Should I Make or Should I Take? Bilateral Trade with a Reference Price*

Simon Jantschgi†

August 22, 2026

Abstract

In financial-market call auctions, trade clears at the feasible price closest to a public reference price. We embed this pricing rule in the canonical model of bilateral trade under incomplete information. The reference price fundamentally alters incentives, price discovery, and welfare compared to canonical pricing rules. Undominated Bayes–Nash equilibria form two Pareto-ranked families of truncated monopsony and monopoly equilibria, with one side making prices above or below the reference price and the other truthfully screening trade. The reference price itself never prints; two-sided price discovery collapses, and pricing power becomes endogenous. The families' frontiers match the buyer's-bid and seller's-ask double auctions up to the reference-imposed cap or floor; when unconstrained, the better guarantees 1/2 of second-best and $1/\pi$ of first-best welfare. Seller price-making is risk-dominant below a unique threshold and buyer price-making above it, with payoff dominance at extreme references. Integrated supply and demand curvature determines whether price-making is a privilege or a burden. Updating the reference to the latest trade generates perpetual role reversal; fixing the price-maker instead drives long-run welfare to zero.

Keywords: Market Design, Bilateral Trade, Reference Prices, Call Auctions, Equilibrium Selection, Price-Discovery

JEL codes: C72, C78, D44, D47, D82.

*I am grateful for comments and suggestions by Dirk Bergemann, Paul Klemperer, Fuhito Kojima, Simon Loertscher, Heinrich Nax, Marco Ottaviani, Bary Pradelski, Marek Pycia, Aldo Rustichini, and Alex Teytelboym.

†Department of Economics and Center for Market Design, University of Zurich; simon.jantschgi@uzh.ch.

1 Introduction

This paper embeds a reference price into the canonical model of bilateral trade under private information. A buyer privately observes her value, a seller privately observes his cost, and the two simultaneously submit a bid and an ask. Trade occurs whenever the bid exceeds the ask. The transaction price is the point in the resulting interval of mutually acceptable prices closest to a publicly known reference price. We show that introducing this seemingly modest price-selection criterion fundamentally alters incentives, price discovery, and welfare relative to the canonical trading rules studied in economic theory.

Reference prices sit at the heart of some of the world's most consequential price-setting mechanisms. Trading on major equity markets begins and ends with a call auction: orders are collected and cleared simultaneously at a single price. These auctions concentrate liquidity at the two decisive moments of the trading day. The opening auction converts the information accumulated while continuous trading was suspended into the first price of the day; the closing auction values portfolios, determines index levels, settles derivatives, and benchmarks vast institutional volume.[1] To set the auction price, exchanges first identify the prices that maximize executable volume and, among these, minimize order imbalance. Yet even then, an interval of market-clearing prices may remain. This residual indeterminacy is resolved by proximity to a publicly known reference price:

> *"If there are two or more prices at which the maximum volume of shares is tradable, the Indicative Match Price will be the price closest to the Auction Reference Price, provided that the Indicative Match Price will not be lower (higher) than the price of an order to buy (sell) [. . .] that was eligible to participate in the applicable auction."*
>
> New York Stock Exchange (2026b, Rule 7.35C)[2]

The appeal of a reference price is that it links current price formation to the recent history of the market.[3] It provides continuity across trading sessions, disciplines price formation when the order book is thin, and avoids jumps between otherwise equivalent clearing prices.

[1] In the first quarter of 2026, the New York Stock Exchange's closing auction executed a record average of over $43 billion per day, peaking at $230.5 billion on a single day; across U.S. equity markets, closing auctions account for almost ten percent of total notional value traded (New York Stock Exchange 2026a).

[2] The London Stock Exchange, Nasdaq, and the Japan Exchange Group employ related algorithmic cascades that combine executable volume, order imbalance, market pressure, and proximity to a reference price (London Stock Exchange 2026, Nasdaq 2026, Tokio Stock Exchange 2026).

[3] The Auction Reference Price is itself inherited from the market's recent history. For the closing auction, it is the last transaction price of the trading day and, absent one, the previous day's official closing price; for the opening auction, it is derived from the pre-opening national best bid and offer, with the previous close as fallback (New York Stock Exchange 2026b, Rules 7.35(a)(8), 7.35C(b)(1)).

In this sense, the reference price enters only after current orders have identified the trades that can be executed and the prices at which they can clear. Yet the choice among these prices is economically consequential: the rule used to resolve price indeterminacy affects traders' incentives, the information reflected in transaction prices, and the division and realization of gains from trade.

Bilateral trade is the natural setting in which to isolate these effects. It has been the canonical model for studying price formation, bargaining, and the limits of trade under private information. Its attraction is precisely that it strips exchange down to the most basic conflict: privately informed parties must determine whether gains from trade exist while simultaneously bargaining over their division. A buyer can improve her terms by understating her willingness to pay, and a seller by overstating his cost, yet the same strategic behavior can prevent mutually beneficial trade. Bilateral trade is also the smallest instance of the institutional mechanism that motivates this paper. With one buyer and one seller, a call auction reduces exactly to a bid, an ask, and a rule for selecting a price whenever the two cross. This case is particularly relevant for thin markets, where individual orders materially affect the market-clearing interval and therefore the price selected from it. The bilateral model thus provides both the cleanest theoretical environment for studying strategic price formation and the limiting case in which reference-price selection is most exposed to individual incentives.

Introducing a reference price into bilateral trade fundamentally alters the strategic incentive of traders. In standard models of bilateral trade, a quote does two jobs at once: it determines which counterparties a trader reaches and, conditional on trade, it enters the transaction price. The reference-price rule separates the two. A buyer bidding below the reference price sets the price with her bid whenever she trades, and so faces the classical monopsony problem of posting a price against an upward-sloping supply curve: a better price against fewer sellers. Once her bid reaches the reference price, further increases affect only which sellers she reaches — the price of every trade already available is set by the reference or the ask, independently of her bid — so her quote serves only to screen which trades are profitable, and truthful reporting is optimal. The seller faces the mirror image, a monopoly problem above the reference and truthful acceptance below it. Each half is familiar: the price-maker solves a textbook monopsony or monopoly problem; the price-taker faces what is effectively a posted price. What no standard trading rule contains is the choice between the two. Under a reference price, what is ordinarily a single continuous bidding problem fractures into a decision every quote must take: *make the price, or take it.*

The reference price rule offers this decision symmetrically to both sides of the market: nothing in it assigns who makes prices and who takes them. The assignment is endogenous. Can price-making and price-taking coexist in equilibrium, and on which side of the market does each settle? What does the answer cost, and whom does it favor? And what becomes of a market in which the reference is itself yesterday's price, so that today's answer sets tomorrow's question? These are the questions this paper answers.

Our Contribution

We consider a canonical bilateral trading environment. A buyer and a seller hold privately known values and costs, drawn independently from commonly known regular distributions; utility is quasilinear, and the outside option of no trade is normalized to zero. The two simultaneously submit a bid and an ask, and trade occurs whenever the bid exceeds the ask, at the price in the interval between ask and bid closest to a publicly known reference price R.

Equilibrium Characterization. The solution concept we consider is Bayes–Nash equilibrium. Bayes–Nash equilibrium alone is too permissive to discipline the maker–taker problem. For any interior price, coordination can sustain an equilibrium in which all trading types quote that price, independently of the reference (Proposition 2). Such equilibria rely on weakly dominated reports that distort acceptance cutoffs without improving the price. We therefore restrict attention to the minimal refinement of undominated strategies.[4]

Our first main result (Theorem 1) shows that under this refinement, either buyers generate all transaction prices below the reference while sellers take prices, or sellers generate all prices above it while buyers take prices. The two channels cannot coexist. The intuition comes from the change in how a trader's quote enters the transaction price at the reference. Once a trader crosses the reference, a more truthful quote can expand the set of trades but cannot improve the price on trades already available, so weak dominance pins the quote to the truthful acceptance cutoff. If both sides attempted to make prices on opposite sides of the reference, these truthful acceptance decisions would create profitable deviations for the opposing price-makers, ruling out coexistence.

[4]This refinement is standard in bilateral-trade models because weakly dominated reports can alter whether trade occurs without ever improving the terms of trade. In particular, at the buyer's-bid and seller's-ask endpoints of the k-double auction, the price-taking side has zero price impact, so truthful reporting of its acceptance cutoff is weakly dominant; eliminating dominated strategies selects the familiar one-sided equilibrium with truthful price-taking.

We then fully characterize the set of undominated Bayesian–Nash equilibria (Theorem 2). Once only one side makes prices, the other truthfully screens trade and the active trader faces the standard monopoly or monopsony problem; only the equilibrium cap or floor remains undetermined. Buyer-side equilibria are exactly truncated monopsony strategies: buyers follow their unconstrained monopsony bids up to an endogenous cap and pool there, while trading sellers report costs truthfully. Seller-side equilibria are the symmetric truncated monopoly strategies indexed by an ask floor. Thus an unrestricted Bayesian game — allowing mixing, pooling, gaps, and non-monotone strategies — reduces to two one-parameter families. The characterization has three striking implications.

- Corollary 1: The reference price itself never prints: every equilibrium transaction price lies strictly on one side of the anchor. The reference matters only indirectly, as a boundary that reorganizes strategic behavior.
- Corollary 2: Price discovery becomes one-sided. The price reflects only the maker's information, up to pooling; the taker's type matters only through acceptance. Trade remains jointly determined, but prices no longer aggregate information from both sides.
- Corollary 3: Frontier outcomes coincide with the buyer's-bid and seller's-ask auctions, which assign pricing power explicitly; the reference-price rule is symmetric and assigns it to neither side ex-ante. Yet equilibrium endogenously recreates the one-sided outcomes.

Welfare. Next, we study the welfare consequences of equilibrium play. Within each regime, equilibria are Pareto ordered (Proposition 3): relaxing the bid cap under buyer control, or the ask floor under seller control, expands trade and makes both sides weakly better off. Each family's Pareto-best outcome is therefore its unconstrained frontier, coinciding with the classical buyer's-bid or seller's-ask auction. Although concentrating price-setting on one side sacrifices trades relative to genuinely two-sided mechanisms, the loss is uniformly bounded (Theorem 3). The better frontier weakly dominates the random-offerer mechanism and therefore guarantees at least one half of Bayesian second-best gains from trade and at least a $\frac{1}{\pi}$ fraction of first-best gains from trade.[5] Reference-price selection therefore eliminates genuinely two-sided price formation, but the surviving one-sided outcomes retain a nontrivial fraction of the welfare benchmarks. Under uniform values, each frontier attains $\frac{8}{9}$ of second-best and $\frac{3}{4}$ of first-best gains from trade.

[5]The random-offerer mechanism chooses uniformly between the buyer- and seller-offering mechanisms. The second-best guarantee follows from Brustle et al. (2017); for the first-best approximation, see Babaioff et al. (2021), Deng et al. (2022), Fei (2022), Jo (2026).

Convention Formation. The equilibrium characterization leaves one multiplicity: which side makes prices. We model this remaining choice as convention formation: the two frontier regimes form a strict 2×2 coordination game between making and taking prices (Proposition 5). We select between them by Harsanyi–Selten risk dominance. The resulting rule is remarkably simple: there exists a unique threshold $R^\dagger$ such that seller control is risk-dominant for $R < R^\dagger$ and buyer control for $R > R^\dagger$ (Theorem 4). Thus the reference price determines the direction of price discovery through a single endogenous switch. The selection also respects unanimous preferences at the boundaries. At sufficiently low references, both sides prefer seller control; at sufficiently high references, both prefer buyer control (Theorem 5). Disagreement is therefore an interior phenomenon.

What that disagreement means, however, depends on the shape of the market. Making prices allows a trader to shade, but only by sacrificing marginal trades. We show that the log-curvature of integrated supply and demand governs whether this trade-off makes pricing power a privilege or a burden (Theorem 6). Under integrated log-concavity, each side earns more when it makes prices, so conflict is *mutual assertion*: both sides seek control. For other smooth regular distributions the ranking reverses, even by an arbitrarily large factor, producing *mutual deference*: each side would rather the other make prices. Risk dominance selects a price-maker in either case.

Self-Reference and Market Dynamics. Finally, we endogenize the reference by updating it to the latest transaction price. Pricing power then becomes self-undermining: buyer-made prices lie below the prevailing reference and seller-made prices above it, so exercising control moves the benchmark toward the region in which the opposite side is selected. We show that this force is universal. Under risk-dominant selection, trades occur infinitely often and pricing power changes sides infinitely often; every run of control by one side has a uniformly geometric tail (Theorem 7). When the risk-dominance threshold separates the two unconstrained pricing programs, the dynamics sharpen. After every trade, the market alternates between the buyer- and seller-offering regimes, and long-run welfare is an interior, liquidity-weighted average of their frontier welfares (Corollary 4). For symmetric environments the weights are $\frac{1}{2}$: long-run average welfare equals the expected gains from trade of the random-offerer mechanism.

However, if the same side is forced to make prices at every reference, each trade pushes the reference price farther toward the edge of the market; long-run average welfare vanishes (Theorem 8). Self-reference therefore sustains trade only through endogenous turnover of pricing power: the market alternates, or it fails.

Related Literature

Because this paper combines the canonical model of bilateral trade under private information with a price-selection rule drawn from financial-market practice, it connects and contributes to multiple strands of literature.

Bilateral Trade and Double Auctions. Bilateral trade is the canonical environment for studying how private information constrains trade and price formation. Myerson and Satterthwaite (1983) provide the mechanism-design benchmark, showing that ex-post efficient trade cannot generally be reconciled with incentive compatibility, individual rationality, and budget balance, and characterizing the second-best mechanism.[6] Rather than optimizing over mechanisms, Chatterjee and Samuelson (1983) study simultaneous bid–ask trade at the midpoint and, for uniform values, derive a linear equilibrium in which both sides shade strategically and some efficient trades are lost. Equilibrium selection is consequential: Leininger et al. (1989) show that the same symmetric mechanism admits uncountably many equilibria whose efficiency ranges from second best to zero. Satterthwaite and Williams (1989) generalize the pricing rule to the k-double auction, $p = kb + (1-k)s$, and establish existence, multiplicity, and generic ex-ante inefficiency. For interior k, both reports affect the price and both sides shade. At the endpoints, price formation becomes one-sided: at $k = 1$ the buyer sets the price and solves a monopsony problem while the seller truthfully reports her acceptance cutoff; at $k = 0$ the seller solves the corresponding monopoly problem and the buyer reports truthfully.[7] At the opposite extreme, Hagerty and Rogerson (1987) show that dominant-strategy incentives essentially force the mechanism toward posted prices that are independent of current reports. Recent work asks how much welfare such simple mechanisms can guarantee when first-best trade is unattainable.[8] Particularly relevant are the buyer- and seller-offering mechanisms corresponding to the two one-sided pricing regimes. Their random mixture guarantees a distribution-free fraction of first-best gains from trade (Babaioff et al. 2021, Deng et al. 2022, Fei 2022, Jo 2026).

Our paper introduces a new, symmetric pricing rule into this canonical environment. Although both traders can in principle influence prices, no undominated equilibrium sustains

[6] Garratt and Pycia (2026) recently show that this impossibility relies importantly on quasilinearity: complementarities or risk aversion can restore efficient trade.

[7] Williams (1987) shows that these one-sided mechanisms implement ex-ante incentive-efficient allocations for a class of distributions, while Gresik (1991) identifies conditions under which linear k-double-auction equilibria are ex-ante efficient.

[8] See, among others, Brustle et al. (2017), Cai and Wu (2023), Colini-Baldeschi et al. (2017), Kang et al. (2022) on simple and fixed-price mechanisms.

two-sided price formation: equilibrium instead collapses to truncated buyer's-bid or seller's-ask regimes, with one side making prices and the other reporting truthfully. Unlike the classical k-double auction, pricing power is therefore not assigned ex ante but becomes an equilibrium-selection problem governed by the reference price. At the frontier, the familiar one-sided mechanisms reappear, allowing us to inherit their welfare guarantees.

Institutional Features and Strategic Trading. A broad market-design literature studies how institutional features omitted from canonical trading models reshape incentives and welfare. Trading frequency and market thickness are themselves design variables: Budish et al. (2015) and Du and Zhu (2017) study how market frequency affects incentives in call auctions, while Loertscher et al. (2022) and Loertscher and Muir (2025) analyse the trade-off between market thickness and delayed exchange. Other work enriches double auction models with features such as transaction costs (Jantschgi et al. 2022), fee transparency (Jantschgi 2024), and multi-unit trade (Loertscher and Marx 2022), again finding that seemingly secondary details of the trading protocol can have first-order effects on incentives and welfare. This perspective is particularly relevant in thin financial markets, where traders have non-negligible price impact (Kyle 1989, Rostek and Weretka 2012, 2015, Vayanos 1999).

Our paper adds the price-selection rule itself to this set of institutional margins. Holding the canonical bilateral environment fixed, we show that introducing the reference-price rule used in financial-market auctions fundamentally changes who has price impact; when the reference is updated to the last transaction price, this institutional detail also creates an endogenous feedback from past prices to future market power.

Reference Prices in Mechanism Design. Reference prices have appeared in mechanism design as a way to select from a set of otherwise admissible prices. Most closely related is Wolitzky (2016). In a model of bilateral trade with maxmin agents, he derives as the outcome of a robust implementation problem the same reference-price mechanism that we take as our starting point. Jantschgi et al. (2024) instead motivates the rule from double-auction theory and financial-market practice, provides a unified definition and axiomatization of the Reference-Price Double Auction, and shows that incentives become approximately truthful as individual price impact vanishes in large markets. Jantschgi et al. (2026) provides a dynamic rationale, showing that selecting the feasible price closest to an inherited benchmark can optimally minimize price adjustment and volatility. In the flow-trading design of Budish et al. (2023), assets clear at discrete intervals, and ties among market-clearing prices are effectively broken toward the price of the previous batch.

The present paper contributes to this literature by including a reference price in the standard Bayesian bilateral-trade game, where bids and asks endogenously determine the feasible price interval. We show that, rather than implementing the reference, equilibrium reorganizes price formation around it, and the reference itself never prints. We then study how the reference selects between pricing regimes, affects welfare, and, when updated to past transaction prices, shapes the dynamics of pricing power and trade.

Call Auctions and Financial-Market Price Discovery. The literature on market microstructure studies how the design of opening and closing call auctions affects price discovery. Early work emphasizes information aggregation in the pre-opening period and at the open (Biais et al. 1999, Madhavan and Panchapagesan 2000), while subsequent studies show that clearing rules themselves affect price efficiency and manipulation (Comerton-Forde and Rydge 2006a,b, Park et al. 2022). More recent work studies strategic responses to auction design, including transparency, time priority, cancellations, and intermediate clearing (Daures et al. 2025, Yamamoto et al. 2026).

Our paper adds a novel feature to this literature: the use of reference prices to select among multiple clearing prices. We study its strategic consequences in the bilateral limit.

Pricing Power, Bargaining, and Surplus Division. One strand of the literature studies how control over the terms of trade affects the division of surplus. Classical bargaining theory makes bargaining power and the protocol for making offers central determinants of who captures the gains from agreement (Nash 1950, Rubinstein 1982). In monopoly and monopsony problems, pricing power instead takes the form of choosing a price while trading off a larger margin against a lower probability or quantity of trade. A central insight is that the value of such pricing power depends on the shape of demand and supply: Bulow and Pfleiderer (1983) link demand curvature to monopoly pass-through, while Weyl and Fabinger (2013) develop a general connection between pass-through, incidence, and the division of surplus under imperfect competition.

Our analysis endogenizes both the assignment and the value of pricing power. Which side makes prices is an equilibrium outcome rather than a feature of the bargaining protocol, and making the price need not be advantageous. We show that the curvature of integrated demand and supply determines whether the price-maker's marginal rent dominates the inframarginal rents earned by price-takers. In our model of bilateral trade with a reference price, pricing power can therefore be a contested privilege that both sides seek or a burden that both would rather relinquish.

2 The Model

A buyer and a seller may trade a single indivisible unit. The buyer's private type is her valuation v, and the seller's private type is his cost c. Private types are drawn independently from commonly known distributions F and G on the common support $[0,1]$, with continuous and strictly positive densities f and g. We assume that the distributions F and G are regular: the buyer's virtual valuation and the seller's virtual cost, $J_B(v) = v - \frac{1-F(v)}{f(v)}$ and $J_S(c) = c + \frac{G(c)}{g(c)}$, are strictly increasing.[9] Preferences are quasilinear and outside options are normalized to zero. Thus, if $x \in \{0,1\}$ indicates trade and p is the transaction price, payoffs from the outcome (x,p) are $x \cdot (v-p)$ for the buyer and $x \cdot (p-c)$ for the seller.

Trade is organized as a one-shot call auction with an exogenous reference price $R \in (0,1)$, commonly known by both traders. This is the bilateral reduction of the price-selection algorithm used in major opening and closing auctions.[10]

Definition 1 (Reference-price double auction)**.** Fix a commonly known reference price $R \in (0,1)$. In the *bilateral reference-price double auction*, the buyer and seller simultaneously submit a *bid* $b \in [0,1]$ and an *ask* $s \in [0,1]$.[11] Their common action space is $\mathcal{A} = [0,1]$. Trade occurs if and only if $s \le b$, making the set of mutually acceptable clearing prices $[s,b]$. Conditional on trade, the transaction price is the clearing price closest to the reference:

$$\varphi_R(s,b) = \arg\min_{p \in [s,b]} |p - R| = \begin{cases} s, & R < s \le b, \\ R, & s \le R \le b, \\ b, & s \le b < R. \end{cases} \qquad \text{(Reference-price rule)}$$

The pricing rule splits the quote space into two regions. Quotes that straddle the reference ($s \le R \le b$) trade at R itself: a posted-price rectangle sits at the anchor. Everywhere else, one quote alone sets the price — the bid below the reference, the ask above it.

[9] Regularity is the natural single-crossing condition for the price-making problems below. It guarantees unique and monotone monopsony and monopoly choices in the characterization of the surviving equilibria. Under our assumptions, J_B and J_S are continuous and strictly increasing, with $J_S(0) = 0$, $J_S(1) > 1$, $J_B(0) < 0$, and $J_B(1) = 1$; the equations $J_S(c) = 1$ and $J_B(v) = 0$ therefore have unique solutions in the interior of $(0,1)$.

[10] Once executable volume and order imbalance have been determined, the exchange selects the market-clearing price closest to the reference price. With one buyer and one seller, all feasible prices generate the same executable volume and imbalance, so the algorithm proceeds immediately to the reference-price criterion.

[11] The boundary quotes $b = 1$ and $s = 0$ represent a market buy and a market sell order, respectively. Abstention may be included as a separate action, but it is weakly dominated by $b = 0$ for every $v > 0$ and by $s = 1$ for every $c < 1$; omitting it leaves all undominated equilibria unchanged.

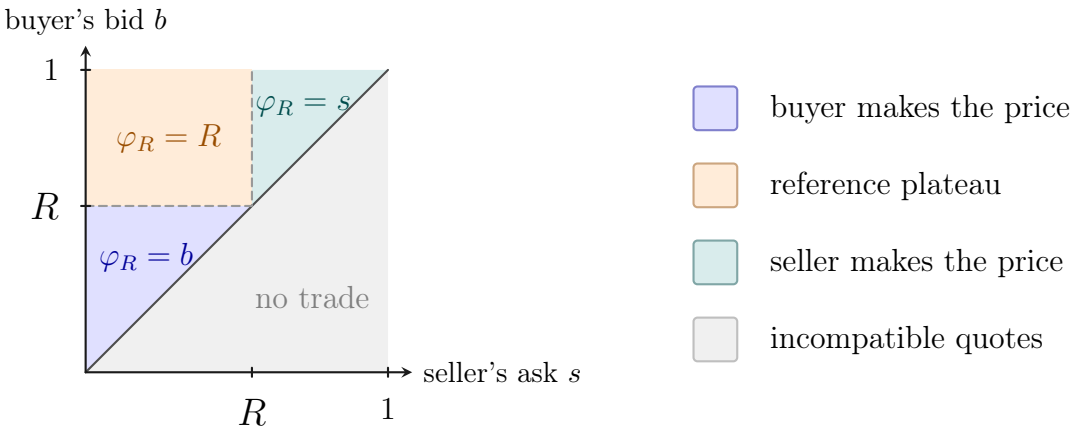


Figure 1: **Geometry of the reference-price rule.** Feasible quote pairs lie weakly above the trade boundary $s = b$. Below the reference, the buyer's bid determines the transaction price; above it, the seller's ask does. When the quotes straddle R, the transaction price equals the reference.

Strategies. A *behavioral strategy* for the buyer is a measurable map $\beta : [0,1] \to \Delta(\mathcal{A})$, where $\beta(\cdot \mid v)$ is the distribution of bids submitted by type v. The seller's strategy $\sigma : [0,1] \to \Delta(\mathcal{A})$ is defined symmetrically:

$$\beta(\cdot \mid v) \in \Delta(\mathcal{A}), \qquad \sigma(\cdot \mid c) \in \Delta(\mathcal{A}). \tag{Strategies}$$

A strategy is *pure* if it assigns probability one to a single quote for every type, in which case we write $b(v)$ and $s(c)$. Type distributions and strategies induce *ex-ante distributions* Σ_B over bids and Σ_S over asks: For every Borel set $B, S \subseteq \mathcal{A}$,

$$\Sigma_B(B) = \int_0^1 \beta(B \mid v)\, dF(v), \quad \Sigma_S(S) \ = \int_0^1 \sigma(S \mid c)\, dG(c). \tag{Quote Distributions}$$

Against these distributions, the interim payoffs from pure quotes are

$$U_B(b; v) = \int_{[0,b]} \big(v - \varphi_R(s,b)\big) \quad \Sigma_S(ds), U_S(s; c) = \int_{[s,1]} \big(\varphi_R(s,b) - c\big)\, \Sigma_B(db). \tag{Interim payoffs}$$

A strategy profile (β, σ) is a *Bayes–Nash equilibrium* if, for almost every type, every quote used with positive probability is an interim best response:

$$\begin{aligned} \operatorname{supp} \beta(\cdot \mid v) &\subseteq \arg\max_{b \in \mathcal{A}} U_B(b; v) \qquad \text{for } F\text{-almost every } v, \\ \operatorname{supp} \sigma(\cdot \mid c) &\subseteq \arg\max_{s \in \mathcal{A}} U_S(s; c) \qquad \text{for } G\text{-almost every } c. \end{aligned} \tag{Bayes–Nash equilibrium}$$

3 Incentives: Make the Price or Take the Price?

The reference price splits a trader's problem of maximizing expected utility into two regimes. Consider the buyer. Bidding below R, she pays her own bid on every trade: in this lower regime, this is a pure monopsony problem, in which shading is priced at the margin. Bidding at or above R, she pays a price determined entirely by the seller and the anchor, which her own bid cannot move. In this upper regime, the bid only *screens* which trades execute, and the best screen is the truthful one. The incentives for the seller are analogous.

A trader's objective is therefore sharply kinked at R, and her decision is not a standard first-order condition but a global choice between two local optima: *make* the price, or *take* the price. For arbitrary distributions Σ_B over bids and Σ_S over asks, we formally define the values of these two regimes for both sides of the market. The *price-maker values* (monopsony and monopoly) and the *price-taker values* (truthfulness) are given by:

$$\begin{aligned} M_B(v) &= \sup_{b\in[0,R)} (v-b)\,\Sigma_S([0,b]), \\ M_S(c) &= \sup_{s\in(R,1]} (s-c)\,\Sigma_B([s,1]). \end{aligned} \qquad \text{(Price-making values)}$$

$$\begin{aligned} T_B(v) &= (v-R)\,\Sigma_S([0,R]) + \int_{(R,\,v]} (v-s)\,\Sigma_S(ds), \\ T_S(c) &= (R-c)\,\Sigma_B([R,1]) + \int_{[c,\,R)} (b-c)\,\Sigma_B(db). \end{aligned} \qquad \text{(Price-taking values)}$$

We show that for any arbitrary distribution of actions — equilibrium play or not — best responses exist for any type and the optimal payoff can be achieved by price-making or price-taking.

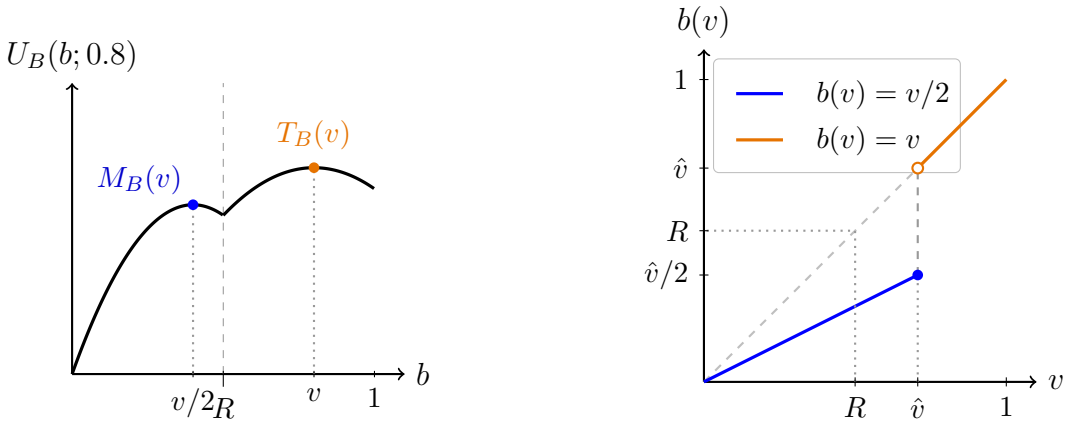
Proposition 1 (Make the price or take the price)**.** *For any quote distributions Σ_B and Σ_S, a best response exists for every type. The optimal payoff satisfies*

$$\sup_{b\in\mathcal{A}} U_B(b;v) \;=\; \begin{cases} \max\{M_B(v),\, T_B(v)\}, & v \geq R, \\ M_B(v), & v < R, \end{cases} \qquad \text{(Buyer Optimization)}$$

$$\sup_{s\in\mathcal{A}} U_S(s;c) \;=\; \begin{cases} \max\{M_S(c),\, T_S(c)\}, & c \leq R, \\ M_S(c), & c > R. \end{cases} \qquad \text{(Seller Optimization)}$$

Proof Outline. Every bid $b < R$ earns $U_B(b; v) = (v - b)\,\Sigma_S([0, b])$, and every ask $s > R$ earns $U_S(s; c) = (s - c)\,\Sigma_B([s, 1])$, whose suprema are $M_B(v)$ and $M_S(c)$. At or across the reference, the quote matters only through the execution set: for any $R \le b \le b'$ it holds that $U_B(b'; v) - U_B(b; v) = \int_{(b,b']}(v - s)\,\Sigma_S(ds)$ and the seller statement mirrors. The increment is nonnegative exactly while the marginal counterparties are profitable, so the truthful quote is optimal in this regime, attaining $T_B(v)$ and $T_S(c)$; for types on the wrong side of the reference ($v < R$ or $c > R$), every straddling or crossing quote yields a non-positive payoff, leaving price-making as the only candidate. The supremum over all quotes is therefore the larger of the two regime optima, and it is attained. Formal proof in Section A.1. □

Example 1 (The regime choice under uniform, truthful play)**.** Let $F = G = U[0, 1]$, let the reference price be $R = \frac{1}{2}$, and let the seller quote truthfully. Then $M_B(v) = v^2/4$, from the monopsony bid $v/2$, while $T_B(v) = \frac{1}{2}(v^2 - \frac{1}{4})$ from the truthful bid v. A price-maker trades with probability $v/2$ at a price she controls; a reference-taker trades with probability v at prices she does not. The values cross at $\hat{v} = \sqrt{2}/2 \approx 0.707$: types below make prices, and types above take the reference by bidding truthfully (Figure 2).



