and carry a known error mode; see §7.4.)

6.3 A commitment that moves through four states — Ford Motor Company (NYSE: F), US

On the Q1 2025 call in May 2025, management quantified the expected tariff impact: *"we estimate a gross adverse EBIT impact of $2.5 billion and a net adverse EBIT impact of about $1.5 billion for full year 2025."*

The net figure then moved twice more: to approximately $2bn at Q2, and back down to approximately $1bn at Q3 in October 2025. The Q4 2025 call reported an actual full-year net tariff impact of $2bn, with management stating: *"That's a $1 billion higher tariff impact than we communicated just in October due to the unexpected and late year change in tariff credits for auto parts."*

One guidance line, four states across four consecutive quarters, closing above the bar originally set. This is the opposite of abandonment and it is worth stating why it belongs in the paper: everything here is on the record. Each revision was made explicitly, on a call, with a stated reason. A reader who tracked only the first and last figures would see a large miss; a reader who tracked the intermediate states sees a company revising in public as facts changed. Our construct classifies this as a revision sequence, not as abandonment, and the distinction is the whole point. **Guidance that moves is not the problem this paper identifies. Guidance that stops is.**

(Quotations verbatim from the corpus, independently confirmed against contemporaneous reporting from CNBC covering the May 2025 and July 2025 calls, and against trade reporting of the Q4 2025 call in February 2026.)

6.4 Withdrawal and reset — The Estée Lauder Companies (NYSE: EL), US

On the Q4 FY24 call, management guided that *"full fiscal year organic net sales are forecasted to range between a decrease of 1% and an increase of 2%."*

On the Q1 FY25 call the outlook was withdrawn in full: *"we are solely issuing an outlook for the second quarter and withdrawing our fiscal year 2025 outlook."* It was reset on the Q3 FY25 call to a range of −9% to −8%. The fiscal year closed at −8% organic — inside the reset band, and far outside the original one. Nearly two-thirds of the decline came from travel retail, which fell 28%.

This is the firm-level cessation phenomenon of Houston et al. (2010) and Aaron et al. (SSRN No. 3794964) occurring at the level of a single commitment: an explicit, dated withdrawal followed by an explicit, dated reset. It is highly visible and easy to audit, and it is precisely the disclosure behaviour that the existing literature is equipped to study. Abandonment is what happens when none of these events occur and the commitment leaves the record without any of them.

(Quotations verbatim from the corpus, independently confirmed against the company's FY2025 results release of 20 August 2025 and its FY2025 Form 10-K.)

---

7. Limitations and threats to validity

We state these in full rather than in summary. Several are material, and a reader who takes the headline figure without them will misuse it.

7.1 Extraction is model-based and not audited at corpus scale

Commitments are extracted and states are assigned by a large language model operating under a fixed written specification with a closed output vocabulary. No human audit of extraction accuracy at corpus scale has been performed and **we report no inter-rater reliability statistic**. Verification in this paper is case-level (§4.6) and therefore establishes that the construct is coherent on the examples we checked; it does not establish a corpus-wide error rate.

The direction of any bias is not obvious a priori. Extraction could over-generate commitments from hedged language, which would inflate the abandonment count, since spurious commitments have nothing to resolve into. It could equally under-generate from oblique phrasing, which would deflate the total. We cannot presently bound either. A human-coded validation sample against a stratified draw is the obvious next step and this paper does not substitute for it.

Two extraction models were used across one corpus. 1,469 of the 1,648 company records were extracted with `gemini-2.5-pro`; the remaining 179 (10.9%) with `gemini-3.1-pro-preview`. The two were not run against a common holdout, so no statement can be made about whether they assign states identically. If they do not, a share of the between-company variation reported here is model artefact rather than firm behaviour. Because assignment of a company to one model or the other was driven by processing date rather than by any firm characteristic, we have no reason to expect the split to correlate with outcomes — but we have not tested this, and a reader who wishes to be conservative should treat any between-company comparison as carrying an unquantified extraction component. The corpus-level rates in §5 are affected only to the extent that the two models differ systematically in aggregate, which is likewise untested.

Additionally, the extraction specification and the internal schema are proprietary and undisclosed, though the models themselves are named above. Exact replication from this paper is therefore not possible. Independent replication of the *construct* is possible — the definitions in §4.1 and §4.4 are complete and

self-contained — and we would regard an independent estimate of the abandonment rate on a separately constructed corpus as the strongest available test of this paper's central claim.

