The Limits of Voluntary Transparency in Competitive Markets*

Terrence Hendershott
Haas School of Business
University of California, Berkeley

Dan Li
Board of Governors of the Federal Reserve System

Dmitry Livdan
Haas School of Business
University of California, Berkeley

Norman Schürhoff
Faculty of Business and Economics and Swiss Finance Institute
University of Lausanne

March 21, 2026

*Terrence Hendershott is with the Walter A. Haas School of Business, University of California, Berkeley, Berkeley CA 94720. Tel.: (510) 643-0619; fax: (510) 643-1412; e-mail: hender@haas.berkeley.edu. He provides expert witness services to a variety of clients. He has taught a course for a financial institution that engages in liquidity provision and high-frequency trading activity. He gratefully acknowledges support from the Norwegian Finance Initiative. Dan Li is with the Board of Governors of the Federal Reserve System, Washington, DC. E-mail: dan.li@frb.gov. She has no conflict of interest to disclose. The views of this paper do not represent the views of the Federal Reserve Board. Dmitry Livdan is with the Walter A. Haas School of Business, University of California, Berkeley, Berkeley CA 94720. fax: (510) 207-3759; e-mail: livdan@haas.berkeley.edu. Livdan is also a Research Fellow of the CEPR and has no conflict of interest to disclose. Norman Schürhoff is with the Faculty of Business and Economics and Swiss Finance Institute, University of Lausanne, CH-1015 Lausanne, Switzerland. Tel.: +41 (21) 692-3447; fax: +41 (21) 692-3435; e-mail: norman.schuerhoff@unil.ch. Schürhoff is also a Research Fellow of the CEPR. He gratefully acknowledges research support from the Swiss Finance Institute and has no conflict of interest to disclose.

The Limits of Voluntary Transparency in Competitive Markets

Abstract

In markets for structured credit, dealers voluntarily share price estimates with competitors before sealed-bid auctions. Using 53,000 auctions, we document that these price talks are informative, yet silent dealers win 55% of auctions. To reconcile these facts, we show theoretically that voluntary communication before competition produces anti-unraveling. Dealers with favorable information remain silent to preserve their competitive advantage in bidding, while dealers with unfavorable information disclose to soften competition. Dealer talk improves price efficiency, but strategic silence generates allocative inefficiency as winner's curse-induced shading causes some nontalkers to underbid talkers.

JEL Classification. G12, G14, G23, D44, D82

Key words. Over-the-counter markets, dealer intermediation, pre-trade disclosure, transparency, auctions, CLOs

1 Introduction

We document that voluntary disclosure in competitive markets generates the opposite of the standard unraveling prediction. In the canonical framework of Grossman (1981) and Milgrom (1981), agents with favorable information disclose to separate from low types, silence signals bad news, and the resulting unraveling leads to full revelation. We show that in competitive auctions this logic reverses. Bidders with the most favorable private information stay silent, while those with less favorable information disclose. The mechanism is straightforward. In an auction, the audience for disclosure is a set of strategic rivals. Revealing an optimistic signal strengthens competitors' bids and erodes the discloser's surplus. The result is an anti-unraveling equilibrium in which silence reflects positive rather than negative information.

Our setting is the U.S. collateralized loan obligation (CLO) market, where secondary trading occurs through "bids-wanted-in-competition" (BWIC) auctions. Before submitting binding bids, dealers voluntarily share non-binding price estimates, known as "price talks," with all market participants. These disclosures are unverifiable and shared directly with rivals who compete in the subsequent sealed-bid auction. We develop a common-value auction model with a pre-auction disclosure stage that explains why such communication arises and why it is strategically selective. In the model, each dealer receives a noisy private signal about the asset's value and chooses whether to disclose it before bidding. The equilibrium features a threshold rule in which dealers with signals below a cutoff talk and those above remain silent. We prove that this equilibrium exhibits a three-region structure. Low-signal dealers disclose and bid a common amount based on the public consensus. Intermediate-signal dealers stay silent but are not competitive because the winner's curse pushes their bids below the talker bid. High-signal dealers stay silent and bid aggressively above the talker bid, winning the auction whenever their private information is sufficiently favorable.

This three-region structure generates a distinctive crossing pattern that serves as our identifying empirical prediction. When a talker wins, the transaction price is anchored by the public consensus. When a nontalker wins, the price reflects favorable private information that was deliberately withheld. The gap between the transaction price and the pre-auction consensus should therefore be larger when nontalkers win, and the gap between the transaction price and the highest individual

talk should narrow. Moreover, these effects should increase with fundamental uncertainty, because the informational advantage of silence is greatest when valuation is hardest.

We test these predictions using the universe of BWIC auctions for U.S. CLOs from 2012 through 2022. The data, sourced from Empirasign, contain over 53,000 auctions and 430,000 dealer-level price talks. We match auction data with transaction prices from TRACE and security characteristics from Bloomberg. Price talk is nearly universal, with 94% of auctions receiving at least one dealer estimate and an average of eight dealers participating.

The empirical evidence strongly supports anti-unraveling. Nontalkers win 55% of auctions at prices systematically above the pre-auction consensus. Our identifying test compares transaction prices to the consensus and to the highest individual talk across winner types. The price premium (price minus consensus) is 8.1 basis points larger when a nontalker wins ($p < 0.01$), and the max-talk crossing (price minus highest talk) narrows by 26.2 basis points ($p < 0.01$). Both effects increase monotonically with tranche risk, from 2.9 and 5.9 basis points for AAA/AA to 13.2 and 19.6 basis points for equity. Within-dealer tests confirm that the same dealer achieves a price premium 9.9 basis points larger when winning as a nontalker than as a talker ($p < 0.01$), ruling out dealer-type sorting. The seniority gradient persists within dealer, rising from 2.6 basis points for AAA/AA to 18.8 basis points for equity. We also directly test the model's three-region structure by exploiting the list structure of BWIC auctions to identify dealers who are present but silent. Among these present-and-silent dealers, talkers win 19% of the time, a sharp rejection of the two-region alternative. Cover bid classification and dealer-level analysis confirm that the same dealer talks on some securities, is non-competitively silent on others, and wins silently on yet others.

Supporting evidence shows that price talk is genuinely informative despite being strategically selective. Consensus surprises predict price surprises with a coefficient of 0.95 and an R^2 of 0.96. Prices in talker-won auctions load primarily on others' signals (73%) rather than the winner's own (27%), consistent with information aggregation. Bid dispersion is roughly half of talk dispersion, indicating substantial Bayesian learning from disclosed signals. These patterns confirm that talk aggregates dispersed private information effectively, even though the most favorable signals are systematically withheld.

Our market-quality analysis reveals a nuanced pattern. Nontalker wins generate prices 8–13

basis points further above public benchmarks, confirming that hidden signals inflate valuations. Yet this price inflation does not translate into proportionally larger winner's surplus, because the cover bid (the second highest bid) adjusts upward alongside the transaction price. The auction mechanism limits the intensive margin of rent extraction even when information asymmetry moves the level of prices. The cost of strategic silence thus falls primarily on price efficiency rather than on the spread between the top two bids.

An important implication of the model, observed in the data, is that price talk can distort allocations by shifting winning probabilities away from the highest-value bidders. In equilibrium, talkers win precisely in states where all silent dealers have signals below the crossing threshold; because high-signal dealers optimally remain silent, talker wins occur when no sufficiently strong private information overturns the public consensus. As a result, a silent dealer with a higher private valuation may not bid aggressively due to the winner's-curse adjustment, allowing a talker with a lower, publicly based valuation to win the asset. This inefficiency reflects a fundamental tradeoff between informational efficiency and market quality. By aggregating dispersed information, price talk reduces uncertainty, lowers trade failure rates, and increases execution probability, but it also weakens the incentives of privately informed dealers to compete aggressively. Consequently, price talk improves market functioning while attenuating the link between allocation and underlying valuations, highlighting a central tension in decentralized OTC markets.

Related literature. Our paper contributes to four literatures. To the literature on OTC market structure (Duffie, Gârleanu, and Pedersen, 2005; Lagos and Rocheteau, 2007; Hugonnier, Lester, and Weill, 2020; Di Maggio, Kermani, and Song, 2017; Li and Schürhoff, 2019; Hollifield, Neklyudov, and Spatt, 2017), we show that dealers provide *information services* beyond balance sheet intermediation. Pre-trade disclosure aggregates dispersed private signals and improves price discovery, a function distinct from and complementary to the liquidity provision role emphasized in prior work.

To disclosure theory (Grossman, 1981; Milgrom, 1981; Verrecchia, 1983; Dye, 1986; Matthews and Postlewaite, 1985), we provide the first large-scale empirical evidence of anti-unraveling. The mechanism arises from competitive exposure rather than disclosure costs or uncertainty about infor-

mation endowment, the two standard exceptions to unraveling. In our setting, favorable information is a liability to disclose because the audience consists of strategic rivals rather than passive investors.

To auction theory (Milgrom and Weber, 1982; Persico, 2000; Ganuza and Penalva, 2010; Campbell and Levin, 2000; Kagel and Levin, 2015), we confirm the linkage principle empirically and provide novel evidence on bidder-initiated information sharing. While prior work studies information disclosure by auctioneers, our setting features voluntary disclosure by bidders themselves, creating strategic incentives that limit full revelation.

To transparency policy (Bessembinder, Maxwell, and Venkataraman, 2006; Edwards, Harris, and Piwowar, 2007; Goldstein and Hotchkiss, 2020; Hendershott and Madhavan, 2020; Boyarchenko, Lucca, and Veldkamp, 2021), we show that voluntary pre-trade transparency can emerge endogenously but with efficiency costs. Regulators have focused primarily on mandated post-trade transparency. Our findings suggest that pre-trade disclosure is important for market quality, but voluntary mechanisms cannot fully substitute for mandates when competitive incentives create selective disclosure.

2 Institutional Setting and Data

2.1 BWIC auctions

Collateralized Loan Obligations (CLOs) are structured credit vehicles that pool leveraged loans and issue tranched securities. The U.S. CLO market exceeded $1 trillion in outstanding securities by 2022. The secondary market is characterized by opacity and infrequent trading. Dealers play a central role in intermediating trades.

The primary mechanism for secondary market trading is the bids-wanted-in-competition (BWIC) auction. A BWIC is an email-based first-price sealed-bid auction with a distinctive two-stage structure. In the first stage, before the bidding deadline, dealers voluntarily provide price talks, non-binding estimates of fair value expressed as a percentage of par. The broker disseminates all talks to market participants. Providing talk incurs no direct cost but reveals information to competitors. In the second stage, dealers submit binding bids by the deadline (typically the same

day). The highest bid wins and determines the transaction price. The second-highest bid, known as the cover, is sometimes disclosed after the auction.

We define notation as follows. For BWIC i with N dealers and $N_i \leq N$ providing talks, let S_{ij} denote the price talk of dealer j expressed as a percentage of par, and let the average talk (consensus) be $\bar{S}_i = N_i^{-1} \sum_{j=1}^{N_i} S_{ij}$. The transaction price is P_i and the cover is C_i. Let $S_i^{1:N_i}$ denote the highest talk among the N_i talkers.

2.2 Data

We combine data from multiple sources. Empirasign provides the universe of BWIC auctions, including dealer-level price talks, auction outcomes, and cover prices when disclosed. TRACE provides transaction prices for executed trades. Bloomberg provides CLO characteristics such as ratings and collateral features. Market conditions data, including VIX and credit spreads, come from FRED.

Our sample covers 2012 through 2022 and includes 53,052 BWIC auctions and over 430,000 dealer-level price talks. Table 1 reports summary statistics. Panel A shows that the average auction features approximately eight talking dealers, and 93.6% of auctions receive at least one price talk. Transaction prices average 94.17% of par with sizable dispersion reflecting heterogeneity in tranche quality and market conditions. The average difference between the transaction price and the cover is 24 basis points, consistent with competitive bidding, while the failure rate of roughly 20% indicates that many auctions do not result in execution. Panel B shows substantial heterogeneity across tranche seniority. Price talk becomes more prevalent for riskier tranches, while auction failure rates increase monotonically with risk, from 8.2% for AAA to over 34% for B-equity tranches.

2.3 Two motivating facts

We start by isolating dealer learning from price talk. We construct a machine-learning benchmark price $\widehat{P}_i^{\mathrm{ML}}$ that predicts transaction prices using only observable security characteristics, excluding any information from dealer disclosures. We train a neural network on features including tranche seniority, vintage, age, yield-to-maturity, deal size, and market conditions (VIX, credit spreads).

Table 1: Summary Statistics

This table reports summary statistics for the sample of BWIC auctions from 2012–2022.

	Mean	Median	SD	P25	P75	Obs
Panel A. Auction-Level						
Number of talkers (N_i)	8.01	8	4.50	5	11	53,052
Has price talk (%)	93.6	100	24.5	100	100	53,052
Transaction price P_i (% par)	94.17	98.57	11.44	94.00	99.96	39,284
Cover disclosed (%, cond. on trade)	77.9	100	41.5	100	100	42,611
Price − Cover (pp of par)	0.24	0.10	0.42	0.02	0.31	30,451
BWIC failure rate (%)	19.7	0	39.8	0	0	53,052
List size (CUSIPs)	8.37	7	6.98	4	10	53,052
Panel B. By Tranche Seniority						
	AAA	AA	A	BBB	BB	B-Eq
Has price talk (%)	88.1	93.7	95.0	96.4	97.2	97.0
Failure rate (%)	8.2	12.6	15.3	20.2	25.9	34.2
Obs	12,039	4,745	4,902	8,170	11,354	5,869

The model is estimated using five-fold cross-validation to generate out-of-sample predictions. The ML benchmark achieves an out-of-sample R^2 of 0.695. This benchmark allows us to define price and consensus *surprises* relative to fundamentals,

$$\tilde{P}_i \equiv P_i - \widehat{P}_i^{\mathrm{ML}}, \qquad \tilde{\bar{S}}_i \equiv \bar{S}_i - \widehat{P}_i^{\mathrm{ML}}. \tag{1}$$

Table 2 presents two motivating facts. Panel A shows that nontalkers win the majority of auctions (55.0%), and transaction prices are significantly higher when nontalkers win, with an average difference of 0.40 percentage points of par. The hidden information premium, the gap between the transaction price and the consensus, is 9.3 basis points larger in nontalker-won auctions ($p < 0.01$). The difference is especially pronounced when few dealers participate in talk (17 bps vs. 8 bps with many talkers), indicating that limited disclosure amplifies the role of private information.