The objective function. Action Space (Fixed type $v = 0.8$): The pricing rule divides the objective function at $b = R$. Bidding below the anchor (price-making) yields a monopsony problem peaking at $M_B(v)$. Bidding above the anchor (price-taking) yields an execution-screening problem peaking at the truthful bid $T_B(v)$.

The best response. Type Space (All types $v \in [0, 1]$): The global best response maps the type space into the two discrete roles. Types below $\hat{v}$ strictly prefer to make the price, shading their bids. At exactly $\hat{v} = \sqrt{2}/2$, the global maximum shifts, and the optimal bid jumps across the plateau to the truthful reference-taking strategy.

Figure 2: **The Fracture of the Continuous Bidding Problem.** By severing trade execution from price determination, the reference-price rule forces traders to choose globally between two distinct local optima, creating a discontinuous jump in the strategic response.

Proposition 1 is a pointwise statement: against any fixed distribution of opposing quotes, every type can attain her optimal payoff by either making the price or taking it. Equilibrium requires these quote distributions to be generated endogenously by traders who are simultaneously solving the same problem. Moreover, Proposition 1 identifies the value of a best response, but not necessarily a unique best-response action; payoff-equivalent quotes may therefore survive and support additional fixed points. Equilibrium is thus a fixed point of a set-valued maker–taker problem: each side's division into price-makers and price-takers must be a best response to the division the other side generates. The rest of the paper asks which such configurations are self-consistent—and first among them, whether price-makers and reference-takers can coexist in equilibrium, on the two sides of the market at once.

4 Equilibrium Analysis

The problem with equilibria in bilateral trade is not existence, but excess. In standard bilateral double auctions, Bayesian–Nash equilibria are notoriously permissive, admitting a continuum of solutions (Leininger et al. 1989). The reference-price double auction inherits this pathology in full: for any candidate price, all trading types can coordinate on quoting it exactly, and the convention sustains itself—quoting less aggressively forfeits trade, while quoting more aggressively cannot improve execution and can only worsen the transaction price.

Proposition 2 (Every price can be an equilibrium price)**.** *Fix any reference price $R \in (0,1)$. For every $p \in (0,1)$, the strategy profile in which a buyer with value $v \geq p$ bids p and a buyer with $v < p$ bids 0, while a seller with cost $c \leq p$ asks p and a seller with $c > p$ asks 1, is a Bayes–Nash equilibrium. Whenever trade occurs, it clears at price p.*

Proof Outline. Consider a buyer with value $v \geq p$. Bidding p trades exactly against sellers with cost $c \leq p$ and yields price p. A lower bid forfeits these trades, while a higher bid reaches no additional profitable seller and can only weakly increase the transaction price. Hence p is optimal. If $v < p$, any bid that induces trade requires paying at least $p > v$, so bidding 0 is optimal. The seller's argument is symmetric: a seller with $c \leq p$ optimally asks p, while a seller with $c > p$ cannot profitably trade at the common price and optimally asks 1. Thus the proposed strategies are mutual best responses for every R. The full proof is relegated to Section A.2.

The proposition is not intended to characterize all Bayes–Nash equilibria. Its point is that even this simple class sustains every interior transaction price, independently of the reference. This mirrors the well-known multiplicity of the classical sealed-bid double auction (Leininger et al. 1989, Satterthwaite and Williams 1989). Bayes–Nash equilibrium alone therefore places little discipline on price discovery.

4.1 A Minimal Refinement: Undominated Equilibrium

The equilibria of Proposition 2 survive only because the convention is never tested. A buyer with a value below p bids 0 and loses nothing, but only because every ask she meets sits at p, above her value: her bid is never payoff-relevant on the path of play. Against any belief that puts positive probability on asks below the convention, a bid below her value is strictly better: it trades at a profit when it trades and never at a loss, while the bid 0 forfeits those trades for nothing. Her prescribed quote is weakly dominated, and the seller quoting 1 above p is in the same position. The one-price equilibria are held together entirely by such quotes: best responses to one perfectly foreseen distribution of opposing play, errors against every other.

We therefore consider the minimal refinement of undominated equilibrium strategies. This rules out only quotes that are never better and sometimes worse against the opponent's behavior; it imposes no purity, monotonicity, or continuity restrictions.[12]

Formally, a bid $b \in \mathcal{A}$ is *weakly dominated* for buyer type v if there exists a mixed action $\mu \in \Delta(\mathcal{A})$, and an ask $s \in \mathcal{A}$ is *weakly dominated* for seller type c if there exists a mixed action $\nu \in \Delta(\mathcal{A})$, such that

$$\int_{\mathcal{A}} U_B(\widetilde{b}, s; v)\, \mu(d\widetilde{b}) \geq U_B(b, s; v) \qquad \text{for every } s \in \mathcal{A}, \qquad \text{(Buyer dominance)}$$

$$\int_{\mathcal{A}} U_S(b, \widetilde{s}; c)\, \nu(d\widetilde{s}) \geq U_S(b, s; c) \qquad \text{for every } b \in \mathcal{A}, \qquad \text{(Seller dominance)}$$

respectively, with strict inequality for at least one opposing quote. A Bayes–Nash equilibrium is *undominated* if the prescribed mixed strategy assigns probability one to undominated quotes for almost every type.

[12] The refinement is standard in bilateral double auctions. At the one-sided endpoints $k \in \{0, 1\}$, one trader determines the transaction price while the other has zero price impact and therefore has a weakly dominant truthful acceptance rule. Indeed, Kadan (2007), Satterthwaite and Williams (1989), Williams (1991) shows that, restricting attention to non-dominated strategies yields a unique equilibrium at these endpoints.

In the reference-price auction, this weak refinement has unusually sharp force. On the price-making side of the reference, a trader's quote determines the transaction price and strategic shading may be optimal. Once the quote crosses the reference, however, it can affect only whether trade occurs, not improve the price on trades already available. Truthful reporting therefore becomes the unique undominated acceptance rule.

Lemma 1 (Dominance beyond the reference)**.** *The following statements hold.*

(i) For every buyer type $v > R$, each $b \in [R, 1] \setminus \{v\}$ is weakly dominated by $b = v$. For every $v < R$, each $b \in [R, 1]$ is weakly dominated.

(ii) For every seller type $c < R$, each $s \in [0, R] \setminus \{c\}$ is weakly dominated by $s = c$. For every $c > R$, each $s \in [0, R]$ is weakly dominated.

Consequently, except for the null boundary types $v = R$ and $c = R$, every undominated quote on the far side of the reference is truthful, and quoting exactly R is weakly dominated.

Proof Outline. Above the reference, a buyer's bid selects which sellers she reaches but never the price of a retained trade. For $v > R$, replacing any $b \geq R$ by v thus leaves existing prices unchanged, adds only profitable trades, and removes only unprofitable ones: truthful bidding weakly dominates. For $v < R$, any trade won by a bid above the reference prices at least at $R > v$, so crossing the reference is weakly dominated. The seller side is symmetric; full proof in Section A.3.

Lemma 1 immediately eliminates the one-price conventions of Proposition 2. If $p < R$, a positive measure of sellers ask p rather than their lower costs; if $p > R$, a positive measure of buyers bid p rather than their higher values; and at $p = R$, both forms of posturing occur.

4.2 The Collapse to One-Sided Price-Making

Once dominated quotes are eliminated, the surviving quotes sort into two kinds. A buyer either bids below the reference—and then, by the price rule, any trade she wins clears at her bid—or bids truthfully above it and accepts a price formed elsewhere; a seller either asks above the reference, pricing any trade he wins, or asks truthfully below it. Nothing in the rule assigns these roles. Both remain open to both sides at once, and the natural conjecture is that equilibrium mixes them—some trades priced by bids below the reference, others by asks above it, with occasional trades at the reference itself whenever quotes straddle it.

Our first main result overturns this conjecture entirely. In every undominated equilibrium, price discovery belongs entirely to one side of the market: whichever side prices one trade prices them all. And the price the rule is built to select is the one price that never realizes.

Theorem 1 (One-sidedness)**.** *In every undominated Bayes–Nash equilibrium, price discovery is one-sided: either almost no seller asks below the reference price, or almost no buyer bids above it. Consequently, within any such equilibrium, either every transaction executes at the seller's ask, strictly above the reference, or every transaction executes at the buyer's bid, strictly below it. In particular, no transaction ever occurs at the reference price itself.*

Proof outline. The formal proof, in Section A.4, accommodates arbitrary mixed and non-monotone strategies. The intuitive argument goes as follows: Suppose both channels are active: the seller asks below the reference with positive probability, and the buyer bids above it with positive probability.

1. *Truthfulness.* By Lemma 1, all such quotes are truthful: an ask below the reference equals the seller's cost, a bid above it equals the buyer's value.

2. *No gap near the reference.* If no ask fell in some interval (b, R), the buyer could lower any truthful bid to b: she would keep every trade against asks below b and pay her own bid on each, rather than at least the reference. Truthful bidding above R would then not be optimal, so the participating asks must reach up to the reference.

3. *The marginal seller defects.* At a cost just below the reference, truthful asking earns almost nothing: every trade pays at most R, so the margin vanishes as $c \to R$. Asking strictly above the reference instead, the seller trades with the buyer's truthful bids at a margin bounded away from zero. For costs close enough to R, the deviation is strictly profitable, contradicting equilibrium.

Each channel thus destabilizes the other: price-making below the reference tempts the truthful buyer to stop taking prices, and truthful bids above it tempt the marginal seller to stop asking below. No undominated equilibrium sustains both. The final clause follows because quotes at the reference itself are dominated (Lemma 1): with one side of the reference unpopulated, every transaction is priced by the other side's quotes, strictly away from R. □

The implication of Theorem 1 is striking. The reference price rule is designed to select the reference whenever the submitted quotes permit it. However, equilibrium ensures that they never do.

4.3 Equilibrium Characterization: Truncated Monopsony & Monopoly

Theorem 1 narrows the possible equilibria to two candidate regimes. For buyer-made prices to support trade, sellers must take those prices by asking below R; for seller-made prices, buyers must take them by bidding above it. Beyond no-trade equilibria, the bilateral double auction with a reference price therefore admits only two asymmetric configurations: buyers make prices below the reference while sellers respond through cost cutoffs, or sellers make prices above the reference while buyers respond through value cutoffs.

This section completes the analysis: we characterize the set of undominated Bayes–Nash equilibria in full. The characterization is simple. Conditional on which side makes prices, every equilibrium outcome belongs to a one-parameter family of truncated monopsony or monopoly strategies, indexed by a single scalar. This simplicity is not built into the refinement. Undominatedness disciplines quotes across the reference, where they determine only execution, but it places almost no direct restriction on the price-making side: price-makers may in principle shade, pool, leave gaps, create atoms, or mix over quotes, and the passive side may participate or abstain. Equilibrium eliminates nearly all of this freedom.

The passive side turns into a truthful demand/supply curve. Once one side makes prices, equilibrium strips the other side of any strategic pricing role. On the range where trade occurs, the passive side simply reveals which types are willing to accept the quoted price. As a result, the price-maker faces the primitive distribution of counterpart types rather than an endogenous distribution of strategic quotes.

- *Buyer-made prices.* Suppose no buyer bids at or above the reference price. Every trade therefore executes at the buyer's bid, and asks above R never trade. Let $\bar{b}$ denote the supremum of the bid distribution. A seller with cost $c < \bar{b}$ trades with positive probability by asking c. Undominatedness leaves him two relevant options: report his true cost below the reference, or submit an inactive ask above it. The latter yields zero, whereas truthful asking earns strictly positive expected surplus against the equilibrium bid distribution. Hence almost every seller with $c < \bar{b}$ asks truthfully; sellers with $c \geq \bar{b}$ never trade, and their quotes are outcome-irrelevant. Throughout the active range, buyers therefore face the primitive supply distribution G itself: a price-making buyer who bids b trades with probability $G(b)$ and pays her own bid.

- *Seller-made prices.* The argument is symmetric. Suppose no seller asks at or below the reference price, so every trade executes at the seller's ask and bids below R never trade.

Let $\underline{s}$ denote the infimum of the ask distribution. Buyers with values $v > \underline{s}$ strictly prefer truthful participation above the reference to an inactive bid below it, while lower-value buyers never trade and their quotes are outcome-irrelevant. Thus, throughout the active range, sellers face the primitive demand curve $1 - F$: a price-making seller who asks s trades with probability $1 - F(s)$ and receives her own ask.

The active side turns into a monopsonist/monopolist. In either regime, strategic interaction therefore collapses to a single price-maker choosing against a truthful acceptance curve. A price-making buyer hence solves a standard monopsony problem; a price-making seller solves its monopoly mirror:

$$\max_b (v - b)\, G(b) \quad \text{and} \quad \max_s (s - c)\big(1 - F(s)\big). \qquad \text{(Monopsony / Monopoly Problem)}$$

Under regularity of distributions, each objective is strictly single-peaked. The optima are characterized by the seller's virtual cost $J_S(c) = c + \frac{G(c)}{g(c)}$ and the buyer's virtual value $J_B(v) = v - \frac{1-F(v)}{f(v)}$; inverting them yields the monopsony bid and monopoly ask:

$$\begin{aligned} b^M(v) &= \arg\max_b (v - b)G(b) = J_S^{-1}(v), \\ s^M(c) &= \arg\max_s (s - c)\big(1 - F(s)\big) = J_B^{-1}(c). \end{aligned} \qquad \text{(Monopsony Bid / Monopoly Ask)}$$

Caps and Ceilings. The most aggressive unconstrained quotes are those of the highest-value buyer and the lowest-cost seller:

$$\bar{b}^* = b^M(1) \quad \text{and} \quad \underline{s}^* = s^M(0). \qquad \text{(Absolute Bid Ceiling / Ask Floor)}$$

The only remaining degree of freedom is where the bid distribution ends. Suppose its highest bid is $\bar{b}$. Types whose unconstrained monopsony bids lie below $\bar{b}$ choose those bids, and choose them uniquely. Types who would prefer to bid above $\bar{b}$ reach no additional active sellers by doing so; they would only pay more on the same set of trades. They therefore pool at $\bar{b}$. Mixing, gaps, and additional atoms disappear, and the buyer strategy takes the truncated form

$$m_{\bar{b}}(v) = \min\{b^M(v), \bar{b}\}.$$

The cap itself cannot exceed the unconstrained ceiling $\bar{b}^*$, and it may approach the reference from below but cannot equal it unless $\bar{b}^*$ already lies weakly below the reference.

The seller-side mirror is governed by an ask floor $\underline{s}$. The admissible parameters are therefore

$$\mathcal{B}(R) = \begin{cases} (0, \bar{b}^*], & R \geq \bar{b}^* \\ (0, R), & R < \bar{b}^* \end{cases} \quad \text{and} \quad \mathcal{S}(R) = \begin{cases} [\underline{s}^*, 1), & R \leq \underline{s}^* \\ (R, 1), & R > \underline{s}^*. \end{cases} \qquad \text{(Admissible Cap / Floor)}$$

Once the price-making side is known, equilibrium is thus indexed by a single scalar: a bid cap under buyer control, or an ask floor under seller control.

Theorem 2 (Equilibrium Characterization)**.** *Consider undominated Bayes–Nash equilibria with an active buyer side ($\beta_+ = 0$) and positive trade probability.*

1. ***Necessity.*** *In every such equilibrium, the highest bid is a cap $\bar{b} = \sup \operatorname{supp} \Sigma_B \in \mathcal{B}(R)$. Almost every buyer plays the pure bid $m_{\bar{b}}(v) = \min\{b^M(v), \bar{b}\}$, and almost every seller with $c < \bar{b}$ asks truthfully. If $\bar{b} < \bar{b}^*$, the types $v > J_S(\bar{b})$ pool in an atom at $\bar{b}$. The equilibrium outcome—the trade set, prices, and payoffs—is determined by $\bar{b}$ alone.*

2. ***Sufficiency.*** *Conversely, for every $\bar{b} \in \mathcal{B}(R)$, the profile in which buyers play $m_{\bar{b}}$, sellers with $c < \bar{b}$ ask truthfully, and all other sellers ask 1, is an undominated Bayes–Nash equilibrium.*

The seller-side mirror replaces the bid cap with an ask floor $\underline{s} \in \mathcal{S}(R)$ and the strategies with $s_{\underline{s}}(c) = \max\{s^M(c), \underline{s}\}$.

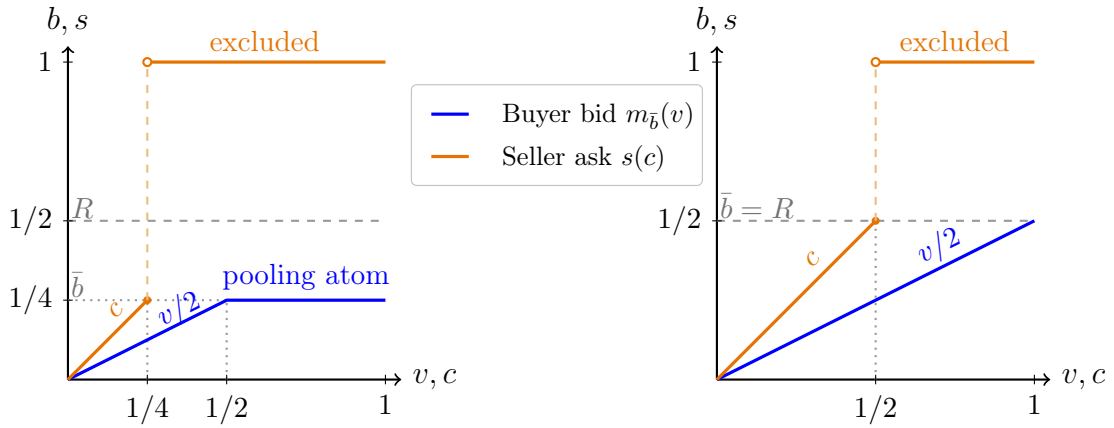
Proof outline. Neccessity is established in three parts:

1. Any seller with cost below $\bar{b}$ earns strictly positive surplus from a truthful ask and zero from an inactive one, so truthful reporting is his unique undominated best response.

2. The buyer therefore faces the supply curve G, and regularity makes each type's monopsony objective strictly single-peaked; every type whose unconstrained optimum is feasible plays it as a unique pure best response, and every type whose optimum exceeds $\bar{b}$ strictly prefers $\bar{b}$ to any lower bid, producing the pooling atom.

3. The cap cannot exceed $\bar{b}^*$: no type's monopsony objective increases beyond the unconstrained optimum of the highest type, so bids above $\bar{b}^*$ are never best responses. And the cap cannot equal the reference when $R < \bar{b}^*$: it would then bind for a positive mass of types, forcing them to pool at R, a quote that Lemma 1 shows is dominated.

Sufficiency follows by verifying that no type on either side has a profitable deviation against the truncated profile. The formal proof is relegated to Section A.5. □



Binding cap ($\bar{b} = 1/4 < \bar{b}^*$): Buyers cap their bids at an arbitrary level, causing high types $v \geq 1/2$ to pool. Sellers with costs $c > 1/4$ retreat to the non-trading ask of 1.

Unconstrained cap ($\bar{b} = \bar{b}^* = 1/2$): The cap is set to the maximal admissible level. Buyers bid strictly monotonically across the entire type space, and no pooling occurs.

Figure 3: **The Rigidity of Truncated Monopsony Equilibria.** Bidding profiles for the uniform benchmark at reference $R = \frac{1}{2}$. The active side (buyers) determines the extent of market participation through the self-fulfilling placement of the bid cap $\bar{b}$. The reference anchor acts strictly as an off-path bound.

Because the surviving play is the classical monopsony program truncated at the cap, we call the resulting outcomes *truncated monopsony equilibria*, and their seller-side mirrors *truncated monopoly equilibria*. The equilibria at the tightest admissible bounds play a central role in the remainder of the paper. For a given reference, consider the *frontier cap* and *floor*

$$q(R) = \min\{R, \bar{b}^*\} \qquad \text{and} \qquad t(R) = \max\{R, \underline{s}^*\}. \qquad \text{(Frontier Cap / Floor)}$$

The *buyer-side frontier equilibrium* is the truncated monopsony equilibrium with cap $q(R)$, the *seller-side frontier equilibrium* the truncated monopoly equilibrium with floor $t(R)$. The frontier bound is admissible exactly when the reference does not constrain the active side: $q(R) = \bar{b}^*$ whenever $R \geq \bar{b}^*$, and $t(R) = \underline{s}^*$ whenever $R \leq \underline{s}^*$; we then call the frontier *unconstrained*. When instead $R < \bar{b}^*$, the frontier cap $q(R) = R$ is not itself admissible: a cap at the reference would force a positive mass of buyers to pool at R, a dominated quote. In this case an *ε-frontier equilibrium* is any truncated monopsony equilibrium with cap in $[R - \varepsilon, R)$.[13] The seller-side mirror replaces the cap with a floor in $(R, R + \varepsilon]$.

[13] Payoffs and outcomes are continuous in the cap, so statements about a constrained frontier are understood as limits as $\varepsilon \to 0$.

Three consequences follow immediately from Theorems 1 and 2.

Corollary 1 (Existence and the equilibrium price set)**.** *Equilibria with positive trade exist on both sides for every reference price $R \in (0,1)$. However, no equilibrium, on either side, generates a transaction price in the interval between the frontier cap $q(R)$ and the frontier floor $t(R)$. Outside this gap, prices are unrestricted: every price below the cap and every price above the floor arises in some equilibrium.*

For every reference, both regimes support equilibria with trade. Buyer-made prices lie strictly below the reference and seller-made prices strictly above it. The reference is not an equilibrium price; it is a boundary around which equilibrium leaves a gap. Theorems 1 and 2 also determines how much information prices reveal.

Corollary 2 (One-sided information revelation)**.** *In a buyer-side equilibrium with cap $\bar{b}$, a traded price $p < \bar{b}$ fully reveals the buyer's type, $V = J_S(p)$, while a trade at the cap reveals only $V \geq J_S(\bar{b})$. The seller's type remains distributed as $G(\cdot \mid C \leq p)$, independently of the buyer's type conditional on trade and price. The seller-side statement is symmetric.*

The price therefore reveals the maker's type, up to pooling at the cap or floor, but reveals about the taker only that trade was acceptable.[14] Price formation is not an aggregation of two strategic reports: one side chooses an optimal quote against the primitive acceptance curve of the other. Finally, at an unconstrained frontier, the mechanism is outcome-equivalent with the classical buyer's bid and seller's ask double auction.[15]

Corollary 3 (Classical one-sided equivalence)**.** *If $R \geq \bar{b}^*$, the buyer-side frontier equilibrium is outcome-equivalent to the unique undominated Bayesian–Nash equilibrium from the buyer's-bid double auction: If $R \leq \underline{s}^*$, the seller-side frontier equilibrium is outcome-equivalent to the unique undominated Bayesian–Nash equilibrium from the seller's-ask double auction.*

The equivalence isolates what the reference price does—and what it does not do. Classical one-sided auctions assign pricing power by design; the reference-price auction is formally symmetric and assigns it to neither side. Yet equilibrium endogenously reproduces the same one-sided outcomes.

[14] This contrasts the k-double auction for $k \in (0,1)$, where the price $kb + (1-k)s$ combines both reports.

[15] These are the $k = 1$ and $k = 0$ endpoints of the k-double auction (Chatterjee and Samuelson 1983, Kadan 2007, Satterthwaite and Williams 1989, Williams 1987).

5 Welfare Analysis

Theorems 1 and 2 reduce the reference-price auction to two families of one-sided equilibria. This section evaluates what equilibrium play costs. Two questions organize it: how the equilibria within each family rank, and how much of the attainable gains from trade the best of them retains.

Gains from Trade. For any strategy profile (β, σ), expected gains from trade are

$$W(\beta, \sigma) \;=\; \mathbb{E}\big[(v - c)\,\mathbf{1}\{s \le b\}\big], \qquad \text{(Expected Gains from Trade)}$$

where the quotes b and s are drawn from $\beta(\cdot \mid v)$ and $\sigma(\cdot \mid c)$. Because utility is quasilinear, transfers cancel between the parties: W is the sum of the two sides' ex-ante payoffs, and welfare depends on the profile only through the trades it executes, not the prices at which it executes them.

By Theorem 2, the outcome of a truncated monopsony equilibrium is determined by its cap alone, and the outcome of a truncated monopoly equilibrium by its floor; we accordingly write $W_B(\bar{b})$ and $W_S(\underline{s})$ for their expected gains from trade.

Proposition 3 (Pareto ranking)**.** *Fix R and admissible caps $\bar{b} < \bar{b}'$ in $\mathcal{B}(R)$. Moving from the $\bar{b}$-cap to the $\bar{b}'$-cap equilibrium leaves no type of either trader worse off and makes a positive measure of both strictly better off (every buyer with $v > J_S(\bar{b})$, every seller with $c < \bar{b}'$). Thus, expected gains from trade are strictly increasing in the cap. The mirror statements hold for the truncated monopoly family, with all orders reversed in the floor $\underline{s}$.*

Proof outline. A higher cap enlarges every buyer's feasible set and leaves the supply curve she faces unchanged, so it helps exactly the types whose bids were capped. The equilibrium bid function rises pointwise, so every seller trades weakly more often at weakly higher prices, strictly so where the uncapped bids reach him. Aggregating over types gives the welfare and trade-probability claims. Formal proof in Section A.9. □

Unanimity is the notable feature. In most mechanisms, expanding trade requires compensating someone; here the cap is a pure constraint, binding on both sides at once, so relaxing it improves the terms of trade for one side and the extent of trade for both. Within each family, coordination is therefore unproblematic: every force pushes toward the frontier, with cap $q(R)$ or floor $t(R)$.

Frontier Welfare. When a frontier is unconstrained, Corollary 3 identifies it with the classical buyer's-bid or seller's-ask auction: the maker bids $b^M(V)$ or asks $s^M(C)$, and trade occurs exactly when the maker's virtual quote clears the taker's type,

$$\begin{aligned} W_B(\bar{b}^*) &= \mathbb{E}\big[(v-c)\,\mathbf{1}\{J_S(c) \le v\}\big], \\ W_S(\underline{s}^*) &= \mathbb{E}\big[(v-c)\,\mathbf{1}\{c \le J_B(v)\}\big]. \end{aligned} \qquad \text{(Frontier Welfares)}$$

These are the best outcomes the two families admit. Which of them the reference makes available is a property of R alone.

Proposition 4 (The optimal reference)**.** *For $R \in (0,1)$, let $W^*(R)$ denote the supremum of expected gains from trade over all undominated Bayes–Nash equilibria at reference R. Then*

$$W^*(R) \;=\; \max\big\{W_B(q(R)),\, W_S(t(R))\big\}, \qquad \sup_{R\in(0,1)} W^*(R) \;=\; \max\big\{W_B(\bar{b}^*),\, W_S(\underline{s}^*)\big\},$$

and the right-hand supremum is attained by an undominated equilibrium at reference R if and only if R leaves a welfare-maximal side unconstrained: $R \ge \bar{b}^*$ and $W_B(\bar{b}^*) \ge W_S(\underline{s}^*)$, or $R \le \underline{s}^*$ and $W_S(\underline{s}^*) \ge W_B(\bar{b}^*)$.

Proof outline. By Theorem 2, every equilibrium belongs to one of the two families, and by Proposition 3 each family's welfare is strictly increasing toward its frontier bound; continuity of W_B and W_S in the bound gives the value $W^*(R)$, whether or not it is attained. The frontier bound is admissible exactly when the reference does not constrain that side (Corollary 3), so attainment holds precisely under the stated condition; when a welfare-maximal side is constrained, every admissible bound on it falls strictly short. Formal proof in Section A.10. □

The proposition is a statement about availability, not prediction: it says which references put the best outcome on the table, with both regimes still candidates. Which regime the market coordinates on is the subject of Section 6, and what a designer who controls only R can extract from the answer is taken up at its end.

A Universal Lower Bound. Write $W_{\mathrm{FB}} = \mathbb{E}\big[(V-C)^+\big]$ for the *first-best expected gains from trade*, realized when every profitable trade executes, and let W_{SB} denote the *Bayesian second best*: the maximum expected gains from trade of any interim individually rational, Bayesian incentive-compatible, and ex-ante weakly budget-balanced mechanism. By Myerson and Satterthwaite (1983), $W_{\mathrm{SB}} < W_{\mathrm{FB}}$: no feasible mechanism trades efficiently. The question is where the reference-price auction falls relative to both.

Theorem 3 (Universal welfare guarantee)**.** *For every regular pair* (F, G)*,*

$$\sup_{R\in(0,1)} W^*(R) \;=\; \max\{W_B(\bar{b}^*),\, W_S(\underline{s}^*)\} \;\geq\; \max\Big\{\tfrac{1}{2}\,W_{\mathrm{SB}},\;\tfrac{1}{\pi}\,W_{\mathrm{FB}}\Big\}. \qquad \text{(Welfare Bounds)}$$

Proof outline. By (Frontier Welfares), the frontier welfares are exactly the gains from trade of the buyer-offering and seller-offering mechanisms of the approximation literature, and their equal mixture is the *random-offerer* mechanism. Brustle et al. (2017) show that the random offerer retains at least half of the Bayesian second best; Deng et al. (2022) established a distribution-free guarantee against first best, sharpened by Fei (2022) and most recently by Jo (2026) to $1/\pi$. Both results hold for arbitrary independent priors, hence under regularity, and the better frontier weakly dominates the average. Formal proof in Section A.11. □

The two lower bounds answer different questions: the $W_{\mathrm{SB}}/2$ bound measures the loss from one-sided pricing against the best any feasible mechanism could do, the W_{FB}/π bound against fully efficient trade, and neither dominates the other across environments. Both are guarantees on the frontier, not on every equilibrium: by Proposition 4, the market inherits them exactly at references that leave a welfare-maximal side unconstrained. The collapse to one-sided pricing is therefore costly but boundedly so.

