Pre-empting Competitive Sales: Private Offers, Early Resolution, and Seller Information

Emil M. S. Halseth*

October 13, 2026

Latest version

[Preliminary: Do not cite or distribute]

Abstract

A private offer can end a competitive sale before other buyers bid. We study costly offers by a buyer with private value and resolution benefits to a seller with private continuation prospects. Accepted trades must finance foregone competition and unsuccessful attempts. Some buyers offer prices that screen seller states; others pay more for certain acceptance. An effective ban preserves later competition and saves offer costs, but sacrifices resolution gains. Its welfare effect depends on the surplus denied to excluded buyers. We establish probing and knockout equilibria and give numerical examples with opposite welfare rankings.

Keywords: auctions, private offers, early contracting, seller information.
JEL Codes: D44, D82, R31.

*Oslo Economics; esh@osloeconomics.no

1 Introduction

A seller preparing a competitive sale may receive a private offer to end the process early. Accepting gives the parties certainty, but prevents other buyers from bidding. A homeowner may sell before an announced bidding round; a business may accept an offer before its equipment auction. Consider also someone reselling a concert ticket: an offer today competes with the possibility of stronger demand close to the event. These settings share a decision about whether to agree now or preserve the opportunity for later competition.

We ask two questions. When will a buyer attempt to pre-empt a planned auction with a private offer, and when will the seller accept? When does a ban on such agreements increase total welfare? The questions are linked because early resolution can benefit the contracting parties while excluding buyers who value the object more.

We study this decision when the seller has private information about the value of continuing the sale. That information may concern prospective buyers, the value of retaining the object, or the benefit of reaching an early agreement. A seller who receives enquiries, for example, may know more about likely participation than an individual buyer does. The buyer privately knows its match value and its own benefit from early agreement. Such benefits may arise from coordinating a move, arranging equipment delivery, or committing to complementary plans.

We model the continuation as an English auction. Ascending bidding is an established way of selling homes and other valuable assets; Norwegian bidding rounds and New Zealand property auctions provide two relevant settings (Finansdepartementet, 2021; Real Estate Authority, nd). The English auction gives a tractable account of how competition determines prices and allocation in such sales. One designated buyer can pay a cost to make a private binding offer and remains free to bid after rejection; the other buyers enter only in the auction. Its offer therefore screens the seller's willingness to continue and reveals information about a buyer who will remain in the sales process.

The model explains why buyers make offers that may be rejected. A relatively low offer purchases early agreement only when the seller's continuation prospects are weak. We call this a *probe*. Paying more can secure acceptance even from a seller with strong prospects: a *knockout* offer. The buyer chooses without observing the seller's state; these terms describe acceptance across states, while different buyer types can choose the same offer. Buyers with little to gain from early resolution wait. Those with intermediate gains may probe, while buyers with larger gains pay for certain acceptance. A seller's willingness to accept also depends on what the offer reveals about the early buyer's valuation, because the buyer can still bid in the auction after rejection.

Our first result identifies what finances early contracting. The seller and early buyer jointly sacrifice a nonnegative *competition wedge* by abandoning the auction. When the

early buyer's value exceeds the seller's fallback, this wedge comes from the possibility that at least two later buyers bid above that value. In every perfect Bayesian equilibrium, the resolution gains from accepted trades must cover this wedge and the cost of all offer attempts, including rejected ones. No restriction on the number of offer prices or on strategy mixing is needed.

For two ordered seller states, we then construct equilibria with waiting and probing, and give sufficient conditions for knockout to coexist with them. Continuously distributed buyer values and resolution benefits determine who selects each action. Locally passive beliefs (defined in Appendix C.1) yield prices that leave the marginal accepting seller indifferent. We also give conditions under which these outcomes admit D1-compatible completions in the contingent-plan representation defined in Appendix C (Cho and Kreps, 1987; Banks and Sobel, 1987).

Finally, we compare discretion to accept early offers with a ban on pre-emption. The early buyer values the option to offer, and the seller accepts only when doing so is worthwhile given its information. Their agreement imposes a negative externality on later buyers by removing their opportunity to acquire the object and earn auction surplus. In our auction, when the early buyer values the object above the seller's fallback, one higher late value can change the efficient allocation, while two are needed to raise the price above the early buyer's value. This difference between allocative value and the contracting parties' opportunity cost drives the welfare result. A ban saves offer costs and preserves late-buyer surplus, but removes gains from early resolution. Two examples with the same mean buyer value give opposite welfare rankings.

Housing-market application. Norwegian homes are commonly sold through broker-mediated ascending bidding after advertised viewings. A buyer can try to pre-empt this process by approaching the seller directly with an offer to conclude the sale early, a practice known as *kupping* (Finansdepartementet, 2021, Section 6.6). It combines the three features central to the model: an offer before open competition, seller information about the alternative to agreement, and a buyer who can still join the bidding round if rejected. Sellers may know more about enquiries or intended participation, while their plans for moving and the value of keeping the property need not be observable to buyers.

Flått et al. (2022, Section 3) identify 2,018 early sales among 84,203 Oslo transactions in 2007–2017, or about 2.4%, using sales fewer than five days after listing as their proxy for kupping. NOU 2021:7 reports an estimate of around 5% from a different brokerage sample and classification (Finansdepartementet, 2021, Section 6.6.3). These figures do not capture failed pre-emption attempts.[1] Whether such agreements should be permitted has also been

[1] Flått et al. (2022, Section 5.2 and Table III) estimate sale prices about 4.6% higher for completed early sales than for auction sales, controlling for property and market characteristics. This conditional difference concerns selected transactions; they suggest that buyers' aversion to auction uncertainty may help explain it.

debated in Norway: NOU 2021:7 considers restrictions on kupping (Finansdepartementet, 2021, Section 6.6.3). The reform effective 1 July 2025 barred brokers from transmitting secret bids and required them to urge sellers to consult the broker about direct offers; it did not ban direct buyer–seller agreements (Finansdepartementet, 2025b,a).

The English auction captures open ascending competition, but imposes a committed allocation and payment rule that Norwegian bidding rounds need not follow exactly. Pre-auction sales are also possible in New Zealand: a seller receiving an early offer may bring the auction forward or withdraw the property and agree a private sale (Real Estate Authority, nd). Beyond housing, industrial-equipment sellers invite pre-auction offers (Prestige Equipment, 2026), and private acquisition offers may be conditional on a target ceasing to entertain competing bidders (Liu et al., 2023, pp. 7–8). These applications motivate the decision to interrupt competition. Less structured sales mechanisms, such as the sequence of unsolicited offers in the ticket example, raise the same question but require a different continuation model.

Related literature. Bulow and Klemperer (1996) show that, under their assumptions, an auction with an additional bidder yields higher expected seller revenue than optimal negotiation with fewer bidders. We study when gains from early resolution can make ending a planned auction jointly attractive to the seller and an early buyer. The welfare analysis includes the surplus lost by buyers excluded through early agreement.

A close procedural comparison is Grebe et al. (2016), where one buyer makes a binding Sell-It-Now offer and enters a second-price auction with a rival after rejection. Their risk-neutral benchmark has a publicly known seller value of zero, fixed participation, no attempt cost, and no direct resolution benefit; under their tie-breaking rule every offer is rejected. We study how private seller continuation prospects interact with voluntary costly offers and heterogeneous resolution benefits.[2]

Seller-posted buyout prices can generate early trade through time sensitivity or risk aversion (Mathews, 2004; Gallien and Gupta, 2007; Budish and Takeyama, 2001; Mathews and Katzman, 2006). In Khezr and Menezes (2018), buyers accept an asking price or submit counteroffers; rejected counteroffers do not trigger an auction. Public preemptive bids can deter rival entry or information acquisition (Fishman, 1988; Cramton and Schwartz, 1991). Here the buyer submits a private binding offer that the seller can accept or reject, and remains a competitor after rejection. Information conveyed by its offer therefore changes the seller's assessment of continuing the sale; in our model, it changes expected auction revenue. Bulow and Klemperer (2009) compare a simultaneous auction with a sequential sale in which public jump bids can deter further costly entry.

[2]The working-paper version refers to Grebe's dissertation for an analysis with risk-averse buyers and/or sellers (Grebe et al., 2015, fn. 7). Our comparison concerns the published risk-neutral benchmark.

We instead study whether a private agreement interrupts a given auction, holding late participation fixed.