Panel B shows that price talk is not cheap talk. The regression

$$\tilde{P}_i = \beta \cdot \tilde{\bar{S}}_i + \alpha_t + \epsilon_i \tag{2}$$

yields $\hat{\beta} = 0.947$ with $R^2 = 0.958$, indicating near-unity pass-through from consensus to prices. The coefficient is lower when nontalkers win ($\beta = 0.935$ vs. 0.962), and R^2 drops from 0.971 to 0.952,

Table 2: Motivating Facts

Panel A compares auction outcomes by winner type. The hidden information premium is the difference between the transaction price and the average talk (in percentage points of par). Panel B reports regressions of price surprises on consensus surprises, both measured relative to the ML benchmark price. Standard errors clustered by security in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Auction Outcomes by Winner Type

	Talker wins	Nontalker wins	Difference
Share of auctions (%)	45.0	55.0	
Transaction price (% of par)	93.94	94.34	0.40***
			(0.08)
Price premium (Price – Consensus)	0.14	0.23	0.09***
			(0.01)
Price premium, few talkers ($N_i \leq 5$)	0.10	0.27	0.17***
			(0.02)
Price premium, many talkers ($N_i > 5$)	0.14	0.22	0.08***
			(0.01)

Panel B. Consensus Predicts Transaction Prices, Less So When Nontalker Wins

	Dep. var. Price Surprise $\tilde{P}_i$			
	All	Talker wins	Nontalker wins	Difference
Consensus surprise $\tilde{\tilde{S}}_i$	0.947***	0.962***	0.935***	-0.027***
	(0.003)	(0.004)	(0.004)	(0.006)
R^2	0.958	0.971	0.952	
Obs	37,026			

consistent with private information entering prices through bidding rather than through talk.

These two facts pose a puzzle. If disclosure is informative and benefits price discovery, why do informed dealers stay silent? And if silence is so profitable, why does anyone voluntarily disclose their valuation to competitors? The following section develops a model that resolves this puzzle.

3 Model

We develop a model of selective disclosure in BWIC auctions that explains the motivating facts and generates testable predictions. The key economic insight is that in competitive auctions, voluntary disclosure reverses the standard unraveling logic because revealing favorable information strengthens rivals' bids and erodes the discloser's surplus.

3.1 Setup

Consider a single BWIC auction with $N \geq 2$ risk-neutral dealers indexed by $j \in \{1, \ldots, N\}$. The payoff to dealer j from winning the security at price P is $v - P$, where v denotes the common value of the asset. CLOs exhibit substantial common-value characteristics because all dealers are ultimately trying to assess the same fundamental, i.e., the present value of cash flows from the underlying loan pool adjusted for credit risk and market conditions. While dealers may have idiosyncratic preferences related to inventory management or client demand, the dominant source of uncertainty is the security's fundamental value, which is common across market participants. This motivates our focus on the common-value paradigm.

The common value is drawn from a normal prior $v \sim \mathcal{N}(\mu_0, \sigma_v^2)$,[1] and each dealer j privately observes a noisy signal

$$s_j = v + \varepsilon_j, \qquad \varepsilon_j \overset{iid}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2), \tag{3}$$

where the noise terms $\{\varepsilon_j\}_{j=1}^N$ are mutually independent and independent of v. Let $\tau_v \equiv 1/\sigma_v^2$ and $\tau_\varepsilon \equiv 1/\sigma_\varepsilon^2$.

The game proceeds in two stages. In the disclosure stage, each dealer j simultaneously chooses whether to disclose her signal, selecting $T_j \in \{\text{Talk}, \text{Silent}\}$. If $T_j = \text{Talk}$, dealer j sends a public message $m_j \in \mathbb{R}$. Let $\mathcal{T} \subseteq \{1, \ldots, N\}$ denote the set of talkers with $n_T = |\mathcal{T}|$, and let $\mathcal{S} = \{1, \ldots, N\} \setminus \mathcal{T}$ denote the set of silent dealers with $n_S = N - n_T$.

In the bidding stage, after observing all disclosed messages $\mathcal{M} = \{m_j : j \in \mathcal{T}\}$, each dealer j submits a sealed bid $b_j \in \mathbb{R}$. The highest bid wins at price $P = \max_j b_j$. Define the consensus, for $n_T \geq 1$, as

$$\bar{S} \equiv \frac{1}{n_T} \sum_{j \in \mathcal{T}} s_j. \tag{4}$$

A talker $j \in \mathcal{T}$ has no private information at the bidding stage since her signal is public. A silent dealer $j \in \mathcal{S}$ retains private information in the form of her signal s_j.

We solve for Perfect Bayesian Equilibrium (PBE) in symmetric threshold strategies. A symmet-

[1] Normality is assumed for closed-form posteriors and the linear decompositions used in the empirical specifications. The equilibrium results (Propositions 1–3) require only that $(v, s_1, \ldots, s_N)$ are affiliated, that posterior expectations are well defined and continuous, and that the monotonicity condition of Lemma 2 holds. These conditions are satisfied by any location family with affiliated signals.

ric threshold equilibrium consists of a disclosure threshold $\bar{s} \in \mathbb{R} \cup \{+\infty\}$ such that dealer j talks if and only if $s_j \leq \bar{s}$, truthful disclosure with $m_j = s_j$ for all $j \in \mathcal{T}$, bidding strategies that are symmetric within talkers and within silent dealers, and beliefs consistent with Bayes' rule on the equilibrium path.

Standard conjugate updating yields the following posteriors (all proofs are in Appendix A). A talker conditions only on public information and has posterior $v \mid \mathcal{M} \sim \mathcal{N}(\mu_T, \sigma^2_{v|\mathcal{M}})$ with

$$\mu_T = \frac{\tau_v \mu_0 + n_T \tau_\varepsilon \bar{S}}{\tau_v + n_T \tau_\varepsilon}, \qquad \sigma^2_{v|\mathcal{M}} = \frac{1}{\tau_v + n_T \tau_\varepsilon}. \tag{5}$$

When $n_T = 0$, no signal is publicly disclosed and $\mathcal{M} = \emptyset$; in that case we interpret $\mu_T = \mu_0$ and $\sigma^2_{v|\mathcal{M}} = 1/\tau_v$. A silent dealer with signal s additionally conditions on her private signal and has a posterior mean

$$\mu_S(s, \mathcal{M}) = \frac{\tau_v \mu_0 + n_T \tau_\varepsilon \bar{S} + \tau_\varepsilon s}{\tau_v + (n_T + 1)\tau_\varepsilon} \equiv \mu_S(0, \mathcal{M}) + \gamma' s, \tag{6}$$

where $\gamma' = \tau_\varepsilon / [\tau_v + (n_T + 1)\tau_\varepsilon]$ measures the loading on the private signal. The consensus receives weight $\beta = n_T \tau_\varepsilon / (\tau_v + n_T \tau_\varepsilon) < 1$ in the talker posterior and a smaller weight $\beta' < \beta$ in the silent posterior. Both weights increase with n_T and converge to one as $n_T \to \infty$.

3.2 Equilibrium

We solve for the equilibrium using backward induction starting with the bidding stage and then finishing with the disclosure stage. The bidding stage is a first-price common-value auction with asymmetric information. Talkers share the same public information $\mathcal{M}$ and submit a common bid $b^T(\mathcal{M})$. Silent dealers retain private signals and use a strictly increasing bid function $b^S(s\,; \mathcal{M})$. Both sides shade their bids below their conditional valuations to account for the winner's curse and strategic competition. We characterize the continuation equilibrium through bidding functions (Appendix A.2).

The silent dealers bid

$$b^S(s;\mathcal{M}) = \mu_S(0,\mathcal{M})\big(1 - I(s \mid \mathcal{M})^{-1}\big) + I(s \mid \mathcal{M})^{-1} b^S(\bar{s};\mathcal{M}) + \\ + \big(1 - I(s \mid \mathcal{M})^{-1}\big)\, \mathbb{E}_X^{h_s}\Big[\gamma' X - \widetilde{\Lambda}_v(X;\mathcal{M})\Big], \tag{7}$$

where $I(s \mid \mathcal{M})$ is the integrating factor given by (26) and $\mathbb{E}_X^{h_s}[\cdot]$ denotes the expectation over X under kernel h_s given by (30). Equation (7) provides a transparent decomposition of the silent bidding strategy. The first term is the propagated contribution of public information. $\mu_S(0,\mathcal{M}) = \mathbb{E}[v \mid s = 0, \mathcal{M}]$ is the posterior mean based on public talks alone, evaluated at zero private signal, and the factor $1 - I(s \mid \mathcal{M})^{-1}$ determines how strongly this public benchmark enters the bid of type s. The second term carries the boundary condition at the disclosure threshold. The bid of the lowest silent type, $b^S(\bar{s};\mathcal{M})$, fixes the overall level of the silent bidding schedule, and the factor $I(s \mid \mathcal{M})^{-1}$ transmits that boundary value across higher types.

The final term is an average net informational surplus. The component $\gamma' X$ captures the incremental contribution of private information, while

$$\widetilde{\Lambda}_v(s;\mathcal{M}) \equiv \mathbb{E}[v \mid s, \mathcal{M}] - \mathbb{E}[v \mid s, \mathcal{M}, S_{(1)} = s], \tag{8}$$

is the exact boundary-state $s = S_{(1)}$, where $S_{(1)}$ is the highest signal among other silent dealers, winner's-curse correction. Their difference is averaged under the endogenous kernel h_s, which is induced by the first-order condition. This kernel is not, in general, literally the conditional distribution of the top rival signal given winning. Rather, it is the first-order-condition-induced weighting measure that summarizes how different rival-signal states affect equilibrium bid shading.

Note that $b^S(\bar{s};\mathcal{M})$ is a price level, not a surplus, and it is pinned down endogenously by the continuation equilibrium at the lowest silent type. The lowest silent dealer need not bid zero, and need not be indifferent between bidding and not bidding. In particular, in the anti-unraveling equilibrium developed below, the disclosure cutoff lies strictly below the crossing signal, so the marginal silent type need not be competitive against talkers. Thus $b^S(\bar{s};\mathcal{M})$ must be kept as an explicit boundary term rather than normalized away.

As with silent dealers, the exact first-order condition for talkers uses the boundary state $B_{(1)} = b^T$, where $B_{(1)}$ is the maximum bid among all silent dealers defined by (33). Define the talker's win-event valuation

$$V^T(b; \mathcal{M}) \equiv \mathbb{E}[v \mid \mathcal{M}, B_{(1)} < b],$$

then talkers bid

$$b^T(\mathcal{M}) = \mu_T - \underbrace{\Big(\mu_T - V^T(b^T(\mathcal{M}); \mathcal{M})\Big)}_{\text{winner's-curse / selection adjustment}} - \underbrace{\frac{F_B(b^T(\mathcal{M}) \mid \mathcal{M})}{f_B(b^T(\mathcal{M}) \mid \mathcal{M})}}_{\text{strategic shading}}, \tag{9}$$

where $F_B(b \mid \mathcal{M})$ and $f_B(b \mid \mathcal{M})$ are the CDF and density of the highest silent bid, respectively. The second term is an economic interpretation rather than an exact restatement of the talkers' first-order condition (34). Talkers share the same public information, but winning is bad news because it means silent dealers did not submit high enough bids to beat them.

The continuation equilibrium summarizes the trading stage of the model.

Proposition 1 (Continuation equilibrium)**.** *Fix $\mathcal{M}$ and (n_T, n_S). Under standard regularity conditions, there exists a Bayes–Nash equilibrium of the bidding stage in which talkers bid a common amount $b^T(\mathcal{M})$ and silent dealers use a strictly increasing strategy $b^S(\cdot; \mathcal{M})$ characterized by (21) and (34).*

We next characterize the disclosure equilibrium, which generates our central result.

Proposition 2 (Selective disclosure equilibrium)**.** *There exists a threshold $\bar{s}^* \in \mathbb{R}$ such that in a symmetric threshold equilibrium, dealers with $s_j \leq \bar{s}^*$ talk by revealing $m_j = s_j$ while dealers with $s_j > \bar{s}^*$ remain silent. The threshold $\bar{s}^*$ satisfies $\Pi^S(s = \bar{s}^*; \bar{s}^*) = \Pi^T(s = \bar{s}^*; \bar{s}^*)$, where Π^S (given by (37)) and Π^T (given by (36)) are the expected profits of the silent dealer and a talker with a signal s. In equilibrium, high-signal dealers remain silent and low-signal dealers talk,*

$$\mathbb{E}[s_j \mid T_j = \text{Talk}] < \mathbb{E}[s_j] < \mathbb{E}[s_j \mid T_j = \text{Silent}].$$

The proof (Appendix A.3) shows that the net benefit of silence, $\Delta(s) \equiv \Pi^S(s; \bar{s}) - \Pi^T(s; \bar{s})$, satisfies single crossing. For low signals, the informational rent from silence is negligible and $\Delta(s) <$

0. For high signals, the rent dominates and $\Delta(s) > 0$. Continuity implies a unique cutoff at which the dealer is indifferent. The second key result establishes a three-region structure.

