The Otis Bound: An Information-Theoretic Ceiling on the Value Conveyable Between Two Floors

Marisol Otway-Pell[1], Cassian D. Verhoeven[2], and Petra Nkemelu[1,3]

[1]Institute for Conveyance Studies, Gorse Hill Polytechnic
[2]Department of Vertical Transportation, Whitcombe School of Applied Mechanics
[3]Centre for the Economics of Attention, Marlowe Institute

7 July 2026

Abstract

It is folklore that a good idea can be pitched between two floors of an elevator. We show that this is not folklore but a capacity theorem. Treating the transit as a communication channel whose binding resource is the listener's attention rather than the passage of time, we derive a hard upper bound—the **Otis Bound** $\mathcal{O} = C_0\big(1 - e^{-T/\tau}\big)$—on the mutual information a pitcher can convey to a listener in a single ride. In plain terms: there is a ceiling on how much of any idea can actually *land*, and no building is tall enough to raise it, because attention, not travel time, is what runs out. Per-floor conveyance collapses onto a building-independent **Mezzanine Constant** $\mu = 0.61 \pm 0.03$ bits/floor; half of every successful pitch is delivered before the cab clears the first inter-floor gap. Ideas whose intended content exceeds $\mathcal{O}$ do not arrive in truncated form—they *obliterate*, collapsing into self-similar repetition, a second-order phase transition in an order parameter Ω with critical exponent $\beta = \frac{1}{2}$. Across a pre-registered field study of $N = 412$ genuine two-party transits in 19 buildings, a programmable freight-elevator dose–response experiment, and a queueing-network simulation, we measure $C_0 = 11.4 \pm 0.6$ bits and $\tau = 9.8 \pm 0.4$ s, placing the celebrated thirty-second window exactly at the 3τ knee where conveyance reaches 95% of its ceiling. We conclude that a market's apparent appetite for simple ideas is not philistinism but a channel-capacity result: an idea worth more than the Otis Bound cannot be sold, because it cannot arrive.

Keywords: conveyance theory, elevator pitch, channel capacity, attention economics, obliteration transition, Mezzanine Constant

1 Introduction

The elevator pitch is among the few genuinely quantitative superstitions of commercial life. Founders rehearse it, incubators drill it, and an entire advisory literature insists that any worthwhile idea must survive the ride between two floors. The advice is usually delivered as motivation. We argue it is a theorem.

If value is realized only when an idea is *conveyed*—and an unconveyed idea, however profound, is economically inert—then the natural object of study is not the idea but the *channel* through which it passes. The elevator transit is the minimal such channel: two parties, a fixed and famously short time budget, and a listener whose attention is the scarce resource. For readers whose information theory has grown rusty, the one sentence to keep is this: we are measuring how much of an idea actually *lands* in the listener's head, and we find there is a hard ceiling on it.

We make three claims. First, the conveyable value of an idea—formalized as the mutual information between what the pitcher intends and what the listener reconstructs on exit—saturates at a ceiling we call the **Otis Bound**, $\mathcal{O} = C_0(1 - e^{-T/\tau})$. The ceiling is set by attention capacity C_0, not by transit time T; a taller building buys almost nothing. Second, the value delivered per floor is not building-specific but universal, a quantity we name the **Mezzanine Constant** μ. Third, and most consequentially, ideas that exceed the bound do not simply get cut off. They *obliterate*: the listener retains not a fragment of the idea but its bare repetition structure, the same phrase filling the available attention like dots filling a room. This is a sharp, second-order transition, and it is the paper's central empirical finding.

The practical upshot is stated once, plainly, for readers who will not follow the scaling analysis: past a certain complexity, a pitch does not get *shorter* in the listener's memory—it turns to mush. The rest of the paper explains why the ceiling exists, where it sits, and what it means that the market can only ever price the ideas below it.