Example 2 (The uniform benchmark)**.** For $F = G = U[0,1]$, the first best trades whenever $v \geq c$ and yields $W_{\mathrm{FB}} = \frac{1}{6}$; the second best, attained by the linear equilibrium of the $k = \frac{1}{2}$ auction of Chatterjee and Samuelson (1983), trades whenever $v \geq c + \frac{1}{4}$ and yields $W_{\mathrm{SB}} = \frac{9}{64}$. The reference-price auction reaches neither at any reference: its best outcome is the frontier, where the buyer-monopsonist trades on the virtual set $v \geq 2c$ and

$$W_B(\bar{b}^*) \;=\; W_S(\underline{s}^*) \;=\; \tfrac{1}{8} \;=\; \tfrac{8}{9}\cdot W_{\mathrm{SB}} \;=\; \tfrac{3}{4}\cdot W_{\mathrm{FB}}.$$

The lost ninth has a name: the burden of unilateral rent extraction. In the interior auction both sides shade and the cost of information rents is shared; the surviving equilibria concentrate the entire burden on one side, whose optimal response is to truncate exactly the marginal trades—those with $c + \frac{1}{4} \leq v < 2c$—that shared shading preserves. Theorem 3 is slack here: $\frac{1}{8}$ sits strictly inside the sandwich $\frac{9}{128} \leq \frac{1}{8} \leq \frac{9}{64}$, and well above $W_{\mathrm{FB}}/\pi = \frac{1}{6\pi}$.

The frontier welfares are generically unequal, $W_B(\bar{b}^*) \neq W_S(\underline{s}^*)$, so the two available outcomes are not interchangeable—and which of them the market reaches is not a parameter anyone sets. It is a coordination problem between the two sides of the market, and resolving it is the task of the next section.

6 Convention Formation: Which Side Makes Prices?

The analysis so far characterizes the set of undominated Bayesian Nash equilibria, but it does not determine which side ultimately assumes the price-making role. We therefore add a simple reduced-form model of *convention formation*. Before private values and costs are realized, the two sides of the market coordinate over whether they will act as price makers or price takers. Each action represents a complete type-contingent trading convention to be followed once private information is realized.[16] Its purpose is to study which of the two one-sided equilibrium regimes is more robust to failures of coordination. The resulting convention-formation problem is a 2×2 coordination game, to which we apply the criteria of risk dominance and payoff dominance (Harsanyi and Selten 1988).

6.1 The Role-Assignment Game

We restrict attention to the Pareto-frontier convention of each one-sided equilibrium family. This selects, conditional on assigning price-making authority to a given side, the equilibrium that both sides weakly prefer within that family. Under *buyer control*, buyers make prices with the capped monopsony strategy $\beta^B(v) = \min\{b^M(v), q(R)\}$ and sellers take prices by reporting costs truthfully; under *seller control*, sellers make prices with the floored monopoly strategy $\sigma^S(c) = \max\{s^M(c), t(R)\}$ and buyers report values truthfully. When a frontier is constrained, the profile and all payoffs below are understood in the ε-frontier limit of Section 4.3. The ex-ante payoffs of the two sides under buyer control are

$$\begin{aligned} \Pi_B^B &= \int_0^1 \big(v - \beta^B(v)\big)\, G\big(\beta^B(v)\big)\, dF(v), \\ \Pi_S^B &= \int_0^1 \int_0^{\beta^B(v)} \big(\beta^B(v) - c\big)\, dG(c)\, dF(v), \end{aligned} \qquad \text{(Buyer Control: Buyer / Seller Payoff)}$$

and under seller control

$$\begin{aligned} \Pi_B^S &= \int_0^1 \int_{\sigma^S(c)}^1 \big(v - \sigma^S(c)\big)\, dF(v)\, dG(c), \\ \Pi_S^S &= \int_0^1 \big(\sigma^S(c) - c\big)\big[1 - F\big(\sigma^S(c)\big)\big]\, dG(c). \end{aligned} \qquad \text{(Seller Control: Buyer / Seller Payoff)}$$

[16] This convention-formation stage is an additional modeling assumption rather than an equilibrium refinement of the underlying Bayesian game. In particular, a deviation at this stage represents a coordinated change in the convention adopted by one side of the market, not a unilateral deviation by an individual type in the trading game.

Two further outcomes arise if the sides pair the roles differently.

- *Both sides make prices.* Their quotes never cross. Bids lie weakly below $q(R) \leq R$ and asks weakly above $t(R) \geq R$. So no trade occurs and both earn zero.
- *Both sides take prices.* If both sides take prices, the buyer bids truthfully above $t(R)$ and the seller asks truthfully below $q(R)$; the quotes straddle the reference. Every such trade executes at R. The resulting *concession payoffs* are

$$X_B = G\big(q(R)\big) \int_{t(R)}^{1} (v - R)\, dF(v),$$
$$X_S = \big[1 - F\big(t(R)\big)\big] \int_{0}^{q(R)} (R - c)\, dG(c). \qquad \text{(Buyer / Seller Concession)}$$

Definition 2 (Role-assignment game)**.** The *role-assignment game* is the two-player normal-form game in which the buyer side and the seller side simultaneously choose to *make* or *take* prices, adopting the corresponding strategies above, with expected payoffs

	Seller takes	Seller makes
Buyer makes	$\Pi_B^B,\ \Pi_S^B$	$0,\ 0$
Buyer takes	$X_B,\ X_S$	$\Pi_B^S,\ \Pi_S^S$

The next proposition shows that the role assignment game is a strict coordination game:

Proposition 5 (Coordination Game)**.** *Buyer control and seller control are strict Nash equilibria, and neither off-diagonal profile is a Nash equilibrium.*

Proof outline. If both sides make prices, either side strictly gains by taking the other's prices, since both regimes trade with positive probability. If both take prices, a deviating buyer can trade with the same participating sellers at her capped monopsony bids rather than at the reference, and shading is strictly profitable for a positive mass of types; the seller mirror is identical. The strict inequalities $\Pi_B^B > X_B$, $\Pi_S^S > X_S$, $\Pi_S^B > 0$, and $\Pi_B^S > 0$ make both diagonal profiles strict equilibria. At constrained frontiers, all four inequalities hold with slack uniform over caps and floors near the frontier bound, so they pass to the ε-frontier limit and the game is well defined at every reference. Formal proof in Section A.12. □

The role-assignment game is an equilibrium-selection device, not a new model: its two actions per side are complete type-contingent strategies of the Bayesian game, and each cell records the exact ex-ante payoff of the corresponding profile. Payoffs are evaluated ex ante because that is the information a selection can use. A convention about who makes prices forms before quotes are submitted, so it can condition only on commonly known objects—the reference, not the realized types.[17]

Coordination games pose a selection question, and game theory offers two classical criteria (Harsanyi and Selten 1988): payoff dominance, which selects an equilibrium both players strictly prefer, and risk dominance, which selects the equilibrium more robust to strategic uncertainty. We take them in turn.

6.2 Risk Dominance: The Threshold of Control

Payoffs alone do not force the choice between two strict equilibria: coordination must survive strategic uncertainty about which regime the opposing side expects. The classical criterion for equilibria robust to exactly this uncertainty is *risk dominance* (Harsanyi and Selten 1988). Evolutionary models of convention formation, in which populations of players adapt by precedent and occasional error, select the risk-dominant equilibrium in the long run whether or not it is the one both players prefer (Kandori et al. 1993, Young 1993), and global games select it as the unique outcome surviving vanishing payoff uncertainty (Carlsson and Van Damme 1993).

Risk dominance compares the losses that miscoordination inflicts. If the market targets buyer control and coordination fails, a buyer side that abandoned its role has fallen back to joint price-taking, earning X_B instead of Π_B^B, while a seller side that abandoned its role has landed at incompatible quotes, earning 0 instead of Π_S^B. The *deviation losses* supporting buyer control are therefore $\Pi_B^B - X_B$ and Π_S^B; those supporting seller control are, in mirror, $\Pi_S^S - X_S$ and Π_B^S. The equilibrium whose product of losses is larger is the harder one to destabilize.

Definition 3 (Risk dominance)**.** Define the *risk-dominance index*

$$\nabla(R) \;=\; \left(\Pi_B^B - X_B\right)\Pi_S^B \;-\; \Pi_B^S\left(\Pi_S^S - X_S\right). \qquad \text{(Risk-Dominance Index)}$$

Buyer control is *strictly risk-dominant* if $\nabla(R) > 0$, and seller control if $\nabla(R) < 0$.

[17] A type-contingent assignment of roles would itself be a strategy, and strategies are already characterized by Theorem 2.

The index inherits a direction from the reference. Raising R relaxes the buyer side's cap, tightens the seller side's floor, and moves the joint-price-taking fallback in the buyer side's favor: both losses supporting buyer control rise, both supporting seller control fall, without any curvature assumption. On the frozen interval, where the regime payoffs are flat in R, the movement of the fallbacks alone keeps the monotonicity strict.

Theorem 4 (The threshold of control)**.** *For every regular pair (F, G) there exists a unique threshold $R^\dagger \in (0,1)$ such that seller control is strictly risk-dominant at every $R < R^\dagger$ and buyer control at every $R > R^\dagger$. For every reflection-symmetric pair, $R^\dagger = \frac{1}{2}$.*

Proof outline. The four deviation losses are continuous in R; the two supporting buyer control are nondecreasing and the two supporting seller control nonincreasing, by an envelope computation combined with a decomposition of the fallback integrands at the virtual quotes $J_S(R)$ and $J_B(R)$. At every reference at least one product moves strictly—through the cap below $\bar{b}^*$, the fallbacks on the frozen interval, and the floor above $\underline{s}^*$—so ∇ is strictly increasing; its boundary limits are $-\Pi_B^S(\Pi_S^S - X_S) < 0$ at 0 and the mirror positive value at 1, so it has exactly one zero. Under reflection symmetry, $R \mapsto 1 - R$ exchanges the regimes and the fallbacks, so $\nabla(1-R) = -\nabla(R)$ and $R^\dagger = \frac{1}{2}$. Formal proof in Section A.13. □

In general, the risk-dominance threshold is not the Walrasian equilibrium price.[18]

6.3 Payoff Dominance: The Limits of Unanimity

Risk dominance selects at every reference, but it is not the only classical criterion. Its companion in Harsanyi and Selten (1988) is payoff dominance: where one equilibrium is better for everyone, a selection should not have to appeal to strategic uncertainty at all—unanimity speaks for itself. This subsection asks when the traders are unanimous.

Definition 4 (Payoff dominance)**.** In the role-assignment game, buyer control *strictly payoff-dominates* seller control if it yields both sides a strictly higher expected payoff,

$$\Pi_B^B > \Pi_B^S \qquad \text{and} \qquad \Pi_S^B > \Pi_S^S;$$

seller control strictly payoff-dominates buyer control if both inequalities are reversed.

[18]The *Walrasian price* p^W equates the mass of buyer types who gain from trading at it to the mass of seller types who do, $1 - F(p^W) = G(p^W)$. Under reflection symmetry both $R^\dagger$ and p^W equal $\frac{1}{2}$, but the coincidence is an artifact of symmetry: for F uniform and $G(c) = c + \frac{1}{2}c(1-c)$, the market clears at $p^W = (5 - \sqrt{17})/2 \approx 0.438$, while $R^\dagger \approx 0.473$. The switch of pricing power is located by deviation losses and fallbacks, not by market clearing, and the two calculations generically disagree.

At extreme references, there exists a payoff dominant equilibrium:

Theorem 5 (Payoff dominance at the extremes)**.** *For every regular pair (F, G) there exist $0 < \underline{r} < \bar{r} < 1$ such that seller control strictly payoff-dominates for all $R < \underline{r}$, and buyer control strictly payoff-dominates for all $R > \bar{r}$. In each case the payoff-dominant regime is also the welfare-superior one.*

Proof outline. As $R \downarrow 0$, the frontier cap $q(R) \to 0$, so under buyer control the probability of trade and both sides' payoffs vanish, while under seller control the floor rests at its unconstrained optimum and both sides' payoffs are bounded away from zero. Both sides therefore strictly prefer seller control, which also generates strictly higher gains from trade, their sum. The mirror applies as $R \uparrow 1$. Formal proof in Section A.14. □

In the interior, unanimity reaches its limit. The welfare of buyer control rises with the reference and that of seller control falls, so the two cross at a balanced reference, where the sides' regime preferences are driven entirely by the maker–taker gap and generically point in opposite directions; by continuity, they do so on an open *conflict region* around it. There, payoff dominance selects nothing, and the threshold of Theorem 4 decides alone. The two criteria are thus mapped: risk dominance points to sellers below the threshold and buyers above it, everywhere; payoff dominance agrees at the extremes and is silent on the conflict region. What neither criterion reveals is what the roles they assign are worth.

6.4 Making the Price: Privilege or Burden?

Is making prices desirable? The presumption behind the phrase "price making" says yes — the maker solves an unconstrained monopoly/monopsony problem, the taker accepts/rejects. The presumption is wrong in general.

Comparing the Maker's and the Takers' Surplus. The comparison reduces to a pointwise calculation. Fix a bid b in the buyer regime and consider the buyer type $v = J_S(b)$ whose unconstrained monopsony bid is b. Her first-order condition equates her markdown to the reverse hazard, $v - b = G(b)/g(b)$, so her expected surplus, and the aggregate expected surplus of the sellers who take her price, are

$$\begin{aligned} M_B(b) = (v - b)\, G(b) \;&=\; \frac{G(b)^2}{g(b)}, \\ T_S(b) = \mathbb{E}\big[(b - C)^+\big] \;&=\; \int_0^b G(c)\, dc \;\equiv\; L_G(b), \end{aligned} \qquad \text{(Maker's / Takers' Surplus)}$$

where L_G is the integrated supply curve. Dividing both surpluses by the probability of trade $G(b)$ makes the economics transparent: the maker side outearns the taker side if and only if

$$\underbrace{\frac{G(b)}{g(b)}}_{\text{marginal rent}} \quad > \quad \underbrace{b - \mathbb{E}[\,C \mid C \le b\,]}_{\text{average inframarginal rent}} .$$

The maker's rent per transaction is her markdown, a marginal quantity; the takers' is the mean surplus of the inframarginal sellers below the bid, an average one. Pricing power is valuable where rents sit at the margin and costly where they sit in the inframarginal mass: if the supply density rises sharply just below the bid, shading down sacrifices too much trade, so the maker bids aggressively, compressing her own margin while the takers below her retain theirs. This comparison is a statement about the curvature of L_G. Since $L_G'' = g$, the integrated supply curve is twice differentiable, and the surplus difference is an exact identity in its log-curvature:

$$M_B(b) - T_S(b) \;=\; -\frac{L_G(b)^2}{g(b)} \left(\log L_G\right)''(b). \qquad \text{(Surplus Differential)}$$

The maker outearns the takers at b exactly where integrated supply is strictly log-concave. Say that the pair (F, G) satisfies *integrated log-concavity* if L_G and the integrated demand curve $L_F(s) = \mathbb{E}[(V - s)^+]$ are strictly log-concave on the relevant ranges:

$$\left(\log L_G\right)''(b) < 0 \text{ and } \left(\log L_F\right)''(s) < 0 \quad \forall b \in (0, \bar{b}^*],\ s \in [\underline{s}^*, 1) \qquad \text{(Integrated Log-Concavity)}$$

The condition is weaker than the familiar log-concavity of the primitives: a log-concave density implies a log-concave distribution and, in turn, a log-concave integral (Bagnoli and Bergstrom 2005). This condition holds for the standard families of distributions–uniform, normal, exponential, logistic.

The identity (Surplus Differential) turns the ranking of maker and taker surplus into a property of the curves L_F and L_G alone. If (F, G) is integrated log-concave, the maker's surplus exceeds the takers' at every quote, so making prices is a privilege on both sides of the market. When the condition fails, the ranking can reverse: there are regular markets in which the takers outearn the maker by an arbitrarily large factor, on an open interval of references around the risk-dominance threshold.

Theorem 6 (Price-making: privilege or burden)**.** *Let (F, G) be the distributions of values and costs.*

(i) If (F, G) satisfy integrated log-concavity , the price-making side strictly outearns the price-taking side in each regime, at every reference: $\Pi_B^B > \Pi_S^B$ and $\Pi_S^S > \Pi_B^S$.

(ii) For every $M > 1$ there is a smooth, strictly regular pair (F, G) — necessarily violating integrated log-concavity — and an open interval of references around $R^\dagger$ on which the price-taking side strictly outearns the price-making side in each regime by more than the factor M:

$$\Pi_S^B > M\,\Pi_B^B \qquad \text{and} \qquad \Pi_B^S > M\,\Pi_S^S.$$

Proof outline. For (i), the identity (Surplus Differential) delivers strict pointwise dominance at every unconstrained bid; integrating over maker's types, and verifying that the comparison is preserved for types pooled at a binding cap, yields the ranking within each regime.

For (ii), the construction exploits the classical link between curvature and pass-through (Bulow and Pfleiderer 1983, Weyl and Fabinger 2013): how much of a change in the maker's willingness to pay is passed into her price is governed by the slope of the virtual cost, and pass-through is exactly what separates the maker's margin from the takers' surplus. The rent function G/g is a free design object—any positive profile with slopes above -1 arises from a strictly regular distribution—so we build a supply schedule whose virtual cost rises with a small constant slope κ over the maker's active range, making pass-through $1/\kappa$. There, each unit of value moves the maker's optimal bid by $1/\kappa$ units: she bids more than one-for-one into a thickening mass of inframarginal sellers, so competition against her own book pins her markdown at the local rent level, of order δ, while the sellers she trades with retain the integrated surplus below her bid, of order δ/κ. Every trade thus splits $1/\kappa$ to the takers for each unit the maker keeps. Concentrating almost all seller mass at a high cost—and, by reflection, almost all buyer mass just above the maker's range—makes this trade the whole of both regimes' payoffs, so the aggregate ratio inherits the pointwise one: $\Pi_S^B \geq (1/\kappa)(1 - o(1))\,\Pi_B^B$, with the seller-regime mirror by reflection. Reflection symmetry places the threshold at the balanced reference $R^\dagger = \frac{1}{2}$, continuity of the payoff maps extends both reversals to an open interval around it, choosing $\kappa < 1/(2M)$ drives the ratio beyond M, and mollification delivers smooth densities preserving every strict inequality. Formal proof in Section A.15. □

The two parts differ in what the conflict is about. Under integrated log-concavity, disagreement is *mutual assertion*: each side prefers the regime in which it makes prices, and risk dominance awards a contested privilege. In the markets of part (ii), disagreement is *mutual deference*: each side prefers that the other make prices, and risk dominance conscripts. The conscript serves because refusing is worse—given that the opponent takes prices, making them strictly beats joint price-taking (Proposition 5)—so someone must set the price, and risk dominance selects the side for which failing to serve is riskier.

Example 3 (The uniform benchmark)**.** For $F = G = U[0,1]$, the unconstrained programs meet exactly at the center, $\bar{b}^* = \underline{s}^* = \frac{1}{2}$: at every reference other than the center, exactly one frontier is unconstrained. As functions of the frontier cap $q \leq \frac{1}{2}$, the payoffs are

$$\Pi_B^B(q) = \tfrac{q}{2} - q^2 + \tfrac{2q^3}{3}, \qquad \Pi_S^B(q) = \tfrac{q^2}{2} - \tfrac{2q^3}{3},$$

with $\Pi_S^S(t) = \Pi_B^B(1-t)$ and $\Pi_B^S(t) = \Pi_S^B(1-t)$ by reflection, for the frontier floor $t \geq \frac{1}{2}$.

- *Risk dominance.* Symmetry places the threshold at the center, $R^\dagger = \frac{1}{2}$ (Theorem 4); the index crosses zero with slope $\nabla'(\frac{1}{2}) = \frac{1}{64}$.

- *Payoff dominance.* The seller side prefers its own regime throughout the lower half, since $\Pi_S^B \leq \frac{1}{24} < \frac{1}{12} = \Pi_S^S(\frac{1}{2})$, so the binding question is the buyer's: capped at R, she prefers seller control once making prices yields less than the $\frac{1}{24}$ she earns taking them. Her indifference condition $\Pi_B^B(R) = \frac{1}{24}$ reduces to $2(2R-1)^3 = -1$, so the agreement boundaries are, in closed form and by symmetry,

  $$R_B = \tfrac{1}{2}\big(1 - 2^{-1/3}\big) \approx 0.103, \qquad R_S = \tfrac{1}{2}\big(1 + 2^{-1/3}\big) \approx 0.897 :$$

  payoff dominance applies to roughly a fifth of the unit interval; on the region between the boundaries, of width $2^{-1/3} \approx 0.79$, there is conflict.

- *Price-Making as a Privilege.* Integrated supply is $L_G(b) = b^2/2$, strictly log-concave, so Theorem 6(i) applies—and with an exact constant: $M_B(b) = b^2 = 2\,T_S(b)$ at every quote, aggregating to the $\frac{1}{12}$-versus-$\frac{1}{24}$ split of Example 2. The conflict is mutual assertion.

7 The Self-Referential Market

The analysis so far takes the reference as given. On the exchanges that motivate it, no one gives it: the reference is a previous market outcome. The closing auction's reference is the day's last transaction price; the opening auction falls back on the previous close whenever the pre-open book is too thin to supply a quote-based benchmark; and on venues for less liquid securities that trade *only* in auctions — the London Stock Exchange's SETSqx segment is the leading example — each auction is anchored to the price of the auction before it. The rule is recursive: the reference for today's trade is the price of yesterday's. Today's equilibrium therefore selects tomorrow's game, and this section studies the market that results.

The Dynamic Model. Trade occurs in periods $t = 1, 2, \ldots$, each period an auction of the form studied above, with the initial reference $R_1 \in (0, 1)$ given. In period t, a buyer and a seller arrive with types drawn independently from F and G, independently across periods; they play the auction at reference R_t; and the reference updates to the transaction price,

$$R_{t+1} = \begin{cases} p_t, & \text{if a trade executes at price } p_t, \\ R_t, & \text{if no trade occurs.} \end{cases}$$

Independence across periods reflects the population these institutions serve: successive auctions draw different traders from an anonymous flow, so no participant expects to face her own past play. Each period is therefore the one-shot game of the preceding sections, and the reference is the sole intertemporal link—the only way today's market touches tomorrow's is through the price it prints.

7.1 Dynamics under Risk-Dominant Equilibrium Selection

First, we assume the market resolves the coordination problem of Section 6 by risk dominance: at every reference $R_t \neq R^\dagger$, play is the risk-dominant frontier equilibrium of Theorem 4. We call this maintained hypothesis the *risk-dominant equilibrium selection* and $R^\dagger$ the *risk-dominance threshold*. Two conventions close the model: at $R_t = R^\dagger$, where neither regime is risk-dominant, control is assigned by an arbitrary fixed tie rule—an event that will be seen to occur with probability zero along the path—and at references where the selected side's frontier is constrained, the market plays an ε-frontier equilibrium for any sufficiently small ε. Every result below holds for all such conventions.

Under the threshold selection, the recursion acquires a direction. Buyer control is selected only above the risk-dominance threshold, and every price it prints is a bid, strictly below the reference that selected it; seller control is selected only below the threshold, and prints only asks strictly above it. Each side, when it prices, drags the reference toward—and with positive probability across—the threshold. The first theorem formalises this dynamic.

Theorem 7 (Alternating Price-Making)**.** *Under the risk-dominant equilibrium selection, for every regular pair (F, G) and every initial reference, almost surely:*

(i) trades occur infinitely often;

(ii) every transaction prints strictly below the current reference under buyer control and strictly above it under seller control; and

(iii) price-making changes sides infinitely often, every run of consecutive trades under one side being stochastically dominated by a geometric distribution whose parameter depends on (F, G).

Proof outline. Prints lie strictly below (above) the reference because the selected side bids below (asks above) it. Trades recur because the selection keeps the selected side viable: buyer control is selected only at $R_t > R^\dagger$, so its cap is bounded away from zero uniformly over references and conventions, and with it the trade probability; the seller mirror is identical. Control flips because the separating types quote independently of the frontier bound: under buyer control, every type $v < J_S(R^\dagger)$ bids $b^M(v) < R^\dagger$ at any cap, so each trade hands control to the sellers with probability at least $a_B = \int_0^{J_S(R^\dagger)} G\big(b^M(v)\big)\, dF(v) > 0$; with the mirror bound a_S, runs are dominated by a geometric distribution with parameter $\min\{a_B, a_S\}$. Formal proof in Section A.16. □

The reference price rule written to provide price continuity thus manufactures perpetual regime change. Pricing power under self-reference is self-undermining: exercising it moves the very benchmark on which it rests, so no side holds price discovery for more than a few trades in expectation, from any starting point and in every regular market.

When, in addition, the risk-dominance threshold separates the two pricing programs, the rotation becomes clockwork: every buyer print lands strictly below $\bar{b}^* \le R^\dagger$, every seller print strictly above $\underline{s}^* \ge R^\dagger$, and every print dethrones its maker. Write $\lambda_B, \lambda_S \in (0, 1)$ for the trade probabilities and μ_B, μ_S for the transaction-price laws of the two unconstrained regimes.

Corollary 4 (Reversal at every trade)**.** *Suppose $\bar{b}^* \le R^\dagger \le \underline{s}^*$. Then, almost surely, after the first trade:*

(i) the selected side is unconstrained at every reference the process visits, and play in each regime coincides with the corresponding one-sided auction (Corollary 3);

(ii) control changes sides at every trade, and the price never lies in $[\bar{b}^*, \underline{s}^*]$;

(iii) the buyer side makes the price a long-run fraction $\lambda_S/(\lambda_B+\lambda_S)$ of periods, the empirical laws of transaction prices and of the reference converge to $\frac{1}{2}\mu_B + \frac{1}{2}\mu_S$ and $(\lambda_B\mu_B + \lambda_S\mu_S)/(\lambda_B+\lambda_S)$, and long-run average welfare per period is

$$\overline{W}_\infty = \frac{\lambda_S\, W_B(\bar{b}^*) + \lambda_B\, W_S(\underline{s}^*)}{\lambda_B + \lambda_S},$$

a strictly interior convex combination of the two frontier welfares;

1. for reflection-symmetric (F, G), $\lambda_B = \lambda_S$ and $\overline{W}^\infty = \frac{1}{2}\big(W_B(\bar{b}^*) + W_S(\underline{s}^*)\big)$, the expected gains from trade of the random-offerer mechanism; in particular, $\overline{W}^\infty \ge \frac{1}{\pi}\, W^{\mathrm{FB}}$.

Proof outline. Buyer control is selected only at $R > R^\dagger \ge \bar{b}^*$, where the cap is the unconstrained $\bar{b}^*$; the seller mirror gives (i), and with it strategies, trade probabilities, and price laws that never depend on the reference. The buyer-control price law is atomless with support in $(0, \bar{b}^*)$, so every buyer print falls strictly below the threshold and flips control; the mirror completes (ii), and atomlessness makes the tie rule irrelevant. By (i), the process is an alternating renewal process—a print by one side is held as the reference for a geometric number of periods while the other side waits to trade—and renewal-reward theory delivers (iii), with interiority from $\lambda_B, \lambda_S \in (0, 1)$. For (iv), reflection symmetry equates λ_B with λ_S, the average is the welfare of the random-offerer mechanism by (Frontier Welfares). Formal proof in Section A.17. □

All of this, however, presumes that the market re-coordinates each period on the risk-dominant regime. The next subsection asks what happens when it does not.

7.2 Dynamics under Fixed Equilibrium Selection

The alternation of pricing power is built into the threshold selection; a natural question is whether it is also an artifact of it. Suppose instead the market settles on a single side. Call a

buyer-side convention any rule—possibly depending on time and on history—that selects, at every reference, an undominated buyer-control equilibrium; a *seller-side convention* mirrors.

The convention is otherwise unrestricted: it may change the cap every period, react to the path of prices, or hold still. Note that below the risk-dominance threshold, a buyer-side convention plays the risk-dominated equilibrium, so the following theorem evaluates a convention, not a refinement.

Theorem 8 (Long-Run Market Failure)**.** *Under every buyer-side convention, almost surely:*

(i) the reference price is nonincreasing, and strictly decreasing at every trade;

(ii) either only finitely many trades occur and the reference freezes at a positive level, or $R_t \to 0$; under every stationary convention—the cap depending only on the current reference—the second branch obtains;

(iii) long-run average welfare per period is zero.

The mirror holds for every seller-side convention, with $R_t \to 1$.

Proof outline. A permanent price-maker moves the reference monotonically toward her own side: every buyer-made price lies strictly below the prevailing reference, and each falling reference caps the next price lower still. The descent cannot stall at a positive level: conditional on a trade at any reference above δ, the price falls below δ with probability at least $\lambda(\delta) = \Pr\big(V < J_S(\delta),\, C \le b^M(V)\big) > 0$, because a low-value buyer separates at a bid below δ whatever the cap. Trades can cease only if the convention shrinks its caps summably fast; a stationary convention cannot, since at a frozen reference the trade probability is a fixed positive number. Welfare vanishes with the reference, since per-period gains are at most $G(R_t)$. Formal proof in Section A.18. □

The convention consumes itself. Each period the price-maker prices optimally against the current market; each optimal price moves the reference toward her own side; and the accumulating references shrink the market she prices in. The short-run optimal exercise of pricing power is, in the long run, the elimination of its object. And the fatal ingredient is the combination: with a fixed reference, a convention is harmless—the market repeats one one-sided equilibrium and earns its welfare every period; under the threshold selection, self-reference is harmless—the market alternates and earns $\overline{W}_\infty$. Only self-reference joined to a frozen role destroys the market. Under a self-referential reference, the market alternates, or it fails.

Example 4 (The uniform benchmark)**.** For $F = G = U[0,1]$, we can solve both dynamics explicitly:

- *Risk-dominant equilibrium selection:* The two regimes trade with equal probability $\lambda_B = \lambda_S = \frac{1}{4}$, and the clockwork of Corollary 4 runs at its simplest: transaction prices alternate between the laws with densities $8p$ on $(0, \frac{1}{2})$ and $8(1-p)$ on $(\frac{1}{2}, 1)$, the reference's long-run law is the symmetric tent $4p \wedge 4(1-p)$ with $\mathbb{E}|R - \frac{1}{2}| = \frac{1}{6}$, and by part (iv) the market trades every fourth period in perpetuity and earns $\overline{W}_\infty = \frac{1}{8}$, the common frontier welfare. The cost is perpetual price movement: consecutive prices lie on opposite sides of the center, with $\mathrm{Var}(p) = \frac{1}{24}$ and $\mathbb{E}|p_{t+1} - p_t| = \frac{1}{3}$.

- *Buyer-side convention:* With frontier caps, Theorem 8 collapses the same market, at an explicit rate. Separating trades arrive with probability R^2, each scaling the reference by a factor of mean logarithm $-\frac{1}{2}$; the resulting drift approximation $\dot{R} = -R^3/2$ gives

$$R_t \;\approx\; \left(R_0^{-2} + t\right)^{-1/2},$$

  and simulated paths track this $t^{-1/2}$ law over five orders of magnitude.