Proposition 3 (Three-region structure)**.** *In the symmetric threshold equilibrium, the disclosure cutoff $\bar{s}^*$ is strictly below the crossing signal $s^\times(\mathcal{M})$ at which $b^S(s^\times(\mathcal{M}); \mathcal{M}) = b^T(\mathcal{M})$.*[2] *This implies three distinct regions of dealer behavior.*

1. ***Talkers*** *($s \leq \bar{s}^*$). Reveal their signal and bid $b^T(\mathcal{M})$.*

2. ***Silent, non-competitive*** *($\bar{s}^* < s < s^\times$). Stay silent but bid below b^T, so they cannot win against talkers.*

3. ***Silent, competitive*** *($s \geq s^\times$). Stay silent and bid above b^T, winning the auction.*

The proof (Appendix A.3) uses Lemma 3, which shows the maximal expected payoff of a bidder without private information is weakly decreasing in public information. At the crossing signal $s^\times$, a dealer strictly prefers silence over talking because staying silent preserves her private information advantage, while talking would expose it to competitors. Since indifference between talking and silence occurs at $\bar{s}^*$, and silence becomes relatively more attractive as s increases, the cutoff must satisfy $\bar{s}^* < s^\times$.

The intermediate region is a distinctive implication of the model. Dealers in this range of signals optimally choose silence in order to avoid revealing unfavorable information to competitors, yet their private signals are not strong enough to overcome the winner's-curse adjustment embedded in equilibrium bids. As a result, these dealers remain silent but submit bids below the talker side and therefore rarely win the auction. In effect, they are "hiding in the crowd": silence protects them from revealing weak signals, but the same signals prevent them from bidding competitively. This region arises from the interaction of two forces that operate simultaneously in common-value auctions with endogenous disclosure: the informational rent from withholding a signal and the winner's-curse penalty associated with aggressive bidding. Dealers with intermediate signals prefer

[2] Because $b^S(\cdot; \mathcal{M})$ is strictly increasing, talkers win whenever the highest signal among silent dealers is below the crossing signal, and silent dealers with signals strictly above $s^\times(\mathcal{M})$ bid above the talker side. In the knife-edge case in which the highest signal among silent dealers is equal to the crossing signal, the winner is determined by tie-breaking.

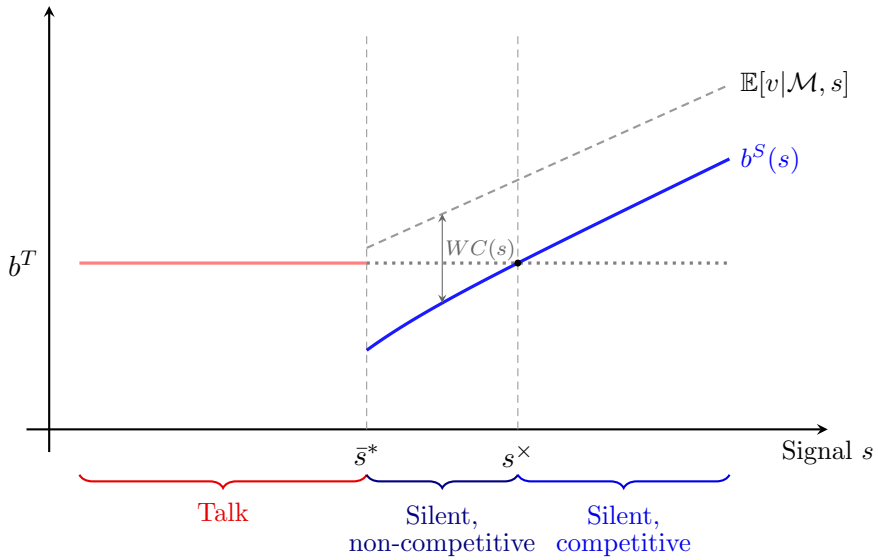


Figure 1: **Equilibrium bid schedule.** Dealers with $s \leq \bar{s}^*$ disclose their signal and bid the common talker bid b^T. Dealers with $s > \bar{s}^*$ remain silent and bid according to the increasing schedule $b^S(s) = \mathbb{E}[v|\mathcal{M}, s] - WC(s)$. For intermediate signals $\bar{s}^* < s < s^\times$, silent dealers bid below the talker bid and therefore do not win. Silent dealers bid above talkers only when $s \geq s^\times$, so nontalkers win only if at least one silent dealer draws $s \geq s^\times$. Talkers win precisely when all silent signals lie in $(\bar{s}^*, s^\times)$.

to keep their information private, but the adverse selection inherent in winning discourages them from bidding enough to beat talkers.

Figure 1 illustrates the equilibrium. The horizontal axis represents the dealer's private signal, and the vertical axis shows the submitted bid. Talkers submit the flat bid b^T, determined by the public consensus. Silent dealers bid according to the increasing schedule $b^S(s)$, which lies below b^T for intermediate signals and above b^T only for signals exceeding the crossing level $s^\times$. The gap $WC(s)$ between the conditional valuation $\mathbb{E}[v|\mathcal{M}, s]$ and the bid reflects the winner's curse, which is most severe near the disclosure threshold (where winning conveys the strongest adverse signal) and diminishes for high signals. Talkers win only when all silent signals fall in the intermediate region $(\bar{s}^*, s^\times)$ and no silent dealer submits a competitive bid. Nontalkers win only when at least one silent dealer draws $s \geq s^\times$ and bids aggressively above b^T. This asymmetry in the allocation mechanism generates the crossing pattern that we test empirically.

This three-region structure differs sharply from the predictions of classic static disclosure models such as Dye (1985) and the unraveling logic of Grossman (1981) and Milgrom and Roberts (1986). In those models, agents with favorable private information voluntarily disclose it in order to sep-

arate themselves from lower types, which often leads to full or near-full disclosure in equilibrium. The logic is that nondisclosure is interpreted pessimistically by the market, inducing higher types to reveal their information. In the BWIC auction environment, this reasoning reverses because disclosure occurs before a competitive bidding stage in which rivals directly incorporate the disclosed information into their bids. Revealing a favorable signal therefore strengthens competitors and erodes the discloser's informational rent. Consequently, high-signal dealers optimally remain silent, while lower-signal dealers reveal their information. The intermediate region captures the fact that some dealers prefer silence even though they are unlikely to win, simply to avoid revealing unfavorable information to the market.

The mechanism also differs from dynamic disclosure models such as Daley and Green (2012). In their setting, privately informed agents optimally delay disclosure in order to preserve informational rents, but information eventually unravels as time passes and uncertainty resolves. As a result, the equilibrium path features gradual revelation of private information: agents strategically postpone disclosure, yet higher types ultimately disclose as the market becomes more informed. In contrast, the BWIC auction environment studied here compresses the disclosure decision into a single pre-auction stage in which competition is immediate. Because rivals can directly use disclosed information to bid more aggressively, revealing a favorable signal immediately intensifies competition and reduces the bidder's expected payoff. This strategic exposure effect prevents unraveling: high-signal dealers strictly prefer to remain silent even in equilibrium. Thus, whereas information is gradually revealed in Daley and Green (2012), the presence of auction competition in our model sustains persistent nondisclosure of favorable information and generates the three-region structure described above.

3.3 Predictions

The model generates five empirical predictions.

Consensus predictiveness. The consensus $\bar{S}_i$ predicts the transaction price P_i with a coefficient $\beta < 1$, and this coefficient is lower when a nontalker wins ($\beta' < \beta$) because the winning bid reflects private information absent from the consensus. Both coefficients increase with n_T as more signals

reduce selection bias, and both are higher for senior tranches where valuation uncertainty is lower.

Price decomposition. When a talker wins, the transaction price loads on the winner's own signal and others' signals with a ratio approximately $1 : (n_T - 1)$, reflecting the wisdom-of-crowds structure of the consensus. Each talker's signal enters symmetrically, so the winner's own signal receives weight $1/n_T$.

Belief convergence. More talk reduces disagreement among dealers. The posterior variance $\text{Var}(v \mid \mathcal{M}) = 1/(\tau_v + n_T\tau_\varepsilon)$ is decreasing in n_T. Empirically, bid dispersion should be smaller than talk dispersion, with the ratio reflecting the degree of learning.

Crossing pattern. The price premium $P_i - \bar{S}_i$ is larger when nontalkers win. The gap between the transaction price and the highest talk, $P_i - S_i^{1:N}$, is more negative when talkers win. Both effects increase with fundamental uncertainty (junior tranches) and decrease with the number of talkers.

Comparative statics. The crossing gap increases for junior tranches (more common-value uncertainty), decreases with more talkers (more public information), increases with few talkers, and increases with higher VIX (more aggregate uncertainty).

4 Empirical Evidence

This section tests the model's predictions. We begin with the crossing test, which provides the sharpest evidence of anti-unraveling, then present within-dealer identification. We directly test the three-region structure of Proposition 3 and conclude with estimates of the information rent and supporting evidence on information aggregation.

4.1 The crossing test

The crossing test provides the paper's identifying prediction. The model implies that the gap between the transaction price and publicly observed talk should be systematically larger when nontalkers win, because nontalker victories reveal favorable private information that was deliberately withheld. We construct two measures that apply uniformly to both winner types. The *price premium* is $P_i - \bar{S}_i$. The *max-talk crossing* is $P_i - S_i^{1:N_i}$, the gap between the transaction price and the most optimistic disclosed signal.

Table 3: The Crossing Test

This table tests whether transaction prices deviate differently from the pre-auction consensus and from the highest individual talk depending on winner type. The price premium is $P_i - \bar{S}_i$. The max-talk crossing is $P_i - S_i^{1:N}$. Panel A reports summary statistics. Panel B reports regressions on a nontalker winner indicator with day fixed effects and standard errors clustered by security. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Summary Statistics

	Mean	Median	% Above Zero	Std. Dev.	Obs
Price − Consensus ($P_i - \bar{S}_i$)					
Talker wins	0.135***	0.101	60.9%	1.095	16,718
	(0.008)				
Nontalker wins	0.227***	0.120	63.3%	1.213	20,423
	(0.008)				
Difference	0.093***	0.019	2.4 pp		
	(0.012)				
Price − Highest Talk ($P_i - S_i^{1:N}$)					
Talker wins	−1.188***	−0.450	13.2%	1.995	16,718
	(0.015)				
Nontalker wins	−0.844***	−0.268	22.3%	1.913	20,423
	(0.013)				
Difference	0.344***	0.182	9.1 pp		
	(0.020)				

Panel B. Regression Evidence

	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Dep. var. Price − Consensus					
Nontalker wins	0.081***	0.029***	0.098***	0.118***	0.132**
	(0.012)	(0.008)	(0.020)	(0.031)	(0.060)
Dep. var. Price − Highest Talk					
Nontalker wins	0.262***	0.059***	0.178***	0.212***	0.196**
	(0.021)	(0.011)	(0.033)	(0.049)	(0.094)
Day FE	Yes	Yes	Yes	Yes	Yes
R^2 (premium)	0.191	0.423	0.359	0.327	0.380
R^2 (crossing)	0.207	0.400	0.405	0.385	0.409
Obs	37,026	12,906	9,361	9,742	4,061

Table 3 presents the results. Panel A reports summary statistics. The average transaction price exceeds the consensus by 22.7 basis points in nontalker-won auctions compared to 13.5 basis points in talker-won auctions, a difference of 9.3 basis points ($t = 7.64$). The max-talk crossing provides a sharper pattern. In talker-won auctions, the average price falls 1.19 percentage points below the highest talk. In nontalker-won auctions, this gap narrows by 34.4 basis points ($t = 16.9$). Nontalker winners are nearly twice as likely to push the price above the highest talk (22.3% versus 13.2%).

Panel B confirms these patterns in a regression framework with day fixed effects. The coefficient on Nontalker Winner is 0.081 ($p < 0.01$) for the price premium and 0.262 ($p < 0.01$) for the max-talk crossing. Both effects increase monotonically with credit risk. For the price premium, the coefficient rises from 2.9 basis points for AAA/AA to 13.2 basis points for equity. For the max-talk crossing, the range is 5.9 to 19.6 basis points. This seniority gradient is consistent with the model. Junior tranches have greater valuation uncertainty and larger information rents from silence, so the hidden signal of a nontalker winner produces a larger deviation between the consensus and the transaction price.

4.2 Within-dealer identification

The crossing results could reflect persistent dealer heterogeneity rather than strategic disclosure if certain dealers systematically win at higher premiums regardless of whether they talk. To address this concern, we restrict the sample to the 95% of auctions whose winning dealer also wins other auctions in the opposite talk status, ensuring within-dealer variation in disclosure behavior. We then re-estimate the crossing regressions with dealer fixed effects, so that identification comes exclusively from comparing the same dealer's outcomes across talk statuses.

Table 4 reports the results. Panel A shows that the same dealer achieves a price premium 9.9 basis points larger when winning as a nontalker ($p < 0.01$). The effect increases monotonically with credit risk, from 2.6 basis points for AAA/AA to 18.8 basis points for equity (all significant at $p < 0.01$). Panel B shows that the same dealer's winning price sits 17.6 basis points closer to the highest talk when winning as a nontalker ($p < 0.01$). The seniority gradient mirrors the pattern without dealer fixed effects, confirming that the strategic value of silence depends on the informativeness of the private signal rather than on dealer identity.

4.3 Direct test of the three-region structure

The crossing test (Section 4.1) and within-dealer identification (Section 4.2) establish that nontalker winners pay more than talker winners, consistent with anti-unraveling. However, these tests do not distinguish a two-region model (talkers below a cutoff, competitive silent dealers above) from the

Table 4: Within-Dealer Evidence

This table re-estimates the crossing regressions from Table 3 adding dealer fixed effects. The sample is restricted to auctions won by dealers who also win at least one auction in the opposite talk status (95% of auctions). All specifications include dealer and day fixed effects. Standard errors clustered by security in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Price Premium ($P_i - \bar{S}_i$)					
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Nontalker$_i$	0.099***	0.026***	0.071***	0.150***	0.188***
	(0.014)	(0.009)	(0.023)	(0.034)	(0.068)
R^2	0.198	0.435	0.371	0.334	0.398
Obs	34,922	12,065	8,859	9,213	3,786
Panel B. Max-Talk Crossing ($P_i - S_i^{1:N}$)					
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Nontalker$_i$	0.176***	0.068***	0.146***	0.230***	0.128
	(0.024)	(0.012)	(0.037)	(0.056)	(0.107)
R^2	0.228	0.413	0.419	0.394	0.425
Obs	34,922	12,065	8,859	9,213	3,786
Dealer FE, Day FE in all specifications					

three-region structure of Proposition 3. In a two-region model, every silent dealer bids above b^T and wins whenever she is the sole nontalker. In the three-region model, a range of silent dealers bid *below* b^T and never win. These dealers are "hiding in the crowd." We exploit the structure of BWIC lists to directly test for the existence of this intermediate, non-competitive region.