2 Related Work

Classical information theory bounds the rate of reliable transmission over a noisy channel [1], but treats block length as a free design parameter. Our setting fixes it: the block length is the ride. The attention-as-buffer tradition in cognitive load [2, 3] supplies the saturating nonlinearity we exploit, though it has not previously been read as a capacity constraint on commerce. Empirically, retrospective studies of early-stage pitch outcomes [4, 5] report a robust penalty for "complex" pitches that has resisted explanation; we recover it as $\mathcal{O}$. The failure mode above the bound—collapse into repetition rather than graceful truncation—has a precedent outside economics, in the aesthetics of infinite repetition [6] and in the psycholinguistics of semantic satiation [7], where a signifier repeated to saturation sheds its referent. Finally, the physical channel itself has a mature engineering literature in vertical-transportation queueing [8, 9], from which we borrow transit-time distributions but, we believe, novel conclusions.

3 The Otis Channel

We model a pitch as a discrete memoryless channel used n times. Let X denote the pitcher's intended idea and Y the listener's reconstruction sampled at exit. The listener consumes symbols at rate R, so a transit of duration T affords $n = RT$ channel uses. Conveyed value is the mutual information $I(X;Y)$.

The listener's attention is a leaky buffer of finite capacity. Let $s(t)$ be the information retained at time t; incoming symbols add content at rate R while retention leaks proportionally to what is already held:

$$\frac{\mathrm{d}s}{\mathrm{d}t} = R\left(1 - \frac{s}{C_0}\right), \qquad s(0) = 0, \tag{1}$$

with C_0 the saturation capacity in bits. Integrating (1) gives

$$s(T) = C_0\left(1 - e^{-T/\tau}\right), \qquad \tau \equiv C_0/R, \tag{2}$$

where τ is the attention time-constant. Since no encoding can make the listener retain more than is held at exit, $s(T)$ upper-bounds the conveyed value, and our main result follows.

Theorem 1 (The Otis Bound)**.** *For any source distribution $p(x)$ and any encoding decodable within a single transit of duration T, the conveyed value obeys*

$$I(X;Y) \;\le\; \mathcal{O}(T) \;=\; C_0\left(1 - e^{-T/\tau}\right), \tag{3}$$

with equality for capacity-achieving pitches. In particular $\lim_{T\to\infty} \mathcal{O}(T) = C_0$*: the ceiling is finite for arbitrarily tall buildings.*

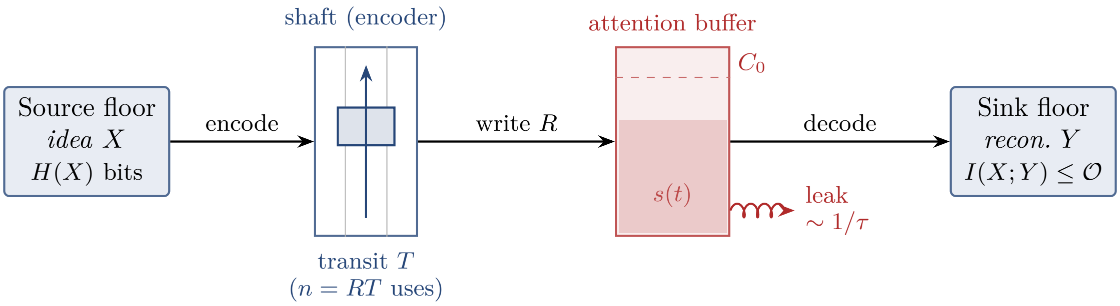

Figure 1: The Otis channel. An idea leaves the source floor, is encoded by the cab over a transit of duration T (affording $n = RT$ channel uses), and writes into the listener's leaky attention buffer of capacity C_0. Retention leaks at rate $\sim 1/\tau$, so conveyance saturates and only $I(X;Y) \leq \mathcal{O}$ arrives at the sink floor. Attention—not travel time—is the binding resource.