8 Concluding Remarks

Call auctions break ties by price proximity to a reference, and this paper has argued that the clause is not a technicality but fundamentally alters incentives and welfare. In the canonical model of bilateral trade, the reference splits every undominated equilibrium into a side that makes prices and a side that takes them, and a unique risk-dominance threshold determines which side is which. When the reference is updated from past prices, as the rulebooks prescribe, pricing power becomes transient: the market alternates makers and attains, without any designer, the welfare guarantee of the random-offerer mechanism. A market that instead freezes the pricing role ratchets its own reference into the tail and extinguishes trade.

The model deliberately isolates one force, and its idealizations mark the openings for theory. Our traders live for one period; a long-lived trader would price today partly to position tomorrow's reference, and the transience result gives that problem an unusual twist—since every print dethrones its maker, a patient price-maker may trade *less* to stay in power, turning the crossed occupancy of the renewal dynamics into a strategic choice and connecting the model to the manipulation of benchmarks. Our primitives are fixed across periods; if

the reference instead shifts the arriving distributions—through participation, anchoring, or information—the feedback loop closes twice over, and the reference becomes an instrument that shapes the very fundamentals it summarizes.

The equilibrium benchmark is bilateral and exact, and its sharpest statements—no prints at the reference, reversal at every trade—should be read as directions rather than point predictions. But the forces behind them are observable in principle, and the natural laboratory exists: auction-only segments for illiquid securities, where the reference recursion is the entire price process and the thinnest instances approach the bilateral case. Whether prints thin out near the reference, whether the aggressive side of the book depends on the reference price, and whether the less active side's prices anchor longer are questions the data can answer, and little is known about strategic behavior in these smallest auctions.

References

Moshe Babaioff, Shahar Dobzinski, and Ron Kupfer. 2021. A Note on the Gains from Trade of the Random-Offerer Mechanism. (2021). https://doi.org/10.48550/arXiv.2111.07790 arXiv:2111.07790 [cs.GT]

Mark Bagnoli and Ted Bergstrom. 2005. Log-concave probability and its applications. *Economic theory* 26, 2 (2005), 445–469.

Bruno Biais, Pierre Hillion, and Chester Spatt. 1999. Price Discovery and Learning during the Preopening Period in the Paris Bourse. *Journal of Political Economy* 107, 6 (1999), 1218–1248. https://doi.org/10.1086/250095

Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. 2017. Approximating Gains from Trade in Two-Sided Markets via Simple Mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. Association for Computing Machinery, 589–590. https://doi.org/10.1145/3033274.3085148

Eric Budish, Peter Cramton, Albert S. Kyle, Jeongmin Lee, and David Malec. 2023. *Flow Trading*. Working Paper 31098. National Bureau of Economic Research.

Eric Budish, Peter Cramton, and John Shim. 2015. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics* 130, 4 (2015), 1547–1621. https://doi.org/10.1093/qje/qjv027

Jeremy I. Bulow and Paul Pfleiderer. 1983. A Note on the Effect of Cost Changes on Prices. *Journal of Political Economy* 91, 1 (1983), 182–185. https://doi.org/10.1086/261135

Yang Cai and Jinzhao Wu. 2023. On the Optimal Fixed-Price Mechanism in Bilateral Trade. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. Association for Computing Machinery, 737–750. https://doi.org/10.1145/3564246.3585171

Hans Carlsson and Eric Van Damme. 1993. Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society* (1993), 989–1018.

Kalyan Chatterjee and William Samuelson. 1983. Bargaining under Incomplete Information. *Operations Research* 31, 5 (1983), 835–851. https://doi.org/10.1287/opre.31.5.835

Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. 2017. Fixed Price Approximability of the Optimal Gain from Trade. In *Web and Internet Economics: 13th International Conference, WINE 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10660)*. Springer, 146–160. https://doi.org/10.1007/978-3-319-71924-5_11

Carole Comerton-Forde and James Rydge. 2006a. Call Auction Algorithm Design and Market Manipulation. *Journal of Multinational Financial Management* 16, 2 (2006), 184–198. https://doi.org/10.1016/j.mulfin.2005.06.002

Carole Comerton-Forde and James Rydge. 2006b. The Influence of Call Auction Algorithm Rules on Market Efficiency. *Journal of Financial Markets* 9, 2 (2006), 199–222. https://doi.org/10.1016/j.finmar.2006.02.001

Laurence Daures, Sophie Moinas, and Selma Boussetta. 2025. Click First or Last? Strategic Order Submission During the Euronext Preopening Session. *Management Science* 72, 3 (2025), 2656–2679. https://doi.org/10.1287/mnsc.2023.03998

Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. 2022. Approximately Efficient Bilateral Trade. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. Association for Computing Machinery, 718–721. https://doi.org/10.1145/3519935.3520054

Songzi Du and Haoxiang Zhu. 2017. What Is the Optimal Trading Frequency in Financial Markets? *The Review of Economic Studies* 84, 4 (2017), 1606–1651. https://doi.org/10.1093/restud/rdx006

Yumou Fei. 2022. Improved Approximation to First-Best Gains-from-Trade. In *Web and Internet Economics: 18th International Conference, WINE 2022, Proceedings (Lecture Notes in Computer Science, Vol. 13778)*. Springer, 204–218. https://doi.org/10.1007/978-3-031-22832-2_12

Rodney Garratt and Marek Pycia. 2026. Efficient Bilateral Trade. (2026). Manuscript, June 2026.

Thomas A. Gresik. 1991. The Efficiency of Linear Equilibria of Sealed-Bid Double Auctions. *Journal of Economic Theory* 53, 1 (1991), 173–184. https://doi.org/10.1016/0022-0531(91)90147-V

Kathleen M. Hagerty and William P. Rogerson. 1987. Robust Trading Mechanisms. *Journal of Economic Theory* 42, 1 (1987), 94–107. https://doi.org/10.1016/0022-0531(87)90104-9

John C Harsanyi and Reinhard Selten. 1988. A general theory of equilibrium selection in games. *MIT Press Books* 1 (1988).

Simon Jantschgi. 2024. Transaction Cost (In)Transparency in Repeated Call Auctions. (2024). https://doi.org/10.2139/ssrn.4861066

Simon Jantschgi, Heinrich H. Nax, Bary S. R. Pradelski, and Marek Pycia. 2022. *Double Auctions and Transaction Costs*. Working Paper 405. University of Zurich, Department of Economics. https://doi.org/10.2139/ssrn.4043576

Simon Jantschgi, Heinrich H. Nax, Bary S. R. Pradelski, and Marek Pycia. 2024. *Foundations of Double Auctions from Theory and Practice*. Working Paper 404. University of Zurich, Department of Economics. https://doi.org/10.2139/ssrn.4043532

Simon Jantschgi, Heinrich H. Nax, Bary S. R. Pradelski, and Marek Pycia. 2026. Minimizing Volatility: Optimal Adjustment with Evolving Feasibility Constraints. (2026). https://doi.org/10.48550/arXiv.2602.15686 arXiv:2602.15686 [econ.TH]

Sunghyeon Jo. 2026. Tighter Bounds for the Random-Offerer Mechanism in Bilateral Trade. (2026). https://doi.org/10.48550/arXiv.2607.13959 arXiv:2607.13959 [cs.GT]

Ohad Kadan. 2007. Equilibrium in the Two-Player, k-Double Auction with Affiliated Private Values. *Journal of Economic Theory* 135, 1 (2007), 495–513. https://doi.org/10.1016/j.jet.2006.06.004

Michihiro Kandori, George J Mailath, and Rafael Rob. 1993. Learning, mutation, and long run equilibria in games. *Econometrica: Journal of the Econometric Society* (1993), 29–56.

Zi Yang Kang, Francisco Pernice, and Jan Vondrák. 2022. Fixed-Price Approximations in Bilateral Trade. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2964–2985. https://doi.org/10.1137/1.9781611977073.115

Albert S. Kyle. 1989. Informed Speculation with Imperfect Competition. *The Review of Economic Studies* 56, 3 (1989), 317–355. https://doi.org/10.2307/2297551

Wolfgang Leininger, Peter B. Linhart, and Roy Radner. 1989. Equilibria of the Sealed-Bid Mechanism for Bargaining with Incomplete Information. *Journal of Economic Theory* 48, 1 (1989), 63–106. https://doi.org/10.1016/0022-0531(89)90120-8

Simon Loertscher and Leslie M. Marx. 2022. Bilateral Trade with Multiunit Demand and Supply. *Management Science* 69, 2 (2022), 1146–1165. https://doi.org/10.1287/mnsc.2022.4399

Simon Loertscher and Ellen V. Muir. 2025. The Matching Benefits of Market Thickness. *Games and Economic Behavior* 153 (2025), 42–66. https://doi.org/10.1016/j.geb.2025.05.010

Simon Loertscher, Ellen V. Muir, and Peter G. Taylor. 2022. Optimal Market Thickness. *Journal of Economic Theory* 200 (2022), 105383. https://doi.org/10.1016/j.jet.2021.105383

London Stock Exchange. 2026. Guide to the Trading System, Issue 15.6. https://docs.londonstockexchange.com/sites/default/files/documents/mit201-guide-to-the-trading-system-15_6_20240429.pdf Last accessed on: August 13, 2026.

Ananth Madhavan and Venkatesh Panchapagesan. 2000. Price Discovery in Auction Markets: A Look Inside the Black Box. *The Review of Financial Studies* 13, 3 (2000), 627–658. https://doi.org/10.1093/rfs/13.3.627

Roger B. Myerson and Mark A. Satterthwaite. 1983. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory* 29, 2 (1983), 265–281. https://doi.org/10.1016/0022-0531(83)90048-0

Nasdaq. 2026. Options Trading Rules, Section 8. Opening and Halt Cross. https://listingcenter.nasdaq.com/RuleBook/Nasdaq/rules/Nasdaq%20Options%203 Last accessed on: August 13, 2026.

Jr. Nash, John F. 1950. The Bargaining Problem. *Econometrica* 18, 2 (1950), 155–162. https://doi.org/10.2307/1907266

New York Stock Exchange. 2026a. Behind the Record Volumes: A Hidden Opportunity. https://www.nyse.com/research/insights/behind-the-record-volumes-a-hidden-opportunity NYSE Research Insights, April 2026. Last accessed on: August 13, 2026.

New York Stock Exchange. 2026b. Rulebook. https://nyseguide.srorules.com/rules Last accessed on: August 13, 2026.

Seongkyu Park, Wing Suen, and Kam-Ming Wan. 2022. Call Auction Design and Closing Price Manipulation: Evidence from the Hong Kong Stock Exchange. *Journal of Financial Markets* 58 (2022), 100700. https://doi.org/10.1016/j.finmar.2021.100700

Marzena Rostek and Marek Weretka. 2012. Price Inference in Small Markets. *Econometrica* 80, 2 (2012), 687–711. https://doi.org/10.3982/ECTA9573

Marzena Rostek and Marek Weretka. 2015. Dynamic Thin Markets. *The Review of Financial Studies* 28, 10 (2015), 2946–2992. https://doi.org/10.1093/rfs/hhv027

Ariel Rubinstein. 1982. Perfect Equilibrium in a Bargaining Model. *Econometrica* 50, 1 (1982), 97–109. https://doi.org/10.2307/1912531

Mark A. Satterthwaite and Steven R. Williams. 1989. Bilateral Trade with the Sealed Bid k-Double Auction: Existence and Efficiency. *Journal of Economic Theory* 48, 1 (1989), 107–133. https://doi.org/10.1016/0022-0531(89)90121-X

Tokio Stock Exchange. 2026. Trading Methods - Itayose Method. https://www.jpx.co.jp/english/derivatives/rules/trading-methods/index.html Last accessed on: August 13, 2026.

Dimitri Vayanos. 1999. Strategic Trading and Welfare in a Dynamic Market. *The Review of Economic Studies* 66, 2 (1999), 219–254. https://doi.org/10.1111/1467-937X.00086

E. Glen Weyl and Michal Fabinger. 2013. Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy* 121, 3 (2013), 528–583. https://doi.org/10.1086/670401

Steven R. Williams. 1987. Efficient Performance in Two Agent Bargaining. *Journal of Economic Theory* 41, 1 (1987), 154–172. https://doi.org/10.1016/0022-0531(87)90010-X

Steven R Williams. 1991. Existence and convergence of equilibria in the buyer's bid double auction. *The Review of Economic Studies* 58, 2 (1991), 351–374.

Alexander Wolitzky. 2016. Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade. *Theoretical Economics* 11, 3 (2016), 971–1004. https://doi.org/10.3982/TE2089

Ryuichi Yamamoto, Xin Fang, and Yukihiko Funaki. 2026. Opening Call Auction Designs and Price Manipulation for Price Discovery: An Experimental Study. (2026). Working paper, January 2026.

H Peyton Young. 1993. The evolution of conventions. *Econometrica: Journal of the Econometric Society* (1993), 57–84.

A Proofs

A.1 Proof of Proposition 1

Proof. Fix Borel probability measures Σ_B, Σ_S on $\mathcal{A} = [0,1]$; no assumption beyond measurability is used. Write $D_S(b) = \Sigma_S([0,b])$, a nondecreasing, right-continuous function with left limits $D_S(b-) = \Sigma_S([0,b))$. We prove the buyer statements (Buyer Optimization); the seller statements follow by the same arguments with the roles of the two half-lines exchanged, and we indicate the required modifications at the end.

Step 1: Payoff formulas. If $b < R$, then every $s \in [0,b]$ satisfies $s \le b < R$, so $\varphi_R(s,b) = b$ by (Reference-price rule), and

$$U_B(b;v) = (v-b)\, D_S(b), \qquad b \in [0,R). \tag{1}$$

If $b \ge R$, decompose $[0,b] = [0,R] \cup (R,b]$ (the second set empty when $b = R$). For $s \in [0,R]$ we have $s \le R \le b$, so $\varphi_R(s,b) = R$; for $s \in (R,b]$ we have $R < s \le b$, so $\varphi_R(s,b) = s$. Hence

$$U_B(b;v) = (v-R)\, D_S(R) + \int_{(R,b]} (v-s)\, \Sigma_S(ds), \qquad b \in [R,1], \tag{2}$$

and consequently, for $R \le b \le b'$,

$$U_B(b';v) - U_B(b;v) = \int_{(b,b']} (v-s)\, \Sigma_S(ds). \tag{3}$$

By (1) and the definition (Price-making values), $\sup_{b \in [0,R)} U_B(b;v) = M_B(v)$.

Step 2: The supremum over $[R,1]$. Suppose first $v \ge R$. For $b \in [R,v]$, the integrand in (3) satisfies $v - s \ge 0$ on $(b,v]$, so $U_B(v;v) \ge U_B(b;v)$; for $b \in (v,1]$, it satisfies $v - s < 0$ on $(v,b]$, so $U_B(v;v) \ge U_B(b;v)$ again. Hence $b = v$ maximizes U_B over $[R,1]$, with value $U_B(v;v) = (v-R)D_S(R) + \int_{(R,v]}(v-s)\,\Sigma_S(ds) = T_B(v)$ by (2) and (Price-taking values). Suppose instead $v < R$. Then for every $b \in [R,1]$, the first term of (2) is nonpositive and the integrand satisfies $v - s \le v - R < 0$ on $(R,b]$, so $U_B(b;v) \le 0$.

Step 3a: The case $v < R$. By Step 2, $U_B \le 0$ on $[R,1]$, while $U_B(0;v) = v\, D_S(0) \ge 0$ by (1). Hence $\sup_{b \in \mathcal{A}} U_B(b;v) = \sup_{b \in [0,R)} U_B(b;v) = M_B(v)$, which proves the value formula. For attainment, write $h(b) = (v-b)D_S(b)$ on $[0,R)$. For $b \in (v,R)$ we have $h(b) \le 0 \le$

$h(0)$, so $M_B(v) = \sup_{b\in[0,v]} h(b)$. On $[0, v]$ the factor $v - b$ is nonnegative, so h is upper semicontinuous there: at any $b_0 \in [0, v]$, right-continuity of D_S gives $\lim_{b\downarrow b_0} h(b) = h(b_0)$, and $\limsup_{b\uparrow b_0} h(b) = (v - b_0)D_S(b_0-) \le (v - b_0)D_S(b_0) = h(b_0)$. An upper semicontinuous function on the compact set $[0, v]$ attains its maximum, at some $b^* \in [0, v] \subseteq [0, R)$; thus $M_B(v)$ is attained and b^* is a best response. (Upper semicontinuity can fail on (v, R): at an atom of Σ_S the factor $v - b$ is negative and h jumps downward. The restriction to $[0, v]$ is therefore essential.)

Step 3b: The case $v \ge R$. Combining Step 2 with $\sup_{[0,R)} U_B = M_B(v)$ yields $\sup_{b\in\mathcal{A}} U_B(b; v) = \max\{M_B(v), T_B(v)\}$, which proves the value formula; it remains to show the supremum is attained. If $T_B(v) \ge M_B(v)$, the bid $b = v$ attains it by Step 2. Otherwise $M_B(v) > T_B(v)$, and we claim $M_B(v)$ is attained on $[0, R)$. Suppose not, and pick $b_n \in [0, R)$ with $h(b_n) \to M_B(v)$; by compactness of $[0, R]$, pass to a subsequence with $b_n \to b_\infty \in [0, R]$. If $b_\infty < R$: along any further subsequence approaching b_∞ from the right, right-continuity of D_S gives $h(b_n) \to h(b_\infty)$; along any approaching from the left or equal to b_∞, $\limsup h(b_n) \le (v - b_\infty)D_S(b_\infty-) \le (v - b_\infty)D_S(b_\infty) = h(b_\infty)$, using $v \ge R > b_\infty$. In either case $h(b_\infty) = M_B(v)$, contradicting non-attainment. Hence $b_\infty = R$, and

$$M_B(v) = \lim_n h(b_n) \le \lim_n (v - b_n)\, D_S(R-) = (v - R)\, D_S(R-) \le (v - R)\, D_S(R) \le T_B(v),$$

where the first inequality uses $D_S(b_n) \le D_S(R-)$ and $v - b_n \ge 0$, and the last uses that the integral in (Price-taking values) is nonnegative for $v \ge R$. This contradicts $M_B(v) > T_B(v)$, so $M_B(v)$ is attained, and the maximizer is a best response.

Step 4: The seller. The seller statements follow by the symmetric argument, with the ask in place of the bid, the tail function $\bar{D}_B(s) = \Sigma_B([s, 1])$ in place of D_S, and the two half-lines exchanged: asks strictly above R price at the ask, so $\sup_{s\in(R,1]} U_S(s; c) = M_S(c)$; asks at or below R affect only the execution set, so the truthful ask is optimal there, with value $T_S(c)$ for $c \le R$ and nonpositive payoff for every $c > R$; and attainment follows as in Steps 3 and 4, noting that $\bar{D}_B$ is left-continuous with right limits, so the roles of the one-sided limits in the semicontinuity arguments are exchanged. □

A.2 Proof of Proposition 2

Proof. Fix $R \in (0,1)$ and $p \in (0,1)$. Since F and G are atomless, the profile induces the quote distributions

$$\Sigma_B = F(p)\,\delta_0 + \big(1 - F(p)\big)\,\delta_p, \qquad \Sigma_S = G(p)\,\delta_p + \big(1 - G(p)\big)\,\delta_1,$$

and because $f, g > 0$ on $[0,1]$, all four masses $F(p)$, $1 - F(p)$, $G(p)$, $1 - G(p)$ are strictly positive. The profile is pure, so the equilibrium condition (Bayes–Nash equilibrium) reduces to pointwise optimality of the prescribed quote for every type against these distributions, which we now verify for the buyer. The case for the seller is analogous.

Buyer, $v \geq p$. The prescribed bid $b = p$ executes exactly against the ask atom at p (the atom at 1 lies outside $[0,p]$ since $p < 1$); the clearing interval is the singleton $\{p\}$, so $\varphi_R(p,p) = p$ for every R, and $U_B(p;v) = G(p)\,(v - p) \geq 0$. We compare all alternative bids. (a) For $b \in [0,p)$: the set $[0,p)$ carries no ask, so $U_B(b;v) = 0 \leq v - p$. (b) For $b \in (p,1)$: the execution set is still only the ask atom at p, with clearing interval $[p,b]$ and price

$$\varphi_R(p,b) = \begin{cases} p, & R < p, \\ R, & p \leq R \leq b, \\ b, & b < R, \end{cases} \qquad \text{i.e.} \qquad \varphi_R(p,b) = \begin{cases} p, & R < p, \\ \min\{R,b\} \;\geq p, & R \geq p. \end{cases}$$

Hence $U_B(b;v) = G(p)\big(v - \varphi_R(p,b)\big) \leq G(p)(v-p) = U_B(p;v)$. (c) For $b = 1$: the bid additionally reaches the ask atom at 1, where $\varphi_R(1,1) = 1$ and the increment is $\big(1 - G(p)\big)(v-1) \leq 0$, while the trade against the atom at p is priced as in (b) with $b = 1$. Hence $U_B(1;v) \leq G(p)(v-p) = U_B(p;v)$. The bid p is therefore optimal for every $v \geq p$.

Buyer, $v < p$. The prescribed bid 0 earns $U_B(0;v) = 0$, since $\Sigma_S(\{0\}) = 0$. Bids $b \in (0,p)$ likewise earn 0 by (a). For $b \in [p,1]$, the trade against the ask atom at p occurs with probability $G(p) > 0$ at a price $\varphi_R(p,b) \geq p$ by the display in (b), contributing at most $G(p)(v-p) < 0$, and the additional trade at $b = 1$ contributes $\big(1 - G(p)\big)(v-1) < 0$. Hence $U_B(b;v) \leq G(p)(v-p) < 0$ for every $b \geq p$, and the bid 0 is optimal.

Trade at price p. Every executed trade in the profile pairs the bid p with the ask p and therefore clears at p.

□

A.3 Proof of Lemma 1

Proof. Recall the ex-post payoffs $u_B(b, s; v) = \mathbf{1}\{s \le b\}\big(v - \varphi_R(s, b)\big)$ and $u_S(s, b; c) = \mathbf{1}\{s \le b\}\big(\varphi_R(s, b) - c\big)$, and that a quote is weakly dominated for a type if some mixed action yields a weakly higher ex-post payoff against every opposing pure quote and a strictly higher one against at least one. All dominating actions exhibited below are pure. We prove the statement for the buyer. The statement by the seller can be proven in analogy.

Step 1: Price invariance above the reference. For all $b, b' \in [R, 1]$ and every $s \le \min\{b, b'\}$,

$$\varphi_R(s, b) = \varphi_R(s, b') = \begin{cases} R, & s \le R, \\ s, & s > R. \end{cases} \tag{4}$$

Indeed, if $s \le R$ then $s \le R \le b$ and $s \le R \le b'$, so by (Reference-price rule) both prices equal R ; if $s > R$ then $R < s \le b$ and $R < s \le b'$, so both prices equal s. A bid at or above the reference therefore determines only the execution set $\{s \le b\}$, never the price of an executed trade.

Step 2a: Buyer types $v > R$. Fix $v > R$ and $b \in [R, 1] \setminus \{v\}$; we show b is weakly dominated by the truthful bid v, comparing ex-post payoffs against every ask $s \in \mathcal{A}$.

- *Case $b < v$.* For $s \le b$, both bids trade, and by (4) at the same price, so the payoffs coincide. For $s \in (b, v]$, only the truthful bid trades; since $s > b \ge R$, the price is s by (4), and the payoff gain is $v - s \ge 0$, strictly positive for every $s \in (b, v)$, a nonempty interval. For $s > v$, neither bid trades. Hence $u_B(v, s; v) \ge u_B(b, s; v)$ for every s, strictly for $s \in (b, v)$.

- *Case $b > v$.* For $s \le v$, both bids trade at the same price by (4), and payoffs coincide. For $s \in (v, b]$, only the bid b trades; since $s > v > R$, the price is s and the retained payoff is $v - s < 0$, so dropping these trades gains $s - v > 0$, strictly, on the nonempty set $(v, b]$. For $s > b$, neither bid trades. Hence $u_B(v, s; v) \ge u_B(b, s; v)$ for every s, strictly for $s \in (v, b]$.

In both cases the truthful bid weakly dominates b.

Step 2b: Buyer types $v < R$. Fix $v < R$ and $b \in [R, 1]$; we show b is weakly dominated by the truthful bid v. Against any $s \leq v$, the truthful bid trades with clearing interval $[s, v] \subset [0, R)$, so the price is the bid itself, $\varphi_R(s, v) = v$, and the payoff is 0; against $s > v$ it does not trade. Thus $u_B(v, \cdot\,; v) \equiv 0$. The bid b, by contrast, trades against every $s \leq b$, at price $\varphi_R(s, b) \geq R > v$ by (4), for an ex-post payoff $v - \varphi_R(s, b) \leq v - R < 0$; against $s > b$ its payoff is 0. Hence $u_B(v, s; v) \geq u_B(b, s; v)$ for every s, strictly for every $s \in [0, b]$, a set that is nonempty since $b \geq R > 0$.

Step 3: Conclusion. By Steps 1 and 2, for every type with $v \neq R$ or $c \neq R$ (by analogous proof), every undominated quote in $[R, 1]$ (buyer) or $[0, R]$ (seller) equals the type's own valuation or cost; the exempted boundary types carry F- and G-measure zero since F and G are atomless. Finally, the quote R itself lies in the dominated set of Step 1 for every $v > R$ (as $R \neq v$), of Step 2 for every $v < R$, and of the mirror statements for every $c \neq R$; hence quoting exactly R is weakly dominated for every type except the null boundary type whose valuation or cost equals R. □

A.4 Proof of Theorem 1

Proof. Write $\sigma_- = \Sigma_S([0, R))$ and $\beta_+ = \Sigma_B((R, 1])$. Suppose, for contradiction, that $\sigma_- > 0$ and $\beta_+ > 0$.

Step 1: Structure of undominated quotes. The engine of the argument is a structural consequence of Lemma 1: beyond the reference, the quote distributions are not arbitrary measures but truthful images of the type distributions, and neither side quotes exactly R.

Claim 1. *In every undominated Bayes–Nash equilibrium, (a) $\Sigma_B(\{R\}) = \Sigma_S(\{R\}) = 0$; and (b) for every Borel $A \subseteq (R, 1]$ and $A' \subseteq [0, R)$,*

$$\Sigma_B(A) = \int_A \beta(\{v\} \mid v)\, dF(v), \qquad \Sigma_S(A') = \int_{A'} \sigma(\{c\} \mid c)\, dG(c);$$

In particular, the restrictions $\Sigma_B|_{(R,1]}$ and $\Sigma_S|_{[0,R)}$ are absolutely continuous with respect to F and G, hence atomless.

Proof of the Claim. By Lemma 1 and the definition of undominated equilibrium, for F-almost every v the support of $\beta(\cdot \mid v)$ contains no dominated quote: for $v > R$, its intersection with $[R, 1]$ is contained in $\{v\}$; for $v < R$, this intersection is empty. Hence, for

Borel $A \subseteq (R, 1]$,

$$\Sigma_B(A) = \int_0^1 \beta(A \mid v)\, dF(v) = \int_{\{v>R\}} \beta\big(A \cap \{v\} \mid v\big)\, dF(v) = \int_A \beta(\{v\} \mid v)\, dF(v),$$

the middle equality because types $v < R$ place no mass in $(R, 1]$ and the exceptional sets (the type $v = R$ and the F-null set of types permitted to play dominated quotes) do not affect the integral. Absolute continuity with respect to F is immediate from the last expression, and F atomless makes the restriction atomless. For (a), the quote R is weakly dominated for every type except the F-null type $v = R$ (Lemma 1), so $\Sigma_B(\{R\}) = \int \beta(\{R\} \mid v)\, dF(v) = 0$. The seller statements follow by the symmetric argument with σ, G, and $[0, R)$ in place of β, F, and $(R, 1]$. △

By the Claim, the restriction of Σ_S to $[0, R)$ is the truthful pushforward of the sellers with $c < R$ who play their costs, the restriction of Σ_B to $(R, 1]$ is the truthful pushforward of the buyers with $v > R$ who play their values, and both restrictions are atomless. Let

$$a^* = \sup \operatorname{supp}\big(\Sigma_S|_{[0,R)}\big) \le R, \qquad v^* = \inf \operatorname{supp}\big(\Sigma_B|_{(R,1]}\big) \ge R.$$

Then $v^* < 1$: otherwise the support of $\Sigma_B|_{(R,1]}$ would be the singleton $\{1\}$, making the positive mass β_+ an atom at 1 and contradicting atomlessness. By the definition of a^* and v^* as extremes of supports,

$$\Sigma_S\big((a^*, R)\big) = 0 \qquad \text{and} \qquad \Sigma_B\big((R, v^*)\big) = 0. \tag{5}$$

Throughout, we use that in equilibrium, for G-almost every c (resp. F-almost every v), every quote in the support of the prescribed mixture is a best response (Bayes–Nash equilibrium).

Step 2: The corridor above the reference carries no asks, $\Sigma_S\big((R, v^*)\big) = 0$. If $v^* = R$ there is nothing to prove, so let $v^* > R$ and fix any $s \in (R, v^*)$ and any type c. The ask s trades exactly against bids $b \ge s$; every such trade has $R < s \le b$, hence price $\varphi_R(s, b) = s$; and by (5) together with part (a) of the Claim, $\Sigma_B([s, 1]) = \Sigma_B([v^*, 1]) = \beta_+$. Hence $U_S(s; c) = (s - c)\,\beta_+$. The alternative ask $s' = \frac{s+v^*}{2} \in (s, v^*)$ earns, by the same computation, $U_S(s'; c) = (s' - c)\,\beta_+ = U_S(s; c) + (s' - s)\,\beta_+ > U_S(s; c)$. Thus no ask in (R, v^*) is a best response for any type; by (Bayes–Nash equilibrium), for G-almost every c the support of $\sigma(\cdot \mid c)$ is disjoint from the open set (R, v^*), so $\Sigma_S\big((R, v^*)\big) = 0$.

Step 3: The below-reference asks reach the reference, $a^* = R$. Suppose $a^* < R$. Since v^* lies in the support of $\Sigma_B|_{(R,1]}$, for every $n \in \mathbb{N}$,

$$\Sigma_B\Big(\big[v^*, v^* + \tfrac{1}{n}\big) \cap (R, 1]\Big) > 0,$$

and by part (b) of the Claim this mass is carried by a positive-F-measure set of types $v \in [v^*, v^* + \frac{1}{n}) \cap (R, 1]$ who play the truthful bid with positive probability.