Identification through list-level dealer presence. The key challenge is identifying dealers who are present but choose to remain silent on a given security. We cannot observe the full set of bidders, but the list structure of BWIC auctions provides a natural solution. Each BWIC list contains multiple securities offered simultaneously. A dealer who submits a price talk on any security on a given list has demonstrably received that list and is active in that auction round. For every other security on the same list where this dealer does not talk, we can infer that the dealer was present (she saw the security) but chose to remain silent.

Consider dealer j who talks on CUSIP A but is silent on CUSIP B, where both appear on the same list. If a *talker* wins CUSIP B, then dealer j was present, silent, and outbid by a talker. In a two-region model this cannot happen. Silence implies $s_j > \bar{s}^*$, which implies $b^S(s_j) \geq b^T$ in a

two-region model, so a silent dealer should always outbid the talker bid. But in the three-region model, dealer j could have a signal in the intermediate range $(\bar{s}^*, s^\times)$, where she bids below b^T despite being silent. The frequency of such events provides a direct measure of the size of region 2.

We construct the test as follows. From the dealer-level talk data (430,921 observations), we identify 84,479 unique (dealer, list) pairs. For each pair, we expand to all securities on that list, producing 444,999 (dealer, security, list) observations. Of these, 14,093 are cases where the dealer was active on the list but did not talk on the security (present and silent). Restricting to securities that received at least one talk ($n_T \geq 1$), we have 11,604 present-and-silent observations for the core test.

The presence test. Table 5, Panel A reports the fraction of present-and-silent dealer-security pairs in which a talker won the auction. If only two regions exist, this fraction should be zero. We find that talkers win 19.1% of the time ($N = 11{,}604$), a sharp rejection of the two-region model.

We progressively tighten the definition of "present" to address the concern that a dealer active on senior tranches might not have information relevant for junior tranches on the same list. In the same-seniority restriction, we require that the dealer talked on another security of the same seniority group on the same list. Even within this restriction, talkers win 17.0% of the time ($N = 6{,}720$). In the same-deal restriction, we require that the dealer talked on another tranche of the same CLO deal. The rate is 16.2% ($N = 388$). All three estimates are of similar magnitude, suggesting that the presence of region 2 dealers is not driven by informational distance between securities.

Panel A also reports the test by seniority. For the baseline specification, the talker-won rate ranges from 17.5% (BB/B) to 22.4% (AAA/AA). For the same-deal restriction, it is highest for senior tranches (29.4% for AAA/AA) and lowest for BB/B (6.5%), though the sample is small. The fact that region 2 behavior is prevalent across the seniority spectrum rules out the possibility that non-competitive silence is confined to a particular asset class.

Runner-up classification. Panel B provides independent corroboration using the cover bid, the second-highest bid in each auction. The cover bid reveals information about the runner-up's type. If the cover falls below every individual talk price, the runner-up bid less than the most pessimistic

Table 5: Direct Test of the Three-Region Equilibrium

This table tests for the existence of silent, non-competitive dealers (region 2 of Proposition 3). Panel A uses list-level dealer activity to identify present-and-silent dealer-security pairs and reports the fraction in which a talker won the auction. In a two-region model this fraction should be zero. The any-activity specification identifies a dealer as present if she talked on any security on the same list. The same-seniority specification requires a talk on a security of the same seniority group. The same-deal specification requires a talk on another tranche of the same CLO deal. Panel B classifies runners-up (cover bidders) by whether their bid falls below all talks, within the talk range, or above all talks. Panel C reports dealer-level region appearances, where a dealer is classified as region 1 if she ever talks, region 2 if she is ever present-and-silent when a talker wins, and region 3 if she ever wins as a nontalker.

Panel A. Talker-Won Rate Among Present-and-Silent Pairs

	All	AAA/AA	A/BBB	BB/B	Equity
Any activity on list	0.191	0.224	0.179	0.175	0.198
	[11,604]	[1,610]	[2,266]	[3,447]	[4,281]
Same seniority	0.170	0.231	0.147	0.156	0.172
	[6,720]	[1,143]	[1,521]	[2,402]	[1,654]
Same deal	0.162	0.294	0.316	0.065	0.143
	[388]	[51]	[19]	[46]	[272]

Panel B. Runner-Up Classification by Cover Bid Position

	Talker-won	Nontalker-won	All
Below all talks (region 2)	10.0%	14.6%	12.6%
Within talk range	81.9%	71.0%	75.7%
Above all talks (region 3)	8.1%	14.4%	11.7%
Obs	9,258	12,205	21,463

Panel C. Dealer-Level Region Appearances

	Dealers	Share
Total dealers in sample	69	
Ever region 1 (talker)	69	100%
Ever region 2 (present, silent, talker won)	42	60.9%
Ever region 3 (won as nontalker)	22	31.9%
In all three regions	22	31.9%

talker, consistent with a silent, non-competitive dealer (region 2). If the cover exceeds every talk, the runner-up outbid all talkers without talking, consistent with a competitive silent dealer (region 3). We classify the runner-up into three categories. The first is below all talks, the second within the talk range, and the third above all talks.

In talker-won auctions, the cover falls below all talks 10.0% of the time. These are auctions where the runner-up was almost certainly a region 2 dealer who bid non-competitively despite being present. When the cover does fall below all talks, the median shortfall is 18 basis points and the mean is 47 basis points, indicating that these runners-up are well below the talk range rather

than marginally below.

In nontalker-won auctions, the cover exceeds all talks 14.4% of the time. In these auctions, both the winner and the runner-up outbid every talker without talking. Both are region 3 dealers. These patterns are consistent across seniority groups.

Same dealer in multiple regions. Panel C addresses whether the three-region structure reflects dealer-level heterogeneity (some dealers always talk, others always stay silent and compete) or signal-dependent behavior by the same dealer. We classify each dealer-security observation into one of the three regions. A dealer is in region 1 if she talked, region 2 if she was present and silent while a talker won, and region 3 if she won as a nontalker. We then ask how many dealers appear in multiple regions across different securities.

All 69 dealers in the sample talk at least once (region 1). Of these, 42 (60.9%) appear at least once in region 2, meaning they were present and silent on a security where a talker won. Twenty-two dealers (31.9%) appear in region 3 as winning nontalkers. Strikingly, every dealer who ever wins as a nontalker (region 3) also appears in region 2 on other securities. That is, 22 dealers, nearly a third of the sample, exhibit behavior in all three regions. The same dealer talks on some securities, is non-competitively silent on others, and wins silently on yet others. This is precisely the signal-dependent behavior predicted by Proposition 3. The dealer's action depends on her private signal relative to the cutoffs $\bar{s}^*$ and $s^\times$, not on a fixed dealer characteristic.

4.4 Estimating the information rent

The crossing coefficient measures the average gap between b^S and b^T at the winning signal. Because the winning nontalker's signal is drawn from the upper tail above $s^\times$, this average gap exceeds the threshold discontinuity Δ in Figure 1. Two tests probe how the estimated coefficient relates to Δ.

First, we interact the nontalker indicator with the number of talkers. When more dealers talk, fewer remain silent, so the winning silent signal is drawn closer to $s^\times$ and the premium should converge toward Δ. Table 6, Panel A confirms this prediction. The interaction is negative and significant for both the price premium (-0.009, $p < 0.05$) and the max-talk crossing (-0.022, $p < 0.01$). The nontalker coefficient falls from 22.9 basis points with one to three talkers to 5.4

Table 6: Estimating the Information Rent

This table estimates the information rent of winning silent dealers. Panel A interacts the nontalker indicator with the number of talkers. Panel B sorts nontalker-won auctions by winner's surplus decile and reports the average price premium within each bin. All regressions include day fixed effects with standard errors clustered by security. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Nontalker Effect by Number of Talkers

	Interaction		By bin			
	(1) Premium	(2) Crossing	(3) 1–3	(4) 4–6	(5) 7–9	(6) 10+
Nontalker_i	0.176***	0.382***	0.229***	0.108***	0.091***	0.054***
	(0.038)	(0.057)	(0.069)	(0.038)	(0.024)	(0.015)
N_i	0.019***	−0.063***				
	(0.004)	(0.006)				
$\text{Nontalker}_i \times N_i$	−0.009**	−0.022***				
	(0.004)	(0.006)				
Day FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.192	0.217	0.384	0.343	0.237	0.238
Obs	37,026	37,026	4,678	5,669	9,354	16,670

Panel B. Price Premium by Winner's Surplus Decile (Nontalker Wins Only)

Surplus decile	1 (low)	2–3	4–5	6–7	8–9	10 (high)
Price premium (bps)	14.8	10.7	18.3	26.0	38.6	57.3
Avg. surplus (bps)	−21.2	1.6	5.7	15.2	43.4	122.9
Obs	1,689	3,385	3,369	3,380	3,375	1,688

basis points with ten or more.

Second, Panel B sorts nontalker-won auctions by winner's surplus (price minus cover). Small surplus means the winner barely outbid the runner-up, placing the winning signal close to $s^{\times}$. In the bottom decile, the average price premium is 14.8 basis points, roughly 65% of the overall nontalker premium. The premium rises monotonically to 57.3 basis points in the top decile, confirming that $b^S(s)$ is increasing. The bottom-decile estimate provides a reduced-form lower bound on Δ.

4.5 Informativeness of price talk

The model predicts that despite selective disclosure, price talk aggregates dispersed information effectively. We test two supporting predictions.

Price decomposition. When a talker wins, the equilibrium price reflects the consensus rather than any individual signal. We decompose the price surprise into contributions from the winner's

Table 7: Winner's Talk vs. Others' Talk

This table decomposes the price surprise into contributions from the winner's own talk versus others' average talk. All variables measured relative to the ML benchmark. Sample restricted to BWICs where the winner talked.

	Dep. var. Price Surprise $\tilde{P}_i$				
	(1) All	(2) AAA/AA	(3) A/BBB	(4) BB/B	(5) Equity
Winner's talk $\tilde{S}_i$	0.273***	0.319***	0.174***	0.123***	0.349***
	(0.022)	(0.034)	(0.028)	(0.046)	(0.064)
Others' average $\tilde{\bar{S}}_{-i}$	0.754***	0.673***	0.832***	0.927***	0.679***
	(0.021)	(0.034)	(0.027)	(0.044)	(0.063)
Implied contribution					
Own signal	27%	32%	17%	12%	34%
Others' signals	73%	68%	83%	88%	66%
R^2	0.973	0.994	0.970	0.986	0.996
Obs	16,342	4,872	3,960	4,942	1,484

signal $\tilde{S}_i$ and others' average $\tilde{\bar{S}}_{-i}$,

$$\tilde{P}_i = \beta_{\text{own}} \cdot \tilde{S}_i + \beta_{\text{others}} \cdot \tilde{\bar{S}}_{-i} + \alpha_t + \epsilon_i. \tag{10}$$

Table 7 reports the results for talker-won auctions. Prices load heavily on others' information ($\hat{\beta}_{\text{others}} = 0.754$) compared to the winner's own signal ($\hat{\beta}_{\text{own}} = 0.273$). Roughly three-quarters of price variation reflects the information of other dealers, consistent with the model's wisdom-of-crowds prediction. The seniority splits reveal meaningful heterogeneity. Mezzanine tranches (A/BBB) show the lowest own-signal weight at 17%, while equity tranches show the highest at 34%. The decomposition explains nearly all price surprise variation, with R^2 exceeding 0.97 across all groups.

Belief convergence. If price talk facilitates Bayesian learning, bid dispersion should be smaller than talk dispersion, with dealers revising their bids toward the emerging consensus. Table 8, Panel A confirms this prediction. Dealers whose talk is below consensus revise upward by 47 bps on average, while those above consensus revise downward by 51 bps, a roughly symmetric pattern consistent with Bayesian updating. The seniority gradient reflects the model. Junior tranches have greater valuation uncertainty, so learning produces larger belief updates. For equity, the

Table 8: Evidence of Learning from Price Talk

Panel A reports average bid revisions ($P_i - S_i$, in percentage points of par) for winning dealers, conditional on whether the winner's talk is below or above consensus. Panel B reports average dispersion (in percentage points of par).

Panel A. Bid Revisions Toward Consensus								
	All	AAA	AA	A	BBB	BB	B	Equity
All revisions	−0.18	−0.03	−0.05	−0.04	−0.19	−0.19	−0.62	−0.84
Talk below consensus	0.47	0.14	0.27	0.49	0.54	0.74	0.48	0.75
Talk above consensus	−0.51	−0.12	−0.23	−0.32	−0.55	−0.69	−1.25	−1.53
Panel B. Dispersion Reduction								
Talk dispersion	0.482	0.122	0.258	0.428	0.545	0.679	0.970	1.690
Bid dispersion	0.226	0.066	0.148	0.185	0.270	0.354	0.429	0.563
Reduction (%)	53	46	43	57	50	48	56	67

downward revision (−1.53 bps) is roughly double the upward revision (0.75 bps), consistent with anti-unraveling (talkers in junior tranches disproportionately hold above-consensus views because the most optimistic dealers stay silent).

Panel B shows that bid dispersion of 0.226% of par is roughly half of talk dispersion of 0.482%, confirming the belief convergence prediction. The percentage reduction increases from 43–46% for senior tranches to 67% for equity, indicating that learning is strongest precisely where prior disagreement is greatest.