7.2 The corpus boundary is earnings calls only

This is the limitation most likely to produce a false inference from our results, and it deserves the emphasis.

Our corpus contains earnings-call transcripts and nothing else. A commitment addressed in a press release, at an investor day, in an annual report, in a regulatory filing, or in a media interview — but not on a subsequent earnings call — will be recorded here as abandoned. **"Abandoned" therefore means "abandoned on the earnings-call record," never "the company never addressed it."**

This is not a hypothetical concern. In preparing the illustrative cases we identified a candidate abandonment at a large listed technology company that our corpus coded correctly by its own definition — the commitment was never closed out on any earnings call — but which the company had in fact addressed publicly outside the call channel. Our construct was right about the call record and would have been badly wrong as a statement about the company. We discarded the case. Any reader applying this construct to a single named firm must perform the same check, and should assume the corpus cannot perform it for them. We did not perform it for the Indian case in §6.1 either, and say so there.

The construct is defensible only because the earnings call is where the commitment was made. A firm that guides on a call and resolves elsewhere has still left the call record incomplete, which is a meaningful observation about the channel. It is not, however, an observation about the firm's total disclosure, and it must not be cited as one.

7.3 Right-censoring

Mean coverage is 4.18 quarters per company. Commitments with horizons longer than the observation window cannot resolve within it. Excluding the 7,486 `new` commitments removes the most obviously censored cases, but as set out in §4.3 it does not make the abandonment rate conservative: a commitment becomes eligible for the abandoned state after a single silent subsequent call, so a long-horizon commitment that has simply not come up again yet is coded abandoned rather than pending. **We regard this as the most serious threat to the headline figure and we take the most plausible direction of the bias to be upward.** We cannot quantify it here, because the correction requires reliable period labels and horizon dating, which §7.4 explains we do not have.

The censoring is also not uniform across companies. A company with two quarters in the corpus and one with eight contribute very differently to the pooled rate, and firms with shallow coverage will contribute disproportionately to the abandoned count for the reason just given. The closeness of the pooled (57.8%) and per-company mean (56.7%) kept rates is mildly reassuring on the weighting question but says nothing about the depth question.

7.4 Period labels are derived and carry a known error

Quarter labels in the Indian corpus are derived rather than taken from an authoritative field, and a known off-by-one error exists in that derivation. **We therefore report no per-quarter decomposition, no time series, and no event-time analysis anywhere in this paper.** This is a design consequence of a data-quality defect, stated so that readers understand why an obvious analysis is absent rather than assuming it was run and suppressed.

The same rule governs the illustrative cases. In §6.1 and §6.2 the Indian sequences are presented in order rather than by quarter label, with one exception: the Eternal origin call is labelled Q1 FY25 because the quotations were verified directly against the company's own Q1 FY25 transcript, which is a primary-source attribution rather than a derived one. Every other Indian period attribution has been removed from this paper rather than carried at a confidence we cannot support. The US cases in §6.3 and §6.4 are unaffected: their period labels come from the companies' own reporting calendars and were confirmed against external sources.

7.5 Selection into the sample

Firms that hold earnings calls and publish transcripts are not representative of Indian listed companies. They skew larger, more widely covered, and more institutionally held. Rates here describe the earnings-call channel, not the market.

7.6 The construct is descriptive, not diagnostic

Because 60.7% of companies in the corpus carry at least one abandoned commitment, the presence of an abandoned commitment carries almost no discriminating information about a company. Any diagnostic use would require intensity, materiality weighting, or persistence — none of which is developed here. We would regard a firm-level red-flag interpretation of this paper as a misreading.

7.7 The figures are a snapshot and will drift

Every number in this paper is a state of the corpus on **4 September 2026**. The corpus grows as new calls are parsed, and existing commitments change state as they resolve. A commitment abandoned today may be revived on a future call; a `new` commitment will resolve into some other state. An earlier snapshot of this corpus differed materially from the figures reported here, which is why we date-stamp rather than describe the figures as standing values. Any citation should carry the date.