By (Bayes–Nash equilibrium), we may pick such a type v_n for which the truthful bid is a best response. Its payoff, using formula (2) from the proof of Proposition 1, part (a) of Claim 1 ($\Sigma_S(\{R\}) = 0$), and Step 1 ($\Sigma_S((R, v^*)) = 0$), is

$$U_B(v_n; v_n) = (v_n - R)\,\sigma_- + \int_{[v^*,\, v_n]} (v_n - s)\,\Sigma_S(ds) \;\le\; (v_n - R)\,\sigma_- + (v_n - v^*),$$

bounding the integrand by $v_n - v^* \ge 0$ and the mass by one. The deviation to the bid $a^* < R$ trades against every ask in $[0, a^*]$, at price $\varphi_R(s, a^*) = a^*$ since $s \le a^* < R$; by (5), $\Sigma_S([0, a^*]) = \sigma_-$, so the deviation earns $(v_n - a^*)\,\sigma_-$. Best-response optimality of the truthful bid yields

$$(v_n - R)\,\sigma_- + (v_n - v^*) \;\ge\; (v_n - a^*)\,\sigma_-, \qquad \text{i.e.} \qquad v_n - v^* \;\ge\; (R - a^*)\,\sigma_-.$$

Letting $n \to \infty$ gives $v_n \to v^*$ and hence $0 \ge (R - a^*)\,\sigma_- > 0$, a contradiction. Therefore $a^* = R$.

Step 3: The marginal sellers jump the reference. Since $\Sigma_B\big((R + \frac{1}{n}, 1]\big) \uparrow \beta_+$ by continuity from below, we may fix $\bar{s} \in (R, 1)$ with $\Sigma_B([\bar{s}, 1]) \ge \beta_+/2$, and set $\kappa = (\bar{s} - R)\,\beta_+/2 > 0$. By Step 2, $a^* = R$, so $\Sigma_S\big((R - \kappa, R)\big) > 0$ (every neighborhood of the support's supremum carries mass), and by part (b) of the Claim this mass comes from a positive-G-measure set of types $c \in (R - \kappa, R)$ playing the truthful ask with positive probability; by (Bayes–Nash equilibrium), for almost every such type the truthful ask is a best response. Fix one such c. Every trade of the ask c pairs it with a bid $b \ge c$ and executes at price $\varphi_R(c, b) = \min\{R, b\} \in [c, R]$, so

$$U_S(c; c) \;\le\; (R - c) \cdot 1 \;<\; \kappa.$$

The deviation to the ask $\bar{s}$ trades against every bid $b \geq \bar{s}$, each priced at $\varphi_R(\bar{s}, b) = \bar{s}$ since $R < \bar{s} \leq b$, and earns

$$U_S(\bar{s}; c) = (\bar{s} - c)\,\Sigma_B([\bar{s}, 1]) \;\geq\; (\bar{s} - c)\,\frac{\beta_+}{2} \;>\; (\bar{s} - R)\,\frac{\beta_+}{2} = \kappa,$$

using $c < R$. Hence the truthful ask is not a best response for this type—contradicting its selection. This contradiction establishes $\sigma_- = 0$ or $\beta_+ = 0$.

Step 5: Conclusion. Suppose $\beta_+ = 0$. With part (a) of Claim 1, all bids lie in $[0, R)$ up to a Σ_B-null set, so almost surely every executed trade satisfies $s \leq b < R$ and prices at the bid, $\varphi_R(s, b) = b < R$. If instead $\sigma_- = 0$, then all asks lie in $(R, 1]$ almost surely, every executed trade satisfies $R < s \leq b$ and prices at the ask, $\varphi_R(s, b) = s > R$. (If both masses vanish, trade occurs with probability zero and both descriptions hold vacuously.) In every case, the transaction price equals R with probability zero, which is also direct: a trade at R requires $s \leq R \leq b$, an event of probability $\Sigma_S([0, R]) \cdot \Sigma_B([R, 1]) = \sigma_- \cdot \beta_+ = 0$ by independence and part (a) of Claim 1. □

A.5 Proof of Theorem 2

Proof of Theorem 2. Consider an undominated Bayes–Nash equilibrium with $\beta_+ = 0$ and positive trade probability. By part (a) of Claim 1 in Step 1 of the proof of Theorem 1, also $\Sigma_B(\{R\}) = 0$, so bids lie in $[0, R)$ up to a Σ_B-null set. Consequently every executed trade satisfies $s \leq b < R$ almost surely and prices at the bid, $\varphi_R(s, b) = b$, and asks in $(R, 1]$ never trade. Write $\bar{b} = \sup \operatorname{supp} \Sigma_B \leq R$.

Step 1: The passive side is a truthful supply curve.

Claim 2. *$\bar{b} \in (0, R]$; for G-almost every seller with $c < \bar{b}$, the prescribed mixture is δ_c; and $\Sigma_S([0, x]) = G(x)$ for every $x \leq \bar{b}$.*

Proof of Claim Suppose $\bar{b} = 0$. Then $\Sigma_B = \delta_0$, and a trade requires an ask at 0. By part (b) of the Claim in Theorem 1, asks in $[0, R)$ are truthful, so an ask at 0 is played only by the G-null type $c = 0$; the trade probability is zero, contradicting the hypothesis. Hence $\bar{b} \in (0, R]$. Fix $c < \bar{b}$. By Lemma 1, the undominated quotes of this type form a subset of

$\{c\} \cup (R, 1]$. Every ask in $(R, 1]$ trades with no bid and earns 0. The truthful ask earns

$$U_S(c; c) = \int_{[c,1]} (b - c)\, \Sigma_B(db) \;\geq\; \int_{(c,\bar{b}]} (b - c)\, \Sigma_B(db) \;>\; 0,$$

using that the price is the bid, that the integrand is strictly positive on $(c, \bar{b}]$, and that $\Sigma_B\big((c, \bar{b}]\big) > 0$ because $\bar{b}$ is the supremum of the support and $\bar{b} > c$. Moreover c is a *global* best response: over $[0, R]$, the increment formula of Step 5 in the proof of Proposition 1 shows the truthful ask maximizes $U_S(\cdot\,; c)$, and over $(R, 1]$ every payoff is $0 < U_S(c; c)$. Any other ask in $[0, R]$ attaining this value would be dominated, hence excluded. Thus for G-almost every $c < \bar{b}$, the support of $\sigma(\cdot \mid c)$ is $\{c\}$, i.e. the mixture is δ_c. For $x \leq \bar{b}$: asks in $[0, x]$ are truthful (part (b) of Claim 1 in the proof of Theorem 1), hence come exactly from types $c \leq x$; conversely almost every type $c < x$ $(\leq \bar{b})$ asks $c \leq x$. The boundary type $c = x$ is G-null, so $\Sigma_S([0, x]) = G(x)$. △

Step 2: The monopsony problem.

Claim 3. *For every $v \in (0, 1]$, the map $b \mapsto (v - b)\, G(b)$ is strictly quasiconcave on $[0, 1]$, with unique maximizer $b^M(v) = J_S^{-1}(v)$; b^M is continuous and strictly increasing with $b^M(0) = 0$ and $0 < b^M(v) < v$ for $v > 0$. Consequently, for every $\bar{b} \in (0, 1]$, the unique maximizer of $(v - b)G(b)$ over $[0, \bar{b}]$ is $m_{\bar{b}}(v) = \min\{b^M(v), \bar{b}\}$, and on $[0, \bar{b}]$ the map is strictly increasing whenever $b^M(v) \geq \bar{b}$.*

Proof of Claim 3. The derivative is $-G(b) + (v - b)g(b) = g(b)\big[v - J_S(b)\big]$, which is strictly positive for $b < J_S^{-1}(v)$ and strictly negative beyond, since $g > 0$ and J_S is strictly increasing (regularity): strict quasiconcavity with the stated maximizer. J_S is continuous and strictly increasing with $J_S(0) = 0$ and $J_S(x) > x$ for $x > 0$ (as $G/g > 0$), so $b^M = J_S^{-1}$ is continuous, strictly increasing, $b^M(0) = 0$, and $b^M(v) < v$; and $b^M(v) > 0$ for $v > 0$. The truncated statement follows from the sign of the derivative on $[0, \bar{b}]$. △

Step 3: Necessity. By Claim 2, $\Sigma_S([0, b]) = G(b)$ for $b \leq \bar{b}$, and since the price is the bid, a bid $b \leq \bar{b}$ earns $(v - b)\, G(b)$ (for $b = \bar{b}$ this uses continuity of G: $\Sigma_S([0, \bar{b}]) = G(\bar{b})$; and if $\bar{b} = R$, the price against every ask $s \leq R$ is $R = \bar{b}$, giving the same expression). Since $\Sigma_B\big((\bar{b}, 1]\big) = 0$, for F-almost every v the support of $\beta(\cdot \mid v)$ lies in $[0, \bar{b}]$, and by (Bayes–Nash equilibrium) each of its elements is a global best response, hence in particular maximizes $(v - b)G(b)$ over $[0, \bar{b}]$. By Claim 3 this maximizer is unique, so for F-almost every v the mixture is $\delta_{m_{\bar{b}}(v)}$: mixing, gaps, and additional atoms disappear.

The cap is then pinned by consistency. Σ_B is the pushforward of F under the continuous nondecreasing map $m_{\bar{b}}$, whose maximum is $m_{\bar{b}}(1) = \min\{\bar{b}^*, \bar{b}\}$; since this pushforward must have supremum of support equal to $\bar{b}$, we get $\bar{b} \leq \bar{b}^*$. Next, $\bar{b} \neq R$ whenever $R < \bar{b}^*$: otherwise the set $\{v > J_S(R)\}$, of positive F-measure because $J_S(R) < J_S(\bar{b}^*) = 1$, plays the bid $\bar{b} = R$ with probability one, and the quote R is weakly dominated for every such type (Lemma 1), contradicting undominatedness. Combining $\bar{b} \in (0, R]$, $\bar{b} \leq \bar{b}^*$, and the exclusion of $\bar{b} = R$ when $R < \bar{b}^*$ yields exactly $\bar{b} \in \mathcal{B}(R)$ as defined in (Admissible Cap / Floor). If $\bar{b} < \bar{b}^*$, then $J_S(\bar{b}) < 1$ and the positive-mass types $v > J_S(\bar{b})$ all bid $\bar{b}$, forming the pooling atom. Finally, almost surely a trade occurs if and only if $c \leq m_{\bar{b}}(v)$, at price $m_{\bar{b}}(v)$; interim payoffs are $(v - m_{\bar{b}}(v))\, G(m_{\bar{b}}(v))$ for the buyer and $\int_{[c,1]} (b - c)\, \Sigma_B(db)$ for the seller, all functions of $\bar{b}$ alone.

Step 4: Sufficiency. Fix $\bar{b} \in \mathcal{B}(R)$ and the profile of the statement; it induces $\Sigma_S = G$ on $[0, \bar{b})$ with the remaining mass $1 - G(\bar{b})$ at 1, and $\Sigma_B = m_{\bar{b}} \# F$, supported in $[0, \bar{b}]$ with $\Sigma_B\big((c, \bar{b}]\big) > 0$ for every $c < \bar{b}$.

- *Buyer optimality.* On $[0, \bar{b}]$ the payoff is $(v-b)G(b)$, uniquely maximized at $m_{\bar{b}}(v)$ (Claim 2), with value $M_{\bar{b}}(v) \geq (v - \bar{b})G(\bar{b})$. Bids $b \in (\bar{b}, R)$ earn $(v - b)\, G(\bar{b}) < (v - \bar{b})\, G(\bar{b})$, since no ask lies in $(\bar{b}, b]$ and $G(\bar{b}) > 0$. Bids $b \in [R, 1)$ trade against exactly the asks in $[0, \bar{b})$, each priced at R (as $s < \bar{b} \leq R \leq b$), earning $(v - R)\, G(\bar{b}) \leq (v - \bar{b})\, G(\bar{b})$, strictly if $\bar{b} < R$. The bid 1 additionally meets the ask atom at 1, adding $(v - 1)\big(1 - G(\bar{b})\big) \leq 0$. Hence $m_{\bar{b}}(v)$ is a global best response for every v.

- *Seller optimality.* All bids lie in $[0, \bar{b}] \subseteq [0, R]$ and the price is the bid, so for $s \leq R$, $U_S(s; c) = \int_{[s,1]} (b - c)\, \Sigma_B(db)$, maximized over $[0, R]$ at $s = c$ by the increment formula (Proof of Proposition 1, Step 5), while asks in $(R, 1]$ earn 0. For $c < \bar{b}$ the truthful value is strictly positive, so $s = c$ is optimal. For $c \geq \bar{b}$: the ask 1 earns 0; any ask $s \leq \bar{b}$ trades only at prices $b \leq \bar{b} \leq c$, a nonpositive payoff; asks in $(\bar{b}, 1]$ earn 0. Hence the ask 1 is optimal.

Claim 4 (Undominatedness of the profile)**.** *(a) For every $v \in (0, 1)$, every bid $b \in (0, \min\{v, R\})$ is undominated; the bid 0 is undominated for $v = 0$. (b) For every $c \in [0, R)$, the truthful ask c is undominated; the ask 1 is undominated for every c.*

Proof of Claim 4. (a) Fix $v \in (0, 1)$ and $b \in (0, \min\{v, R\})$, and test against the ask $s = b$: the bid b trades at price $\varphi_R(b, b) = b$, earning $v - b > 0$; any bid $b'' < b$ earns 0; any

$b'' \in (b, R)$ trades at price b'' and earns $v - b'' < v - b$; any $b'' \geq R$ trades at price R (as $s = b < R \leq b''$), earning $v - R < v - b$. Hence every pure alternative loses strictly against $s = b$, and therefore so does every mixture other than δ_b; no action weakly dominates b. For $v = 0$, every payoff is at most 0 and the bid 0 earns exactly 0 against every ask, so no action can improve on it anywhere, let alone strictly.

(b) Fix $c \in [0, R)$. Below the reference the bid sets the price, so against a fixed bid $b < R$ *all* asks $s \leq b$ are payoff-equivalent; the truthful ask must instead be separated through low bids. Let $\mu \in \Delta(\mathcal{A})$ weakly dominate δ_c. Against any bid $b \in (c, R)$, the ask c earns $b - c > 0$, while μ earns $\mu([0, b])\,(b - c)$ (components $s \leq b$ trade at price b; components $s > b$ earn 0, using $b < R$ so no straddling ask beats it); weak dominance forces $\mu([0, b]) = 1$ for every $b \in (c, R)$, hence $\mu([0, c]) = 1$. If $\mu([0, c)) > 0$, pick $s'' < c$ in its support and test against a bid $b \in (s'', c)$: the ask c earns 0 (no trade), while μ earns $\mu([0, b])\,(b - c) < 0$ since $\mu([0, b]) \geq \mu([0, s'']) > 0$ and $b < c$ — contradicting weak dominance. Hence $\mu = \delta_c$, and the truthful ask is undominated. For the ask 1 and any $c < 1$, test against the bid $b = 1$: the ask 1 trades at price $\varphi_R(1, 1) = 1$, earning $1 - c$, the maximal feasible price, while every ask $s < 1$ trades at price R (if $s \leq R$) or s (if $s \in (R, 1)$), both strictly below 1; so no mixture other than δ_1 weakly dominates it, and none dominates it. For $c = 1$ every payoff is at most 0 and the ask 1 earns 0 everywhere, so it is undominated. △

The prescribed quotes fall under Claim 3: for $v \in (0, 1)$, $m_{\bar{b}}(v) \in (0, \min\{v, R\})$, since $b^M(v) < v$ and, on the pooled region, $v > J_S(\bar{b}) \geq \bar{b}$, while $m_{\bar{b}}(v) \leq \bar{b} < R$ when $\bar{b} < R$; if $\bar{b} = R$ (possible only with $\bar{b}^* = R \leq R$), the sole type reaching the quote R is the F-null type $v = 1$, which may be reassigned its truthful bid without affecting any distribution. The type $v = 0$ bids $m_{\bar{b}}(0) = 0$. Sellers with $c < \bar{b}$ ($< R$ or $= R$ with the null-type caveat) play the undominated truthful ask, and sellers with $c \geq \bar{b}$ play the undominated ask 1. Finally, the trade probability is $\int_0^1 G\big(m_{\bar{b}}(v)\big) dF(v) > 0$, since $m_{\bar{b}}(v) > 0$ for $v > 0$. This verifies every requirement, and the profile is an undominated Bayes–Nash equilibrium.

The seller-side mirror. This follows by the symmetric argument, exchanging bid and ask, buyer and seller, F and G, J_S and J_B, minimum and maximum, the ceiling $\bar{b}^*$ and the floor $\underline{s}^*$, and the sets $\mathcal{B}(R)$ and $\mathcal{S}(R)$; above the reference the ask sets the price, so Claim 1's truthful side is the buyer, Claim 3 becomes the monopoly problem $(s - c)(1 - F(s))$ with maximizer $s^M(c) = J_B^{-1}(c)$, and in Claim 4(b) the payoff-equivalence now runs over asks against a fixed high bid, with separation through high asks. □

A.6 Proof of Corollary 1

Proof. We break up the proof into 3 steps.

Step 1: Existence. For every $R \in (0,1)$, the admissible sets are nonempty: $\mathcal{B}(R)$ contains $\min\{R, \bar{b}^*\}/2 > 0$, and $\mathcal{S}(R)$ contains $(1 + \max\{R, \underline{s}^*\})/2 < 1$. By Theorem 2 (Sufficiency) and its mirror, each admissible cap and floor supports an undominated equilibrium, with trade probability $\int_0^1 G(m_{\bar{b}}(v))\, dF(v) > 0$ and its mirror, respectively.

Step 2: The gap. Note first that $q(R) \leq R \leq t(R)$ always, so the interval $(q(R), t(R))$ is well defined (possibly empty). In every buyer-side equilibrium, Theorem 2 (Necessity) gives $\bar{b} \in \mathcal{B}(R)$, so $\bar{b} \leq \min\{R, \bar{b}^*\} = q(R)$, and almost surely every transaction price equals $m_{\bar{b}}(V) \leq \bar{b} \leq q(R)$. In every seller-side equilibrium, the mirror gives prices at least $\underline{s} \geq t(R)$ almost surely. Hence no equilibrium generates a transaction price in $(q(R), t(R))$ with positive probability; indeed, none generates a price there at all, off null sets.

Step 3: Attainment. Fix $p \in (0, q(R))$. Then $p \in \mathcal{B}(R)$: if $R \geq \bar{b}^*$ this reads $p \in (0, \bar{b}^*]$, and if $R < \bar{b}^*$ it reads $p \in (0, R)$, both true since $p < q(R)$. Consider the equilibrium with cap $\bar{b} = p$. Since $p < \bar{b}^*$, we have $J_S(p) < J_S(\bar{b}^*) = 1$, so the pooled set $\{v > J_S(p)\}$ has F-measure $1 - F(J_S(p)) > 0$, every pooled type bids exactly p, and each such bid trades with probability $G(p) > 0$. Hence the price p occurs with positive probability in this equilibrium. The mirror construction, an ask floor at $p \in (t(R), 1)$, covers the prices above the floor. Finally, the frontier bounds themselves arise with positive probability only when constrained away from the absolute optima: at an unconstrained frontier ($\bar{b} = \bar{b}^*$), the pooled set $\{v > J_S(\bar{b}^*)\} = \{v > 1\}$ is empty, so the price $\bar{b}^*$ is attained only by the F-null type $v = 1$; this is why the attainment claim is stated on the open intervals. □

A.7 Proof of Corollary 2

Proof. By Theorem 2 (Necessity), the equilibrium outcome coincides, off null sets, with that of the profile in which the buyer bids $m_{\bar{b}}(V)$ and sellers with $C < \bar{b}$ ask truthfully; we may therefore compute all conditional laws from this profile. The trade event is $T = \{C \leq m_{\bar{b}}(V)\}$ (the boundary $C = m_{\bar{b}}(V)$ and the excluded sellers $C \geq \bar{b}$ contribute null sets), and on T the price is $P = m_{\bar{b}}(V)$.

For Borel sets $A, B \subseteq [0,1]$ and $\mathcal{P} \subseteq (0, \bar{b}]$, independence of V and C gives

$$\Pr\big(T,\ P \in \mathcal{P},\ V \in A,\ C \in B\big) = \int_{A \cap m_{\bar{b}}^{-1}(\mathcal{P})} G\big(B \cap [0, m_{\bar{b}}(v)]\big)\, dF(v). \tag{6}$$

Split the domain at $J_S(\bar{b})$ (with the convention $J_S(\bar{b}) = 1$ if $\bar{b} = \bar{b}^*$). On $\{v < J_S(\bar{b})\}$, the map $m_{\bar{b}}(v) = b^M(v) = J_S^{-1}(v)$ is a continuous strict increasing bijection onto $(0, \bar{b})$; substituting $v = J_S(p)$ in (6) shows that the part of the measure with $P < \bar{b}$ disintegrates over the price as

$$\Pr\big(T, P \in dp, V \in A, C \in B\big) = \mathbf{1}\{J_S(p) \in A\}\, G\big(B \cap [0,p]\big)\, \nu(dp),$$
$$\text{where} \quad \nu = \big(b^M\big)\# F\big|_{\{v < J_S(\bar{b})\}}.$$

Hence a version of the regular conditional distribution of (V, C) given T and $P = p < \bar{b}$ is $\delta_{J_S(p)} \otimes G(\cdot \mid C \le p)$: the buyer's type is $J_S(p)$ almost surely, the seller's type is distributed as $G(\cdot \mid C \le p)$, and the product form is conditional independence. On the pooled event $\{V > J_S(\bar{b})\}$ (nonempty exactly when $\bar{b} < \bar{b}^*$), every type bids $\bar{b}$, so $P = \bar{b}$ and, crucially, the trade probability given $V = v$ is the *constant* $G(\bar{b})$; (6) with $\mathcal{P} = \{\bar{b}\}$ therefore factorizes as

$$\Pr\big(T,\ P = \bar{b},\ V \in A,\ C \in B\big) = F\big(A \cap (J_S(\bar{b}), 1]\big) \cdot G\big(B \cap [0, \bar{b}]\big),$$

so conditional on T and $P = \bar{b}$, the types are independent, with $V \sim F(\cdot \mid V > J_S(\bar{b}))$—in particular $V \ge J_S(\bar{b})$ almost surely—and $C \sim G(\cdot \mid C \le \bar{b})$, consistent with the statement at $p = \bar{b}$. The seller-side mirror follows by the symmetric computation with $s_{\underline{s}}$, J_B, and the roles of F and G exchanged. □

A.8 Proof of Corollary 3

Proof. Let $R \ge \bar{b}^*$, so the frontier cap is $q(R) = \bar{b}^*$ and the frontier equilibrium prescribes $m_{\bar{b}^*}(v) = b^M(v)$ for every buyer (as $b^M(v) \le b^M(1) = \bar{b}^*$), truthful asks for sellers with $c < \bar{b}^*$, and the ask 1 otherwise. We prove the statement for the buyer's-bid double auction. The seller-side statement follows by the symmetric argument with the seller's-ask auction, the floor $\underline{s}^*$, monopoly asks $s^M(c)$, and truthful bids, for every $R \le \underline{s}^*$.

Step 1: The undominated equilibrium of the buyer's-bid auction. In the buyer's-bid auction, the ex-post seller payoff is $\mathbf{1}\{s \le b\}(b - c)$: the ask affects execution only, never the price. For any type c, the truthful ask weakly dominates every $s \ne c$. Indeed, against any

bid $b \geq \max\{s, c\}$ both asks trade at the same price b, and against any bid $b < \min\{s, c\}$ neither trades, so the payoffs can differ only against bids between the two asks: for $s > c$, the truthful ask additionally trades against every $b \in [c, s)$, at payoff $b - c \geq 0$, strictly positive for $b \in (c, s)$; for $s < c$, it sheds the trades against every $b \in [s, c)$, each at payoff $b - c < 0$. Hence the truthful ask is weakly dominant, and it is itself undominated, since by the same case analysis every alternative ask loses strictly against some bid. In every undominated equilibrium, therefore, almost every seller asks truthfully, so the buyer faces the supply curve G and pays her bid. By the proof of Theorem 2, the unique best response of type v is $b^M(v)$; the bid $b^M(v)$ is undominated, since the test-against-matching-ask argument applies verbatim (for $b < v$ the buyer's-bid price coincides with the reference-price auction's price at bids below R). Hence the profile (truthful asks, monopsony bids) is the essentially unique undominated Bayes–Nash equilibrium of the buyer's-bid auction, unique up to the play of null type sets and of sellers who never trade.

Step 2: Outcome equivalence. In both mechanisms and both profiles, the buyer bids $b^M(V)$ and every executed trade prices at the bid: in the buyer's-bid auction by definition; in the reference-price auction because $b^M(V) \leq \bar{b}^* \leq R$ places every trade in the bid-pricing branch of (Reference-price rule) (bids equal to R occur only for the null type $v = 1$ when $\bar{b}^* = R$). The trade events coincide off null sets: in the frontier equilibrium, trade occurs iff $C \leq b^M(V)$ with $C < \bar{b}^*$, and in the buyer's-bid equilibrium iff $C \leq b^M(V)$; the difference $\{C \geq \bar{b}^*,\, C \leq b^M(V)\} \subseteq \{C = b^M(V) = \bar{b}^*\}$ is null. Allocation and price being identical functions of (V, C) off null sets, the transaction-price distributions coincide; buyer interim payoffs coincide as identical functionals of G; seller interim payoffs coincide, equal to $\int_{[c,1]}(b - c)\, d\big(b^M \# F\big)$ for $c < \bar{b}^*$ and to 0 for $c \geq \bar{b}^*$ in both (in the buyer's-bid equilibrium, a seller with $c \geq \bar{b}^*$ trades only against the null event $b^M(V) \geq c$). Finally, neither profile's strategies depend on R; hence for every $R \geq \bar{b}^*$ the equilibrium strategies, and therefore the induced outcome, are identical. □

A.9 Proof of Proposition 3

Proof of Proposition 3. By Theorem 2 (Necessity), every equilibrium with cap $\bar{b}$ induces, off null sets, the outcome of the canonical profile: buyers bid $m_{\bar{b}}(v) = \min\{b^M(v), \bar{b}\}$, sellers with $c < \bar{b}$ ask truthfully, all others ask 1. All payoffs below are computed from these profiles; write $m = m_{\bar{b}}$ and $m' = m_{\bar{b}'}$, and note

$$m'(v) \geq m(v) \text{ for every } v, \quad \text{with} \quad m'(v) > m(v) \iff v > J_S(\bar{b}), \tag{7}$$

since $m'(v) = m(v) = b^M(v)$ when $b^M(v) \leq \bar{b}$, while for $b^M(v) > \bar{b}$, i.e. $v > J_S(\bar{b})$, we have $m(v) = \bar{b}$ and $m'(v) = \min\{b^M(v), \bar{b}'\} > \bar{b}$. The set $\{v > J_S(\bar{b})\}$ has positive F-measure: $\bar{b} < \bar{b}' \leq \bar{b}^*$ gives $J_S(\bar{b}) < J_S(\bar{b}^*) = 1$.

Step 1: The buyer's payoff. In the cap-$\bar{b}$ equilibrium, type v's payoff is $\max_{b \in [0,\bar{b}]}(v - b)G(b)$, attained at $m(v)$ (Theorem 2, using $\Sigma_S([0, b]) = G(b)$ on the active range). Enlarging the feasible set to $[0, \bar{b}']$ can only raise the maximum, so no buyer type is worse off. For $v > J_S(\bar{b})$: by the strict quasiconcavity of the monopsony objective established in the proof of Theorem 2, $(v - b)G(b)$ is strictly increasing on $[0, b^M(v)]$, and $m(v) = \bar{b} < m'(v) \leq b^M(v)$ by (7), so $(v - m'(v))\, G(m'(v)) > (v - \bar{b})\, G(\bar{b})$: strictly better off.

Step 2: The seller's payoff. In the cap-$\bar{b}$ equilibrium, a seller of type c earns $\int_0^1 \big(m(v) - c\big)^+ dF(v)$: types with $c < \bar{b}$ earn this as the truthful-ask payoff (every trade prices at the bid), and types with $c \geq \bar{b}$ earn 0, which the formula also returns since $m \leq \bar{b} \leq c$. By (7) and the monotonicity of $x \mapsto (x - c)^+$, the integrand rises pointwise, so no seller type is worse off. For strictness, fix $c < \bar{b}'$ and note that for v in a neighborhood of 1, $m'(v) = \min\{b^M(v), \bar{b}'\} \to \min\{\bar{b}^*, \bar{b}'\} = \bar{b}'$ by continuity of b^M, so there is a positive-F-measure set of types with $m'(v) > \max\{c, \bar{b}\} \geq m(v)$; on this set $(m'(v) - c)^+ > (m(v) - c)^+$ (the right side being $(\bar{b} - c)^+$ or, for $c \geq \bar{b}$, zero). Hence every seller with $c < \bar{b}'$ is strictly better off. Together with Step 1 this proves the Pareto claims.

Step 3: Welfare and trade probability. By (Expected Gains from Trade), ex-ante welfare is the sum of the two sides' ex-ante payoffs, so the Pareto claims give $W_B(\bar{b}') > W_B(\bar{b})$. Directly: the trade region expands from $\{(v, c) : c \leq m(v)\}$ to $\{(v, c) : c \leq m'(v)\}$; every added profile satisfies $c \leq m'(v) \leq b^M(v) < v$ and is therefore strictly surplus-creating, and the added region has probability $\int_{\{v > J_S(\bar{b})\}} \big[G(m'(v)) - G(\bar{b})\big] dF(v) > 0$ since G is strictly

increasing. The same display shows the trade probability $\int_0^1 G\big(m_{\bar{b}}(v)\big)dF(v)$ is strictly increasing in the cap.

Step 4: The mirror. The truncated monopoly statements follow by the symmetric argument: floors $\underline{s} > \underline{s}'$ in $\mathcal{S}(R)$, strategies $s_{\underline{s}}(c) = \max\{s^M(c), \underline{s}\}$ falling pointwise as the floor falls, strictly exactly for $c < J_B(\underline{s})$; buyers' payoffs $\int_0^1 \big(v - s_{\underline{s}}(c)\big)^+ dG(c)$ rising pointwise; the demand-side objective $(s-c)(1-F(s))$ strictly decreasing on $[s^M(c), 1]$; and every buyer with $v > \underline{s}'$ strictly gaining. All orders reverse in the floor. □

A.10 Proof of Proposition 4

Proof. We prove the proposition in multiple steps.

Step 1: Welfare as a function of the bound. By Theorem 2 (Necessity), the outcome of every buyer-side equilibrium with cap q is that of the canonical profile, so its expected gains from trade depend on the cap alone:

$$W_B(q) = \int_0^1 \int_0^{m_q(v)} (v-c)\, dG(c)\, dF(v), \qquad W_S(t) = \int_0^1 \int_{a_t(c)}^1 (v-c)\, dF(v)\, dG(c), \quad (8)$$

with $m_q(v) = \min\{b^M(v), q\}$ on $(0, \bar{b}^*]$ and $a_t(c) = \max\{s^M(c), t\}$ on $[\underline{s}^*, 1)$. Both maps are strictly monotone toward their frontier bound: any two caps $q < q' \le \bar{b}^*$ are admissible at every reference $R \ge \bar{b}^*$, so Proposition 3 applies and $W_B(q) < W_B(q')$; the mirror gives W_S strictly decreasing in t. Both maps are continuous: $m_q(v)$ and $a_t(c)$ are 1-Lipschitz in the bound for every type, the inner integrals are bounded by 1, and dominated convergence passes the limit through the outer integral. Finally, both are strictly positive, since $m_q(v) > 0$ for every $v > 0$ and $a_t(c) < 1$ for every $c < 1$, and the densities are strictly positive.