Nontalker information by seniority. Table 9 extends the predictiveness analysis with seniority splits and a nontalker interaction. Panel A shows that the consensus coefficient is high across all seniority groups, ranging from 0.914 for B-equity to 0.997 for A. Panel B adds the nontalker indicator and its interaction with the consensus surprise. Price surprises are 9.3 bps higher when nontalkers win ($p < 0.01$), increasing to 12.7 bps with the interaction included. The interaction term (−0.025, $p < 0.01$) indicates that consensus is 2.5 percentage points less predictive when nontalkers win. The nontalker price premium increases with credit risk, from 3.7 bps for AAA/AA to 15.2 bps for equity, and the interaction likewise grows in magnitude and significance, consistent with the model's prediction that the information rent from silence is greatest where fundamental uncertainty is highest.

Table 9: Nontalker Information by Seniority

This table reports regressions of price surprises on consensus surprises and nontalker winner indicators. All variables measured relative to the ML benchmark. Standard errors clustered by security.

Panel A. Consensus Predicts Prices

	Dep. var. Price Surprise $\tilde{P}_i$						
	All	AAA	AA	A	BBB	BB	B-Eq
Consensus surprise $\tilde{\tilde{S}}_i$	0.947***	0.987***	0.980***	0.997***	0.981***	0.994***	0.914***
	(0.003)	(0.005)	(0.006)	(0.011)	(0.010)	(0.009)	(0.005)
R^2	0.958	0.992	0.979	0.935	0.945	0.966	0.967
Obs	37,026	9,163	3,383	3,361	5,649	7,393	6,378

Panel B. Prices Higher and Talk Less Predictive When Nontalkers Win

	Dep. var. Price Surprise $\tilde{P}_i$					
	(1) All	(2) All	(3) AAA/AA	(4) A/BBB	(5) BB/B	(6) Equity
Consensus surprise $\tilde{\tilde{S}}_i$	0.947***	0.961***	0.985***	0.992***	0.981***	0.932***
	(0.003)	(0.004)	(0.003)	(0.011)	(0.008)	(0.007)
Nontalker$_i$	0.093***	0.127***	0.037***	0.129**	0.124***	0.152**
	(0.015)	(0.020)	(0.011)	(0.052)	(0.041)	(0.073)
$\tilde{\tilde{S}}_i \times$ Nontalker$_i$		−0.025***	−0.002	−0.012	−0.015*	−0.035***
		(0.006)	(0.005)	(0.015)	(0.009)	(0.010)
R^2	0.958	0.958	0.987	0.936	0.967	0.969
Obs	37,026	37,026	12,906	9,361	9,742	4,061

5 Market Quality Implications

The crossing test establishes that nontalker wins generate prices further above public benchmarks. This section asks whether the price inflation translates into larger dealer rents and whether more disclosure causally improves market outcomes.

5.1 Winner's surplus and seller revenue

We use the winner's surplus, defined as $P_i - C_i$, to measure rents extracted by the winning dealer. Under full disclosure, competition among equally informed dealers compresses this gap toward zero. Under strategic silence, hidden information allows the winner to bid above competitors who lack favorable private signals. The difference in surplus across winner types measures the incremental revenue that sellers forgo because of strategic silence.

Table 10, Panel A reveals a nuanced pattern. Unconditionally, talker winners capture 24.7

Table 10: Winner's Surplus

This table reports the winner's surplus (price minus cover, in percentage points of par) by winner type. Panel A reports summary statistics. Panel B reports regressions with day fixed effects and standard errors clustered by security. ***, **, * denote significance at 1%, 5%, 10%.

Panel A. Winner's Surplus by Winner Type

	All	Talker Wins	Nontalker Wins	Difference
Price – Cover (pp of par)	0.240	0.247	0.233	−0.014***
				(0.005)
By seniority				
AAA/AA	0.090	0.081	0.095	0.014***
A/BBB	0.237	0.226	0.246	0.020**
BB/B	0.371	0.366	0.376	0.009
Equity	0.412	0.431	0.398	−0.033*
By talk intensity				
Few talkers ($N_i \leq 5$)	0.253	0.277	0.247	−0.029**
Many talkers ($N_i > 5$)	0.236	0.244	0.226	−0.018***

Panel B. Regression Evidence

	Dep. var. Winner's Surplus (Price – Cover)						
	(1) All	(2) All	(3) All	(4) AAA/AA	(5) A/BBB	(6) BB/B	(7) Equity
Nontalker_i	−0.025***	−0.017***	−0.005	0.007*	0.005	−0.030**	−0.043*
	(0.005)	(0.005)	(0.015)	(0.004)	(0.009)	(0.013)	(0.023)
N_i		0.006***	0.007***	−0.000	−0.000	−0.012***	−0.006
		(0.001)	(0.001)	(0.001)	(0.002)	(0.004)	(0.005)
$\text{Nontalker}_i \times N_i$			−0.001				
			(0.001)				
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.198	0.199	0.199	0.392	0.335	0.285	0.341
Obs	30,322	30,322	30,322	10,590	7,974	7,559	3,216

basis points of surplus compared to 23.3 for nontalker winners, a difference of −1.4 basis points ($p < 0.01$). This comparison confounds winner type with tranche composition, because talker wins are disproportionately concentrated in junior tranches where surplus is mechanically larger. Within seniority groups, nontalker winners capture modestly more surplus for AAA/AA (+1.4 bps, $p < 0.01$) and A/BBB (+2.0 bps, $p < 0.05$). For BB/B, the difference is small and insignificant (0.9 bps). For equity, the sign reverses (−3.3 bps, $p < 0.10$).

Comparing the surplus with the crossing results is instructive. Table 3 shows that nontalker-won auctions generate prices 8–13 basis points further above the consensus and 26 basis points closer to the highest talk. Nontalker wins thus produce systematically higher prices. Yet this price

Table 11: Effects of Price Talk on Market Quality

This table reports instrumental variables estimates of the effect of price talk on market quality. The instrument is the leave-one-out average number of talks on other CUSIPs in the same list. All specifications include controls, day FE, and list FE.

Panel A. First Stage						
			Dep. var. Number of Talks N_i			
	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	AAA/AA	A/BBB	BB/B	Equity
Z_i (leave-one-out mean)	−5.253***	−5.249***	−5.712***	−5.440***	−5.537***	−4.145***
	(0.115)	(0.115)	(0.249)	(0.288)	(0.200)	(0.482)
First-stage F	2,102	2,097	525	358	769	74
Panel B. Trade Failures						
			Dep. var. Trade Failure$_i$			
	OLS	IV	IV. AAA/AA	IV. A/BBB	IV. BB/B	IV. Eq
N_i / $\widehat{N}_i$	−0.060***	−0.008***	−0.002*	−0.013***	−0.016***	−0.020
	(0.002)	(0.001)	(0.001)	(0.002)	(0.003)	(0.018)
Obs	48,277	47,079	16,784	13,072	15,179	2,044
Panel C. Cover Disclosure						
			Dep. var. Public Sale$_i$			
	OLS	IV	IV. AAA/AA	IV. A/BBB	IV. BB/B	IV. Eq
N_i / $\widehat{N}_i$	0.033***	0.028***	0.024***	0.026***	0.030***	−0.020
	(0.004)	(0.001)	(0.001)	(0.002)	(0.003)	(0.025)
Obs	35,581	35,653	14,501	9,905	10,145	1,102
Controls, Day FE, List FE in all specifications						

inflation does not translate into proportionally larger winner's surplus, because the cover bid adjusts upward alongside the transaction price. The auction mechanism absorbs much of the informational advantage into higher overall price levels rather than a wider gap between the top two bids.

Panel B confirms this pattern. The pooled nontalker coefficient on surplus is −2.5 bps ($p < 0.01$), driven by the composition effect. Once controls for the number of talkers are included and the sample is split by seniority, the coefficient is negligible for investment-grade tranches and modestly negative for junior tranches. The cost of strategic silence thus falls primarily on price efficiency, the level at which the market clears, rather than on the spread between winning and runner-up bids.

5.2 Effects of price talk on market quality

The preceding results are correlational and may reflect endogenous dealer behavior. To isolate plausibly exogenous variation in talk intensity, we exploit the batch structure of BWIC lists. When

sellers distribute lists containing multiple securities, dealers who choose to respond typically provide talks on many or all CUSIPs, reflecting administrative convenience and attention allocation rather than purely security-specific information. This institutional feature generates cross-CUSIP dependence in talk provision within a list.

We instrument for the number of talks on CUSIP i using the leave-one-out average number of talks on other CUSIPs in the same list,

$$Z_i = \frac{1}{|L_i| - 1} \left(\sum_{j \in L_i} N_j - N_i \right), \tag{11}$$

where L_i denotes the set of CUSIPs in the same BWIC list. Intuitively, Z_i captures how intensively dealers are providing talks to *other* securities in the same batch. Conditional on list fixed effects, this variation reflects how dealer attention is allocated across CUSIPs within a list, rather than overall list-level participation. We estimate the second stage

$$Y_i = \beta_0 + \beta_1 \widehat{N}_i + \beta_2' \mathbb{X}_i + \alpha_{\ell(i)} + \varepsilon_i, \tag{12}$$

where $\alpha_{\ell(i)}$ denotes list fixed effects and Y_i is an outcome such as trade failure or cover disclosure.

Table 11 reports the results. Panel A shows that the instrument is strongly relevant, with first-stage F-statistics exceeding 350 across all rating groups except equity tranches ($F = 74$). The first-stage coefficient is negative and statistically significant at the 1%-level. With list fixed effects absorbed, the instrument captures within-list substitution in dealer attention, that is, when other CUSIPs in the same list attract more talks, CUSIP i receives fewer talks. This pattern is consistent with dealer attention being a scarce resource that is allocated across securities within a batch. Importantly, the identifying variation is therefore orthogonal to list-level demand or participation, which is absorbed by fixed effects, and instead reflects relative attention shifts across otherwise comparable securities.

Panel B reports estimates for trade failures. The IV estimate (-0.008) is substantially smaller in magnitude than the OLS estimate (-0.060), indicating that OLS likely overstates the effect due to reverse causality or omitted variables. In particular, CUSIPs that are ex ante easier to

trade may both attract more talks and exhibit lower failure rates, biasing OLS downward. The IV estimate implies that an additional talker reduces the probability of trade failure by 0.8 percentage points. The effect is economically larger for junior tranches (e.g., -1.6 percentage points for BB/B), consistent with information frictions being more severe in lower-rated securities.

Panel C shows that price talk increases the probability of cover disclosure. The IV estimate implies that each additional talker raises the likelihood of a public sale by 2.8 percentage points. This finding suggests that pre-trade information aggregation and post-trade transparency are complements, that is, when more information is revealed before the auction, sellers are more likely to disclose outcomes afterward, consistent with improved price discovery and reduced uncertainty.

The IV estimates can be interpreted as local average treatment effects for CUSIPs whose talk intensity responds to within-list variation in dealer attention. These are precisely the securities for which talk provision is marginal and therefore most sensitive to institutional frictions such as time constraints or operational costs. As such, the estimates capture the causal effect of additional talkers in settings where dealer participation is not fully saturated.

We interpret these results as suggestive evidence that price talk improves market quality, while acknowledging limitations of the identification strategy. The exclusion restriction requires that talks on other CUSIPs affect outcomes for CUSIP i only through their impact on N_i. This assumption is plausible under the limited-attention-allocation interpretation, but may be violated if dealer activity on other CUSIPs directly affects competition on CUSIP i. In particular, dealers who talk on other securities in the same list may also be more likely to bid on CUSIP i, creating a direct competition channel.

The three-region test in Section 4.3 provides evidence that such cross-CUSIP participation is empirically relevant since dealers who are active on other securities are indeed present and bidding on CUSIP i. While list fixed effects absorb the overall level of dealer participation at the list level, they cannot fully eliminate this within-list reallocation of competitive pressure. Consequently, the IV estimates should be interpreted with caution as capturing a combination of attention-driven variation in talk and residual shifts in competitive intensity.

Taken together, the evidence suggests that price talk enhances market quality by reducing trade failures and increasing transparency, particularly in settings with greater informational frictions. At

the same time, the identification strategy highlights an important feature of OTC markets which is that dealer attention and participation are jointly determined across securities, implying that information provision and competition are inherently interconnected.

An important implication of the model is that price talk can induce inefficient allocations. In the equilibrium characterized in Section 4.3, talkers win precisely in states where all silent dealers have signals below the crossing threshold. Because high-signal dealers optimally remain silent, talker wins necessarily occur when no sufficiently strong private information is present to overturn the public consensus. As a result, the winning bidder in talker-won auctions need not be the dealer with the highest valuation. In particular, a silent dealer with a moderately higher private valuation may choose not to bid aggressively due to the winner's-curse adjustment, allowing a talker with a lower valuation—based only on public information—to win the asset. Thus, price talk can distort the allocation by shifting winning probabilities away from the highest-value bidders.

This allocation inefficiency reflects a fundamental tradeoff between informational efficiency and market quality. Price talk aggregates dispersed information and reduces uncertainty about asset values, which in turn lowers trade failure rates and increases the likelihood of transaction completion, as documented in Table 11. However, the same mechanism dampens the incentives of privately informed dealers to bid aggressively, since revealing information intensifies competition and reduces informational rents. Consequently, while talk improves market quality along dimensions such as execution probability and transparency, it does so at the cost of allocative efficiency. In this sense, price talk operates as a coordination device that facilitates trade but weakens the link between allocation and the underlying valuations. The empirical findings therefore suggest that improvements in market quality in OTC environments may come at the expense of efficient allocation, highlighting a central tension in the design of decentralized trading mechanisms.

6 Conclusion

This paper shows that voluntary disclosure in competitive markets generates anti-unraveling. In the CLO BWIC market, dealers with the most favorable private information stay silent because disclosure strengthens rivals' bids and erodes informational advantage. Those with less favorable

information disclose at no competitive cost. The result is a selective disclosure equilibrium in which publicly shared information systematically underrepresents the most optimistic assessments of value.