Section 2 gives the model and financing result. Section 3 constructs probing and knockout equilibria. Section 4 studies the ban. The appendices provide the equilibrium proofs, supporting payoff calculations, and details of the numerical examples. The Online Appendix studies further regimes, disclosure of verifiable seller information, risk aversion as a source of gains from early agreement, continued bargaining before the auction, and seller counteroffers after bidding.

2 A Private Offer before an English Auction

2.1 Timing, preferences, and information

A seller owns one indivisible object. An early buyer privately observes its value $b \sim F$ on $[\underline{b}, \bar{b}]$ and benefit from early resolution $d_B \sim G$ on $[0, \bar{d}]$. The two draws are independent and have continuously differentiable positive densities f and g. Write d for a realization of d_B.

The seller privately observes $s \in \{L, H\}$, independent of the buyer's type, with $\Pr(s = H) = \pi \in (0, 1)$. State s determines the distribution Q_s of the number N of late buyers, the seller's no-sale value $v_s \in (\underline{b}, \bar{b})$, and its benefit from early resolution $d_{S,s} \geq 0$. Every late buyer draws independently from the same F as the early buyer. The distributions are common knowledge, but the buyer does not observe the seller's state. Participation is exogenous. We use stronger prospects for late participation as the running interpretation of H; a higher value of retaining the object or a smaller benefit from early agreement can also make the seller more reluctant to accept.

The early buyer either waits or pays a cost $\kappa > 0$ to submit a private, binding price p. The seller observes the price and accepts or rejects. Acceptance ends the sale. Rejection leads to the same auction as waiting, with the early buyer still present. Late buyers arrive and the realized participants are observed. We use the English-auction allocation and payment rule, implemented by a second-price mechanism with a fixed, privately known seller threshold v_s. The highest bidder wins if its bid reaches v_s, paying the larger of v_s and the highest rival bid; otherwise the seller retains the object. The seller cannot revise the threshold after seeing the bids. Bidding one's value is optimal for every realized competing threshold: it secures the object when its value exceeds that threshold and avoids a purchase at a loss. This remains optimal when the threshold is unknown. Figure 2.1 summarizes the sequence.

Before the offer cost, early agreement at p gives the buyer $b - p + d_B$ and the seller $p + d_{S,s}$. If the auction is reached, each receives its auction payoff without these resolution benefits. The buyer additionally pays κ whenever it offers, whether accepted or rejected. Thus

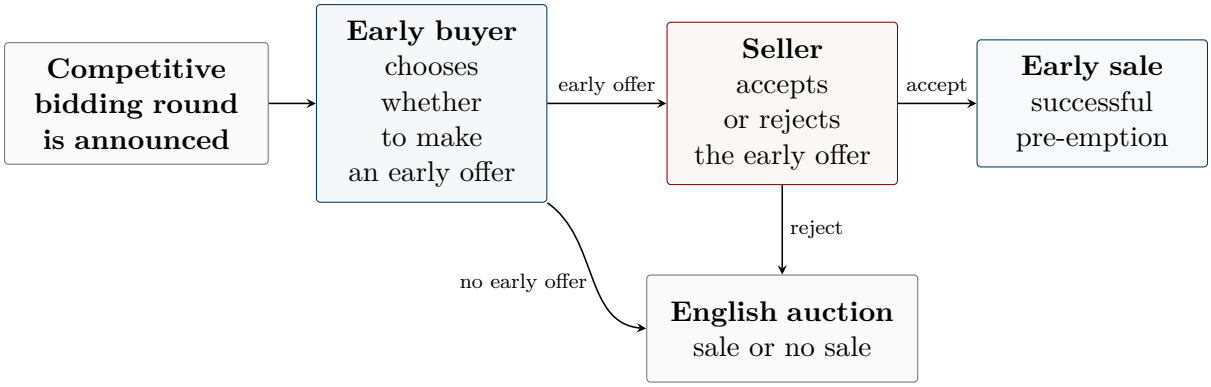


Figure 2.1: Sequence of decisions. Acceptance produces an early sale; rejection and no offer lead to the same auction, in which the early buyer can still bid.

the cost is attached to the attempt, whereas the resolution benefits require a completed early agreement. All agents are risk neutral. The equilibrium concept is perfect Bayesian equilibrium (PBE). Except in the aggregate financing result, we construct pure strategies, with sellers accepting at indifference. Table 2.1 summarizes the model's primitives and their economic roles.

Let $Y_{1N} \geq Y_{2N}$ denote the two highest late values. For $N = 0$, set both to $\underline{b}$; for $N = 1$, do so only for Y_{2N}. The seller's realized payoff $P_s(b)$ equals the sale price when the object is sold and v_s otherwise. Conditional on (b, s), the early buyer's expected auction surplus and the seller's expected auction payoff are

$$u_s(b) = \mathbb{E}_s[\max\{b - \max\{v_s, Y_{1N}\}, 0\}], \tag{2.1}$$

$$R_s(b) = \mathbb{E}_s[P_s(b)], \qquad P_s(b) = \max\{v_s, Y_{2N}, \min\{b, Y_{1N}\}\}. \tag{2.2}$$

Here $\mathbb{E}_s$ averages over the number of late buyers and their valuations. Define the buyer's willingness to pay for early ownership before resolution benefits, $w_s(b) = b - u_s(b)$. The seller's minimum early price conditional on (b, s) is $C_s(b)$, and averaging the buyer's willingness to pay over seller states gives $\bar{w}(b)$:

$$C_s(b) = R_s(b) - d_{S,s}, \qquad \bar{w}(b) = (1 - \pi)w_L(b) + \pi w_H(b). \tag{2.3}$$

2.2 The competition wedge and the financing of attempts

The seller and early buyer can divide b by selling early, before resolution benefits. Together they expect to obtain $R_s(b)+u_s(b)$ in the auction. Their opportunity cost of early agreement is therefore

$$\Gamma_s(b) := R_s(b) - w_s(b) = \mathbb{E}_s[\max\{\max\{v_s, Y_{2N}\} - b, 0\}] \geq 0. \tag{2.4}$$

Table 2.1: Primitives of the baseline model.

Symbol	Name and information	Economic role
b, F, f, $[\underline{b}, \bar{b}]$	Object value; privately known to each buyer.	Values determine the gains from ownership and the auction outcome. F is the common value distribution, f its density, and the endpoints its support.
d_B, d, G, g, $\bar{d}$	Buyer's resolution benefit; privately known to the early buyer.	Measures the gain from completing the purchase before the auction. G has density g on $[0, \bar{d}]$; d is a realization. Heterogeneity changes the incentive to wait, probe, or secure acceptance.
$s \in \{L, H\}$, π	Seller state; privately known to the seller.	Determines continuation prospects and early-resolution benefits. π is the commonly known probability of H, so the buyer faces uncertainty about the seller's acceptance.
N, Q_s	Number of late buyers; drawn conditional on s.	Captures the intensity of later competition. Participation is exogenous; Q_s is common knowledge and realized participants are observed at the auction.
v_s	Seller's fallback value and fixed auction threshold.	Values retaining the object and determines when auction trade occurs. It is known to the seller through s and cannot be revised after bids.
$d_{S,s}$	Seller's resolution benefit; known through s.	Measures the gain from completing the sale early and lowers the seller's minimum acceptable early price.
$\kappa > 0$	Cost of submitting an offer; common knowledge.	Makes an unsuccessful attempt costly relative to waiting. It is paid whether the offer is accepted or rejected.

The distributions and support bounds are common knowledge. The densities f and g are implied by F and G, rather than additional primitives. Differences in Q_s, v_s, and $d_{S,s}$ are alternative or combinable sources of private seller information; all three need not vary across states. Appendix A.2 gives sufficient conditions for each source.

Appendix A proves the identity. If $b \geq v_s$, the wedge is positive precisely when there is positive probability of two late values above b.

> **Example: one late buyer or two.** Take $b = 100$ and a seller fallback below 100. One late buyer with value 140 wins at 100, leaving the seller and early buyer jointly with 100. With two late values of 140 and 120, the seller instead receives 120 and the early buyer receives zero. In the latter realization, the auction produces 20 more than these two parties can divide through an early sale before resolution benefits.