Step 2: The value at a fixed reference. Fix $R \in (0,1)$. By Theorems 1 and 2, every undominated equilibrium either induces no trade, with welfare 0, or belongs to one of the two one-sided families, and both families are nonempty (Corollary 1). Within the buyer-side family, the attainable welfares are $\{W_B(q) : q \in \mathcal{B}(R)\}$. For the admissible caps, $\sup \mathcal{B}(R) = q(R)$ and $\mathcal{B}(R)$ contains a sequence increasing to $q(R)$, so by monotonicity and continuity of W_B,

$$\sup_{q \in \mathcal{B}(R)} W_B(q) = W_B\big(q(R)\big),$$

attained if and only if $q(R) \in \mathcal{B}(R)$, i.e. $R \geq \bar{b}^*$. The mirror yields the seller-side value $W_S\big(t(R)\big)$, attained if and only if $R \leq \underline{s}^*$. Since both family values are strictly positive by Step 1, no-trade equilibria affect neither supremum, and $W^*(R) = \max\big\{W_B\big(q(R)\big),\, W_S\big(t(R)\big)\big\}$.

Step 3: The supremum over references. For every R, $q(R) \leq \bar{b}^*$ and $t(R) \geq \underline{s}^*$, so Step 1 gives $W^*(R) \leq \max\{W_B(\bar{b}^*), W_S(\underline{s}^*)\}$. Conversely, every $R \geq \bar{b}^*$ has $W^*(R) \geq W_B(q(R)) = W_B(\bar{b}^*)$, and every $R \leq \underline{s}^*$ has $W^*(R) \geq W_S(\underline{s}^*)$; both ranges are nonempty since $\bar{b}^*, \underline{s}^* \in (0,1)$. Hence $\sup_R W^*(R) = \max\{W_B(\bar{b}^*), W_S(\underline{s}^*)\} =: W^*$.

Step 4: Attainment. *Sufficiency.* If $R \geq \bar{b}^*$ and $W_B(\bar{b}^*) \geq W_S(\underline{s}^*)$, then $\bar{b}^* \in \mathcal{B}(R)$, and by Theorem 2 (Sufficiency) the truncated monopsony equilibrium with cap $\bar{b}^*$ exists exactly and has welfare $W_B(\bar{b}^*) = W^*$. The seller side mirrors.

Necessity. Suppose an undominated equilibrium at reference R attains W^*. Since $W^* > 0$ by Step 1, the equilibrium has positive trade probability and belongs to one of the two families. If it is buyer-side with cap $q \in \mathcal{B}(R)$, then $W_B(q) = W^* \geq W_B(\bar{b}^*)$, while $W_B(q) \leq W_B(\bar{b}^*)$ by Step 1; hence $W_B(q) = W_B(\bar{b}^*)$, and strict monotonicity forces $q = \bar{b}^*$. Admissibility of $\bar{b}^*$ requires $R \geq \bar{b}^*$, and $W_B(\bar{b}^*) = W^* \geq W_S(\underline{s}^*)$. The seller-side case mirrors, yielding $R \leq \underline{s}^*$ and $W_S(\underline{s}^*) \geq W_B(\bar{b}^*)$. These are exactly the two stated alternatives. □

A.11 Proof of Theorem 3

Proof. The equality is Proposition 4. It remains to prove

$$\max\big\{W_B(\bar{b}^*),\, W_S(\underline{s}^*)\big\} \;\geq\; \max\Big\{\tfrac{1}{2}\,W_{\mathrm{SB}},\; \tfrac{1}{\pi}\,W_{\mathrm{FB}}\Big\}. \tag{9}$$

Step 1: Identification with the one-sided offer mechanisms. In the *buyer-offering* mechanism, the buyer observes v and posts a take-it-or-leave-it price p, which the seller accepts if and only if $c \leq p$; the buyer solves $\max_p (v-p)G(p)$. The derivative $g(p)\big(v - J_S(p)\big)$ and strict regularity make the objective strictly single-peaked with maximizer $b^M(v) = J_S^{-1}(v)$, so trade occurs exactly when $J_S(c) \leq v$. In the *seller-offering* mechanism, the seller solves $\max_p (p-c)(1-F(p))$, with derivative $f(p)\big(c - J_B(p)\big)$ and maximizer $s^M(c) = J_B^{-1}(c)$, so trade occurs exactly when $c \leq J_B(v)$. Comparing with (Frontier Welfares), the expected gains from trade of the two mechanisms are exactly the two frontier welfares:

$$\mathrm{GFT}(\text{buyer-offering}) = W_B(\bar{b}^*) \qquad \text{and} \qquad \mathrm{GFT}(\text{seller-offering}) = W_S(\underline{s}^*). \tag{10}$$

Step 2: The second-best bound. Brustle et al. (2017) prove, for arbitrary independent priors, that the better of the generalized seller-offering and generalized buyer-offering mechanisms attains at least one half of the maximum expected gains from trade over all interim individually rational, Bayesian incentive-compatible, ex-ante budget-balanced mechanisms. Two reductions transfer their bound to (9). First, the generalized mechanisms trade on ironed virtual sets; under regularity no ironing occurs, their allocations reduce to $\{(v,c) : v \geq J_S(c)\}$ and $\{(v,c) : c \leq J_B(v)\}$, and since expected gains from trade depend on a mechanism only through its allocation, (10) identifies their welfares with the frontier welfares. Second, their program imposes exact ex-ante budget balance, whereas W_{SB} is defined with the weak inequality; the two programs have the same value, since any ex-ante budget surplus can be rebated through type-independent transfers, which preserve Bayesian incentive compatibility and interim individual rationality and leave the allocation, hence the gains from trade, unchanged. Together, $\max\{W_B(\bar{b}^*), W_S(\underline{s}^*)\} \geq \frac{1}{2} W_{\mathrm{SB}}$.

Step 3: The first-best bound. The *random-offerer* mechanism runs the two offer mechanisms with equal probability; by (10), its expected gains from trade are $\frac{1}{2}\big(W_B(\bar{b}^*) + W_S(\underline{s}^*)\big)$. Jo (2026), sharpening Fei (2022) and Deng et al. (2022), proves that the random offerer attains at least $\frac{1}{\pi}\, W_{\mathrm{FB}}$ for every pair of independent priors on a common compact interval, hence under our assumptions. Since the maximum of two numbers weakly exceeds their average,

$$\max\big\{W_B(\bar{b}^*),\, W_S(\underline{s}^*)\big\} \;\geq\; \tfrac{1}{2}\big(W_B(\bar{b}^*) + W_S(\underline{s}^*)\big) \;\geq\; \tfrac{1}{\pi}\, W_{\mathrm{FB}}.$$

Combining Steps 2 and 3 yields (9). □

A.12 Proof of Proposition 5

Proof. Throughout, fix $R \in (0,1)$ and write $q = q(R) \leq R \leq t = t(R)$. The four side-strategies are those of the two frontier profiles: the buyer's *making* strategy $\beta^B(v) = \min\{b^M(v), q\}$ and *taking* strategy $\beta^S(v) = v$ for $v > t$, $\beta^S(v) = 0$ otherwise; the seller's *taking* strategy $\sigma^B(c) = c$ for $c < q$, $\sigma^B(c) = 1$ otherwise, and *making* strategy $\sigma^S(c) = \max\{s^M(c), t\}$.

These profiles are exact equilibrium profiles of the quote game whenever the corresponding frontier bound is admissible: the cap $q(R)$ lies in $\mathcal{B}(R)$ if and only if $R \geq \bar{b}^*$, and the floor $t(R)$ lies in $\mathcal{S}(R)$ if and only if $R \leq \underline{s}^*$; both hold exactly on the frozen interval. We therefore

prove the proposition first in the exact case—in fact, uniformly over a set of caps and floors that contains the exact pair—and then show, in a last step closing the proof, that every conclusion passes to the ε-frontier limits that define the game at all remaining references.

To this end, throughout Steps 1–4 we fix an arbitrary

$$q' \in \left[\tfrac{q}{2},\, q\right] \qquad \text{and} \qquad t' \in \left[t,\, \tfrac{1+t}{2}\right],$$

consider the profiles built from cap q' and floor t', and establish every bound with slack uniform over this set; it contains the exact pair (q, t) and every ε-frontier pair for ε small.

Step 1: The cells. We evaluate all four cells.

- (β^B, σ^B) and (β^S, σ^S) are the frontier profiles at cap q' and floor t', with payoffs (Buyer Control: Buyer / Seller Payoff) and (Seller Control: Buyer / Seller Payoff) evaluated there, by Theorem 2 (Necessity).

- Under (β^B, σ^S), every bid is at most $q' \le R$ and every ask at least $t' \ge R$; a crossing requires bid = ask = $q' = R = t'$, which involves at most the null types $v = 1$ (if $q' = \bar{b}^* = R$) and $c = 0$ (if $t' = \underline{s}^* = R$), so trade has probability zero and the payoffs are $(0, 0)$; at ε-frontier pairs with $q' < R$ or $t' > R$, trade is impossible outright.

- Under (β^S, σ^B), trade pairs a buyer $v > t'$ bidding v with a seller $c < q'$ asking c; since $c < q' \le R$ and $v > t' \ge R$, the quotes straddle the reference and $\varphi_R(c, v) = R$; buyers with $v \le t'$ bid 0 and sellers with $c \ge q'$ ask 1, trading only on null events. Independence gives exactly (Buyer / Seller Concession), evaluated at (q', t').

Step 2: A uniform gap for making over taking against the taking opponent. We claim

$$\Pi_B^B \;\ge\; X_B + \delta, \qquad \delta := \Big[F\big(\tfrac{3R}{4}\big) - F\big(\tfrac{R}{2}\big)\Big] \cdot \frac{R}{4} \cdot G\Big(\min\big\{\tfrac{R}{4}, \tfrac{q}{2}\big\}\Big) \;>\; 0, \tag{11}$$

at every (q', t') in the fixed set. Both payoffs face the same opponent strategy σ^B, so (11) follows by integrating a pointwise comparison over buyer types. Consider first the set $A = [\frac{R}{2}, \frac{3R}{4}]$, of F-measure $F(\frac{3R}{4}) - F(\frac{R}{2}) > 0$; since $\frac{3R}{4} < R \le t'$, every type $v \in A$ takes a payoff of 0 under β^S. Under price-making, the bid $b(v) = \min\{v/2,\, q/2\}$ is feasible ($b(v) \le q/2 \le q'$), satisfies $b(v) \ge \min\{R/4,\, q/2\} > 0$ and $v - b(v) \ge v/2 \ge R/4$, and trades against every seller $c \le b(v)$ at price $b(v) < R$; hence the type's making payoff is at least $\frac{R}{4}\, G\big(\min\{\frac{R}{4}, \frac{q}{2}\}\big)$.

Next, for types $v \notin A$ with $v \le t'$, the taking payoff is 0 while the making payoff is nonnegative. Finally, for $v > t'$, the taking payoff is $(v-R)\,\Sigma_S\big([0,v]\big) \le (v-R)\,G(q')$, while the deviation available to the maker of bidding q' itself trades against the same participating sellers—all asks lie in $[0,q')\cup\{1\}$—at price $\min\{q',R\} \le R$, earning at least $(v-R)\,G(q')$; the making payoff, a maximum over feasible bids, is weakly larger. Summing the three cases yields (11). The mirror argument, with the set $A' = \big[\frac{1+3R}{4}, \frac{1+R}{2}\big]$, the ask $s(c) = \max\big\{\frac{1+c}{2}, \frac{1+t}{2}\big\}$, and the bound $s(c) \le s_1 := \max\big\{\frac{3+R}{4}, \frac{1+t}{2}\big\} < 1$, gives

$$\Pi_S^S \;\ge\; X_S + \delta', \qquad \delta' := \Big[G\big(\tfrac{1+R}{2}\big) - G\big(\tfrac{1+3R}{4}\big)\Big] \cdot \frac{1-R}{4} \cdot \big[1 - F(s_1)\big] \;>\; 0,$$

again uniformly over the fixed set of caps and floors.

Step 3: Positivity of the taker payoffs. Set $\beta_0 := \min\big\{b^M(\tfrac{1}{2}),\, q/2\big\} > 0$. For every $v \ge \frac{1}{2}$ and every $q' \ge q/2$, $\beta^B(v) = \min\{b^M(v), q'\} \ge \beta_0$, and each executed trade prices at $\min\{\beta^B(v), R\} \ge \beta_0$; hence

$$\Pi_S^B \;\ge\; \big[1 - F(\tfrac{1}{2})\big] \int_0^{\beta_0} \big(\beta_0 - c\big)\, dG(c) \;>\; 0,$$

uniformly over the fixed set. Symmetrically, Π_B^S is bounded below by a strictly positive constant uniformly over the set.

Step 4: The equilibrium properties. At (β^B, σ^B): the buyer's deviation to β^S yields $X_B < \Pi_B^B$ by (11); the seller's deviation to σ^S yields the payoff of asks at least $t' \ge R$ against bids at most $q' \le R$, namely 0 (probability-zero trade, as in Step 1), which is strictly below $\Pi_S^B > 0$ by Step 3. Both deviations lose strictly, so buyer control is a strict Nash equilibrium of the role-assignment game; seller control mirrors, using $\Pi_S^S > X_S$ and $\Pi_B^S > 0$. At the off-diagonal (β^B, σ^S), the buyer's deviation to β^S moves the profile to seller control and yields $\Pi_B^S > 0$, the cell payoff: a strict improvement, so the profile is not a Nash equilibrium. At (β^S, σ^B), the buyer's deviation to β^B yields $\Pi_B^B > X_B$: again a strict improvement. This proves all four assertions at every admissible exact pair, and in particular on the frozen interval.

Step 5: The ε-frontier limits. We prove the following claim:

Claim 5. *At every reference at which a frontier bound is inadmissible, the payoffs of the ε-frontier profiles converge, as $\varepsilon \downarrow 0$, to the six payoff formulas evaluated at $(q(R), t(R))$, and every conclusion of Steps 1–4 holds for the limit payoffs.*

Hence, the role-assignment game is well defined at every reference, with payoffs (Buyer Control: Buyer / Seller Payoff)–(Buyer / Seller Concession) read at $(q(R), t(R))$, and Proposition 5 holds verbatim.

Proof of the Claim. Convergence of the six payoffs is direct. The frontier strategies are Lipschitz in their bound: for caps q', q'' and every type, $|\min\{b^M(v), q'\} - \min\{b^M(v), q''\}| \le |q' - q''|$, and likewise for floors. Each of the regime payoffs (Buyer Control: Buyer / Seller Payoff)–(Seller Control: Buyer / Seller Payoff) is the integral of a bounded integrand obtained by composing the frontier strategy with the continuous functions G, F, and affine maps of the type; along $q_\varepsilon \to q(R)$ and $t_\varepsilon \to t(R)$ (which satisfy $|q_\varepsilon - q(R)| \le \varepsilon$ and $|t_\varepsilon - t(R)| \le \varepsilon$ by construction), the integrands converge pointwise in the type and are bounded by 1, so dominated convergence gives convergence of all four payoffs to their values at $(q(R), t(R))$. The fallback payoffs $X_B = G(q_\varepsilon) \int_{t_\varepsilon}^1 (v - R)\, dF(v)$ and $X_S = \big[1 - F(t_\varepsilon)\big] \int_0^{q_\varepsilon} (R - c)\, dG(c)$ are compositions of continuous functions of $(q_\varepsilon, t_\varepsilon)$ and converge likewise. The cell values and inequalities then pass to the limit. The mutual-making cell is exactly $(0, 0)$ at every ε (Step 1), hence in the limit; the mutual-taking cell equals (X_B, X_S) at $(q_\varepsilon, t_\varepsilon)$ and converges to its value at $(q(R), t(R))$. For $\varepsilon < \min\{q/2, (1 - t)/2\}$, the pair $(q_\varepsilon, t_\varepsilon)$ lies in the fixed set of Steps 1–4, so the inequalities $\Pi_B^B \ge X_B + \delta$ and $\Pi_S^S \ge X_S + \delta'$, and the two positive lower bounds of Step 3, hold at every such ε with slacks δ, δ' and bounds that do not depend on ε; passing to the limit preserves each of them with its slack, hence strictly. Since every payoff comparison in Step 4 rests on exactly these four strict inequalities, all four assertions of the proposition hold for the limit payoffs. △

In view of Step 5, the role-assignment game is well defined at every reference, its payoffs are (Buyer Control: Buyer / Seller Payoff)–(Buyer / Seller Concession) evaluated at $\big(q(R), t(R)\big)$, and all subsequent results state payoffs and inequalities there directly, with no further reference to the approximation. □

A.13 Proof of Theorem 4

Proof. For $q \in (0, \bar{b}^*]$ and $t \in [\underline{s}^*, 1)$, write $\Pi_B^B(q), \Pi_S^B(q), \Pi_B^S(t), \Pi_S^S(t)$ for the payoffs (Buyer Control: Buyer / Seller Payoff)–(Seller Control: Buyer / Seller Payoff) with cap q and floor t; where a frontier is constrained, all payoffs are the ε-frontier limits, which equal these formulas evaluated at $q(R)$ and $t(R)$, so every identity below applies verbatim at every reference. Define the four *deviation losses*

$$L_1(R) = \Pi_B^B\big(q(R)\big) - X_B(R), \qquad L_2(R) = \Pi_S^B\big(q(R)\big),$$

$$L_3(R) = \Pi_S^S\big(t(R)\big) - X_S(R), \qquad L_4(R) = \Pi_B^S\big(t(R)\big),$$

so that $\nabla = L_1 L_2 - L_3 L_4$ by (Risk-Dominance Index).

Step 1: Regularity of the payoff maps. We start by showing that the four payoff maps are continuously differentiable, with

$$\big(\Pi_B^B\big)'(q) = g(q)\,\mathbb{E}\big[(V - J_S(q))^+\big], \qquad \big(\Pi_S^B\big)'(q) = G(q)\big[1 - F\big(J_S(q)\big)\big],$$
$$\big(\Pi_B^S\big)'(t) = -\big[1 - F(t)\big]\,G\big(J_B(t)\big), \qquad \big(\Pi_S^S\big)'(t) = -f(t)\int_0^{J_B(t)} \big[J_B(t) - c\big]\,dG(c),$$

with the convention that integrals over $[0, J_B(t)]$ vanish when $J_B(t) \le 0$. The first two derivatives are nonnegative, strictly positive if and only if $q < \bar{b}^*$; the last two are nonpositive, strictly negative if and only if $t > \underline{s}^*$. Moreover, for every $q \in (0, \bar{b}^*]$ and $t \in [\underline{s}^*, 1)$,

$$0 < \Pi_B^B(q) \le G(q), \quad 0 < \Pi_S^B(q) \le G(q), \quad 0 < \Pi_B^S(t) \le 1 - F(t), \quad 0 < \Pi_S^S(t) \le 1 - F(t).$$

To prove this, fix v and let $\phi(q, v) = \max_{b \in [0,q]} (v - b)\,G(b)$. By the proof of Theorem 2, the objective is strictly quasiconcave with peak $b^M(v)$. Hence $\phi(q, v) = (v - q)G(q)$ for $q < b^M(v)$, with q-derivative $-G(q) + (v - q)g(q) = g(q)\big(v - J_S(q)\big) > 0$, and $\phi(q, v)$ is constant in q for $q \ge b^M(v)$; at the junction, where $v = J_S(q)$, the one-sided derivatives both vanish. Thus $\partial_q \phi(q, v) = g(q)\big(v - J_S(q)\big)^+$ for every (q, v), jointly continuous and bounded by $\|g\|_\infty < \infty$. Applying dominated convergence to the difference quotients of $\Pi_B^B(q) = \int_0^1 \phi(q, v)\,dF(v)$ yields the first formula, continuous in q; it is strictly positive if and only if $F\big(\{v : v > J_S(q)\}\big) > 0$, i.e. $J_S(q) < 1$, i.e. $q < \bar{b}^*$.

Integrating by parts, $\int_0^m (m-c)\,dG(c) = \int_0^m G(c)\,dc = L_G(m)$, so $\Pi_S^B(q) = \int_0^1 L_G\big(m_q(v)\big)\,dF(v)$ with $m_q(v) = \min\{b^M(v), q\}$. For every $v \neq J_S(q)$, the map $q \mapsto L_G(m_q(v))$ is differentiable at q with derivative $G(q)\,\mathbf{1}\{v > J_S(q)\}$; the exceptional type is F-null, and since $m_\cdot(v)$ and L_G are 1-Lipschitz, the difference quotients are bounded by 1. Dominated convergence yields the second formula, continuous in q, with the same strict-positivity characterization.

By the mirror integration by parts, $\int_s^1 (v-s)\,dF(v) = \int_s^1 \big(1-F(v)\big)dv = L_F(s)$, so $\Pi_B^S(t) = \int_0^1 L_F\big(\sigma_t(c)\big)\,dG(c)$ and $\Pi_S^S(t) = \int_0^1 \psi(t,c)\,dG(c)$, where $\psi(t,c) = \max_{s\in[t,1]}(s-c)\big(1-F(s)\big)$. The objective's s-derivative is $\big(1-F(s)\big) - (s-c)f(s) = f(s)\big[c - J_B(s)\big]$, strictly positive for $s < s^M(c)$ and strictly negative for $s > s^M(c)$ by regularity of J_B, so the maximizer over $[t,1]$ is $\sigma_t(c) = \max\{s^M(c), t\}$. For a floored type ($c < J_B(t)$), $\partial_t\psi(t,c) = f(t)\big[c - J_B(t)\big] < 0$; for an unfloored type, $\psi(\cdot,c)$ is locally constant; the one-sided derivatives vanish together at the junction $c = J_B(t)$. Both integrands have difference quotients bounded by $\max\{\|f\|_\infty, 1\}$, so dominated convergence yields the last two formulas, continuous in t; each is strictly negative if and only if the floored set $\{c : c < J_B(t)\}$ has positive G-mass, i.e. $t > \underline{s}^*$.

Finally, we provide the neccessary bounds. $\Pi_B^B(q) \geq G(q)\int_q^1 (v-q)\,dF(v) > 0$, since the bid q is feasible and $F\big((q,1]\big) > 0$ for $q < 1$; $\Pi_S^B(q) = \int L_G(m_q(v))\,dF(v) > 0$, since $m_q(v) > 0$ for $v > 0$ and $L_G > 0$ on $(0,1]$. For every $c < 1$, regularity gives $s^M(c) < 1$ (as $J_B(1) = 1$) and $s^M(c) > c$ (as $J_B(s) < s$ on $(0,1)$), hence $c < \sigma_t(c) < 1$, so the integrands of $\Pi_S^S(t)$ and $\Pi_B^S(t)$ are strictly positive for G-almost every c. The upper bounds follow from $\Pi_B^B(q) \leq \max_{b\leq q} G(b) = G(q)$, $\Pi_S^B(q) \leq L_G(q) \leq q\,G(q)$, $\Pi_S^S(t) \leq \max_{s\geq t}\big(1-F(s)\big) = 1-F(t)$, and $\Pi_B^S(t) \leq L_F(t) \leq (1-t)\big(1-F(t)\big)$.

Step 2: Absolute continuity and the fallback derivatives. The maps $q(R) = \min\{R, \bar{b}^*\}$ and $t(R) = \max\{R, \underline{s}^*\}$ are nondecreasing and 1-Lipschitz, differentiable except at the two kinks $\bar{b}^*$ and $\underline{s}^*$, with $q' \in \{0,1\}$ and $t' \in \{0,1\}$, and with the identities

$$\big(R - q(R)\big)\,q'(R) = 0 \qquad \text{and} \qquad \big(t(R) - R\big)\,t'(R) = 0 \tag{12}$$

wherever the derivatives exist, since $q' = 1$ only where $q = R$ and $t' = 1$ only where $t = R$. By Step 1 and the product and chain rules, all four losses are Lipschitz on compact subsets of $(0,1)$, hence absolutely continuous, and monotonicity may be read off almost-everywhere derivatives. Differentiating $X_B(R) = G\big(q(R)\big)\int_{t(R)}^1 (v-R)\,dF(v)$ and using (12) to remove

the term $-G(q)\big(t-R\big)f(t)\,t'$,

$$X_B'(R) = q'\,g(q)\int_t^1 (v-R)\,dF - G(q)\big(1-F(t)\big) \qquad \text{a.e.,} \tag{13}$$

and, symmetrically, removing $(R-q)\,g(q)\,q'$ from the differentiation of $X_S(R) = \big(1-F(t(R))\big)\int_0^{q(R)}(R-c)\,dG(c)$,

$$X_S'(R) = -t'\,f(t)\int_0^q (R-c)\,dG + \big(1-F(t)\big)G(q) \qquad \text{a.e.} \tag{14}$$

Step 3: The losses are monotone. By Step 1, (13), and the chain rule, almost everywhere

$$L_1'(R) = q'\,g(q)\Big[\mathbb{E}\big[(V-J_S(q))^+\big] - \int_t^1 (v-R)\,dF\Big] + G(q)\big(1-F(t)\big). \tag{15}$$

Where $q'(R)=0$, i.e. $R>\bar{b}^*$, $L_1'(R) = G(\bar{b}^*)\big(1-F(t)\big) > 0$, since $t(R)<1$. Where $q'(R)=1$, we have $q=R$, and we bound the fallback integral by the decomposition $v-R = \big(v-J_S(R)\big)+\big(J_S(R)-R\big)$:

$$\int_t^1 (v-R)\,dF = \int_t^1 \big(v-J_S(R)\big)dF + \big(J_S(R)-R\big)\big(1-F(t)\big) \le \mathbb{E}\big[(V-J_S(R))^+\big] + \frac{G(R)}{g(R)}\big(1-F(t)\big),$$

since the first integrand is bounded above by its positive part, the domain of a nonnegative integrand may be enlarged, and $J_S(R)-R = G(R)/g(R)$. Substituting into (15),

$$L_1'(R) \;\ge\; -\,g(R)\cdot\frac{G(R)}{g(R)}\big(1-F(t)\big) + G(R)\big(1-F(t)\big) \;=\; 0.$$

Hence L_1 is nondecreasing on $(0,1)$, strictly increasing on $(\bar{b}^*,1)$. The mirror computation for L_3, using (14) and the decomposition $R-c = \big(R-J_B(R)\big)+\big(J_B(R)-c\big)$ with $R-J_B(R) = \big(1-F(R)\big)/f(R)$ and

$$\int_0^q \big(J_B(R)-c\big)dG \le \int_0^1 \big(J_B(R)-c\big)^+ dG = \int_0^{J_B(R)} \big(J_B(R)-c\big)dG,$$

gives, almost everywhere, $L_3'(R)\le 0$ where $t'=1$, and $L_3'(R) = -\big(1-F(\underline{s}^*)\big)G\big(q(R)\big) < 0$ where $t'=0$, i.e. on $(0,\underline{s}^*)$: L_3 is nonincreasing, strictly decreasing on $(0,\underline{s}^*)$. Finally, L_2 is nondecreasing with $L_2'(R) = G(R)\big[1-F(J_S(R))\big] > 0$ almost everywhere on $(0,\bar{b}^*)$, and L_4 is nonincreasing with $L_4'(R) = -\big(1-F(R)\big)G\big(J_B(R)\big) < 0$ almost everywhere on $(\underline{s}^*,1)$, by

Step 1 and the chain rule.

Step 4: The losses are positive. By the proof of Proposition 5, $L_1(R) > 0$ and $L_3(R) > 0$ for every $R \in (0,1)$: the uniform coordination gap established there survives the ε-frontier limit with its slack. By Step 1, $L_2(R) > 0$ and $L_4(R) > 0$, since $q(R) > 0$ and $t(R) < 1$.

Step 5: ∇ is continuous and strictly increasing. Continuity is inherited from the losses. Since all four losses are nonnegative, with L_1, L_2 nondecreasing and L_3, L_4 nonincreasing, the product $L_1 L_2$ is nondecreasing and $L_3 L_4$ nonincreasing, so ∇ is nondecreasing. For strictness, fix $R < R'$ in $(0,1)$ and observe that $[R, R']$ meets at least one of the sets $(0, \bar{b}^*)$, $(\underline{s}^*, 1)$, and—when $\bar{b}^* \le \underline{s}^*$—the frozen interval $(\bar{b}^*, \underline{s}^*)$, in a set of positive length: the three sets cover $(0,1)$ up to at most two points, and when $\bar{b}^* > \underline{s}^*$ the first two alone cover. On $(0, \bar{b}^*) \cap [R, R']$, almost everywhere $(L_1 L_2)' \ge L_1 L_2' > 0$, since L_1 is bounded below on $[R, R']$ by its strictly positive minimum (continuity and Step 4). On $(\underline{s}^*, 1) \cap [R, R']$, almost everywhere $(L_3 L_4)' \le L_3 L_4' < 0$, likewise. On the frozen interval, q and t are constant, so almost everywhere $(L_1 L_2)' = L_1' L_2 = G(\bar{b}^*)\big(1 - F(\underline{s}^*)\big) L_2 > 0$. In every case $\nabla' = (L_1 L_2)' - (L_3 L_4)'$ is strictly positive on a subset of $[R, R']$ of positive measure and nonnegative almost everywhere else, so $\nabla(R') > \nabla(R)$ by absolute continuity.

Step 6: Boundary signs and the threshold. As $R \downarrow 0$: by Step 1, $0 < L_1(R) \le \Pi_B^B(q(R)) \le G(R) \to 0$ and $0 < L_2(R) \le G(R) \to 0$; meanwhile $t(R) = \underline{s}^*$, so $L_4(R) = \Pi_B^S(\underline{s}^*) > 0$ is constant, and since $0 \le X_S(R) \le \int_0^{q(R)} R \, dG \le R \to 0$, also $L_3(R) \to \Pi_S^S(\underline{s}^*) > 0$. Hence $\nabla(R) \to -\Pi_S^S(\underline{s}^*)\,\Pi_B^S(\underline{s}^*) < 0$, so $\nabla < 0$ near 0. Symmetrically, as $R \uparrow 1$: $L_3, L_4 \le 1 - F(R) \to 0$, $q(R) = \bar{b}^*$, $0 \le X_B(R) \le \int_{t(R)}^1 (v - R)\, dF \le 1 - R \to 0$, and $\nabla(R) \to \Pi_B^B(\bar{b}^*)\,\Pi_S^B(\bar{b}^*) > 0$: $\nabla > 0$ near 1. A continuous, strictly increasing function on $(0,1)$ that is negative near 0 and positive near 1 has exactly one zero $R^\dagger \in (0,1)$; by the sign convention of Definition 3, seller control is strictly risk-dominant on $(0, R^\dagger)$ and buyer control on $(R^\dagger, 1)$, proving the first statement.