The crossing test confirms this mechanism empirically. Transaction prices deviate further above the pre-auction consensus when silent dealers win, with the effect increasing monotonically in tranche risk. Within-dealer tests show that the same dealer achieves higher prices when winning as a nontalker, ruling out dealer-type sorting. A direct test of the three-region structure exploits the list structure of BWIC auctions to show that present-and-silent dealers are outbid by talkers 19% of the time, confirming the existence of the intermediate non-competitive region. Supporting evidence shows that price talks are highly informative despite being strategically selective. Consensus explains 96% of price variation, dealers revise bids toward the emerging consensus, and bid dispersion is roughly half of talk dispersion. The cost of strategic silence falls on price efficiency rather than on winner's surplus, because the first-price auction mechanism limits rent extraction even when hidden information moves the level of prices.

These findings have implications for market design. Voluntary transparency can emerge and function effectively absent regulatory mandates, but competitive incentives impose a fundamental limit on its reach. The dealers whose information would contribute most to price efficiency are precisely those with the strongest incentive to remain silent. This creates a wedge between the private optimality of silence and the social value of disclosure that voluntary mechanisms cannot close. The anti-unraveling result demonstrates that full disclosure is not a natural equilibrium outcome when information is shared among competitors, implying that market design choices, not market forces alone, determine the boundary between what is revealed and what remains hidden.

References

Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman (2006). "Market Transparency and the Corporate Bond Market". In: *Journal of Financial Economics* 82.2, pp. 251–288.

Boyarchenko, Nina, David O. Lucca, and Laura Veldkamp (2021). "Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Auctions". In: *Journal of Political Economy* 129.2, pp. 607–645.

Campbell, Colin M. and Dan Levin (2000). "Can the Seller Benefit from an Insider in Common-Value Auctions?" In: *Journal of Economic Theory* 91.1, pp. 76–98.

Daley, Brendan and Brett Green (2012). "Waiting for News in the Market for Lemons". In: *Econometrica* 80.4, pp. 1433–1504.

Di Maggio, Marco, Amir Kermani, and Zhaogang Song (2017). "The Value of Trading Relations in Turbulent Times". In: *Journal of Financial Economics* 124.2, pp. 266–284.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen (2005). "Over-the-Counter Markets". In: *Econometrica* 73.6, pp. 1815–1847. DOI: 10.1111/j.1468-0262.2005.00639.x.

Dye, Ronald A. (1985). "Disclosure of Nonproprietary Information". In: *Journal of Accounting Research* 23.1, pp. 123–145.

— (1986). "Proprietary and Nonproprietary Disclosures". In: *Journal of Business* 59.2, pp. 331–366.

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar (2007). "Corporate Bond Market Transaction Costs and Transparency". In: *Journal of Finance* 62.3, pp. 1421–1451.

Ganuza, Juan-Jose and Jose Penalva (2010). "Signal Ordering Based on Dispersion and the Supply of Private Information in Auctions". In: *Econometrica* 78.3, pp. 1007–1030.

Goldstein, Michael A. and Edith S. Hotchkiss (2020). "Providing Liquidity in an Illiquid Market: Dealer Behavior in US Corporate Bonds". In: *Journal of Financial Economics* 135.1, pp. 16–40. DOI: 10.1016/j.jfineco.2019.05.014.

Grossman, Sanford J. (1981). "The Informational Role of Warranties and Private Disclosure about Product Quality". In: *Journal of Law and Economics* 24.3, pp. 461–483.

Hendershott, Terrence and Ananth Madhavan (2020). "Market Quality and ETF Trading". In: *Journal of Financial and Quantitative Analysis* 55.7, pp. 2201–2229.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt (2017). "Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations". In: *Review of Financial Studies* 30.9, pp. 3048–3085.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill (2020). "Frictional Intermediation in Over-the-Counter Markets". In: *Review of Economic Studies* 87.2, pp. 714–752.

Kagel, John H. and Dan Levin (2015). "Auctions: A Survey of Experimental Research, 1995–2010". In: *Handbook of Experimental Economics*. Vol. 2. Princeton University Press.

Lagos, Ricardo and Guillaume Rocheteau (2007). "Search in Asset Markets: Market Structure, Liquidity, and Welfare". In: *American Economic Review* 97.2, pp. 198–202.

Li, Dan and Norman Schürhoff (2019). "Dealer Networks". In: *Journal of Finance* 74.1, pp. 91–144.

Matthews, Steven and Andrew Postlewaite (1985). "Quality Testing and Disclosure". In: *RAND Journal of Economics* 16.3, pp. 328–340.

Milgrom, Paul and John Roberts (1986). "Relying on the Information of Interested Parties". In: *RAND Journal of Economics* 17.1, pp. 18–32.

Milgrom, Paul and Robert J. Weber (1982). "A Theory of Auctions and Competitive Bidding". In: *Econometrica* 50.5, pp. 1089–1122. DOI: 10.2307/1911865.

Milgrom, Paul R. (1981). "Good News and Bad News: Representation Theorems and Applications". In: *Bell Journal of Economics* 12.2, pp. 380–391.

Persico, Nicola (2000). "Information Acquisition in Auctions". In: *Econometrica* 68.1, pp. 135–148.

Verrecchia, Robert E. (1983). "Discretionary Disclosure". In: *Journal of Accounting and Economics* 5, pp. 179–194. DOI: 10.1016/0165-4101(83)90011-3.

A Proofs

A.1 Useful technical results

Lemma 1 (Bayesian posteriors)**.** *Given the revealed signals $\{s_k\}_{k\in\mathcal{T}}$, the posterior distribution of v for a talker conditioning only on public information is $v \mid \mathcal{M} \sim \mathcal{N}(\mu_T, \sigma^2_{v|\mathcal{M}})$ with*

$$\mu_T \equiv \mathbb{E}[v \mid \mathcal{M}] = \alpha\mu_0 + \beta\bar{S}, \qquad \sigma^2_{v|\mathcal{M}} \equiv \mathrm{Var}(v \mid \mathcal{M}) = \frac{1}{\tau_v + n_T\tau_\varepsilon}, \tag{13}$$

where $\alpha \equiv \tau_v\sigma^2_{v|\mathcal{M}}$ and $\beta \equiv n_T\tau_\varepsilon\sigma^2_{v|\mathcal{M}}$.

The posterior distribution for a silent dealer $j \in \mathcal{S}$ with signal $s_j = s$ is $v \mid (\mathcal{M}, s) \sim \mathcal{N}(\mu_S(s,\mathcal{M}),\ \sigma^2_S(\mathcal{M}))$ with posterior mean

$$\mu_S(s,\mathcal{M}) = \frac{\tau_v\mu_0 + n_T\tau_\varepsilon\bar{S} + \tau_\varepsilon s}{\tau_S(\mathcal{M})} = \alpha'\mu_0 + \beta'\bar{S} + \gamma' s \equiv \mu_S(0,\mathcal{M}) + \gamma' s, \tag{14}$$

and posterior precision

$$\tau_S(\mathcal{M}) \equiv \frac{1}{\sigma^2_S(\mathcal{M})} = \tau_v + (n_T + 1)\tau_\varepsilon. \tag{15}$$

Finally, a generic unrevealed dealer signal satisfies marginally

$$s_k \mid \mathcal{M} \sim \mathcal{N}(\mu_T, \sigma^2_s), \qquad \sigma^2_s \equiv \sigma^2_{v|\mathcal{M}} + \sigma^2_\varepsilon.$$

Proof. Standard conjugate normal updating. □

Order statistics for silent competitors. Throughout the silent-bidder analysis we focus on the interior case $n_T \geq 1$ and $n_S \geq 2$; boundary cases can be handled separately. Fix a silent dealer j with signal $s_j = s > \bar{s}$ and let

$$S_{(1)} \equiv \max_{k\in\mathcal{S}\setminus\{j\}} s_k$$

denote the highest signal among dealer j's silent competitors.

A central feature of the common-value environment is that rivals' silent signals are generally *not* independent of dealer j's own signal after conditioning only on the public information $\mathcal{M}$. The reason is that all signals remain correlated through the common value v. A dealer with a high private signal assigns more probability to high realizations of v, and therefore also assigns more probability to her rivals having high signals. For that reason, the relevant order-statistic object for a silent bidder must keep track of two distinct arguments: the *threshold* at which the top rival signal is evaluated, and the bidder's *own signal* that shifts the conditional distribution of rival signals.

Accordingly, define

$$G(x \mid s, \mathcal{M}) \equiv \Pr\big(S_{(1)} \leq x \,\big|\, s, \mathcal{M},\ \text{equilibrium set of silent rivals}\big), \tag{16}$$

whenever the density exists. Thus, $G(x \mid s, \mathcal{M})$ is the conditional CDF of the highest rival silent signal from the viewpoint of a silent bidder with own signal s, and $g(x \mid s, \mathcal{M}) \equiv \frac{\partial}{\partial x}G(x \mid s, \mathcal{M})$ is the corresponding density. The two arguments play different roles. The variable x is a rival-signal threshold that enters the bidder's winning probability when she deviates to bid as if she had pseudo-signal x. The variable s is the bidder's true private signal and enters because, in a common-value

model, the bidder's own signal is informative about the distribution of rivals' signals through the latent common value.

This distinction matters economically. If dealer j has a higher private signal, she not only values the asset more highly herself; she also infers that rivals are more likely to have strong signals. Hence, a higher own signal raises both the benefit of bidding aggressively and the expected toughness of competition. The object $G(x \mid s, \mathcal{M})$ captures exactly this tradeoff. The event $\{S_{(1)} \leq x\}$ says that the strongest rival silent dealer is weak enough to be beaten by a bid corresponding to pseudo-signal x, while the conditioning on $(s, \mathcal{M})$ incorporates dealer j's private and public information when assessing how likely that event is.

Define also the corresponding reverse hazard rate

$$K(x \mid s, \mathcal{M}) \equiv \frac{g(x \mid s, \mathcal{M})}{G(x \mid s, \mathcal{M})}, \tag{17}$$

whenever $G(x \mid s, \mathcal{M}) > 0$. Economically, $K(x \mid s, \mathcal{M})$ measures the local intensity of competition among silent rivals, from the viewpoint of a bidder with information $(s, \mathcal{M})$. Throughout, we impose the regularity condition that $g(s \mid s, \mathcal{M}) > 0$ on the relevant support so that the boundary-state first-order condition is well defined.

Similarly, let

$$S_{\mathcal{S}}^{\max} \equiv \max_{k \in \mathcal{S}} s_k$$

denote the highest signal among all silent dealers, and define

$$G_0(x \mid \mathcal{M}) \equiv \Pr(S_{\mathcal{S}}^{\max} \leq x \,|\, \mathcal{M}), \qquad g_0(x \mid \mathcal{M}) \equiv \frac{\partial}{\partial x} G_0(x \mid \mathcal{M}). \tag{18}$$

Unlike the rival-specific object $G(x \mid s, \mathcal{M})$, the talker-side object $G_0(x \mid \mathcal{M})$ conditions only on public information because a talker has no private signal at the bidding stage.

A.2 Analysis of bidding stage

We solve the game by backward induction. The continuation game is a first-price common-value auction with asymmetric information: talkers condition on $\mathcal{M}$ only, while silent dealers condition on $(\mathcal{M}, s)$.

We focus on a continuation equilibrium with the following structure: (i) talkers submit a common bid $b^T(\mathcal{M})$; and (ii) silent dealers use a strictly increasing bid function $b^S(\cdot; \mathcal{M})$ in their private signal.

Silent dealers. Fix $\mathcal{M}$ and suppose talkers bid a constant $b^T(\mathcal{M})$ while silent competitors follow a strictly increasing strategy $b^S(\cdot; \mathcal{M})$. Consider a silent dealer j with true signal $s > \bar{s}$ who deviates to bid b. Because $b^S(\cdot; \mathcal{M})$ is strictly increasing, there is a unique pseudo-signal

$$x \equiv (b^S)^{-1}(b)$$

such that $b = b^S(x; \mathcal{M})$. Thus, deviating to bid b is equivalent to bidding as if she had signal x.

Under monotonicity, dealer j wins against silent rivals if and only if $S_{(1)} < x$. Therefore her win probability when she bids $b = b^S(x; \mathcal{M})$ is

$$\Pr(\text{win} \mid s, \mathcal{M}, b) = \mathbb{1}\{b > b^T(\mathcal{M})\} \cdot G(x \mid s, \mathcal{M}). \tag{19}$$

Define the maximum bid among dealer j's silent rivals as

$$B^-_{(1)} \equiv b^S(S_{(1)}; \mathcal{M}).$$

By a change of variables, its conditional density is

$$f_{B^-_{(1)}}(b \mid s, \mathcal{M}) = \frac{g((b^S)^{-1}(b) \mid s, \mathcal{M})}{b^{S\prime}((b^S)^{-1}(b); \mathcal{M})}. \tag{20}$$

The silent dealer's expected payoff from bidding b is

$$\Pi(b; s, \mathcal{M}) = \Pr(\text{win} \mid s, \mathcal{M}, b) \cdot \Big(\mathbb{E}[v \mid s, \mathcal{M}, \text{win}] - b\Big),$$

where, under monotone bidding,

$$\{b > b^T(\mathcal{M})\} \cap \{B^-_{(1)} < b\} \quad \Longleftrightarrow \quad \{b > b^T(\mathcal{M})\} \cap \{S_{(1)} < x\}.$$

The exact first-order characterization is boundary-based.

Proposition 4 (Silent dealer first-order condition)**.** *In a symmetric equilibrium with $b = b^S(s; \mathcal{M})$ and monotone bids, an interior silent bid satisfies*

$$b^S(s; \mathcal{M}) = \mathbb{E}\big[v \;\big|\; s, \mathcal{M},\ S_{(1)} = s\big] - \frac{1}{K(s \mid s, \mathcal{M})} \frac{d}{ds} b^S(s; \mathcal{M}), \qquad s > \bar{s}, \tag{21}$$

provided $g(s \mid s, \mathcal{M}) > 0$.