For a given buyer type (b, d) and seller state s, and setting aside the offer cost, early agreement creates combined surplus $d + d_{S,s} - \Gamma_s(b)$ relative to the auction. A price can give both parties nonnegative gains exactly when

$$R_s(b) - d_{S,s} \leq p \leq w_s(b) + d \quad \Longleftrightarrow \quad d + d_{S,s} \geq \Gamma_s(b). \tag{2.5}$$

When $b < v_s$, the wedge includes the value of retaining the object. When $b \geq v_s$, it comes entirely from late competition. The next result aggregates these gains over actual equilibrium trade and includes the cost of rejected attempts.

Proposition 2.1 (Aggregate financing of early offers)**.** *In every PBE of the one-offer game, without restricting offer prices or requiring pure strategies, let I be the event that an offer is made and $A_c \subseteq I$ the event that it is accepted. Then*

$$\mathbb{E}[\{d_B + d_{S,s} - \Gamma_s(b)\}\mathbf{1}_{A_c}] \geq \kappa \Pr(I). \tag{2.6}$$

The conclusion also holds at $\kappa = 0$.

Proof. The buyer can wait and the seller can reject, so each must obtain a nonnegative expected gain relative to the auction, conditional on its information. Let g_B and g_S denote gains relative to the expected auction payoffs conditional on (b, s):

$$g_B = \mathbf{1}_{A_c}[w_s(b) + d_B - p] - \kappa \mathbf{1}_I, \qquad g_S = \mathbf{1}_{A_c}[p - C_s(b)].$$

Buyer optimality gives $\mathbb{E}[g_B \mid b, d] \geq 0$. For each seller state s and almost every offer p reached in equilibrium, seller optimality and Bayes' rule give $\mathbb{E}[g_S \mid s, p] \geq 0$: acceptance has nonnegative conditional expected gain, and rejection gives zero. The early price cancels when the two gains are added:

$$g_B + g_S = \{d_B + d_{S,s} - \Gamma_s(b)\}\mathbf{1}_{A_c} - \kappa \mathbf{1}_I.$$

Taking expectations over buyer types, the seller's state, and any randomization in offer and acceptance strategies, and using $\mathbb{E}[g_B], \mathbb{E}[g_S] \geq 0$, proves the result. □

Accepted trades must pay for unsuccessful attempts as well as for the competition forgone. Dividing (2.6) by $\Pr(A_c) > 0$ gives $\mathbb{E}[d_B + d_{S,s} - \Gamma_s(b) \mid A_c] \geq \kappa \Pr(I)/\Pr(A_c)$: the required net gain per accepted offer rises with the number of attempts per acceptance. With no resolution gains and positive attempt costs, no buyer offers. Only optimality and Bayes' rule at reached information sets enter the proof; no particular price selection or off-path belief is used.

The accounting idea extends beyond English auctions. For a specified continuation mechanism M, define the pair's opportunity cost as $\Gamma_s^M(b) = R_s^M(b) + u_s^M(b) - b$. This difference need not be nonnegative. With only the winner paying, nonnegativity follows if a late winner pays at least b and the seller retains the object only when $v_s \geq b$. In our auction the early buyer remains present and bids its value b, so a late winner must pay at least b, and the seller cannot retain the object when $b > v_s$. Appendix A.1 gives the general accounting identity and the additional continuation restriction needed for the financing argument.

2.3 Private continuation values

We maintain the following ordering in the equilibrium construction. Appendix A.2 gives sufficient primitive conditions for it, based on participation prospects, the value of retaining the object, or benefits from early agreement.

Assumption 1 (Ordered continuation values)**.** The functions w_s are continuous and strictly increasing; C_s are continuous, nondecreasing, and strictly increasing on a top interval, with $C_L(\underline{b}) > 0$. For every b, $w_L(b) \leq w_H(b)$ and $C_L(b) < C_H(b)$. The gap $w_H(b) - w_L(b)$ is weakly increasing in the buyer's valuation b, and $\Gamma_L(b)$ is weakly decreasing in b.

A higher buyer value raises willingness to pay for early ownership and weakly improves the seller's auction alternative. State H requires a higher early price and leaves the buyer with no more auction surplus than L. As the buyer's value rises, the willingness-to-pay gap across states weakly widens and the pair's competition wedge in L weakly shrinks.

After observing p, the seller updates its belief about the early buyer's value. Let $\mu_B(\cdot \mid p)$ denote this posterior. It is common across seller states at on-path prices because the buyer cannot condition its offer on s, which is independent of its type. Our equilibrium constructions also use common posteriors at unused prices; Appendix C explains this restriction for the D1 test. Seller s accepts exactly when

$$p \geq \mathcal{C}_s(p) := \int C_s(b)\, d\mu_B(b \mid p). \tag{2.7}$$

Since both states use the same posterior and $C_L(b) < C_H(b)$ for every buyer value, their posterior minimum prices satisfy $\mathcal{C}_L(p) < \mathcal{C}_H(p)$. Acceptance is therefore nested: neither

state accepts, only L accepts, or both accept. Under pure responses, two used prices with the same accepting states cannot differ, since every buyer would prefer the lower price. An offer rejected in both states loses κ and is dominated by waiting. At most two positive prices remain relevant.

Private seller information makes rejection possible without making the attempt pointless. An offer can secure a purchase from a seller with weak prospects for later competition or a large benefit from early resolution, yet be rejected by a seller with better alternatives. If the buyer instead observed s before offering, a costly offer certain to be rejected would be dominated by waiting in any pure-strategy PBE of the one-offer game.

3 Probing and Certain Acceptance

3.1 Buyer selection and prices

Denote waiting by W, a probing offer q_L by P, and a knockout offer q_H by K. Their incremental payoffs relative to waiting are

$$\begin{aligned} V_0(b,d) &= 0, \\ V_L(b,d) &= (1-\pi)[w_L(b) + d - q_L] - \kappa, \\ V_H(b,d) &= \bar{w}(b) + d - q_H - \kappa. \end{aligned} \tag{3.1}$$

Probing purchases resolution only in state L; knockout purchases it in both states. Hence the payoff from knockout rises faster with d. The waiting–probing and probing–knockout cutoffs are

$$D_L(b; q_L) = q_L - w_L(b) + \frac{\kappa}{1-\pi}, \tag{3.2}$$

$$D_H(b; q_L, q_H) = \frac{q_H - (1-\pi)q_L}{\pi} - w_H(b). \tag{3.3}$$

For given prices with $D_L < D_H$, the buyer waits below D_L, probes between the cutoffs, and chooses knockout above D_H. Probing is chosen with positive probability when its interval overlaps the support of resolution benefits: $\max\{0, D_L\} < \min\{\bar{d}, D_H\}$. Both cutoffs decline with the early buyer's valuation b. A higher-value buyer needs a smaller resolution benefit to justify either switch. Figure 3.1 shows the three payoffs for a fixed valuation and price pair.

Prices also affect the seller's inference. We select *boundary prices*: a probe makes L indifferent, and knockout makes H indifferent. Local passivity of beliefs after small undercuts justifies this selection. If the marginal accepting state earned positive rent at a used price, a sufficiently small undercut that preserved its posterior would retain

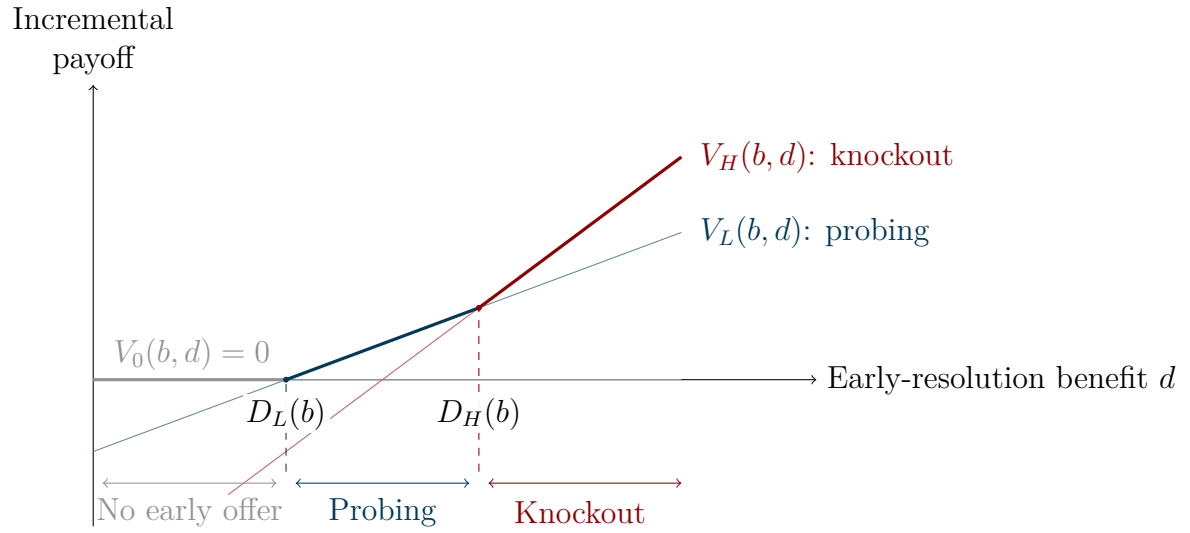


Figure 3.1: Buyer payoffs relative to waiting, holding valuation and prices fixed. The payoff slopes are 0 for waiting, $1-\pi$ for probing, and 1 for knockout. Thick segments trace the optimal action when $D_L < D_H$.