7.8 Why no cross-market comparison is reported

We maintain a parallel corpus of US earnings calls and we deliberately report **no US rates and no India-versus-US comparison**. Three confounders make such a comparison invalid, and we state them so that no reader constructs one:

1. **Coverage depth differs materially.** Mean coverage is 6.99 quarters per company in the US corpus against 4.18 in the Indian corpus. (The US figure is the mean number of distinct quarters in which a

company stated at least one commitment, over a separate US store of 37,529 commitments across 885 companies; it is not computed from the Indian corpus described in §3.2 and the two are not pooled anywhere in this paper.) Since abandonment requires subsequent calls to be observable, and resolution requires time, differing depth alone moves every rate.

2. **The state vocabularies are not identical.** The US corpus populates `not_reported` and `unverifiable` states that are inert in the Indian corpus. These absorb ambiguous cases that, under the Indian vocabulary, must be assigned to a substantive state. This alone mechanically depresses the measured US abandonment rate relative to India's, by construction rather than by conduct.
3. **The corpora were extracted under different specifications.** The two extraction specifications have diverged, so a difference in measured rates cannot be attributed to a difference in managerial behaviour.

Any comparative statement of the form "managements in market A keep more of their guidance than in market B" derived from these two corpora would be a claim about a methodological artefact. We do not make it, and we ask that it not be constructed from this paper.

7.9 Competing interests

The author is affiliated with the commercial platform that maintains the corpus (note 1). The platform sells access to the firm-level records underlying the aggregates reported here, and the aggregate findings are favourable to the commercial case for that product. Readers should weigh this. We have tried to mitigate it by reporting the complete state distribution rather than only the headline, by reporting a kept rate of 57.8% and a counter-example case rather than presenting the corpus as uniformly damning, by declining to publish an available cross-market comparison that would have been more striking and less valid, and by stating limitations at length. None of that eliminates the interest, and we do not claim it does.

---

8. Conclusion

Guidance is usually studied as something a firm either provides or withholds. Measured instead at the level of the individual commitment, a third possibility becomes visible: guidance that is provided, and then neither kept nor broken, but simply left unresolved.

In a corpus of 8,220 earnings-call transcripts from 1,654 Indian listed companies yielding 24,328 tracked management commitments, 2,271 commitments — 13.5% of those old enough to resolve — are never mentioned again on any subsequent call. Sixty per cent of the companies in the corpus have at least one. Over the same corpus, 57.8% of resolvable commitments are kept.

The finding is descriptive and its limitations are real: the extraction is model-based and unaudited at scale, the corpus sees only earnings calls, and the observation window is short. Two of those limitations bear directly on the headline. A commitment becomes eligible for the abandoned state after one silent subsequent call, and mean coverage is 4.18 quarters; and the Indian corpus populates no residual state, so cases the extraction cannot resolve have nowhere to go but a substantive state. Both plausibly push 13.5%

up rather than down. We would therefore ask that the figure be cited as an upper-leaning estimate of a phenomenon whose existence and breadth are the durable findings here, rather than as a settled parameter. We have set the limitations out at length rather than in passing, because the construct is easy to over-read and the most damaging misreading — treating an abandoned commitment as proof that a company never addressed a target — is one our design cannot rule out for any individual firm.

What the measurement does establish is that the earnings call, which is the venue where much operational guidance is made and in many cases the only venue where it is made, does not systematically close what it opens. That is a property of a disclosure channel, and it is measurable. Three directions follow: a human-coded validation of extraction accuracy; a test of whether abandonment predicts anything — subsequent performance, revision behaviour, analyst forecast dispersion, or the information-environment deterioration that Houston et al. (2010) document after cessation; and independent construction of the same measure on an unrelated corpus. We would welcome the third most of all.

Statements

Funding. No external funding was received for this work.

Competing interests. The author is affiliated with the commercial platform that maintains the corpus described in note 1 and that sells access to the firm-level records underlying the aggregate statistics reported here. See §7.9.

Ethics. This study uses only publicly disclosed corporate documents. No human subjects are involved and no ethics approval was required.

Data availability. See §3.5. Aggregate statistics are fully reproducible from Tables 1 and 2: every rate in the text is an arithmetic function of counts printed there. Firm-level records are not distributed, and no public sub-sample is offered. The coding of any individual case discussed in this paper will be supplied on request to the corresponding author, with the underlying transcript passage and the commitment's state history.