Proof. Fix $\mathcal{M}$ and suppress it when no confusion arises. If a type-s silent bidder deviates to the pseudo-signal x, then for $b^S(x) > b^T$ her expected payoff is

$$\Pi(x; s) = \int_{\bar{s}}^{x} \Big(\mathbb{E}[v \mid s, \mathcal{M}, S_{(1)} = t] - b^S(x)\Big)\, g(t \mid s, \mathcal{M})\, dt.$$

Differentiating with respect to x yields

$$\frac{\partial \Pi(x; s)}{\partial x} = \Big(\mathbb{E}[v \mid s, \mathcal{M}, S_{(1)} = x] - b^S(x)\Big) g(x \mid s, \mathcal{M}) - b^{S\prime}(x; \mathcal{M}) \int_{\bar{s}}^{x} g(t \mid s, \mathcal{M})\, dt.$$

Since

$$\int_{\bar{s}}^{x} g(t \mid s, \mathcal{M})\, dt = G(x \mid s, \mathcal{M}),$$

the first-order condition becomes

$$\Big(\mathbb{E}[v \mid s, \mathcal{M}, S_{(1)} = x] - b^S(x; \mathcal{M})\Big) g(x \mid s, \mathcal{M}) - b^{S\prime}(x; \mathcal{M}) G(x \mid s, \mathcal{M}) = 0.$$

Evaluating at the symmetric optimum $x = s$ gives (21). □

Equation (21) is the exact boundary-state first-order condition. The valuation term is the

bidder's posterior mean in the marginal state in which the strongest silent rival is exactly tied with her own type. The shading term is governed by $1/K(s \mid s, \mathcal{M})$.

For interpretation, define the exact boundary-state winner's-curse correction

$$\widetilde{\Lambda}_v(s; \mathcal{M}) \equiv \mathbb{E}[v \mid s, \mathcal{M}] - \mathbb{E}[v \mid s, \mathcal{M}, S_{(1)} = s]. \tag{22}$$

This is the correct winner's-curse object for the exact first-order condition, because first-price optimality is determined by the marginal boundary state. It differs from the more familiar win-event correction

$$\Lambda_v(s; \mathcal{M}) \equiv \mathbb{E}[v \mid s, \mathcal{M}] - \mathbb{E}[v \mid s, \mathcal{M}, S_{(1)} < s],$$

which is often easier to interpret economically, but is not the exact object entering the FOC.

Lemma 2 provides a sufficient condition for the monotonicity of the silent bidding schedule. Economically, the condition requires that the marginal informational benefit of a higher private signal, γ', dominates any increase in the boundary-state winner's-curse correction. Under this condition, higher-signal silent dealers submit strictly higher bids.

Lemma 2 (Monotonicity of the silent bidding strategy)**.** *Fix $\mathcal{M}$ and suppose the boundary-state winner's-curse correction $\widetilde{\Lambda}_v(s; \mathcal{M})$ is continuously differentiable on $(\bar{s}, \infty)$. If*

$$\gamma' - \frac{\partial}{\partial s}\widetilde{\Lambda}_v(s; \mathcal{M}) > 0 \qquad \text{for all } s > \bar{s}, \tag{23}$$

then any continuously differentiable solution $b^S(s; \mathcal{M})$ to the silent first-order condition (21) is strictly increasing in s on $(\bar{s}, \infty)$.

Proof. Using (14) and the definition of $K(s \mid s, \mathcal{M})$, we can rewrite (21) as

$$b^S(s; \mathcal{M}) = \mu_S(0, \mathcal{M}) + \gamma' s - \widetilde{\Lambda}_v(s; \mathcal{M}) - \frac{1}{K(s \mid s, \mathcal{M})}\, b^{S\prime}(s; \mathcal{M}). \tag{24}$$

Rearranging (24) gives

$$b^{S\prime}(s; \mathcal{M}) = K(s \mid s, \mathcal{M})\Big(\mu_S(0, \mathcal{M}) + \gamma' s - \widetilde{\Lambda}_v(s; \mathcal{M}) - b^S(s; \mathcal{M})\Big).$$

Thus $b^{S\prime}(s; \mathcal{M})$ has the same sign as the surplus term in parentheses because $K(s \mid s, \mathcal{M}) > 0$. In an interior first-price equilibrium, the bidder shades strictly below her boundary-state valuation, so

$$b^S(s; \mathcal{M}) < \mu_S(0, \mathcal{M}) + \gamma' s - \widetilde{\Lambda}_v(s; \mathcal{M}).$$

Hence $b^{S\prime}(s; \mathcal{M}) > 0$. Condition (23) ensures that the boundary-state valuation itself is strictly increasing in s, so this strict inequality is preserved along the equilibrium path. Therefore $b^S(s; \mathcal{M})$ is strictly increasing. □

Once again rearranging (24) yields

$$b^{S\prime}(s; \mathcal{M}) + K(s \mid s, \mathcal{M}) b^S(s; \mathcal{M}) = K(s \mid s, \mathcal{M})\big(\mu_S(0, \mathcal{M}) + \gamma' s - \widetilde{\Lambda}_v(s; \mathcal{M})\big). \tag{25}$$

Now define the integrating factor

$$I(x \mid \mathcal{M}) \equiv \exp\left(\int_{\bar{s}}^{x} dt\, K(t \mid t, \mathcal{M}) \right). \tag{26}$$

Multiplying (25) by $I(s \mid \mathcal{M})$ yields

$$\frac{d}{ds}\Big(I(s \mid \mathcal{M}) b^S(s; \mathcal{M}) \Big) = I(s \mid \mathcal{M}) K(s \mid s, \mathcal{M}) \big(\mu_S(0, \mathcal{M}) + \gamma' s - \widetilde{\Lambda}_v(s; \mathcal{M}) \big).$$

Integrating from $\bar{s}$ to s gives the general solution

$$b^S(s; \mathcal{M}) = I(s \mid \mathcal{M})^{-1} \big[b^S(\bar{s}; \mathcal{M}) + \int_{\bar{s}}^{s} dx\, I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M}) \big(\mu_S(0, \mathcal{M}) + \gamma' x - \widetilde{\Lambda}_v(x; \mathcal{M}) \big) \big]. \tag{27}$$

Because $\mu_S(0, \mathcal{M})$ is constant in x, this representation can be simplified by separating the public-information benchmark from the type-dependent term:

$$\begin{aligned} b^S(s; \mathcal{M}) = I(s \mid \mathcal{M})^{-1} \big[b^S(\bar{s}; \mathcal{M}) + \mu_S(0, \mathcal{M}) \int_{\bar{s}}^{s} dx\, I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M}) \big] \\ + I(s \mid \mathcal{M})^{-1} \int_{\bar{s}}^{s} dx\, I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M}) \big(\gamma' x - \widetilde{\Lambda}_v(x; \mathcal{M}) \big). \end{aligned} \tag{28}$$

Since

$$\frac{d}{dx} I(x \mid \mathcal{M}) = K(x \mid x, \mathcal{M}) I(x \mid \mathcal{M}),$$

we have

$$\int_{\bar{s}}^{s} dx\, I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M}) = I(s \mid \mathcal{M}) - 1.$$

Substituting this identity into (28) yields

$$\begin{aligned} b^S(s; \mathcal{M}) = & \mu_S(0, \mathcal{M}) + I(s \mid \mathcal{M})^{-1} \big[b^S(\bar{s}; \mathcal{M}) - \mu_S(0, \mathcal{M}) + \\ & + \int_{\bar{s}}^{s} dx\, I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M}) \big(\gamma' x - \widetilde{\Lambda}_v(x; \mathcal{M}) \big) \big]. \end{aligned} \tag{29}$$

The remaining integral can be written as an expectation under an endogenous weighting measure. Define a density on $(\bar{s}, s)$ by

$$h_s(x \mid \mathcal{M}) \equiv \frac{I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M})}{I(s \mid \mathcal{M}) - 1} \mathbb{1}\{\bar{s} < x < s\}. \tag{30}$$

This is a valid density because

$$\int_{\bar{s}}^{s} dx\, I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M}) = I(s \mid \mathcal{M}) - 1.$$

Hence

$$\int_{\bar{s}}^{s} dx\, I(x \mid \mathcal{M}) K(x \mid x, \mathcal{M}) \big(\gamma' x - \widetilde{\Lambda}_v(x; \mathcal{M})\big) = \big(I(s \mid \mathcal{M}) - 1\big)\, \mathbb{E}_X^{h_s}\Big[\gamma' X \widetilde{\Lambda}_v(X; \mathcal{M})\Big], \tag{31}$$

where X has density $h_s(\cdot \mid \mathcal{M})$. Therefore

$$\begin{aligned} b^S(s; \mathcal{M}) =& \mu_S(0, \mathcal{M})\big(1 - I(s \mid \mathcal{M})^{-1}\big) + I(s \mid \mathcal{M})^{-1} b^S(\bar{s}; \mathcal{M}) + \\ &+ \big(1 - I(s \mid \mathcal{M})^{-1}\big)\, \mathbb{E}_X^{h_s}\Big[\gamma' X - \widetilde{\Lambda}_v(X; \mathcal{M})\Big]. \end{aligned} \tag{32}$$

Talkers. Define the maximum bid among all silent dealers as

$$B_{(1)} \equiv b^S(S_{\mathcal{S}}^{\max}; \mathcal{M}). \tag{33}$$

Let

$$F_B(b \mid \mathcal{M}) \equiv \Pr(B_{(1)} < b \mid \mathcal{M}), \qquad f_B(b \mid \mathcal{M}) \equiv \frac{\partial}{\partial b} F_B(b \mid \mathcal{M}),$$

denote the CDF and density of the highest silent bid. A talker choosing bid b wins iff $B_{(1)} < b$, so

$$\Pi^T(b; \mathcal{M}) = F_B(b \mid \mathcal{M})\Big(\mathbb{E}[v \mid \mathcal{M}, B_{(1)} < b] - b\Big).$$

The exact first-order condition is again boundary-based.

Proposition 5 (Talker first-order condition)**.** *An interior optimal talker bid $b^T(\mathcal{M})$ satisfies*

$$b^T(\mathcal{M}) = \mathbb{E}\big[v \;\big|\; \mathcal{M},\ B_{(1)} = b^T(\mathcal{M})\big] - \frac{F_B(b^T(\mathcal{M}) \mid \mathcal{M})}{f_B(b^T(\mathcal{M}) \mid \mathcal{M})}. \tag{34}$$

Proof. Let

$$m(x) \equiv \mathbb{E}[v \mid \mathcal{M}, B_{(1)} = x].$$

Then

$$\mathbb{E}[v \mid \mathcal{M}, B_{(1)} < b] = \frac{1}{F_B(b)} \int_{-\infty}^{b} m(x) f_B(x)\, dx,$$

so

$$\Pi^T(b) = \int_{-\infty}^{b} m(x) f_B(x)\, dx - b F_B(b).$$

Differentiating gives

$$\frac{d\Pi^T(b)}{db} = m(b) f_B(b) - \big(F_B(b) + b f_B(b)\big).$$

Setting this equal to zero and rearranging yields (34). □

This concludes the construction of the continuation equilibrium and the proof of Proposition 1.

Crossing signal. Define the crossing signal $s^\times(\mathcal{M})$ as the unique signal satisfying

$$b^S(s^\times(\mathcal{M}); \mathcal{M}) = b^T(\mathcal{M}). \tag{35}$$

Because $b^S(\cdot; \mathcal{M})$ is strictly increasing, talkers win whenever

$$S_{\mathcal{S}}^{\max} < s^\times(\mathcal{M}),$$

and silent dealers with signals strictly above $s^\times(\mathcal{M})$ bid above the talker side. At the knife-edge case $S_{\mathcal{S}}^{\max} = s^\times(\mathcal{M})$, the winner is determined by tie-breaking.

A.3 Analysis of disclosure stage

We now characterize the equilibrium disclosure threshold. The key force is competitive exposure: revealing favorable information raises rivals' bids and erodes the discloser's informational advantage.

Fix a dealer j with signal $s_j = s$ contemplating whether to talk. Suppose all other dealers follow the symmetric threshold strategy with cutoff $\bar{s}$:

$$T_k = \text{Talk} \iff s_k \le \bar{s}, \qquad m_k = s_k \text{ if } T_k = \text{Talk}.$$

If the dealer talks. If dealer j talks, her signal becomes public. Let $\mathcal{T}_{-j}$ denote the talkers among the other dealers and $n_{T,-j} = |\mathcal{T}_{-j}|$. Then

$$n_T^{\text{in}} \equiv n_{T,-j} + 1, \qquad \bar{S}^{\text{in}}(s) \equiv \frac{1}{n_T^{\text{in}}} \left(s + \sum_{k \in \mathcal{T}_{-j}} s_k \right), \qquad \mathcal{M}^{\text{in}}(s) \equiv \{s\} \cup \{s_k : k \in \mathcal{T}_{-j}\}.$$

Let $(b^{T,\text{in}}, b^{S,\text{in}})$ denote the continuation equilibrium under $\mathcal{M}^{\text{in}}(s)$, and define $s^{\times,\text{in}} \equiv s^\times(\mathcal{M}^{\text{in}}(s))$. A talker wins iff $S_{\mathcal{S}}^{\max} < s^{\times,\text{in}}$, ignoring the zero-probability tie case under continuity. Conditional on a talker win, symmetry implies each talker wins with probability $1/n_T^{\text{in}}$. Therefore

$$\Pi^T(s; \bar{s}) = \frac{1}{n_T^{\text{in}}} \Pr\left(S_{\mathcal{S}}^{\max} < s^{\times,\text{in}}\right) \cdot \left[\mathbb{E}\left(v \middle| \mathcal{M}^{\text{in}}(s), S_{\mathcal{S}}^{\max} < s^{\times,\text{in}}\right) - b^{T,\text{in}} \right]. \tag{36}$$

If the dealer stays silent. If dealer j stays silent, the public information is

$$n_T^{\text{out}} \equiv n_{T,-j}, \qquad \bar{S}^{\text{out}} \equiv \frac{1}{n_T^{\text{out}}} \sum_{k \in \mathcal{T}_{-j}} s_k, \qquad \mathcal{M}^{\text{out}} \equiv \{s_k : k \in \mathcal{T}_{-j}\}.$$

Let $(b^{T,\text{out}}, b^{S,\text{out}})$ denote the continuation equilibrium under $\mathcal{M}^{\text{out}}$, and define $s^{\times,\text{out}} \equiv s^\times(\mathcal{M}^{\text{out}})$. Dealer j wins iff she has the highest silent signal and her bid beats the talker bid, which under monotonicity is equivalent to

$$\{S_{(1)} < s\} \cap \{s > s^{\times,\text{out}}\}.$$

Hence

$$\Pi^S(s;\bar{s}) = \Pr\big(S_{(1)} < s,\ s > s^{\times,\text{out}}\big) \cdot \Big[\mathbb{E}\big(v\big|s, \mathcal{M}^{\text{out}}, \text{win}\big) - b^{S,\text{out}}(s)\Big]. \tag{37}$$

We can now prove Proposition 2.