acceptance and benefit the buyer. Without this restriction, a small undercut may be attributed to a higher-value buyer, raising the seller's continuation value enough to induce rejection. Appendix C.1 states the belief restriction precisely. The zero-rent condition is relative to the buyer types selecting each offer: changing the price changes that selection and hence the seller's expected continuation value. The existence results below do not establish uniqueness of the resulting fixed point.

3.2 Probing with acceptance and rejection

For a candidate probing price q, buyers with $d \geq D_L(b;q)$ offer. With the convention $G(x) = 0$ below zero and $G(x) = 1$ above $\bar{d}$, seller L's posterior minimum early price is

$$T_L(q) = \frac{\int C_L(b)[1 - G(D_L(b;q))]f(b)\,db}{\int [1 - G(D_L(b;q))]f(b)\,db}. \tag{3.4}$$

Unspecified value integrals run over $[\underline{b}, \bar{b}]$. Under the candidate strategy, $1 - G(D_L(b;q))$ is the probability of offering q at value b, and the denominator is the total offer probability. Normalizing these weights times $f(b)$ gives the seller's posterior density conditional on observing q; $T_L(q)$ averages $C_L(b)$ under that posterior. A fixed point $q_L = T_L(q_L)$ makes L indifferent, while H prefers to continue under the same posterior. Define the top-type entry thresholds and their difference by

$$\begin{aligned} \bar{D}_L &= C_L(\bar{b}) - w_L(\bar{b}) + \frac{\kappa}{1-\pi}, \\ \bar{D}_H &= C_H(\bar{b}) - \bar{w}(\bar{b}) + \kappa, \qquad \Delta_D = \bar{D}_H - \bar{D}_L. \end{aligned} \tag{3.5}$$

For the highest-value buyer, $\bar{D}_L$ is the resolution benefit needed to prefer an L-only offer at $C_L(\bar{b})$ to waiting; $\bar{D}_H$ is the corresponding threshold for acceptance in both states at $C_H(\bar{b})$. A belief concentrated on $\bar{b}$ induces these prices, the highest minimum early prices across posteriors. Comparing the thresholds with $\bar{d}$ gives conditions for waiting and probing equilibria; D1 compatibility is checked separately.

Proposition 3.1 (Probing equilibrium)**.** *Under Assumption 1, a PBE in which every buyer waits exists if and only if*

$$\bar{d} \leq \min\{\bar{D}_L, \bar{D}_H\}. \tag{3.6}$$

Suppose instead that

$$\Delta_D > 0, \qquad \bar{D}_L < \bar{d} \leq \bar{D}_L + \Delta_D/\pi, \qquad C_L(\underline{b}) - w_L(\underline{b}) + \kappa/(1-\pi) > 0. \tag{3.7}$$

Then T_L has an interior fixed point $q_L \in (C_L(\underline{b}), C_L(\bar{b}))$. There is a PBE in which the buyer waits below $D_L(b; q_L)$ and offers q_L otherwise; L accepts and H rejects. Waiting, rejected offers, and accepted offers each have positive probability. If in addition $\bar{D}_L \geq 0$, this on-path outcome admits a D1-compatible belief completion as defined in Appendix C.

At the lower candidate price, the seller's expected minimum early price exceeds the offer; at the upper candidate price it is lower. Continuity therefore gives a price at which the seller is indifferent. The lower bound in (3.7) makes offering worthwhile for some buyers; the upper bound ensures that even the highest-value buyer with the largest resolution benefit cannot gain by paying for acceptance in both states. Appendix B.2 gives the complete argument.

For a simple participation example, let $F = U[0,1]$, $v_L = v_H = .4$, $d_{S,L} = d_{S,H} = 0$, $(n_L, n_H) = (2, 6)$, $\pi = .5$, and $\kappa = .02$. Equation (A.2) gives $\bar{D}_L = .04$ and $\bar{D}_H \simeq 0.10469$. For example, $\bar{d} = .1$ satisfies the probing bounds, and the resulting on-path outcome admits D1 beliefs. Private information about later participation alone therefore suffices for accepted and rejected attempts; seller preferences need not differ across states.

3.3 Paying for certain acceptance

The gain from choosing knockout rather than probing is

$$V_H(b,d) - V_L(b,d) = \pi[w_H(b) + d] - q_H + (1-\pi)q_L.$$

A higher resolution benefit can justify the extra payment for acceptance in state H. For each candidate price pair, the resulting action probabilities determine the seller's posterior over buyer values.

For $q = (q_L, q_H)$ with $0 < D_L(b) < D_H(b) < \bar{d}$, let $m_a(b; q)$ be the probability of action $a \in \{W, P, K\}$ conditional on valuation b, $M_a(q)$ its unconditional probability, and $h_a(b; q)$ the posterior density of b after that action:

$$\begin{aligned} m_W(b; q) &= G(D_L(b)), \\ m_P(b; q) &= G(D_H(b)) - G(D_L(b)), \\ m_K(b; q) &= 1 - G(D_H(b)), \end{aligned} \tag{3.8}$$

$$M_a(q) = \int m_a(b; q) f(b)\, db, \qquad h_a(b; q) = \frac{m_a(b; q) f(b)}{M_a(q)}. \tag{3.9}$$

The price map is

$$T_P(q) = \int C_L(b) h_P(b; q)\, db, \qquad T_K(q) = \int C_H(b) h_K(b; q)\, db. \tag{3.10}$$

At $q = T(q)$, the seller is indifferent in L after probing and in H after knockout. In state L, knockout leaves a rent: the higher price is needed to secure acceptance in the other state.

Proposition 3.2 (Probing and knockout)**.** *Suppose Assumption 1 and the price-rectangle conditions in Appendix B.3 hold. Then T has a fixed point $q_L < q_H$ supporting a PBE with*

$$p(b, d) = \begin{cases} \varnothing, & d < D_L(b), \\ q_L, & D_L(b) \le d < D_H(b), \\ q_H, & d \ge D_H(b). \end{cases} \tag{3.11}$$

Seller L accepts both prices; seller H accepts only q_H. For each valuation b, positive intervals of resolution benefits lead to waiting, probing, and knockout. The on-path outcome admits both a locally passive PBE completion and a D1-compatible completion in the contingent-plan representation, which need not use the same off-path beliefs.

The conditions bracket both prices: at each lower boundary the relevant posterior minimum price exceeds the offer, and at each upper boundary it falls below it. Continuity yields a pair of consistent prices. These are sufficient conditions for existence, not global uniqueness. They require $0 < D_L(b) < D_H(b) < \bar{d}$ at every valuation, which is stronger than positive aggregate probabilities of the three actions. Appendix B.3 states the inequalities, illustrates their geometry, and checks deviations to every unused price.

The observable probabilities are M_W for no offer, πM_P for rejection, and $(1 - \pi)M_P + M_K$ for successful pre-emption. Knockout adds certain acceptance to a mechanism that already generates attempts and rejections through probing. The Online Appendix also characterizes equilibria in which all submitted offers are accepted.

4 Restricting Pre-emption: Allocation and Welfare

Compare seller discretion, denoted D, with a ban on early agreement, denoted B. The Norwegian inquiry NOU 2021:7 considered restrictions on kupping, including possible prohibitions, motivated in part by the loss of open competition (Finansdepartementet, 2021, Section 6.6.3 and Chapter 22). It also emphasized the difficulty of preventing direct agreements. Our counterfactual is an effective ban that sends every buyer to the same terminal auction, holding participation, values, and costs fixed and imposing no enforcement cost. The comparison therefore applies to sales with an identifiable early-offer stage and an enforceable restriction on early agreement.