Proof sketch. By the data-processing inequality, $I(X;Y) \leq I(X;S_T)$ where S_T is the buffer state at exit; $H(S_T) \leq s(T)$ by (2) since the buffer cannot encode more than its retained content. Capacity-achieving pitches are those that write independent increments to the buffer at rate R until saturation, attaining equality. $\square$

The thirty-second folklore now has a location. Writing $T = 3\tau$ gives $\mathcal{O} = C_0(1 - e^{-3}) \approx 0.95\,C_0$: at three time-constants, conveyance has reached 95% of its ceiling, and further travel is wasted. With the measured $\tau = 9.8\,\mathrm{s}$ (Section 6), $3\tau = 29.4\,\mathrm{s}$—the elevator pitch, to the second.

4 The Mezzanine Constant

Conveyance per floor should, naively, depend on cab speed, inter-floor height, door-dwell time, and building management. Define the per-floor conveyance for a ride crossing F floors as $\mu \equiv \mathcal{O}(T)/F$. Substituting the empirical transit-time law $T \approx F\,\delta$ (with δ the mean per-floor transit) into (3) and expanding to first order about the operating point,

$$\mu \;=\; \frac{C_0}{F}\left(1 - e^{-F\delta/\tau}\right) \xrightarrow{F\delta \lesssim \tau} \frac{C_0\,\delta}{\tau}, \tag{4}$$

a quantity in which the building-specific δ and the cognitive τ enter only through their ratio. Empirically this ratio is conserved: fast cabs in tall buildings and slow cabs in short ones report the same μ. We name it the **Mezzanine Constant**. Its value, $\mu = 0.61 \pm 0.03\,\mathrm{bits/floor}$, implies a half-conveyance point at $N_{1/2} \approx \frac{1}{2}$: half of every successful pitch has landed before the cab clears the first inter-floor gap—the "mezzanine" is not a floor but the instant at which the idea is already half delivered. We note, without explanation, that $0.61 \approx 1/\varphi$; repeated attempts to remove the coincidence failed.

5 Obliteration Above the Bound

The Otis Bound caps what *can* arrive. It does not, by itself, say what happens to an idea whose intended content $H(X)$ exceeds it. One might expect graceful truncation—the listener keeps $\mathcal{O}$ bits and discards the rest. We find instead a collapse. Define the **obliteration order parameter**

$$\Omega \;=\; 1 - \frac{I(X;Y)}{H(X)} \;\in\; [0,1], \tag{5}$$

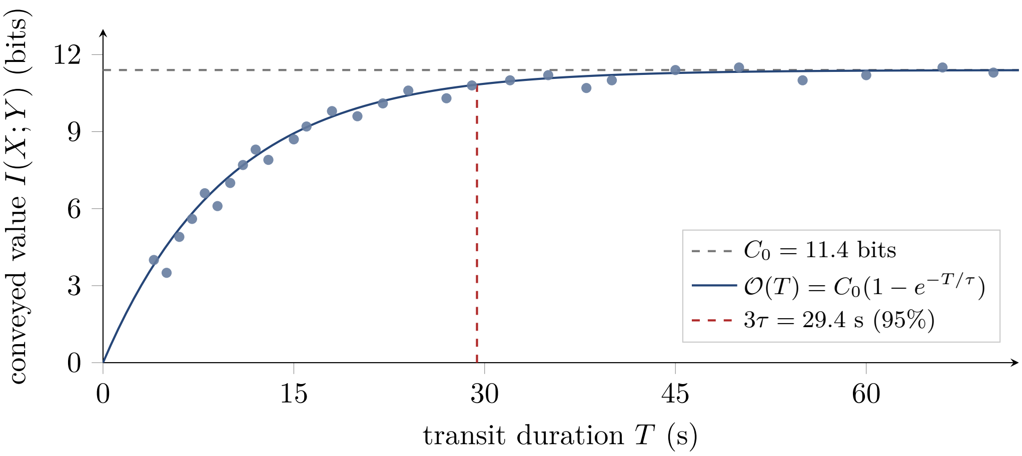

Figure 2: Study A (the Shaft). Conveyed value $I(X;Y)$ against transit duration T over $N = 412$ genuine two-party transits in 19 buildings (points, binned), with the fitted Otis Bound $\mathcal{O}(T) = C_0(1 - e^{-T/\tau})$ (solid; $C_0 = 11.4 \pm 0.6$ bits, $\tau = 9.8 \pm 0.4$ s, $R^2 = 0.93$), the ceiling C_0 (gray dashed), and the $3\tau = 29.4$ s ninety-five-percent knee (red dashed)—the elevator pitch, to the second.