Generative AI disclosure. During the preparation of this work the author used **Google Gemini 2.5 Pro** (`gemini-2.5-pro`, 1,469 of 1,648 company extraction runs) and **Google Gemini 3.1 Pro Preview** (`gemini-3.1-pro-preview`, the remaining 179 runs) in order to extract management commitments from earnings-call transcripts and to assign each commitment a state from a closed, pre-specified vocabulary, as described in §4. This is a substantive component of the study's method rather than a writing aid; its limitations, including the consequences of using two extraction models across one corpus, are stated in §7.1. The author additionally used **Anthropic Claude Opus 5** in order to draft and edit portions of the manuscript text. After using these tools, the output was reviewed, edited and verified by the author, who takes full responsibility for the content. No data, results, citations or references were generated by an AI tool; every figure was derived by direct query against the corpus database and every reference was checked against the publishing journal or repository record.

References

Aaron, A., J. Kang, J. Ng, and T. O. Rusticus. *Withdrawal of Management Guidance During the COVID-19 Pandemic.* SSRN Working Paper No. 3794964.

Chen, S., D. A. Matsumoto, and S. Rajgopal. 2011. "Is silence golden? An empirical analysis of firms that stop giving quarterly earnings guidance." *Journal of Accounting and Economics* 51 (1–2): 134–150.

Gow, I. D., D. F. Larcker, and A. A. Zakolyukina. 2021. "Non-Answers During Conference Calls." *Journal of Accounting Research* 59 (4): 1349–1384. https://doi.org/10.1111/1475-679X.12371

Hollander, S., M. Pronk, and E. Roelofsen. 2010. "Does Silence Speak? An Empirical Analysis of Disclosure Choices During Conference Calls." *Journal of Accounting Research* 48 (3): 531–563. https://doi.org/10.1111/j.1475-679X.2010.00365.x

Houston, J. F., B. Lev, and J. W. Tucker. 2010. "To Guide or Not to Guide? Causes and Consequences of Stopping Quarterly Earnings Guidance." *Contemporary Accounting Research* 27 (1): 143–185. https://doi.org/10.1111/j.1911-3846.2010.01005.x

Lansford, B., J. Lee, and J. W. Tucker. *Disclosure of Management Guidance in Conference Calls: Materiality, Determinants, and Consequences.* SSRN Working Paper No. 1270974.

Securities and Exchange Board of India. 2021. *SEBI (Listing Obligations and Disclosure Requirements) (Second Amendment) Regulations, 2021.* Notified 6 May 2021.

Appendix A. The state vocabulary

Each tracked commitment holds exactly one state at any point. The vocabulary is closed and pre-specified; extraction cannot emit a state outside it. Twelve of its values are populated in the Indian corpus and are listed first; the final row gives the two values that exist in the vocabulary but are not populated here.

State	Meaning
`new`	Made too recently to be resolvable. Excluded from all rate denominators.
`on_track`	Reiterated or confirmed as still standing on a subsequent call, without resolution.
`achieved`	Delivered on the terms originally stated.
`achieved_diluted`	Delivered, but against a scope or measurement base narrower than originally stated.
`achieved_delayed`	Delivered, but later than the stated horizon.
`missed`	The horizon passed and the commitment was not met; management acknowledged this on a call.
`ghosted`	**Abandoned. At least one subsequent call exists, and no subsequent call resolves, revises, flags, reiterates, or otherwise references the commitment.**
`at_risk`	Explicitly flagged on a subsequent call as unlikely to be met on current terms.
`delayed`	The horizon was explicitly pushed out.
`revised_up`	The target level was explicitly raised.
`revised_down`	The target level was explicitly lowered.
`diluted`	The scope or measurement base was explicitly narrowed without the level changing.
(`not_reported`, `unverifiable`)	Present in the vocabulary but not populated in the Indian corpus. Excluded from rate denominators where they occur.

The distinction that carries the paper is between `ghosted` and every other non-`new` state. Every other state requires management to have *said something* on a subsequent call. `ghosted` is the state that requires them to have said nothing — which is also why, in a corpus with no populated residual category, it is the state most exposed to absorbing cases the extraction could not resolve (§4.6).

Working paper. Not peer reviewed. Comments welcome.