The early buyer weakly prefers discretion because waiting remains available. The seller also weakly prefers it at the offer stage, accepting only when its conditional expected payoff warrants acceptance. In the boundary-price equilibrium with positive knockout mass, its ex ante gain is strictly positive: a knockout offer pays for state H even when the seller is in L. Later buyers weakly prefer the ban, which restores their opportunity to win. Appendix D derives these payoff comparisons.

The allocative benefit of preserving the auction after a proposed early sale to value b is

$$L_s(b) = \mathbb{E}_s[\max\{\max\{v_s, Y_{1N}\} - b, 0\}]. \tag{4.1}$$

Count v_s as the social value of the seller retaining the object, rather than a later transfer; count d_B and $d_{S,s}$ at their monetary utility values and κ as a real effort cost. The highest late value determines whether the auction improves allocation. By contrast, the second-highest late value determines the contracting parties' wedge. Their difference is the expected surplus accruing to late buyers, which is positive only when a late buyer wins,

$$\Lambda_s(b) = \mathbb{E}_s[\max\{Y_{1N} - \max\{v_s, b, Y_{2N}\}, 0\}], \qquad L_s(b) = \Gamma_s(b) + \Lambda_s(b). \tag{4.2}$$

The auction allocates the object efficiently and divides its value among the seller, the early buyer, and late buyers. Subtracting b gives $L_s(b) = [R_s(b) + u_s(b) - b] + \Lambda_s(b)$. Early agreement eliminates both the opportunity cost borne by its participants and the surplus left to an excluded buyer.

Proposition 4.1 (Welfare effect of a ban)**.** *Holding the terminal environment fixed, the welfare difference is*

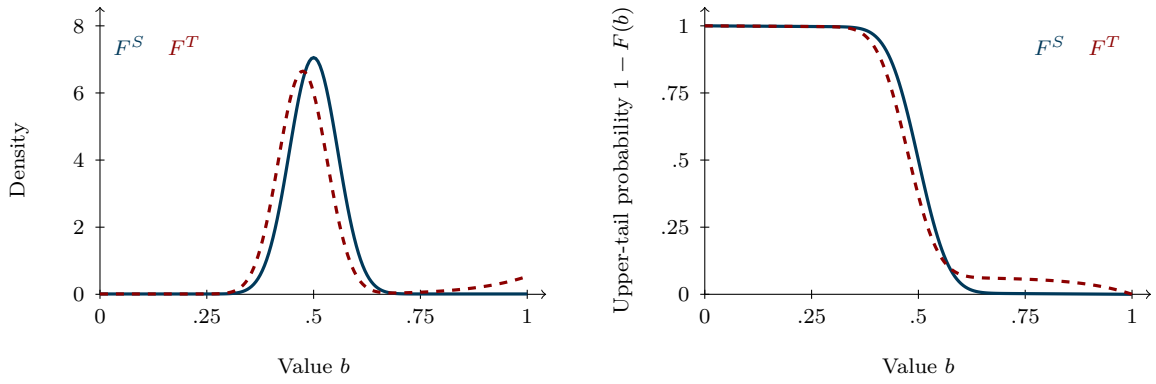
$$W^B - W^D = \mathbb{E}[\kappa \mathbf{1}_I + \mathbf{1}_{A_c}\{L_s(b) - d_B - d_{S,s}\}]. \tag{4.3}$$

A ban raises surplus precisely when the allocative gains and avoided offer costs exceed the resolution gains from accepted trades.

Proof. After a rejected offer, both regimes reach the same auction, but the ban saves κ. After an accepted offer, it also replaces an allocation to b by one worth $\max\{v_s, b, Y_{1N}\}$, while forgoing the resolution benefits $d_B + d_{S,s}$. No offer leaves the regimes identical. The purchase price is a transfer from the winning buyer to the seller. Summing payoffs and averaging over the three outcomes gives (4.3). □

The financing result concerns the contracting parties: gains from accepted offers must cover rejected attempts in equilibrium. The welfare result also counts $\Lambda_s(b)$, the surplus of excluded late buyers. Larger resolution benefits make early contracting easier to sustain; larger excluded surplus makes preserving competition more valuable socially.

Table 4.1 gives two examples with identical parameters except for the common value distribution of early and late buyers. Both distributions have mean .5. The first is concentrated around the mean; the second adds probability near the upper end while preserving that mean. In both cases, buyers wait, probe, or choose knockout. Appendix D.2 specifies the distributions and supplies bounds certifying equilibrium existence and the welfare signs. Figure 4.1 shows the two distributions.



(a) Densities **(b)** Survival functions

Figure 4.1: Value distributions in the ban comparison. Both have mean .5; F^T places more probability near the upper end. The panels show the density and upper-tail probability. The distribution of resolution benefits is held fixed.

Table 4.1: The Welfare Effect of a Ban with Mean-Matched Values

	Concentrated F^S	Upper tail F^T
Rejected attempt ($\Pr(I \setminus A_c)$)	.05631	.06674
Successful pre-emption ($\Pr(A_c)$)	.15435	.06839
Avoided attempt costs ($\kappa \Pr(I)$)	.00421	.00270
Allocative gain ($\mathbb{E}[\mathbf{1}_{A_c} L_s(b)]$)	.00388	.00276
Forgone resolution gains ($\mathbb{E}[\mathbf{1}_{A_c}(d_B + d_{S,s})]$)	.01162	.00516
$W^B - W^D$	-.00353	.00030

Notes: The first two rows are probabilities; welfare entries are per object. The seller's resolution benefit is zero. Reported point estimates are rounded numerical approximations; interval bounds in Appendix D.2 establish the signs, rather than all displayed decimal places.

Discretion raises welfare under F^S, while the ban raises it under F^T. A single exceptional late value can increase the allocative value of waiting without raising the auction price proportionately. Excluded late-buyer surplus per accepted early sale rises from .01909 to .03689 between the examples.[3] The distributions also yield different equilibrium prices and action probabilities: successful pre-emption occurs with probability .15435 under F^S and .06839 under F^T.

The ban keeps late participation fixed. If cancellation saves preparation costs for late buyers, these savings are an additional benefit of early agreement and must be subtracted from the measured benefit of a ban. Endogenous participation would also change the set of later competitors; assessing that response requires a model of entry.

5 Conclusion

A private offer before a competitive sale purchases early resolution at the cost of possible competition. Seller information makes this purchase uncertain: a probing price is accepted when continuation prospects are weak and rejected when they are strong. Buyers with sufficiently large resolution benefits may instead pay for acceptance in both states. The offer also reveals information about a buyer who returns after rejection, so prices and continuation expectations must be jointly consistent.

The same mechanism creates a policy tradeoff. Pre-emption imposes a negative externality on late buyers: the contracting parties internalize the auction revenue they forgo, but not all the surplus available to excluded buyers. Holding the auction environment fixed, an effective ban protects that surplus and saves attempt costs while removing gains from early agreement. English auctions make these components explicit. Studying other

[3]Section 3 of the Calculation Workbook reports the unconditional surplus numerators and completed-sale probabilities used in these ratios.

competitive-sale and search mechanisms would clarify how the tradeoff changes when the continuation process differs.

References

Banks, J. S. and Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica*, 55(3):647–661.

Budish, E. B. and Takeyama, L. N. (2001). Buy prices in online auctions: Irrationality on the internet? *Economics Letters*, 72(3):325–333.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.

Bulow, J. and Klemperer, P. (2009). Why do sellers (usually) prefer auctions? *American Economic Review*, 99(4):1544–1575.

Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2):179–221.

Cramton, P. and Schwartz, A. (1991). Using auction theory to inform takeover regulation. *The Journal of Law, Economics, and Organization*, 7(1):27–53.

Finansdepartementet (2021). NOU 2021: 7: Trygg og enkel eiendomsmegling: Evaluering av eiendomsmeglingsloven og forslag til en fremtidsrettet regulering. Norges offentlige utredninger 2021: 7, Finansdepartementet, Oslo.

Finansdepartementet (2025a). Endringer i eiendomsmeglingsloven og forskrift om eiendomsmegling trer i kraft 1. juli. Government announcement. 10 June 2025.