Proof. Define $\Delta(s) \equiv \Pi^S(s;\bar{s}) - \Pi^T(s;\bar{s})$. Under monotone continuation strategies, higher signals increase the marginal informational rent from remaining silent and thus raise the relative benefit of silence. For low s, the informational rent from silence is negligible and $\Delta(s) < 0$; for high s, the informational rent dominates and $\Delta(s) > 0$. By continuity, there exists a cutoff $\bar{s}^*$ such that $\Delta(\bar{s}^*) = 0$. □

The proof of Proposition 3 follows from Lemma 3 and Proposition 6.

Lemma 3. *Fix the disclosure cutoff $\bar{s}$ followed by all other dealers and fix the identity and number of opponents (n_T, n_S). Consider two public information sets in the bidding stage,*

$$\mathcal{M} \subset \mathcal{M}', \qquad \mathcal{M}' = \mathcal{M} \cup \{s\},$$

where s is an additional publicly revealed signal. Assume:

1. *The joint distribution of $(v, \{s_k\}_k)$ is affiliated.*
2. *In the continuation game, silent bids are strictly increasing and the talker bid is well defined.*
3. *For any b, the CDF of the highest opposing bid under $\mathcal{M}'$ is weakly lower than under $\mathcal{M}$:*

$$\Pr(B^{\text{opp}}_{(1)}(\mathcal{M}') < b) \le \Pr(B^{\text{opp}}_{(1)}(\mathcal{M}) < b).$$

Then the maximal expected payoff of a bidder with no private information is weakly decreasing in public information:

$$\sup_b \Pi^{\text{no priv}}(b;\mathcal{M}') \le \sup_b \Pi^{\text{no priv}}(b;\mathcal{M}), \tag{38}$$

where

$$\Pi^{\text{no priv}}(b;\mathcal{M}) = \Pr(B^{\text{opp}}_{(1)}(\mathcal{M}) < b) \cdot \Big(\mathbb{E}[v \mid \mathcal{M}, B^{\text{opp}}_{(1)}(\mathcal{M}) < b] - b\Big).$$

Proof. The assumption on stochastic dominance implies that, for every bid b, winning becomes weakly harder as public information becomes more favorable to opponents. Under affiliation and increasing strategies, this weakly lowers the bidder's optimized payoff because the higher public signal both reduces win probability and tightens the conditioning event implicit in winning. Taking the supremum over b yields the claim. □

Proposition 6. *Fix a symmetric threshold profile for the other dealers with cutoff $\bar{s}$ and truthful talks. Let $s^{\times}(\mathcal{M}^{\text{out}})$ satisfy*

$$b^S(s^{\times}(\mathcal{M}^{\text{out}}); \mathcal{M}^{\text{out}}) = b^T(\mathcal{M}^{\text{out}}),$$

where $\mathcal{M}^{\text{out}}$ is the public information set when dealer j stays silent. Let $\bar{s}^*$ denote a symmetric equilibrium disclosure cutoff. Under the maintained monotone continuation equilibrium and Lemma 3, we have

$$\bar{s}^* < s^{\times}(\mathcal{M}^{\text{out}}).$$

Proof. Fix the realized public talks from the other dealers and define

$$\mathcal{M}^{\text{out}} = \{s_k : k \in \mathcal{T}_{-j}\}, \qquad \mathcal{M}^{\text{in}}(s) = \mathcal{M}^{\text{out}} \cup \{s\}.$$

Let

$$\Delta(s) \equiv \Pi^S(s; \bar{s}) - \Pi^T(s; \bar{s}).$$

Consider type $s = s^\times \equiv s^\times(\mathcal{M}^{\text{out}})$. By definition,

$$b^S(s^\times; \mathcal{M}^{\text{out}}) = b^T(\mathcal{M}^{\text{out}}).$$

Because $b^S(\cdot; \mathcal{M}^{\text{out}})$ is strictly increasing, any type just above $s^\times$ bids above the talker bid and wins with positive probability. Moreover, if

$$G(s^\times \mid s^\times, \mathcal{M}^{\text{out}}) > 0,$$

then type $s^\times$ itself has strictly positive win probability against silent rivals at the boundary. Since the silent first-order condition implies strict bid shading whenever $g(s^\times \mid s^\times, \mathcal{M}^{\text{out}}) > 0$ and $K(s^\times \mid s^\times, \mathcal{M}^{\text{out}})$ is finite, we have

$$b^S(s^\times; \mathcal{M}^{\text{out}}) < \mathbb{E}[v \mid s^\times, \mathcal{M}^{\text{out}}, S_{(1)} = s^\times],$$

so type $s^\times$ earns strictly positive conditional surplus in the boundary state. Combining positive boundary-state surplus with positive win probability yields

$$\Pi^S(s^\times; \bar{s}) > 0. \tag{39}$$

If the same type talks, she enters the bid stage with no private information under the richer public information set $\mathcal{M}^{\text{in}}(s^\times)$. By Lemma 3,

$$\Pi^T(s^\times; \bar{s}) \leq \sup_b \Pi^{\text{no priv}}(b; \mathcal{M}^{\text{out}}).$$

But under $\mathcal{M}^{\text{out}}$, a silent bidder of type $s^\times$ has at least as much information as any bidder who conditions only on $\mathcal{M}^{\text{out}}$. Since the feasible bid set is the same, the value of the bidding problem is weakly increasing in information:

$$\sup_b \Pi^{\text{no priv}}(b; \mathcal{M}^{\text{out}}) \leq \Pi^S(s^\times; \bar{s}).$$

Therefore

$$\Pi^T(s^\times; \bar{s}) \leq \Pi^S(s^\times; \bar{s}),$$

so $\Delta(s^\times) \geq 0$, and under strict exposure the inequality is strict.

In symmetric threshold equilibrium, $\Delta(\bar{s}^*) = 0$. Since $\Delta(s)$ satisfies single crossing and $\Delta(s^\times) \geq 0$, the zero of Δ must occur below $s^\times$. Hence

$$\bar{s}^* < s^\times.$$

□

A.4 Testable predictions

Result 1 (Consensus predictiveness)**.** *Transaction prices load positively on the talk consensus, but the loading is stronger when a talker wins than when a nontalker wins. Moreover, the consensus becomes more informative as more dealers talk.*

Proof. When a talker wins, the winning bid is pinned down by the public-information side of the auction, so prices are determined primarily by the public posterior formed from disclosed talks. When a nontalker wins, the winning bid reflects both public information and the winner's retained private signal, so the consensus explains a smaller share of price variation. As the number of talkers increases, the public posterior aggregates more information and becomes more precise, so the consensus becomes more informative in both cases. □

Result 2 (Price decomposition)**.** *Transaction prices reflect posterior beliefs about the asset's value. When a talker wins, the transaction price depends on the public posterior mean* $\mathbb{E}[v \mid \mathcal{M}]$*, which aggregates all disclosed talks. When a nontalker wins, the transaction price depends on the private posterior mean* $\mathbb{E}[v \mid s, \mathcal{M}]$*, because the winning bidder conditions on both the public information and her retained private signal. The public components of these posterior means admit linear decompositions in the talk consensus:*

$$\mathbb{E}[v \mid \mathcal{M}] = \frac{\beta}{n_T} s_i + \frac{\beta(n_T - 1)}{n_T} \bar{S}_{-i} + \alpha \mu_0, \tag{40}$$

$$\mathbb{E}[v \mid s = 0, \mathcal{M}] = \frac{\beta'}{n_T} s_i + \frac{\beta'(n_T - 1)}{n_T} \bar{S}_{-i} + \alpha' \mu_0. \tag{41}$$

Proof. The posterior means follow directly from Lemma 1. The decompositions in (40)–(41) use the identity

$$\bar{S} = \frac{s_i + (n_T - 1)\bar{S}_{-i}}{n_T}.$$

When a talker wins, the winning bid depends only on public information because talkers have no private signal at the bidding stage. The transaction price therefore depends on the public posterior mean $\mathbb{E}[v \mid \mathcal{M}]$. When a nontalker wins, the winning bid depends on the full private posterior mean $\mathbb{E}[v \mid s, \mathcal{M}] = \mu_S(0, \mathcal{M}) + \gamma' s$, together with equilibrium shading and winner's-curse adjustments. The displayed decomposition isolates the public component $\mu_S(0, \mathcal{M}) = \mathbb{E}[v \mid s = 0, \mathcal{M}]$ of that private posterior. □

Result 3 (Belief convergence)**.** *More talk (higher* n_T*) increases the precision of the public posterior about* v *and reduces disagreement about the common-value component. In particular, conditional on* $\mathcal{M}$*, the posterior variance* $\mathrm{Var}(v \mid \mathcal{M}) = 1/(\tau_v + n_T \tau_\varepsilon)$ *is decreasing in* n_T*.*

Proof. By Lemma 1, $\mathrm{Var}(v \mid \mathcal{M}) = 1/(\tau_v + n_T \tau_\varepsilon)$, which is strictly decreasing in n_T. □

Corollary 1 (Comparative statics of the disclosure cutoff)**.** *Under the monotonicity and regularity conditions stated below, the disclosure cutoff* $\bar{s}^*$ *is decreasing in the intensity of competition and in the importance of the common-value component, and increasing in signal noise. In particular,*

$$\frac{\partial \bar{s}^*}{\partial N} < 0, \qquad \frac{\partial \bar{s}^*}{\partial \sigma_\varepsilon^2} > 0, \qquad \frac{\partial \bar{s}^*}{\partial \sigma_v^2} < 0.$$

Proof. Let

$$\Delta(s; N, \sigma_\varepsilon^2, \sigma_v^2) \equiv \Pi^S(s; \bar{s}) - \Pi^T(s; \bar{s})$$

denote the net benefit of remaining silent rather than talking for a dealer with signal s, given that all other dealers use the cutoff rule with threshold $\bar{s}$. In a symmetric equilibrium, the disclosure cutoff $\bar{s}^*$ solves

$$\Delta(\bar{s}^*; N, \sigma_\varepsilon^2, \sigma_v^2) = 0. \tag{42}$$

Assume that Δ is continuously differentiable in $(s, N, \sigma_\varepsilon^2, \sigma_v^2)$ in a neighborhood of the equilibrium and that single crossing holds:

$$\Delta_s(\bar{s}^*; \cdot) > 0. \tag{43}$$

Then the Implicit Function Theorem implies

$$\frac{\partial \bar{s}^*}{\partial \theta} = -\frac{\Delta_\theta(\bar{s}^*; \cdot)}{\Delta_s(\bar{s}^*; \cdot)}, \qquad \theta \in \{N, \sigma_\varepsilon^2, \sigma_v^2\}. \tag{44}$$

Hence the sign of each comparative static is the opposite of the sign of the corresponding effect on Δ.

The key object is the information rent from silence. A dealer who talks gives up her private signal before bidding and competes only on the basis of public information. A dealer who stays silent retains private information and can use it strategically in the auction. The value of silence is therefore governed by how much private information matters in the bidding stage. Parameters that increase the marginal informational rent from remaining silent raise $\Delta(s; \cdot)$ and therefore lower the cutoff; parameters that weaken that informational rent reduce $\Delta(s; \cdot)$ and therefore raise the cutoff.

First, increasing the number of competing dealers N makes silence more attractive. More dealers intensify competition for favorable assets and raise the value of preserving a private informational advantage. In particular, when there are more silent rivals, a high private signal is more valuable because it is precisely what allows a dealer to outbid the talker side and other silent competitors. Thus an increase in N raises the relative payoff from remaining silent, so $\Delta_N(\bar{s}^*; \cdot) > 0$, which implies

$$\frac{\partial \bar{s}^*}{\partial N} < 0.$$

Second, increasing signal noise σ_ε^2 makes silence less attractive. Noisier signals reduce the precision of private information and weaken the bidder's informational advantage from withholding it. They also weaken the informational content of winning, which dampens the strategic value of conditioning on a private signal in the first place. As a result, the gain from remaining silent falls, so $\Delta_{\sigma_\varepsilon^2}(\bar{s}^*; \cdot) < 0$, which implies

$$\frac{\partial \bar{s}^*}{\partial \sigma_\varepsilon^2} > 0.$$

Third, increasing the prior variance σ_v^2 makes silence more attractive. A larger common-value component raises the payoff relevance of private information because a favorable signal becomes more informative about the asset's fundamental value. This increases the informational rent from keeping the signal private and using it in the bidding stage rather than disclosing it ex ante. Hence $\Delta_{\sigma_v^2}(\bar{s}^*; \cdot) > 0$, which implies

$$\frac{\partial \bar{s}^*}{\partial \sigma_v^2} < 0.$$

This establishes the stated comparative statics. □