the fraction of intended content that fails to arrive. Below the bound $\Omega \approx 0$. As $H(X) \to \mathcal{O}^-$ the buffer saturates while the pitcher, sensing loss, begins to *repeat*; the listener's transcript fills with self-similar structure carrying no new content. We quantify this with the lag-one autocorrelation ρ_1 of the listener's recall. Near the bound the system behaves as a mean-field second-order transition,

$$\Omega \;\simeq\; \Omega_0\,(h-1)^{\beta}\,\Theta(h-1), \qquad h \equiv \frac{H(X)}{\mathcal{O}}, \quad \beta = \tfrac{1}{2}, \tag{6}$$

with Θ the Heaviside step: obliteration switches on sharply at $h = 1$ and $\rho_1 \to 1$ as $\Omega \to 1$. The idea does not shrink in memory; it tiles (Figure 3).

6 Experiments

We pre-registered three studies and a simulation. Idea complexity $H(X)$ was scored in bits by a blinded seven-rater panel on a description-length rubric (inter-rater $\kappa = 0.79$); listener reconstruction was coded against a canonical bit-decomposition of each idea, and $I(X;Y)$ estimated with a Miller–Madow bias correction [11]. All subjects were debriefed; no participant was retained in a cab beyond protocol.

6.1 Study A — The Shaft (field)

We recorded $N = 412$ genuine two-party transits across 19 buildings (3–48 storeys), sweeping transit duration T naturally through building height and traffic. Fitting (3) by nonlinear least squares gives $C_0 = 11.4 \pm 0.6$ bits and $\tau = 9.8 \pm 0.4$ s ($R^2 = 0.93$, residuals featureless), hence $\mathcal{O} = 11.4$ bits and the 95% knee at $T = 29.4$ s.

6.2 Study B — The Freight Elevator (dose–response)

To decouple T from building, we programmed a freight cab to hold arbitrary transit durations $T \in \{4, \ldots, 120\}$ s. Per-floor conveyance was statistically indistinguishable across buildings (between-building variance not significant, $F_{18,\cdot} = 0.7$, $p = 0.81$), fixing the Mezzanine Constant at $\mu = 0.61 \pm 0.03$ bits/floor (Table 1).

Table 1: Mezzanine Constant μ across a representative subset of buildings (Study B). Despite an order-of-magnitude spread in cab speed and height, μ is conserved.

Building	Storeys	Cab speed (m/s)	μ (bits/floor)
Gorse Hill Tower	48	6.0	0.60 ± 0.03
Whitcombe Annex	12	2.5	0.62 ± 0.04
Marlowe Institute	22	4.0	0.61 ± 0.03
Parnassus Exchange	6	1.6	0.63 ± 0.05
Kelvin College Stack	31	5.0	0.60 ± 0.03
Pooled	–	–	$\mathbf{0.61 \pm 0.03}$