Finansdepartementet (2025b). Prop. 41 L (2024–2025): Endringer i eiendomsmeglingsloven (budgivning, egnethetskrav og sanksjoner mv.). Proposition to the Storting.

Fishman, M. J. (1988). A theory of preemptive takeover bidding. *The RAND Journal of Economics*, 19(1):88–101.

Flått, M., Olaussen, J. O., Oust, A., and Sønstebø, O. J. (2022). Hijacking the auction: seller's or buyer's curse? *Journal of European Real Estate Research*, 15(3):387–404.

Gallien, J. and Gupta, S. (2007). Temporary and permanent buyout prices in online auctions. *Management Science*, 53(5):814–833.

Grebe, T., Ivanova-Stenzel, R., and Kröger, S. (2015). "Buy-It-Now" or "Sell-It-Now" Auctions: Effects of changing bargaining power in sequential trading mechanisms. IZA Discussion Paper 9566, Institute for the Study of Labor (IZA).

Grebe, T., Ivanova-Stenzel, R., and Kröger, S. (2016). "Buy-It-Now" or "Sell-It-Now" Auctions: Effects of changing bargaining power in sequential trading mechanisms. *Economics Letters*, 142:27–30.

Khezr, P. and Menezes, F. M. (2018). Auctions with an asking price. *International Journal of Game Theory*, 47(4):1329–1350.

Liu, T., Officer, M. S., and Tu, D. (2023). Negotiation, auction, or negotiauction?! evidence from the field. Working paper, December. American Finance Association conference version.

Mathews, T. (2004). The impact of discounting on an auction with a buyout option: A theoretical analysis motivated by eBay's "buy-it-now" feature. *Journal of Economics*, 81(1):25–52.

Mathews, T. and Katzman, B. E. (2006). The role of varying risk attitudes in an auction with a buyout option. *Economic Theory*, 27(3):597–613.

Prestige Equipment (2026). Surplus to continuing operations of Johnson Controls: Two-day online auction. Auction brochure. Federalsburg, Maryland, 3–4 March; pre-auction offers invited on major assets, p. 1; accessed 18 September 2026.

Real Estate Authority (n.d.). Auctions. New Zealand Real Estate Authority website. Accessed 28 July 2026.

A Auction Payoffs and Sources of Seller Information

With n independent late buyers, the probability that all their valuations are at most t is $F(t)^n$. Averaging over the state-dependent participation distribution gives $\sum_{n\geq 0} Q_s(n)F(t)^n = \phi_s(F(t))$, where $\phi_s(x) = \mathbb{E}_{Q_s}[x^N]$, with $x^0 = 1$, is the probability-generating function of late participation. Integrating the early buyer's winning probability gives its willingness to pay,

$$w_s(b) = \begin{cases} b, & b \leq v_s, \\ v_s + \displaystyle\int_{v_s}^{b} [1 - \phi_s(F(t))]\, dt, & b > v_s. \end{cases} \tag{A.1}$$

For $b \leq v_s$, the early buyer has no auction surplus, so its willingness to pay is $w_s(b) = b$. The no-sale probability is $\mathbf{1}\{b < v_s\}\phi_s(F(v_s))$ conditional on (b, s). The unconditional probability in state s is $F(v_s)\phi_s(F(v_s))$; stronger participation lowers it, while a higher fallback raises it.

To prove (2.4), set $Z = \max\{v_s, Y_{2N}\}$ and $M = \min\{b, \max\{v_s, Y_{1N}\}\}$. The terminal price is $\max\{Z, M\}$, and $w_s(b) = \mathbb{E}_s[M]$. If $Z \leq b$, then $Z \leq M$; otherwise $M = b$. Thus the price less M equals $\max\{Z - b, 0\}$ in both cases. The wedge is weakly decreasing in b, and $\Gamma_s(\bar{b}) = 0$, so $C_s(\bar{b}) = w_s(\bar{b}) - d_{S,s}$.

The seller knows the distribution Q_s, but need not know eventual participation. When its state instead reveals a fixed count $n \geq 1$, and values are uniform on $[0, 1]$, direct integration gives

$$\begin{aligned} w_{n,v}(b) &= \begin{cases} b, & b \leq v, \\ b - \dfrac{b^{n+1} - v^{n+1}}{n+1}, & b > v, \end{cases} \\ R_{n,v}(b) &= \begin{cases} \dfrac{n-1}{n+1} + v^n - \dfrac{n-1}{n+1}v^{n+1}, & b \leq v, \\ \dfrac{n-1}{n+1} + b^n - \dfrac{n}{n+1}b^{n+1} + \dfrac{v^{n+1}}{n+1}, & b > v. \end{cases} \end{aligned} \tag{A.2}$$

A.1 The wedge under other continuation mechanisms

Fix a continuation mechanism M, quasi-linear payoffs, and any history on which the mechanism conditions. Let x_B, x_j, x_S indicate allocation to the early buyer, late buyer j, or the seller, respectively, with $x_B + \sum_j x_j + x_S = 1$. Let t_B, t_j be payments to the seller; assume no transfers to outside parties and set aside resolution benefits and continuation

costs. The seller receives $r = v_s x_S + t_B + \sum_j t_j$, and the early buyer receives $u = bx_B - t_B$. Their combined payoff satisfies

$$r + u - b = (v_s - b)x_S + \sum_j (t_j - bx_j).$$

Taking conditional expectations gives $\Gamma_s^M(b) = R_s^M(b) + u_s^M(b) - b$. If only the winner pays, the realized difference is zero when the early buyer wins; it equals the winning late buyer's payment minus b when a late buyer wins, and $v_s - b$ when the seller retains the object. This proves the sufficient conditions for nonnegativity stated in the main text. Efficient allocation alone does not ensure the payment condition.

The financing argument carries over if, conditional on primitive types, rejecting an offer leaves each party with the same expected continuation payoff as waiting. The definitions of g_B and g_S then apply with the corresponding continuation values, and their sum gives (2.6) with Γ_s^M, which may be signed. If a rejected offer changes who participates, how they bid, or the seller's threshold, the resulting continuation-payoff changes must also enter the accounting. In our model, the early buyer remains in the auction and bids its value after either waiting or rejection; late participation is exogenous, the other buyers also bid their values, and the seller's threshold is fixed. Hence, conditional on (b, s), the expected auction payoffs are the same after the two histories.

A.2 Sources of seller information

Suppose $Q_s(N \geq 1) > 0$ and $d_{S,L} < R_L(\underline{b})$. The following three structures satisfy Assumption 1. Counts are finite almost surely, and the formulas below apply wherever the displayed derivatives exist; they also establish the corresponding monotonicity at kinks.

Participation. Hold $v_s = v_0$ and $d_{S,s} = d_S$ fixed. Let Q_H first-order stochastically dominate Q_L, with a strict tail inequality $\Pr_H(N \geq k) > \Pr_L(N \geq k)$ for some $k \geq 2$. Couple the participation draws so that $N_H \geq N_L$, using the same valuations for buyers present in both states. The highest and second-highest late values can then only increase, giving $w_H \geq w_L$ and $C_H \geq C_L$. For $b > v_0$,

$$w_H(b) - w_L(b) = \int_{v_0}^{b} [\phi_L(F(t)) - \phi_H(F(t))]\, dt,$$

which is strictly increasing in the buyer's value b for $v_0 < b \leq \bar{b}$. The extra tail condition makes the seller inequality strict even below the fallback; its proof follows below.

Fallback. Hold Q and d_S fixed, and let $v_H > v_L$. Then

$$C_H(b) - C_L(b) = \int_{v_L}^{v_H} \Pr(\max\{Y_{2N}, \min\{b, Y_{1N}\}\} \leq t)\ dt > 0.$$

The integrand is weakly decreasing in b. The willingness gap is zero below v_L and equals $\int_{v_L}^{\min\{b,v_H\}} \phi(F(t))\,dt$ above it, increasing to a plateau at v_H.

Early resolution. Hold Q and v fixed, and let $d_{S,L} > d_{S,H} \geq 0$. Then $w_H = w_L$ and $C_H - C_L = d_{S,L} - d_{S,H} > 0$. The seller's resolution benefit changes its minimum early price without changing the buyer's auction payoff. Participation and fallback information can change both.