Step 7: Reflection symmetry. We show that if $F(v) = 1 - G(1 - v)$, then $\underline{s}^* = 1 - \bar{b}^*$, $q(1 - R) = 1 - t(R)$, $t(1 - R) = 1 - q(R)$, and

$$\Pi_B^B(q) = \Pi_S^S(1 - q), \qquad \Pi_S^B(q) = \Pi_B^S(1 - q), \qquad X_B(1 - R) = X_S(R);$$

consequently $L_1(1 - R) = L_3(R)$ and $L_2(1 - R) = L_4(R)$.

It follows from $f(v) = g(1-v)$ and $1 - F(v) = G(1-v)$ that

$$J_B(v) = v - \frac{1-F(v)}{f(v)} = v - \frac{G(1-v)}{g(1-v)} = 1 - \Big[(1-v) + \frac{G(1-v)}{g(1-v)}\Big] = 1 - J_S(1-v).$$

So $J_B(v) = 0$ if and only if $J_S(1-v) = 1$, giving $\underline{s}^* = 1 - \bar{b}^*$, and the two frontier-bound identities follow from the definitions of q and t. The change of variables $v = 1-c$, $b = 1-s$ transforms the floored monopoly objective under (F, G) into the capped monopsony objective under the same pair: $(s-c)\big(1-F(s)\big) = \big((1-c)-(1-s)\big)\,G(1-s)$ with the constraint $s \geq t$ becoming $1-s \leq 1-t$, so $\psi(t,c) = \phi\big(1-t,\,1-c\big)$ pointwise; integrating against $dG(c) = dF(1-c)$ yields $\Pi_S^S(t) = \Pi_B^B(1-t)$. The same substitution in the taker integrals gives $L_F(s) = L_G(1-s)$ and hence $\Pi_B^S(t) = \Pi_S^B(1-t)$. For the fallbacks, $G\big(q(1-R)\big) = G\big(1-t(R)\big) = 1 - F\big(t(R)\big)$ and, substituting $v = 1-c$,

$$\int_{t(1-R)}^{1} \big(v - (1-R)\big)\,dF(v) = \int_{1-q(R)}^{1} \big(v - 1 + R\big)\,dF(v) = \int_{0}^{q(R)} \big(R - c\big)\,dG(c),$$

so $X_B(1-R) = X_S(R)$. The loss identities follow by assembling: $\nabla(1-R) = L_1(1-R)L_2(1-R) - L_3(1-R)L_4(1-R) = L_3(R)L_4(R) - L_1(R)L_2(R) = -\nabla(R)$, so ∇ is antisymmetric about $\frac{1}{2}$; its unique zero is therefore $R^\dagger = \frac{1}{2}$. □

A.14 Proof of Theorem 5

Proof. For $q \in (0, \bar{b}^*]$ and $t \in [\underline{s}^*, 1)$, write again $\Pi_B^B(q), \Pi_S^B(q), \Pi_B^S(t), \Pi_S^S(t)$ for the payoffs with cap q and floor t, and define the *stakes* $\Delta_B(R) = \Pi_B^B(q(R)) - \Pi_B^S(t(R))$ and $\Delta_S(R) = \Pi_S^B(q(R)) - \Pi_S^S(t(R))$, each side's ex-ante gain from buyer control over seller control.

Where a frontier is constrained, the regime payoffs are the ε-frontier limits, which equal the displayed formulas evaluated at $q(R)$ and $t(R)$; the identities below therefore apply verbatim at every reference.

Step 1: Regularity of the payoff maps. We start with the following claim:

Claim 6. *The four payoff maps are continuously differentiable, with*

$$\big(\Pi_B^B\big)'(q) = g(q)\,\mathbb{E}\big[(V - J_S(q))^+\big], \qquad \big(\Pi_S^B\big)'(q) = G(q)\big[1 - F\big(J_S(q)\big)\big],$$
$$\big(\Pi_B^S\big)'(t) = -\big[1 - F(t)\big]\,G\big(J_B(t)\big), \qquad \big(\Pi_S^S\big)'(t) = -f(t)\int_0^{J_B(t)} \big[J_B(t) - c\big]\,dG(c),$$

with the convention that integrals over $[0, J_B(t)]$ *vanish when* $J_B(t) \leq 0$. *The first two derivatives are nonnegative, strictly positive if and only if* $q < \bar{b}^*$; *the last two are nonpositive, strictly negative if and only if* $t > \underline{s}^*$. *Moreover, for every* $q \in (0, \bar{b}^*]$ *and* $t \in [\underline{s}^*, 1)$,

$$0 < \Pi_B^B(q) \leq G(q), \quad 0 < \Pi_S^B(q) \leq G(q), \quad 0 < \Pi_B^S(t) \leq 1 - F(t), \quad 0 < \Pi_S^S(t) \leq 1 - F(t).$$

Proof of Claim 6. We distinguish between the buyer and the seller regime.

- **Buyer regime.** Fix v and let $\phi(q, v) = \max_{b \in [0,q]} (v - b)\, G(b)$. By the proof of Theorem 2, the objective is strictly quasiconcave with peak $b^M(v)$. Hence $\phi(q, v) = (v - q)G(q)$ for $q < b^M(v)$, with q-derivative $-G(q) + (v - q)g(q) = g(q)\big(v - J_S(q)\big) > 0$, and $\phi(q, v)$ is constant in q for $q \geq b^M(v)$; at the junction, where $v = J_S(q)$, the one-sided derivatives both vanish. Thus $\partial_q \phi(q, v) = g(q)\big(v - J_S(q)\big)^+$ for every (q, v), jointly continuous and bounded by $\|g\|_\infty < \infty$. Applying dominated convergence to the difference quotients of $\Pi_B^B(q) = \int_0^1 \phi(q, v)\, dF(v)$ yields the first formula, and its continuity in q is inherited from that of g, J_S, and $q \mapsto \mathbb{E}[(V - J_S(q))^+]$. It is strictly positive if and only if $F\big(\{v : v > J_S(q)\}\big) > 0$, i.e. $J_S(q) < 1$, i.e. $q < \bar{b}^*$. Integrating by parts, $\int_0^m (m - c)\, dG(c) = \int_0^m G(c)\, dc = L_G(m)$, so $\Pi_S^B(q) = \int_0^1 L_G\big(m_q(v)\big)\, dF(v)$ with $m_q(v) = \min\{b^M(v), q\}$. For every $v \neq J_S(q)$, the map $q \mapsto L_G(m_q(v))$ is differentiable at q with derivative $G(q)\, \mathbf{1}\{v > J_S(q)\}$; the exceptional type is F-null, and since $m_\cdot(v)$ and L_G are 1-Lipschitz, the difference quotients are bounded by 1. Dominated convergence yields the second formula, continuous in q, with the same strict-positivity characterization.

- **Seller regime.** By the mirror integration by parts, $\int_s^1 (v - s)\, dF(v) = \int_s^1 \big(1 - F(v)\big) dv = L_F(s)$, so $\Pi_B^S(t) = \int_0^1 L_F\big(\sigma_t(c)\big)\, dG(c)$ and $\Pi_S^S(t) = \int_0^1 \psi(t, c)\, dG(c)$, where $\psi(t, c) = \max_{s \in [t,1]} (s - c)\big(1 - F(s)\big)$. The objective's s-derivative is $\big(1 - F(s)\big) - (s - c) f(s) = f(s)\big[c - J_B(s)\big]$, strictly positive for $s < s^M(c)$ and strictly negative for $s > s^M(c)$ by regularity of J_B; hence the maximizer over $[t, 1]$ is $\sigma_t(c) = \max\{s^M(c), t\}$. For a floored type ($c < J_B(t)$), $\partial_t \psi(t, c) = f(t)\big[c - J_B(t)\big] < 0$; for an unfloored type, $\psi(\cdot, c)$ is locally constant; the one-sided derivatives vanish together at the junction $c = J_B(t)$. Both integrands have difference quotients bounded by $\max\{\|f\|_\infty, 1\}$, so dominated convergence yields the two formulas, continuous in t; each is strictly negative if and only if the floored set $\{c : c < J_B(t)\}$ has positive G-mass, i.e. $J_B(t) > 0$, i.e. $t > \underline{s}^*$.

We finally prove the bounds in Claim 6.

- **Lower bound:** $\Pi_B^B(q) \geq G(q)\int_q^1 (v-q)\,dF(v) > 0$, since the bid q is feasible and $F\big((q,1]\big) > 0$ for $q < 1$; $\Pi_S^B(q) = \int L_G(m_q(v))\,dF(v) > 0$, since $m_q(v) > 0$ for every $v > 0$ and $L_G > 0$ on $(0,1]$. On the seller side, for every $c < 1$ we have $s^M(c) < 1$ (as $J_B(1) = 1$) and $s^M(c) > c$ (as $J_B(s) < s$ on $(0,1)$), hence $c < \sigma_t(c) < 1$; thus the integrand of $\Pi_S^S(t)$ is strictly positive for G-almost every c, and $L_F\big(\sigma_t(c)\big) > 0$ likewise, giving $\Pi_S^S(t), \Pi_B^S(t) > 0$ for $t < 1$.

- **Upper bound:** $\Pi_B^B(q) \leq \max_{b\leq q} G(b) \leq G(q)$ and $\Pi_S^B(q) \leq L_G(q) \leq q\,G(q) \leq G(q)$; $\Pi_S^S(t) \leq \max_{s\geq t}\big(1-F(s)\big) \leq 1-F(t)$ and $\Pi_B^S(t) \leq L_F(t) \leq (1-t)\big(1-F(t)\big) \leq 1-F(t)$.

Step 2: Both stakes are negative near 0 **and positive near** 1**.** By Claim 6 and the chain rule, composed with the nondecreasing 1-Lipschitz maps $q(\cdot)$ and $t(\cdot)$, the stakes Δ_B and Δ_S are continuous on $(0,1)$. As $R \downarrow 0$: $q(R) = R \to 0$, so by the upper bounds of Claim 6, $\Pi_B^B\big(q(R)\big)$ and $\Pi_S^B\big(q(R)\big)$ are bounded by $G(R) \to 0$; meanwhile $t(R) = \underline{s}^*$ for every $R \leq \underline{s}^*$, so the seller-regime payoffs are constant at $\Pi_B^S(\underline{s}^*) > 0$ and $\Pi_S^S(\underline{s}^*) > 0$. Hence

$$\lim_{R\downarrow 0} \Delta_B(R) = -\,\Pi_B^S(\underline{s}^*) < 0, \qquad \lim_{R\downarrow 0} \Delta_S(R) = -\,\Pi_S^S(\underline{s}^*) < 0,$$

and by continuity there exists $r \in (0,1)$ with $\Delta_B < 0$ and $\Delta_S < 0$ on $(0,r)$: both sides earn strictly more under seller control, so seller control strictly payoff-dominates there (Definition 4). Symmetrically, as $R \uparrow 1$: $t(R) = R \to 1$, so the seller-regime payoffs are bounded by $1 - F(R) \to 0$, while $q(R) = \bar{b}^*$ for every $R \geq \bar{b}^*$, so the buyer-regime payoffs are constant at $\Pi_B^B(\bar{b}^*) > 0$ and $\Pi_S^B(\bar{b}^*) > 0$; hence both stakes have strictly positive limits at 1, and there exists $\bar{r} \in (0,1)$ with $\Delta_B > 0$ and $\Delta_S > 0$ on $(\bar{r},1)$: buyer control strictly payoff-dominates there. Finally, no reference lies in both regions, since the two dominance configurations are mutually exclusive; hence $r \leq \bar{r}$, and replacing r by $r/2$ gives $r < \bar{r}$.

Step 3: Welfare superiority. By quasilinearity, the expected gains from trade of a regime equal the sum of the two sides' ex-ante payoffs, so at every reference

$$W_B\big(q(R)\big) - W_S\big(t(R)\big) = \Delta_B(R) + \Delta_S(R).$$

For $R < r$ the right side is strictly negative and for $R > \bar{r}$ strictly positive: in each case, the payoff-dominant regime also generates strictly higher expected gains from trade. □

A.15 Proof of Theorem 6

Proof of Theorem 6 Part (i). Fix $R \in (0,1)$; where a frontier is constrained, the regime payoffs are the ε-frontier limits, which equal the formulas (Buyer Control: Buyer / Seller Payoff)–(Seller Control: Buyer / Seller Payoff) evaluated at $q = q(R) \in (0, \bar{b}^*]$ and $t = t(R) \in [\underline{s}^*, 1)$, so it suffices to prove the two strict inequalities at every such q and t. As the ε-frontier payoffs converge to these values, the theorem's inequalities hold at every reference.

Step 1: The buyer regime. By the parts-integration representation in the proof of Theorem 4, $\Pi_S^B(q) = \int_0^1 L_G\big(m_q(v)\big)\, dF(v)$ with $m_q(v) = \min\{b^M(v), q\}$, so

$$\Pi_B^B(q) - \Pi_S^B(q) = \int_0^1 \Big[\big(v - m_q(v)\big)\, G\big(m_q(v)\big) - L_G\big(m_q(v)\big)\Big]\, dF(v),$$

and we show the integrand is strictly positive for every $v \in (0,1]$. If $v \le J_S(q)$ (a separating type), then $m_q(v) = b^M(v) =: b \in (0,q]$, and the first-order condition $v = J_S(b)$ gives $v - b = G(b)/g(b)$, so the integrand equals

$$\frac{G(b)^2}{g(b)} - L_G(b) \;>\; 0$$

by integrated log-concavity at $b \in (0, \bar{b}^*]$. If $v > J_S(q)$ (a capped type), then $m_q(v) = q$ and

$$(v-q)\, G(q) \;>\; \big(J_S(q) - q\big)\, G(q) = \frac{G(q)^2}{g(q)} \;>\; L_G(q),$$

the first inequality strict because $v > J_S(q)$ and $G(q) > 0$, the second by integrated log-concavity at q. The remaining type $v = 0$ is F-null. Hence $\Pi_B^B(q) > \Pi_S^B(q)$ for every $q \in (0, \bar{b}^*]$, and in particular at $q(R)$.

Step 2: The seller regime. By the mirror representation, $\Pi_B^S(t) = \int_0^1 L_F\big(\sigma_t(c)\big)\, dG(c)$ with $\sigma_t(c) = \max\{s^M(c), t\}$, so

$$\Pi_S^S(t) - \Pi_B^S(t) = \int_0^1 \Big[\big(\sigma_t(c) - c\big)\big(1 - F(\sigma_t(c))\big) - L_F\big(\sigma_t(c)\big)\Big]\, dG(c),$$

with $\sigma_t(c) \in [\underline{s}^*, 1)$ for every $c < 1$, since $s^M(c) \in [\underline{s}^*, 1)$ and $t \ge \underline{s}^*$. For an unfloored type ($c \ge J_B(t)$), $\sigma_t(c) = s^M(c) =: s$, and the first-order condition $c = J_B(s)$ gives $s - c =$

$\big(1-F(s)\big)/f(s)$, so the integrand equals

$$\frac{\big(1-F(s)\big)^2}{f(s)} - L_F(s) \;>\; 0$$

by integrated log-concavity at $s \in [\underline{s}^*, 1)$. For a floored type ($c < J_B(t)$), $\sigma_t(c) = t$ and

$$(t-c)\big(1-F(t)\big) \;>\; \big(t - J_B(t)\big)\big(1-F(t)\big) = \frac{\big(1-F(t)\big)^2}{f(t)} \;>\; L_F(t),$$

the first inequality strict because $c < J_B(t)$ and $F(t) < 1$, the second by integrated log-concavity at t. The remaining type $c = 1$ is G-null. Hence $\Pi_S^S(t) > \Pi_B^S(t)$ for every $t \in [\underline{s}^*, 1)$, and in particular at $t(R)$. □

Proof of Theorem 6 Part(ii). We first explain the intuition behind what must be engineered.

Step -1: Intuition. By Theorem 6, integrated log-concavity forces the maker to outearn the takers, so a burden market must violate it decisively: at the bids the maker actually uses, her *marginal* rent G/g must be far smaller than the takers' *average* inframarginal rent, L_G/G per trade. The construction achieves this through pass-through. We build a supply schedule on which the virtual cost J_S rises with a small constant slope κ over the maker's active range: there, every unit increase in the buyer's value moves her optimal bid up by $1/\kappa$ units. Competition against her own inframarginal book forces her to bid aggressively into a thickening mass of sellers, pinning her markdown at the small constant $\varphi = G/g \approx \delta$, while the seller's types she trades with retain integrated surplus $L_G \approx \delta/\kappa$: the taker's types collect $1/\kappa$ times what the maker's types keep. Around this engine we place the population masses: almost all seller's types in a narrow bump at a high cost $\bar{c}$, so that the buyer-side maker prices only against a thin low-cost fringe, and almost all buyer's types—by reflection—at a value v^* just above the pass-through stretch, so that the bulk of the buyer side actually uses it. Reflection symmetry then makes the seller regime the exact mirror and, by the symmetry part of the proof of Theorem 4, places the threshold at $R^\dagger = \frac{1}{2}$; the estimates are carried out at the balanced reference and extended to an open interval by continuity.

Step 0: Notation. Fix $\kappa \in (0,1)$ with $1/\kappa > 2M$ and $b^* \in (0, \frac{1}{2})$, and let $\delta \downarrow 0$ be the construction parameter. Set $v^* = b^* + \delta$, $\bar{c} = 1 - v^*$, $\omega = \delta^2$, $\varepsilon' = \delta^4$, $b^+ = b^* + 2\omega/\kappa$, and fix

$$h \in \Big(0,\; \min\Big\{\tfrac{\bar{c}-b^*}{4},\; \tfrac{1+\kappa}{2}\big(\bar{c} - \tfrac{1}{2}\big)\Big\}\Big).$$

All order symbols below refer to $\delta \downarrow 0$ at fixed (κ, b^*, h), and "for small δ" means for all δ below a threshold $\delta_0(\kappa, b^*, h)$.

Claim 7. *Let $\varphi : (0,1) \to (0,\infty)$ be continuous and piecewise C^1 with $\varphi' \geq -1+\zeta$ almost everywhere for some $\zeta > 0$, with $\int_0 \varphi^{-1} = \infty$, and with $\varphi(c) = c$ on a right neighborhood of 0. Then*

$$G(c) = \exp\Big(-\int_c^1 \frac{du}{\varphi(u)}\Big)$$

is a strictly regular distribution on $[0,1]$ with continuous, strictly positive density $g = G/\varphi$, satisfying $G/g = \varphi$ and $J_S = \mathrm{id} + \varphi$, so that $J_S' \geq \zeta$ almost everywhere.

Proof of Claim 7. G is strictly increasing with $G(1) = 1$ and $G(0^+) = 0$ by the divergent integral, and $g = G' = G/\varphi > 0$ is continuous on $(0,1)$ with $\int_0^1 g = G(1) - G(0^+) = 1$. Near 0, $\varphi(c) = c$ gives $(\log G)' = 1/c$, so G is linear there, $G(c) = Kc$ with $K > 0$; hence $g \equiv K$ near 0 and extends continuously with $g(0) = K > 0$. Near 1, φ is bounded away from zero, so $g(1) = 1/\varphi(1) > 0$. Finally $J_S' = 1 + \varphi' \geq \zeta > 0$ almost everywhere and J_S is continuous, hence strictly increasing: the pair is strictly regular. $\triangle$

The rent function $\varphi = G/g$ is thus a free design object: any positive profile with slopes above -1 is realizable, and $J_S' = 1 + \varphi'$ translates its slopes into pass-through. We now specify the construction.

Step 1: The construction. Let $c_2 = b^*/2$, $\varphi(c_2) = \delta + (1-\kappa)\, b^*/2$, and $c_1 = \min\{b^*/4,\ \varphi(c_2)\}$. Define φ piecewise:

(z1) $\varphi(c) = c$ on $(0, c_1]$: a uniform-like base, guaranteeing $g(0) > 0$ and the divergent integral;

(z2) affine from (c_1, c_1) to $(c_2, \varphi(c_2))$, of nonnegative slope by the choice of c_1: a neutral join;

(z3) the *pass-through stretch* $\varphi(c) = \delta + (1-\kappa)(b^* - c)$ on $[c_2, b^+]$: here $J_S' \equiv \kappa$, with $J_S(b^*) = v^*$, $J_S(b^+) = v^* + 2\omega$, and $\varphi(b^+) = \delta\big(1 - 2(1-\kappa)\delta/\kappa\big) \geq \delta/2$ for small δ—the maker's rents are pinned near δ throughout, while the stretch is long enough (b^+ exceeds b^* by $2\omega/\kappa$) that every bulk buyer's optimum lands strictly inside it;

(z4) affine from $\varphi(b^+)$ to h on $[b^+, b^+ + \delta]$; the plateau $\varphi \equiv h$ on $[b^+ + \delta, c_3]$; and an affine descent at slope $-\frac{1+\kappa}{2}$ from h to ε' on $[c_3, \bar{c}-\omega]$, with $c_3 = \bar{c}-\omega-2(h-\varepsilon')/(1+\kappa)$: dead space carrying almost no mass, positioned—by the choice of h—so that $\frac{1}{2} \in (b^+ + \delta, c_3)$ for small δ;

(z5) $\varphi \equiv \varepsilon'$ on $[\bar{c} - \omega, \bar{c} + \omega]$: a tiny rent function forces $(\log G)' = 1/\varepsilon'$ to be huge, concentrating essentially all seller mass in this bump;

(z6) affine from ε' to $1/\delta$ on $[\bar{c}+\omega, \bar{c}+\omega+\delta]$, then $\varphi \equiv 1/\delta$ up to 1: a thin tail whose steep φ places $\bar{b}^*$ just above the bump.

Every descent has slope $-(1-\kappa)$ or $-\frac{1+\kappa}{2}$, so $\varphi' \geq -1+\zeta$ almost everywhere with $\zeta = \min\{\kappa, \frac{1-\kappa}{2}\} > 0$, and Claim 1 applies. Let G be the induced distribution, $F(v) = 1-G(1-v)$ its reflection, so the pair is strictly regular and reflection-symmetric; write $\varepsilon_\delta = G(b^*)$ for the (small) mass of the low-cost fringe.

Step 3: Masses, frontiers, and the execution bound. We start with the following claims about our construction:

Claim 8. *For all small δ:*

(a) $G(\bar{c}-\omega) \leq e^{-2/\delta^2}$ and $1 - G(\bar{c}+\omega) \leq C\delta$; by reflection, buyer mass at least $1 - C\delta$ lies in $[v^* - \omega, v^* + \omega]$, mass at most $C\delta$ below it, and at most e^{-2/δ^2} above it;

(b) $\bar{b}^* \in (\bar{c}+\omega,\, \bar{c}+\omega+\delta)$; hence $\bar{b}^* > \frac{1}{2} > 1 - \bar{b}^* = \underline{s}^*$, and $q(R) = t(R) = R$ for every $R \in (\underline{s}^*, \bar{b}^*)$, in particular at $\frac{1}{2}$;

(c) $\sup_{b \leq 1/2} G(b) = G(\frac{1}{2}) \leq C_0\, \varepsilon_\delta$ with $C_0 = e^{1/(2h)+1}$.

Proof of Claim 8. (a) $G(\bar{c}-\omega) \leq \exp\big(-\int_{\bar{c}-\omega}^{\bar{c}+\omega} du/\varepsilon'\big) = e^{-2\omega/\varepsilon'} = e^{-2/\delta^2}$. For the upper tail, $1 - e^{-y} \leq y$ gives $1 - G(\bar{c}+\omega) \leq \int_{\bar{c}+\omega}^{1} \varphi^{-1}$; the (z6) rise is affine with slope $(1/\delta - \varepsilon')/\delta$, so its contribution is $\frac{\delta}{1/\delta - \varepsilon'} \log \frac{1/\delta}{\varepsilon'} \leq 6\delta^2 \log \frac{1}{\delta} \leq \delta$ for small δ, and the constant tail contributes at most $\delta \cdot (1-\bar{c}) \leq \delta$; take $C = 2$. The buyer statements are the reflection $F(v) = 1 - G(1-v)$ applied at $v^* \mp \omega = 1 - (\bar{c} \pm \omega)$. (b) On $(0, \bar{c}+\omega]$, $J_S(c) = c + \varphi(c) \leq \bar{c} + \omega + \max\{\varepsilon', h\} < 1$ for small δ (using $h < \bar{c} - b^* \leq \bar{c} < 1$ and the zone bounds), while $J_S(\bar{c}+\omega+\delta) \geq 1/\delta > 1$; since J_S is strictly increasing and continuous, its crossing of 1—which defines $\bar{b}^*$—lies in the (z6) rise. As $\delta \downarrow 0$, $\bar{b}^* \to 1 - b^* > \frac{1}{2}$, and reflection gives $\underline{s}^* = 1 - \bar{b}^*$; the band identities $q(R) = \min\{R, \bar{b}^*\} = R$ and $t(R) = \max\{R, \underline{s}^*\} = R$ hold on $(\underline{s}^*, \bar{b}^*)$. (c) $G(\frac{1}{2}) = \varepsilon_\delta \exp\big(\int_{b^*}^{1/2} \varphi^{-1}\big)$, and the exponent splits into the stretch tail $[b^*, b^+]$, where $\varphi \geq \delta/2$ over

length $2\omega/\kappa$, contributing at most $4\delta/\kappa$; the (z4) rise, affine from $\varphi(b^+) \geq \delta/2$ to h over length δ, contributing at most $\frac{\delta}{h}\log\frac{2h}{\delta} \leq \frac{2\delta}{h}\log\frac{1}{\delta}$; and the plateau up to $\frac{1}{2}$, contributing at most $1/(2h)$. For small δ the total is below $1/(2h)+1$. $\triangle$

Claim 8 already shows why the reference $\frac{1}{2}$ is balanced and constrained at once: it lies strictly between $\underline{s}^*$ and $\bar{b}^*$, so both frontier bounds equal the reference itself, and all payoffs at $\frac{1}{2}$ are understood in the ε-frontier sense; the estimates below are uniform over caps in $[\frac{1}{2}-\varepsilon, \frac{1}{2}]$ for $\varepsilon < \frac{1}{2} - b^+ - \delta$, so the limits equal the displayed values.

Step 4: Payoffs at the balanced reference. We next bound the payoffs around the balanced reference price.

Claim 9. *At $R = \frac{1}{2}$,*

$$\Pi_B^B = \varepsilon_\delta\,\delta\,\big(1+O(\delta)\big), \qquad \Pi_S^B \geq \frac{\varepsilon_\delta\,\delta}{\kappa}\,\big(1-o(1)\big).$$

Proof of Claim 9. On the stretch, integrating $(\log G)' = \varphi^{-1}$ in closed form,

$$G(c) = \varepsilon_\delta\Big(\frac{\delta}{\delta+(1-\kappa)(b^*-c)}\Big)^{1/(1-\kappa)}, \qquad \varphi(c)\,G(c) = \varepsilon_\delta\,\delta^{\frac{1}{1-\kappa}}\big(\delta+(1-\kappa)(b^*-c)\big)^{-\frac{\kappa}{1-\kappa}},$$

so φG increases along the stretch and is maximized at its top, where it equals $\varepsilon_\delta\delta\big(1+O(\delta)\big)$; and since the antiderivative of G on the stretch is $\frac{1}{\kappa}\varphi G$, evaluating at the endpoints and using that the zones below the stretch contribute $\int_0^{c_2} G \leq G(c_2) = \varepsilon_\delta\,O\big(\delta^{1/(1-\kappa)}\big)$,

$$L_G(b) = \frac{\varepsilon_\delta\,\delta}{\kappa}\,\big(1+o(1)\big) \qquad \text{uniformly over } b \in [\,b^*-\omega/\kappa,\ b^+\,].$$

- A *bulk* buyer $v \in [v^*-\omega, v^*+\omega]$ solves $J_S(b) = v$ inside the stretch, at $b(v) = b^* + (v-v^*)/\kappa \in [b^*-\omega/\kappa,\, b^+]$; because J_S is globally strictly increasing and $b(v) < \frac{1}{2}$, this is her unique maximizer under every cap in $[\frac{1}{2}-\varepsilon, \frac{1}{2}]$, and she earns $\big(v-b(v)\big)G(b(v)) = \varphi(b(v))\,G(b(v)) = \varepsilon_\delta\delta\big(1+O(\delta)\big)$, while the sellers she trades with collect $L_G(b(v)) = \frac{\varepsilon_\delta\delta}{\kappa}\big(1+o(1)\big)$. This is the pass-through mechanism in one line: her margin is the rent function at her bid, pinned near δ, while the fringe's integrated surplus is the same quantity divided by κ.

- An *exceptional* buyer $v < v^*-\omega$ earns at most $C_1\varepsilon_\delta\delta$ from any feasible bid, with $C_1 = \max\{2/(1-\kappa)+1,\, C_0\}$: for bids below the stretch, $(v-b)G(b) \leq G(c_2) = \varepsilon_\delta\,O(\delta^{1/(1-\kappa)}) \leq \varepsilon_\delta\delta$; for bids on the stretch with $b \leq b^*$, $v-b < \delta + (b^*-b) \leq \varphi(b)/(1-\kappa)$,

so $(v-b)G(b) \le \frac{1+O(\delta)}{1-\kappa}\,\varepsilon_\delta\delta$; for $b \in (b^*, b^+]$, $v - b < \delta \le 2\varphi(b)$, so $(v-b)G(b) \le 2\varepsilon_\delta\delta\big(1+O(\delta)\big)$; and for $b \in (b^+, \frac12]$, $v-b<\delta$ and $(v-b)G(b) \le \delta\, G(\frac12) \le C_0\varepsilon_\delta\delta$ by Claim 2(c). Buyers above $v^*+\omega$, of mass at most e^{-2/δ^2}, earn at most $G(\frac12) \le C_0\varepsilon_\delta$.

Assembling with the masses of Claim 8(a): the bulk, of mass $1-O(\delta)$, contributes $\varepsilon_\delta\delta\big(1+O(\delta)\big)$ to Π_B^B and at least $\big(1-O(\delta)\big)\frac{\varepsilon_\delta\delta}{\kappa}\big(1-o(1)\big)$ to Π_S^B; the exceptional mass $O(\delta)$ contributes $O(\varepsilon_\delta\delta^2)$; and the top mass contributes $\varepsilon_\delta\, e^{-2/\delta^2} = o(\varepsilon_\delta\delta)$. △

Step 5: Burden on an open interval. By reflection symmetry and Claim 2 of the proof of Theorem 4, at $R=\frac12$ (where $q=t=\frac12$) we have $\Pi_S^S = \Pi_B^B$ and $\Pi_B^S = \Pi_S^B$. Hence, by Claim 3, for all small δ,

$$\Pi_S^B - \Pi_B^B \;\ge\; \varepsilon_\delta\,\delta\left[\tfrac{1}{\kappa}\big(1-o(1)\big) - \big(1+o(1)\big)\right] \;>\; 0, \qquad \frac{\Pi_S^B}{\Pi_B^B} \;\ge\; \frac{1}{\kappa}\big(1-o(1)\big) \;>\; M,$$

using $1/\kappa > 2M$; and by the reflection identities the same strict inequality and the same ratio hold in the seller regime, $\Pi_B^S > \Pi_S^S$ with $\Pi_B^S/\Pi_S^S = \Pi_S^B/\Pi_B^B$. Fix such a δ. The four payoff maps are continuous (Claim 1 in the proof of Theorem 4) and $q(R)=t(R)=R$ near $\frac12$ by Claim 2(b), so both strict inequalities, and the ratio bound with M in place of $\frac{1}{\kappa}(1-o(1))$, persist on an open interval $I \ni \frac12$. The pair is regular, so Theorem 4 applies, and reflection symmetry gives $R^\dagger = \frac12 \in I$.