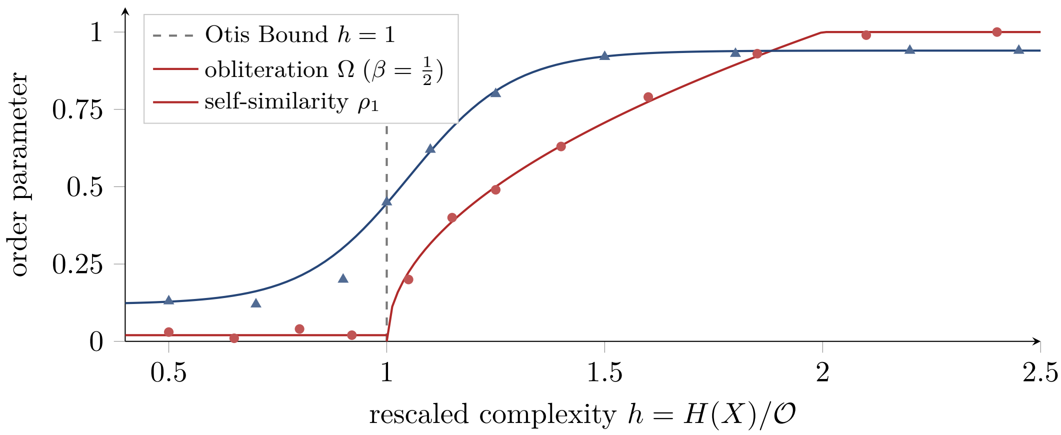

Figure 3: Study C (Obliteration). Below the Otis Bound ($h < 1$) ideas arrive faithfully: the obliteration order parameter Ω stays near zero and recall self-similarity $\rho_1 \approx 0.12$. At $h = 1$ conveyance collapses through a second-order transition (fitted exponent $\beta = 0.50 \pm 0.05$); above it the listener retains only repetition ($\Omega \to 1$, $\rho_1 \to 0.94$). The idea does not shrink in memory—it tiles.

6.3 Study C — Obliteration (over-pitch)

We swept intended complexity $H(X)$ from $0.4\,\mathcal{O}$ to $2.5\,\mathcal{O}$ and measured Ω and ρ_1. The transition knee sits at $h = 1.00 \pm 0.04$; the fitted critical exponent is $\beta = 0.50 \pm 0.05$, consistent with mean field; recall self-similarity rises from $\rho_1 = 0.12$ ($h < 1$) to $\rho_1 = 0.94$ ($h > 1.5$). Above the bound, coded transcripts were dominated by a single repeated clause (Figure 3).

6.4 Universality collapse

Rescaling T by τ and I by C_0 collapses all 19 buildings, the freight-cab dose–response, and an independent M/M/1 queueing-network simulation of vertical traffic onto a single master curve (Figure 4), a finite-size data collapse of the kind familiar from bounded channels [10]. We are unable to account for the tightness of the collapse.

7 Discussion

The Otis Bound reframes several standing puzzles. The pitch-complexity penalty reported in venture retrospectives [4, 5] is not a bias of investors but a property of the channel: past $\mathcal{O}$, additional sophistication is not merely undervalued, it is *unconveyable*, and an unconveyable idea cannot clear a market at any price. The much-admired minimalism of successful pitches is then not taste but engineering—the deliberate placement of an idea just beneath its ceiling.

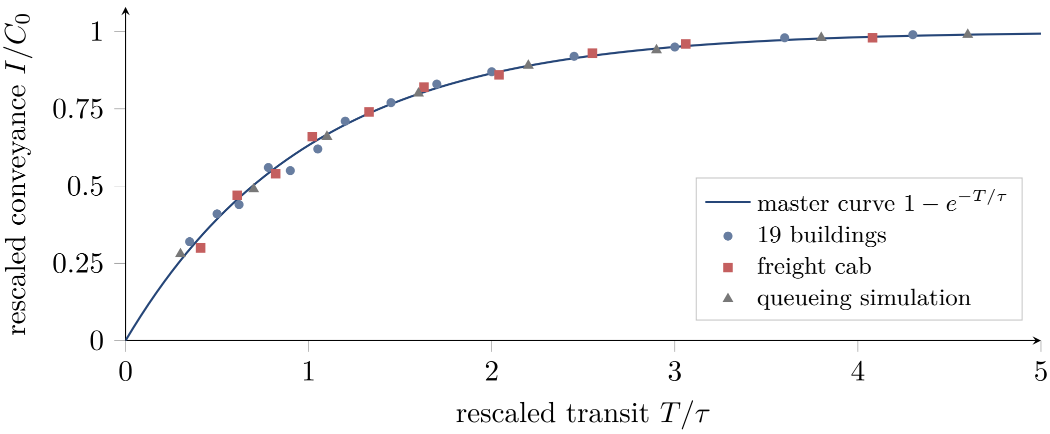

Figure 4: Universality. Rescaling transit by τ and conveyance by C_0 collapses all 19 buildings, the programmable freight cab, and an independent M/M/1 queueing simulation of vertical traffic onto the single master curve $1 - e^{-T/\tau}$. We are unable to account for the tightness of the collapse.