These changes can also be combined in the same direction: stronger participation, a higher fallback, and smaller seller resolution benefits in H. For $v_H \geq v_L$, the willingness difference is

$$w_H(b) - w_L(b) = \begin{cases} 0, & b \leq v_L, \\ \displaystyle\int_{v_L}^{b} \phi_L(F(t))\,dt, & v_L < b \leq v_H, \\ \displaystyle\int_{v_L}^{v_H} \phi_L(F(t))\,dt + \int_{v_H}^{b} [\phi_L(F(t)) - \phi_H(F(t))]\,dt, & b > v_H. \end{cases} \tag{A.3}$$

This is weakly increasing. In each state,

$$R_s'(b) = \begin{cases} 0, & b < v_s, \\ [1 - F(b)]\phi_s'(F(b)) > 0, & v_s < b < \bar{b}. \end{cases} \tag{A.4}$$

Equations (A.1) and (A.4) give regularity and monotonicity. Equation (2.4) also shows that $\Gamma_L(b)$ is weakly decreasing in the buyer's value, as required by Assumption 1. Any of the strict state changes just specified makes $C_H > C_L$ everywhere. Thus the same equilibrium argument covers all three sources of seller information without requiring that they have identical effects on buyer selection.

Under the aligned primitive changes, monotone coupling gives $\Gamma_H(b) \geq \Gamma_L(b)$. Together with $d_{S,H} \leq d_{S,L}$, this orders the bilateral thresholds $\Gamma_s(b) - d_{S,s}$. As d rises, condition (2.5) holds in neither state, in L alone, or in both; strict inequality permits strict gains for both parties. These comparisons hold (b, s) fixed. In the game, the seller evaluates offers using its posterior over buyer values, while the buyer offers without observing s.

Strict ordering under participation

For $x \in (0, 1)$, let $a_n(x)$ be the probability that at least two of n draws exceed quantile x. Since $a_0 = a_1 = 0$ and $a_n - a_{n-1} = (n-1)(1-x)^2 x^{n-2} > 0$ for $n \geq 2$, the tail-sum identity gives

$$\Pr_H(Y_{2N} > t) - \Pr_L(Y_{2N} > t) = \sum_{k=2}^{\infty} (k-1)[1 - F(t)]^2 F(t)^{k-2} \{\Pr_H(N \geq k) - \Pr_L(N \geq k)\}.$$

For $v_0 < t < \bar{b}$, this is strictly positive under the stated tail condition. If $b \leq v_0$, integrating above v_0 gives $C_H(b) > C_L(b)$. If $v_0 < b < \bar{b}$, integrate above b, noting that the difference below b is weakly positive. At $b = \bar{b}$, a nontrivial count shift strictly increases the probability that at least one late value exceeds each $t \in (v_0, \bar{b})$. This establishes strictness at the endpoint as well. A shift only from zero to one late buyer cannot produce strict ordering below the fallback, because the terminal price there remains v_0.

B Proofs of the Equilibrium Results

B.1 Prices and action support

Nested acceptance limits the number of used prices. The following condition also ensures a positive mass of waiting buyers:

$$\chi_L = C_L(\underline{b}) - w_L(\underline{b}) + \frac{\kappa}{1-\pi} > 0, \qquad \chi_H = C_H(\underline{b}) - \bar{w}(\underline{b}) + \kappa > 0. \tag{B.1}$$

These inequalities require the least-valued patient buyer to prefer waiting even at the lowest possible accepted price. Proposition 3.1 uses the first inequality; Proposition 3.2 instead imposes the stronger cutoff interiority below.

Lemma B.1 (On-path prices and action support)**.** *Suppose $\kappa > 0$. Restrict attention to pure-strategy PBE, with seller indifference resolved in favor of acceptance. All on-path probing prices then coincide almost surely, and all on-path knockout prices coincide almost surely. No offer rejected in both states is submitted with positive probability. Under (B.1), every such PBE has positive waiting mass. Thus, denoting the unique probing and knockout price actions by P and K when present, the action support is $\{W\}$, $\{W, P\}$, $\{W, K\}$, or $\{W, P, K\}$.*

Proof. The common posterior and $C_L < C_H$ imply nested acceptance. Any price rejected by both states is dominated by waiting. For two prices inducing the same nonempty acceptance pattern, every buyer strictly prefers the lower; thus the price is constant almost surely within that pattern, even if the candidate price rule has a continuous range. Seller optimality gives $q_L \geq C_L(\underline{b})$ and $q_H \geq C_H(\underline{b})$. At $(\underline{b}, 0)$, offer gains are bounded above by $-(1-\pi)\chi_L$ and $-\chi_H$, respectively. Continuity and positive type densities give a positive neighborhood of waiting types. The two price bounds then leave exactly the four possible supports stated. □

B.2 Probing

Proof of Proposition 3.1. For sufficiency of the no-early-offer condition, assign every off-path price to $b = \bar{b}$. The cheapest prices inducing L-only and both-state acceptance are

then $C_L(\bar{b})$ and $C_H(\bar{b})$. Their payoffs are increasing in (b, d), so the two inequalities in (3.6) deter every buyer. The same posterior survives the D1 test defined in Appendix C, which describes each seller response by an acceptance probability in each state. For every such response optimal under some common posterior, the gain from the unused price is weakly increasing in d and in each $w_s(b)$; rejection in both states loses κ for every type. Hence no other type can D1-dominate $(\bar{b}, \bar{d})$.

For necessity, first suppose $\bar{d} > \bar{D}_H$. A price just above $C_H(\bar{b})$ is accepted in both states under every posterior and is profitable for types near $(\bar{b}, \bar{d})$. If instead $\bar{d} > \bar{D}_L$, a price just above $C_L(\bar{b})$ is accepted by seller L under every posterior and its L-only payoff is positive for the top type. If seller H also accepts under the induced posterior, this only raises that type's payoff: at the limiting price,

$$w_H(\bar{b}) + \bar{d} - C_L(\bar{b}) = w_H(\bar{b}) - w_L(\bar{b}) + \bar{d} + d_{S,L} > 0,$$

where $R_L(\bar{b}) = w_L(\bar{b})$. Thus violation of either inequality creates a profitable deviation even when continuation-value supports overlap.

Now impose (3.7). Its strict lower bound rules out no-early-offer PBE. On $[C_L(\underline{b}), C_L(\bar{b})]$, the probing map is continuous wherever its denominator is positive. At the lower endpoint its posterior mean is strictly above $C_L(\underline{b})$: indeed, $\bar{d} > \bar{D}_L$ and $C_L(\underline{b}) < C_L(\bar{b})$ imply

$$D_L(\bar{b}; C_L(\underline{b})) < \bar{D}_L < \bar{d},$$

so a positive interval of high values on which $C_L(b) > C_L(\underline{b})$ probes. At the upper endpoint, $D_L(\bar{b}; C_L(\bar{b})) = \bar{D}_L < \bar{d}$, so a positive interval of high-b types probes and the posterior mean is strictly below $C_L(\bar{b})$. The intermediate value theorem gives an interior fixed point q_L. At that point the cutoff $D_L(b; q_L)$ implements the stated buyer strategy, Bayes' rule makes seller L indifferent at q_L, and seller H rejects because, under the same probe posterior μ_B,

$$\int C_H(b)\, d\mu_B(b) > \int C_L(b)\, d\mu_B(b) = q_L.$$

For unused offers $p < q_L$ sufficiently close to q_L, assign the seller the on-path posterior following a probe. Seller L's minimum early price is then q_L, so an undercut is rejected by both seller states. Assign every remaining unused price to $b = \bar{b}$. A lower price is rejected, while a higher price that only seller L accepts is dominated by q_L. The cheapest price accepted in both states is $C_H(\bar{b})$. Relative to probing, its payoff satisfies

$$\begin{aligned}
&\bar{w}(b) + d - C_H(\bar{b}) - \kappa - \{(1-\pi)[w_L(b) + d - q_L] - \kappa\} \\
&\quad = \pi[w_H(b) + d] + (1-\pi)q_L - C_H(\bar{b}) \\
&\quad \leq \pi[w_H(\bar{b}) + \bar{d}] + (1-\pi)C_L(\bar{b}) - C_H(\bar{b}) \leq 0,
\end{aligned}$$

where the final inequality is equivalent to the upper bound on $\bar{d}$ in (3.7). Thus no knockout deviation gains, including for a type that waits, since a positive knockout payoff would imply a positive probing payoff.