Step 6: Smoothing. In a final step, we need to smooth the constructed distributions and show that the same bounds still apply.

Claim 10. *There is a C^∞ profile φ_r, coinciding with φ near 0 and 1, with $\varphi_r' \ge -1+\zeta$, such that the induced reflection-symmetric pair (F_r, G_r) has C^∞ densities, is strictly regular, and satisfies every strict inequality of Step 5 on a compact subinterval $K \subset I$ with $\frac12 \in \operatorname{int} K$, for all small r.*

Proof of Claim 10. All kinks of φ lie in a compact subset of $(0,1)$; mollify at scale r smaller than the distance from that compact to $\{0,1\}$, so that $\varphi_r = \varphi * \rho_r$ is C^∞, coincides with φ near the endpoints (where φ is affine), is positive for small r, and satisfies $\varphi_r' = \varphi' * \rho_r \ge -1+\zeta$, a continuous piecewise-C^1 function having no singular part. Claim 1 applies to φ_r; set $F_r = 1 - G_r(1-\cdot)$, preserving reflection symmetry exactly. Uniform convergence $\varphi_r \to \varphi$ gives $J_S^{(r)} \to J_S$ uniformly with slopes bounded below by ζ, hence uniform convergence of the inverses, the strategies, and (via the exact common tail near 0 and uniform convergence

of $\int_c^1 \varphi_r^{-1}$ on compacts of $(0,1]$) of $G_r \to G$ and $F_r \to F$ on $[0,1]$. Every payoff map is a bounded-integrand functional of these objects, so the six payoffs converge uniformly over the relevant compact sets of (q,t), and $\bar{b}^{*(r)} \to \bar{b}^*$. The finitely many strictly positive quantities of Step 5—the two minimal payoff gaps over K and the ratio margin at $\frac{1}{2}$—therefore remain strictly positive for all small r, and $R^{\dagger(r)} = \frac{1}{2} \in K$ by symmetry. Near the endpoints the densities are exactly the piecewise forms of the unsmoothed pair (g_r constant near 0, exponential near 1), so the densities are C^∞ on $[0,1]$. △

For small r, the pair (F_r, G_r) and the interval K satisfy every assertion, which proves part (ii) of Theorem 6. □

A.16 Proof of Theorem 7

Proof. We start by fixing the notation.

Step 0: Framework, conventions, and constants. Let $(V_t, C_t)_{t\geq 1}$ be i.i.d. with law $F \otimes G$ on a probability space carrying the filtration $\mathcal{F}_t = \sigma\big((V_s, C_s)_{s\leq t}\big)$, $\mathcal{F}_0$ trivial; the reference R_t is then $\mathcal{F}_{t-1}$-measurable by the update rule. Play in period t is constructed as follows. The controlling side is $\xi_t = B$ if $R_t > R^\dagger$ and $\xi_t = S$ if $R_t < R^\dagger$, by the risk-dominant equilibrium selection; at $R_t = R^\dagger$, ξ_t is given by the tie rule, an arbitrary $\mathcal{F}_{t-1}$-measurable choice. Thus $\xi_t = B$ only if $R_t \geq R^\dagger$, and $\xi_t = S$ only if $R_t \leq R^\dagger$. If $\xi_t = B$, play is a truncated monopsony equilibrium of Theorem 2, determined by its *cap* κ_t:

$$\kappa_t = \bar{b}^* \ \text{ if } R_t \geq \bar{b}^* \qquad \text{(the frontier cap, admissible here)},$$

and, if $R_t < \bar{b}^*$, so that the frontier is inadmissible, κ_t is an $\mathcal{F}_{t-1}$-measurable choice from $[R_t - \varepsilon, R_t)$, the ε-frontier convention. Buyers bid $\min\{b^M(V_t), \kappa_t\}$, sellers with $C_t < \kappa_t$ ask truthfully and all others ask 1; off null sets, a trade occurs iff $C_t \leq \min\{b^M(V_t), \kappa_t\}$, at the bid. If $\xi_t = S$, play is the mirror truncated monopoly equilibrium, determined by its *floor* φ_t: $\varphi_t = \underline{s}^*$ if $R_t \leq \underline{s}^*$, and otherwise an $\mathcal{F}_{t-1}$-measurable choice from $(R_t, R_t + \varepsilon]$; sellers ask $\max\{s^M(C_t), \varphi_t\}$, buyers with $V_t > \varphi_t$ bid truthfully and all others bid 0, and a trade executes at the ask. A *convention* is the triple of adapted rules just described—the tie rule and the two ε-frontier selections—and ε is fixed with

$$\varepsilon \ \leq \ \varepsilon_0 := \tfrac{1}{2} \min\big\{R^\dagger,\ 1 - R^\dagger,\ 1 - \underline{s}^*\big\}.$$

Every print is a bid $\min\{b^M(V_t), \kappa_t\} > 0$ or an ask $\max\{s^M(C_t), \varphi_t\} < 1$ for a.e. type, so $R_t \in (0,1)$ for all t almost surely and the construction is well defined along the entire path. Define the constants

$$\kappa_0 := \min\{\bar{b}^*,\ R^\dagger - \varepsilon_0\}, \qquad s_1 := \max\{\underline{s}^*,\ R^\dagger\} + \varepsilon_0,$$

$$\alpha_B := \int_0^{J_S(\kappa_0)} G\big(b^M(v)\big)\, dF(v), \qquad \alpha_S := \int_{J_B(s_1)}^1 \big[1 - F\big(s^M(c)\big)\big]\, dG(c), \qquad a := \min\{\alpha_B, \alpha_S\}.$$

Here $\kappa_0 > 0$ since $\varepsilon_0 \le R^\dagger/2$, and $s_1 < 1$ since $\varepsilon_0 \le \frac{1}{2}\min\{1 - R^\dagger, 1 - \underline{s}^*\}$; hence $J_S(\kappa_0) > 0$ and $J_B(s_1) < 1$, both integrals are strictly positive, $\kappa_0 \le R^\dagger$, and $s_1 > R^\dagger$. All three constants depend only on (F, G).

Step 1: Print geometry. Suppose $\xi_t = B$ and a trade occurs. If the frontier is unconstrained, $R_t \ge \bar{b}^*$, the cap is $\bar{b}^*$ and every bid is at most $\bar{b}^* \le R_t$; the print equals R_t only if it equals the cap *and* $\bar{b}^* = R_t$, in which case the pooled set $\{v > J_S(\bar{b}^*)\} = \{v > 1\}$ is empty and the cap is bid only by the F-null type $v = 1$. If the frontier is constrained, $R_t < \bar{b}^*$, the cap lies in $[R_t - \varepsilon, R_t)$ and every print is strictly below R_t outright. In both cases, almost surely every transaction under buyer control prints strictly below the current reference; by the symmetric argument, with the floor in place of the cap and the G-null type $c = 0$ at $R_t = \underline{s}^*$, every transaction under seller control prints strictly above it. This proves part (ii).

Step 2: A uniform crossing bound. Call period t a *crossing* if a trade occurs whose print lies strictly beyond the risk-dominance threshold from the controlling side: below $R^\dagger$ if $\xi_t = B$, above $R^\dagger$ if $\xi_t = S$. We claim that there exists a constant $a > 0$ such that

$$\Pr\big(\text{crossing at } t \,\big|\, \mathcal{F}_{t-1}\big) \ \ge\ a \qquad \text{almost surely, for every } t. \tag{16}$$

Suppose $\xi_t = B$, so $R_t \ge R^\dagger$, and let κ_t be the prevailing cap. If the frontier is unconstrained, $\kappa_t = \bar{b}^* \ge \kappa_0$; if constrained, $\kappa_t \ge R_t - \varepsilon \ge R^\dagger - \varepsilon_0 \ge \kappa_0$. Hence $\kappa_t \ge \kappa_0$ in every case. Every buyer type $v < J_S(\kappa_0)$ has $b^M(v) < \kappa_0 \le \kappa_t$, so she separates and bids $b^M(v)$, independently of the cap, the reference, and the convention; her bid satisfies $b^M(v) < \kappa_0 \le R^\dagger$, and she trades with probability $G\big(b^M(v)\big)$, producing a print strictly below $R^\dagger$. Integrating over these types and using the independence of (V_t, C_t) from $\mathcal{F}_{t-1}$,

$$\Pr\big(\text{crossing at } t \,\big|\, \mathcal{F}_{t-1}\big) \ \ge\ \int_0^{J_S(\kappa_0)} G\big(b^M(v)\big)\, dF(v) \ =\ \alpha_B.$$

Suppose instead $\xi_t = S$, so $R_t \le R^\dagger$, with prevailing floor φ_t. If unconstrained, $\varphi_t = \underline{s}^* \le s_1$; if constrained, $\varphi_t \le R_t + \varepsilon \le R^\dagger + \varepsilon_0 \le s_1$. Every seller type $c > J_B(s_1)$ has $s^M(c) > s_1 \ge \varphi_t$, so he separates and asks $s^M(c) > s_1 > R^\dagger$, trading with probability $1 - F\big(s^M(c)\big)$ and printing strictly above $R^\dagger$; integrating gives the bound α_S. This proves (16).

Step 3: Trades occur infinitely often. A crossing is in particular a trade, so $\Pr\big(T_t \mid \mathcal{F}_{t-1}\big) \ge a$ for every t, where T_t is the trade event. Hence $\sum_t \Pr\big(T_t \mid \mathcal{F}_{t-1}\big) = \infty$ almost surely, and by Lévy's extension of the Borel–Cantelli lemma, T_t occurs for infinitely many t almost surely: part (i). Let $\tau_1 < \tau_2 < \cdots$ denote the trade times, and write $\mathcal{G}_k = \mathcal{F}_{\tau_k}$.

Step 4: Runs are geometrically dominated and price-making changes sides infinitely often. Between consecutive trades the reference, hence the controlling side, is unchanged, so the side at trade $k+1$ is determined by the print of trade k: writing X_k for the event that trade k is a crossing, we have $\xi_{\tau_{k+1}} \ne \xi_{\tau_k}$ on X_k, whatever the tie rule, since the new reference lies strictly beyond the threshold. We first transfer (16) to trade times:

$$\Pr\big(X_{k+1} \,\big|\, \mathcal{G}_k\big) \;\ge\; a \qquad \text{almost surely, for every } k \ge 0. \tag{17}$$

Indeed, decomposing over the date of trade $k{+}1$ and using that $\{\tau_{k+1} \ge t\} \in \mathcal{F}_{t-1}$,

$$\Pr\big(X_{k+1} \mid \mathcal{G}_k\big) = \sum_{t > \tau_k} \mathbb{E}\Big[\mathbf{1}\{\tau_{k+1} \ge t\}\, \Pr\big(\text{crossing at } t \mid \mathcal{F}_{t-1}\big) \;\Big|\; \mathcal{G}_k\Big] \tag{18}$$

$$\ge\; a \sum_{t > \tau_k} \mathbb{E}\Big[\mathbf{1}\{\tau_{k+1} \ge t\}\, \Pr\big(T_t \mid \mathcal{F}_{t-1}\big) \;\Big|\; \mathcal{G}_k\Big], \tag{19}$$

where the inequality holds term by term because a crossing at t has conditional probability at least $a \ge a\ \Pr(T_t \mid \mathcal{F}_{t-1})$; and the final sum equals $\Pr\big(\tau_{k+1} < \infty \mid \mathcal{G}_k\big) = 1$. Now fix k and let $L_k = \min\{j \ge 1 : \xi_{\tau_{k+j}} \ne \xi_{\tau_k}\}$ be the number of further trades the side controlling trade k retains control. Since $\{L_k > m\} \subseteq X_k^c \cap X_{k+1}^c \cap \cdots \cap X_{k+m-1}^c$, iterating (17) backward gives

$$\Pr\big(L_k > m \,\big|\, \mathcal{G}_{k-1}\big) \;\le\; (1-a)^m, \qquad m = 0, 1, 2, \ldots,$$

so every run of consecutive trades under one side is stochastically dominated by a geometric distribution with parameter a, depending only on (F, G). In particular every run is finite almost surely; combined with Step 3, the trade sequence decomposes into infinitely many finite runs, so price-making changes sides infinitely often almost surely. □

A.17 Proof of Corollary 4

Proof. Throughout, adopt the framework and notation of Step 0 of the proof of Theorem 7, and define

$$\lambda_B := \int_0^1 G\big(b^M(v)\big)\, dF(v), \qquad \lambda_S := \int_0^1 \big[1 - F\big(s^M(c)\big)\big]\, dG(c),$$

and μ_B, μ_S as the conditional laws of $b^M(V)$ given $\{C \le b^M(V)\}$ and of $s^M(C)$ given $\{V \ge s^M(C)\}$.

Step 1: The selected side is always unconstrained. Suppose $\xi_t = B$, so $R_t \ge R^\dagger$ by the risk-dominant equilibrium selection (or the tie rule at equality). The hypothesis gives $R_t \ge R^\dagger \ge \bar{b}^*$, so the frontier cap is $q(R_t) = \bar{b}^*$, which is admissible; the selected equilibrium is exact, and by Corollary 3 its play coincides with the buyer's-bid auction: bids $b^M(V_t)$, truthful asks. Symmetrically, $\xi_t = S$ forces $R_t \le R^\dagger \le \underline{s}^*$, floor $t(R_t) = \underline{s}^*$, and the seller's-ask auction. This holds at *every* reference, including the initial one and any reference inside $[\bar{b}^*, \underline{s}^*]$: under the hypothesis, whichever side the selection assigns is unconstrained there. This proves (i), with no ε-frontier approximation anywhere on the path. Consequently, in each regime the strategy profile, and with it the trade event and the transaction price, are fixed functions of the current types alone.

Under buyer control, every bid satisfies $b^M(V_t) \le \bar{b}^* \le R_t$ and every truthful ask lies below it on the trade event, so the price is the bid; the trade probability is λ_B and the print law is μ_B, at every reference with buyer control. The seller mirror gives λ_S and μ_S. Both probabilities are interior: $\lambda_B \ge \int_{1/2}^1 G\big(b^M(\tfrac{1}{2})\big) dF > 0$ and $\lambda_B \le G(\bar{b}^*) < 1$ since $\bar{b}^* < 1$; the mirror bounds use $\underline{s}^* > 0$. Finally, both print laws are atomless with the stated supports: $b^M = J_S^{-1}$ is strictly increasing and continuous, so $b^M(V)$ inherits atomlessness from F, conditioning on a positive-probability event preserves it, and $\mu_B\big(\{\bar{b}^*\}\big) = \Pr\big(V = 1 \mid C \le b^M(V)\big) = 0$, so μ_B concentrates on $(0, \bar{b}^*)$; symmetrically μ_S on $(\underline{s}^*, 1)$.

Step 2: Reversal at every trade and the excluded interval. By Step 1, every print under buyer control lies in $(0, \bar{b}^*)$ almost surely, hence strictly below $R^\dagger$; every print under seller control lies in $(\underline{s}^*, 1)$, strictly above it. In particular no transaction price ever lies in $[\bar{b}^*, \underline{s}^*]$. At every trade, the new reference lies strictly beyond the risk-dominance threshold from the trading side, and the selection assigns control to the opposite side at all subsequent periods until the next trade, since $R_t \notin [\bar{b}^*, \underline{s}^*] \supseteq \{R^\dagger\}$ from then on.

Between trades the reference is constant, so control changes sides exactly at the trades, at every one of them; and since every print lies in $(0, \bar{b}^*) \cup (\underline{s}^*, 1)$, the reference never returns to $[\bar{b}^*, \underline{s}^*]$ after the first trade. This proves (ii). Trades occur infinitely often almost surely: the per-period conditional trade probability is λ_B or λ_S, hence at least $\min\{\lambda_B, \lambda_S\} > 0$, and Lévy's Borel–Cantelli lemma applies as in Step 3 of the proof of Theorem 7. Write $\tau_1 < \tau_2 < \cdots$ for the trade times and P_k for the print of trade k.

Step 3: The cycle decomposition. For $k \geq 1$, call the periods $\{\tau_k + 1, \ldots, \tau_{k+1}\}$ *cycle* $k{+}1$, with duration $D_{k+1} = \tau_{k+1} - \tau_k$; throughout cycle $k{+}1$, the reference equals P_k and control is held by the side opposite the one that traded at τ_k, by Step 2. Since the types $(V_{\tau_k+j}, C_{\tau_k+j})_{j\geq 1}$ are i.i.d. and independent of $\mathcal{F}_{\tau_k}$, while the controlling side of cycle $k{+}1$ is $\mathcal{F}_{\tau_k}$-measurable and its strategies are reference-free by Step 1, the pair (D_{k+1}, P_{k+1}) is, conditionally on $\mathcal{F}_{\tau_k}$, distributed as $\mathrm{Geom}(\lambda_\xi) \otimes \mu_\xi$ for the controlling side ξ: the duration is the first success time in i.i.d. trials of success probability λ_ξ, and the print is the successful trial's price, whose law is μ_ξ independently of which trial succeeded. Because sides alternate deterministically after the first trade, the cycles $(D_k, P_k)_{k\geq 2}$ form an independent sequence with laws alternating between $\mathrm{Geom}(\lambda_B) \otimes \mu_B$ and $\mathrm{Geom}(\lambda_S) \otimes \mu_S$.

Step 4: Renewal–reward limits. Group the cycles from the second onward into consecutive pairs, one of each side; the pairs are i.i.d., with integrable durations $D^{(BS)}$ of mean $1/\lambda_B + 1/\lambda_S$. The strong law of large numbers for the paired sums, together with the almost-sure finiteness of the initial segment before the first trade, yields, for every ratio of cumulative rewards to cumulative periods, the ratio of the pair means.

Buyer control prevails exactly during the cycles that end in a buyer trade, of mean duration $1/\lambda_B$; the long-run fraction of periods under buyer control is therefore

$$\frac{1/\lambda_B}{1/\lambda_B + 1/\lambda_S} = \frac{\lambda_S}{\lambda_B + \lambda_S}.$$

For a fixed Borel set A, the reward $\mathbf{1}\{P_k \in A\}$ per trade gives, since prints alternate between independent draws from μ_B and μ_S, the trade-average $\frac{1}{2}\mu_B(A) + \frac{1}{2}\mu_S(A)$; the reward $D_{k+1}\mathbf{1}\{P_k \in A\}$—the number of periods the print P_k is held as the reference, with D_{k+1} independent of P_k and of mean $1/\lambda_S$ after a buyer print, $1/\lambda_B$ after a seller print—gives the period-average

$$\frac{\mu_B(A)/\lambda_S + \mu_S(A)/\lambda_B}{1/\lambda_B + 1/\lambda_S} = \frac{\lambda_B\,\mu_B(A) + \lambda_S\,\mu_S(A)}{\lambda_B + \lambda_S}.$$

Applying both simultaneously over a countable class of intervals with rational endpoints, and using that the limit laws are atomless, the empirical laws of prints and of the reference converge weakly almost surely to $\frac{1}{2}\mu_B + \frac{1}{2}\mu_S$ and $(\lambda_B\mu_B + \lambda_S\mu_S)/(\lambda_B + \lambda_S)$.

Realized gains accrue only at trades, and the expected realized gain of a cycle equals $\mathbb{E}\big[(V - C)\mathbf{1}\{\text{trade}\}\big]/\lambda_\xi = W_B(\bar{b}^*)/\lambda_B$ or $W_S(\underline{s}^*)/\lambda_S$ by (Frontier Welfares); gains are bounded by 1, so renewal–reward gives, almost surely,

$$\overline{W}_\infty = \frac{W_B(\bar{b}^*)/\lambda_B + W_S(\underline{s}^*)/\lambda_S}{1/\lambda_B + 1/\lambda_S} = \frac{\lambda_S\, W_B(\bar{b}^*) + \lambda_B\, W_S(\underline{s}^*)}{\lambda_B + \lambda_S},$$

a convex combination with both weights in $(0,1)$, since $\lambda_B, \lambda_S \in (0,1)$ by Step 1.

Step 5: Reflection symmetry. Suppose $F(v) = 1 - G(1-v)$. By the reflection identities in the proof of Theorem 4, $J_B(v) = 1 - J_S(1-v)$, hence $s^M(c) = J_B^{-1}(c) = 1 - b^M(1-c)$. Since $1 - F(s) = G(1-s)$ and the map $c \mapsto 1-c$ pushes G forward to F,

$$\lambda_S = \int_0^1 \big[1 - F\big(s^M(c)\big)\big]\, dG(c) = \int_0^1 G\big(b^M(1-c)\big)\, dG(c) = \int_0^1 G\big(b^M(v)\big)\, dF(v) = \lambda_B.$$

Substituting $\lambda_B = \lambda_S$ into the welfare expression of (iii) collapses it to $\overline{W}_\infty = \frac{1}{2}\big(W_B(\bar{b}^*) + W_S(\underline{s}^*)\big)$. Finally, the average $\frac{1}{2}\big(W_B(\bar{b}^*) + W_S(\underline{s}^*)\big)$ equals the expected gains from trade of the random-offerer mechanism, and the bound $\overline{W}^\infty \geq \frac{1}{\pi}W^{\mathrm{FB}}$ follows as in Step 3 of the proof of Theorem 3. □

A.18 Proof of Theorem 8

Proof. We prove the buyer-side statements; the seller-side mirror is Step 6.

Step 0: Framework and conventions. Adopt the filtration framework of Step 0 of the proof of Theorem 7. By Theorem 2 (Necessity), every undominated buyer-control equilibrium at reference R is a truncated monopsony equilibrium indexed by its cap in $\mathcal{B}(R)$; a buyer-side convention is therefore an adapted cap process (κ_t) with $\kappa_t \in \mathcal{B}(R_t)$, each κ_t being $\mathcal{F}_{t-1}$-measurable, possibly with the help of randomization independent of the current types. All bounds below hold conditionally on the realized cap and are uniform over $\mathcal{B}(R_t)$ where needed, so randomization is harmless and we treat κ_t as $\mathcal{F}_{t-1}$-measurable. In period t, buyers bid $m_{\kappa_t}(v) = \min\{b^M(v), \kappa_t\}$, sellers with $c < \kappa_t$ ask truthfully and all others ask 1; off null sets, the trade event is $T_t = \{C_t \leq m_{\kappa_t}(V_t)\}$, with print $p_t = m_{\kappa_t}(V_t)$.

Step 1: Monotone descent. Note that $\kappa_t \le \bar{b}^*$ and $\kappa_t \le R_t$, with $\kappa_t < R_t$ whenever $R_t < \bar{b}^*$. Hence $p_t \le \kappa_t \le R_t$ on every trade. Equality $p_t = R_t$ requires $\kappa_t = R_t$, hence $R_t \ge \bar{b}^*$ and $\kappa_t = \bar{b}^* = R_t$, together with a print at the cap $\bar{b}^*$; the pooled set $\{v > J_S(\bar{b}^*)\} = \{v > 1\}$ is empty, so the cap $\bar{b}^*$ is bid only by the F-null type $v = 1$. Almost surely, therefore, every trade prints strictly below the prevailing reference, simultaneously over all t; between trades the reference is unchanged. This proves (i). Consequently R_t is nonincreasing and converges to a limit $R_\infty \ge 0$; moreover $R_t > 0$ for every t almost surely, since every print $m_{\kappa_t}(V_t)$ is strictly positive for $V_t > 0$.

Step 2: The descent bound. For $\delta \in (0, 1)$ define

$$\lambda(\delta) := \Pr\big(V < J_S(\delta),\ C \le b^M(V)\big) = \int_0^{J_S(\delta)} G\big(b^M(v)\big)\, dF(v) \ > \ 0,$$

strictly positive since $J_S(\delta) > 0$ and $G\big(b^M(v)\big) > 0$ for $v > 0$. We claim that, for every t and every δ, almost surely

$$\Pr\big(T_t \cap \{p_t < \delta\} \,\big|\, \mathcal{F}_{t-1}\big) \ \ge \ \lambda(\delta)\ \Pr\big(T_t \,\big|\, \mathcal{F}_{t-1}\big), \tag{20}$$

with no condition on the reference. Split on the $\mathcal{F}_{t-1}$-measurable event $\{\kappa_t < \delta\}$. If $\kappa_t < \delta$, every print satisfies $p_t \le \kappa_t < \delta$, so the left side equals $\Pr(T_t \mid \mathcal{F}_{t-1})$ and (20) holds since $\lambda(\delta) \le 1$. If $\kappa_t \ge \delta$, every type $v < J_S(\delta)$ has $b^M(v) < \delta \le \kappa_t$, so she separates and bids $b^M(v) < \delta$, independently of the cap, trading exactly against $\{C_t \le b^M(v)\}$ at a print below δ; integrating over these types, the left side is at least $\lambda(\delta) \ge \lambda(\delta) \Pr(T_t \mid \mathcal{F}_{t-1})$.

Step 3: The dichotomy. Let $N = \sum_t \mathbf{1}\{T_t\}$ be the number of trades. On $\{N < \infty\}$, the reference is constant after the last trade and equals either R_1 or the last print; by Step 1 it is positive, so the reference freezes at a positive level. Now fix $\delta > 0$ and work on $\{N = \infty\}$. By Lévy's extension of the Borel–Cantelli lemma—for adapted events, $\{A_t \text{ i.o.}\}$ and $\{\sum_t \Pr(A_t \mid \mathcal{F}_{t-1}) = \infty\}$ coincide up to null sets—the event $\{N = \infty\} = \{T_t \text{ i.o.}\}$ coincides almost surely with $\{\sum_t \Pr(T_t \mid \mathcal{F}_{t-1}) = \infty\}$. On this event, (20) gives $\sum_t \Pr\big(T_t \cap \{p_t < \delta\} \mid \mathcal{F}_{t-1}\big) = \infty$, so by the same lemma a trade with print below δ occurs almost surely on $\{N = \infty\}$; by Step 1 the reference then remains below δ forever. Intersecting over $\delta = 1/n$, $n \in \mathbb{N}$, yields $R_t \to 0$ almost surely on $\{N = \infty\}$.

The first branch is not vacuous: under the adapted convention

$$\kappa_t = \min\left\{2^{-t},\ \frac{R_t}{2},\ \frac{\bar{b}^*}{2}\right\} \in \mathcal{B}(R_t),$$

we have $\Pr(T_t) \le G(\kappa_t) \le G(2^{-t})$. Since G has a bounded density on $[0,1]$, $\sum_t G(2^{-t}) < \infty$, and hence Borel–Cantelli implies that only finitely many trades occur almost surely. By Step 1, the reference therefore freezes at a positive level.

Step 4: Stationary conventions. Let the convention be stationary: $\kappa_t = k(R_t)$ for a fixed measurable selection $k(r) \in \mathcal{B}(r)$ (under stationary randomization, a fixed kernel over $\mathcal{B}(r)$; the argument below applies conditionally on the realized cap). On $\{N < \infty\}$, let τ be the last trade time ($\tau = 0$ if none) and $r = R_\infty > 0$ the frozen reference. For every $t > \tau$, $\Pr\big(T_t \mid \mathcal{F}_{t-1}\big) = \pi(r)$, where $\pi(r) = \int_0^1 G\big(m_{k(r)}(v)\big)\, dF(v) > 0$ pointwise, since $k(r) > 0$ and $m_{k(r)}(v) > 0$ for $v > 0$; no uniformity in r is needed, as r is a fixed (random) positive number on the event. Hence $\sum_t \Pr(T_t \mid \mathcal{F}_{t-1}) = \infty$ on $\{N < \infty\}$, which by Lévy's lemma forces infinitely many trades there—a contradiction unless $\Pr(N < \infty) = 0$. Under every stationary convention, therefore, the second branch obtains.

Step 5: Welfare. The realized per-period gain is $g_t = (V_t - C_t)\mathbf{1}\{T_t\} \in [0,1]$, and on T_t, $C_t \le \kappa_t \le R_t$. Fix $\eta > 0$. On $\{R_t \to 0\}$ there is almost surely a finite t_0 with $R_t < \eta$ for all $t \ge t_0$, so $g_t \le \mathbf{1}\{C_t \le \eta\}$ for $t \ge t_0$, and by the strong law of large numbers for the i.i.d. sequence (C_t),

$$\limsup_{T\to\infty} \frac{1}{T}\sum_{t\le T} g_t \ \le\ \limsup_{T\to\infty} \frac{1}{T}\sum_{t\le T} \mathbf{1}\{C_t \le \eta\} \ =\ G(\eta) \qquad \text{almost surely,}$$

the finite prefix $t < t_0$ vanishing in the average. On $\{N < \infty\}$, the sum of gains is finite and the average converges to zero. The two events cover almost surely by Step 3, and letting $\eta \downarrow 0$ along a sequence, $G(\eta) \to G(0) = 0$, so long-run average welfare per period is zero almost surely. This proves (iii).

Step 6: The mirror. For a seller-side convention, substitute throughout: floors $\varphi_t \in \mathcal{S}(R_t)$ for caps, the truncated monopoly profile $a_{\varphi_t}(c) = \max\{s^M(c), \varphi_t\}$ for the truncated monopsony profile, prints $p_t = a_{\varphi_t}(C_t) \ge \varphi_t \ge R_t$ with strictness almost surely (at $R_t = \underline{s}^*$ the floor is asked only by the G-null type $c = 0$, since the pooled set $\{c < J_B(\underline{s}^*)\} = \{c < 0\}$ is empty),

so the reference is nondecreasing and strictly increasing at every trade, converging to $R_\infty \le 1$ with $R_t < 1$ throughout. The ascent bound replaces $\lambda(\delta)$ by $\tilde{\lambda}(\delta) = \Pr\big(C > J_B(\delta),\, V \ge s^M(C)\big) > 0$ for $\delta \in (0,1)$: types $c > J_B(\delta)$ separate at asks $s^M(c) > \delta$ whenever the floor lies below δ, and every print exceeds the floor otherwise. The dichotomy, the summable-floors witness

$$\varphi_t = \max\left\{1 - 2^{-t},\ \frac{1+R_t}{2},\ \frac{1+\underline{s}^*}{2}\right\} \in \mathcal{S}(R_t),$$

, and the stationary exclusion are verbatim, giving $R_t \to 1$ on infinitely many trades; and welfare vanishes since on T_t, $V_t \ge \varphi_t \ge R_t$, so $g_t \le \mathbf{1}\{V_t \ge 1-\eta\}$ eventually, with $1 - F(1-\eta) \to 0$ as $\eta \downarrow 0$. □