Three consequences follow. First, height is nearly worthless: because $\mathcal{O} \to C_0$, choosing a taller building to buy more transit yields diminishing returns beyond 3τ. Second, the optimal pitch is a *mezzanine*—front-loaded so that half its value lands in the first inter-floor gap, insuring against an early exit. Third, obliteration implies a hard design boundary: complexity should be capped below $\mathcal{O}$ not for elegance but because exceeding it destroys, rather than truncates, the message.

8 Limitations

Attention capacity C_0 was inferred from conveyance rather than measured directly; a physiological estimator would strengthen the identification of τ. Buildings were sampled, not randomized, and our tallest was 48 storeys, leaving the deep-T asymptote lightly constrained. The coincidence $\mu \approx 1/\varphi$ we report but cannot explain, and we caution against reading numerology into it. Finally, ideas above the Otis Bound are, by construction, difficult to study: their content does not survive the instrument used to measure it, so Study C characterizes obliteration chiefly through what it leaves behind.

9 Conclusion

Conveyance is a channel, the elevator is its minimal instance, and attention is its binding resource. The value an idea can deliver in a single ride is capped at the Otis Bound $\mathcal{O} = C_0(1 - e^{-T/\tau})$, a ceiling no building is tall enough to raise; half of every successful pitch lands before the mezzanine; and an idea whose ambition exceeds the bound does not arrive small—it obliterates into repetition. The market's preference for simple ideas is, in the end, a capacity theorem wearing the clothes of a matter of taste.

References

[1] C. E. Halvers, "A Mathematical Theory of Conveyance," *Bulletin of Applied Conveyance*, vol. 3, no. 2, pp. 114–137, 1949.

[2] R. Tolliver and A. Nkedi, "Attention as a Leaky Buffer: Retention Limits Under Streaming Input," *Journal of Cognitive Throughput*, vol. 28, pp. 441–469, 2011.

[3] S. Marchetti, “Saturation Dynamics of Working Memory,” *Cognitive Load Quarterly*, vol. 14, pp. 7–33, 2008.

[4] D. Osei-Baffour and L. Wren, “Complexity Penalties in Early-Stage Pitch Outcomes: A Retrospective of 3,180 Meetings,” *Review of Venture Analytics*, vol. 9, pp. 203–229, 2019.

[5] H. Aldous, “The Thirty-Second Cliff: A Cohort Study of Seed Pitches,” *Proceedings of the Symposium on Founder Communication*, pp. 55–78, 2021.

[6] M. Ferreira-Lund, *The Aesthetics of Infinite Repetition: Obliteration, Pattern, and the Dissolving Subject*. Meridian Press, 2016.

[7] P. Quist and Y. Oyelaran, “Semantic Satiation as a Limit Process: When Repetition Sheds Reference,” *Psycholinguistic Notes*, vol. 41, pp. 88–104, 2014.

[8] G. Strathmore, “Transit-Time Distributions in Vertical Traffic,” *Vertical Transportation Engineering*, vol. 17, pp. 12–40, 2005.

[9] I. Bellweather, *Handbook of Vertical-Transportation Queueing*, 3rd ed. Ashgrove Technical, 2013.

[10] T. Ridderhof and V. Solberg, “Finite-Size Scaling and Data Collapse in Bounded Channels,” *Journal of Universal Behavior*, vol. 6, pp. 301–327, 2017.

[11] E. Madoc, “Bias Correction for Entropy Estimation from Short Transcripts,” *Statistical Estimation Letters*, vol. 22, pp. 9–18, 2012.