The final condition in (3.7) gives $D_L(\underline{b}; q_L) > 0$, while $\bar{d} > \bar{D}_L > D_L(\bar{b}; q_L)$. Both waiting and probing therefore have positive probability. Since both continuation states have positive probability, probing generates both rejection and acceptance with positive probability.

□

B.3 The full menu

Boundary-price fixed points lie in

$$\mathcal{Q} = [C_L(\underline{b}), C_L(\bar{b})] \times [C_H(\underline{b}), C_H(\bar{b})]. \tag{B.2}$$

Let $\mathcal{Q}_0 = [\ell_L, u_L] \times [\ell_H, u_H] \subseteq \mathcal{Q}$, with both intervals nondegenerate. The price-rectangle conditions in Proposition 3.2 require, for every $q \in \mathcal{Q}_0$ and every b,

$$0 < D_L(b; q_L) < D_H(b; q) < \bar{d}, \tag{B.3}$$

and the inward face inequalities

$$\begin{aligned} &T_P(\ell_L, q_H) > \ell_L, && T_P(u_L, q_H) < u_L && \text{for all } q_H \in [\ell_H, u_H], \\ &T_K(q_L, \ell_H) > \ell_H, && T_K(q_L, u_H) < u_H && \text{for all } q_L \in [\ell_L, u_L]. \end{aligned} \tag{B.4}$$

For the deviation comparisons define

$$\widetilde{D}_H(b; q_H) = q_H - \bar{w}(b) + \kappa, \tag{B.5}$$

$$Z(b; q) = q_H - q_L - \pi[w_H(b) - w_L(b)] - \frac{\pi\kappa}{1-\pi}. \tag{B.6}$$

These give

$$D_H - D_L = Z/\pi, \quad \widetilde{D}_H - D_L = Z, \quad D_H - \widetilde{D}_H = (1-\pi)Z/\pi. \tag{B.7}$$

Thus (B.3) implies $Z > 0$. By the monotonicity of w_L, w_H and $w_H - w_L$, interiority is equivalent to the three endpoint checks

$$\begin{aligned} \ell_L - w_L(\bar{b}) + \frac{\kappa}{1-\pi} &> 0, \\ \frac{u_H - (1-\pi)\ell_L}{\pi} - w_H(\underline{b}) &< \bar{d}, \\ \ell_H - u_L - \pi\big[w_H(\bar{b}) - w_L(\bar{b})\big] - \frac{\pi}{1-\pi}\kappa &> 0. \end{aligned} \tag{B.8}$$

These inequalities ensure all three action intervals at every match value, uniformly over $\mathcal{Q}_0$. Seller acceptance remains nested by the common-posterior ordering in Section 2.

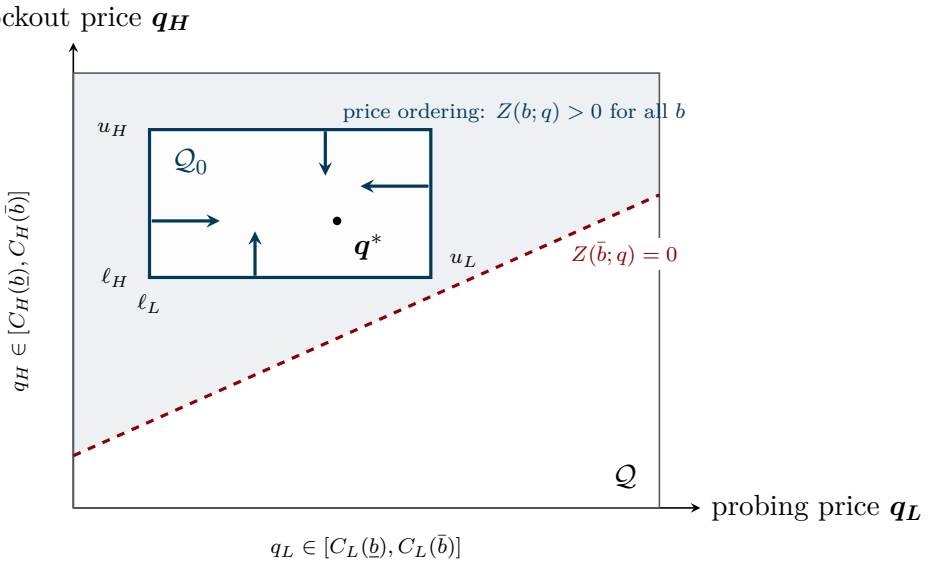


Figure B.1: Price-space geometry for the three-action fixed point. The outer rectangle $\mathcal{Q}$ contains all boundary-price candidates. Since $w_H(b) - w_L(b)$ is weakly increasing, $Z(b;q) > 0$ for every b is the region above $Z(\bar{b};q) = 0$. The proof places $\mathcal{Q}_0$ inside this region. The horizontal arrows show the residual $T_P(q) - q_L$ on the two vertical faces, and the vertical arrows show $T_K(q) - q_H$ on the two horizontal faces. Their inward directions imply a fixed point $\boldsymbol{q}^*$. Other interiority conditions are omitted.

Proof of Proposition 3.2. For every $q \in \mathcal{Q}_0$, condition (B.3) and the positive density of G imply that all three actions have positive probability for every b. The conditional action probabilities in (3.8) and their integrals $M_a(q)$ are continuous in q. Since each $M_a(q) > 0$, the posterior densities in (3.9) and the price map T are continuous as well. Let $\Psi(q) = T(q) - q$. Condition (B.4) makes the first coordinate of Ψ positive and negative on the lower and upper q_L-faces, respectively, and does the same for the second coordinate on the corresponding q_H-faces. The Poincaré–Miranda theorem therefore gives a zero of Ψ in $\mathcal{Q}_0$; the strict face inequalities put it in the interior.

Fix such a point. At q_L, seller L's posterior minimum early price is

$$\int C_L(b) h_P(b;q)\, db = q_L,$$

so it accepts. Seller H rejects because

$$\int C_H(b)h_P(b;q)\,db > \int C_L(b)h_P(b;q)\,db = q_L.$$

At q_H, seller H accepts at equality, while seller L accepts strictly because

$$\int C_L(b)h_K(b;q)\,db < \int C_H(b)h_K(b;q)\,db = q_H.$$

Finally, $Z(b;q) > 0$ implies $q_H - q_L > \pi[w_H(b) - w_L(b)] + \pi\kappa/(1-\pi) > 0$, so $q_L < q_H$.

For each b, equation (B.7) and $Z(b;q) > 0$ imply

$$D_L(b) < \widetilde{D}_H(b) < D_H(b).$$

Below $D_L(b)$, both submitted-offer payoffs are negative. Between $D_L(b)$ and $\widetilde{D}_H(b)$, only probing is profitable. Between $\widetilde{D}_H(b)$ and $D_H(b)$, both offers are profitable and probing yields the larger payoff. Above $D_H(b)$, knockout yields the larger payoff. These comparisons prove (3.11).

It remains to rule out off-path deviations. First consider sufficiently small, disjoint left neighborhoods of the two on-path prices. After an undercut of q_L, retain the on-path probing posterior. Seller L's minimum early price is then $q_L > p$, so both states reject. After an undercut of q_H, retain the on-path knockout posterior. Seller H's minimum early price is $q_H > p$, so H rejects; if seller L accepts, the deviation induces the same L-only response as the strictly cheaper on-path probe q_L. No local undercut is profitable, and these posteriors converge to their on-path counterparts as required by Assumption 2.

At every remaining off-path price, concentrate the posterior on $b = \bar{b}$. A price below $C_L(\bar{b})$ is rejected by both sellers, a price in $[C_L(\bar{b}), C_H(\bar{b}))$ is accepted only by seller L, and a price at least $C_H(\bar{b})$ is accepted by both. A rejected deviation yields $-\kappa$. Since

$$q_L = \int C_L(b)h_P(b;q)\,db \le C_L(\bar{b}),$$

any off-path price accepted only in state L is weakly more expensive than probing. Likewise,

$$q_H = \int C_H(b)h_K(b;q)\,db \le C_H(\bar{b}),$$

so any off-path price accepted in both states is weakly more expensive than knockout. No buyer can improve on its best on-path action. Seller actions are sequentially rational by (2.7), completing the proof.

The D1-compatible completion is established in Appendix C. □

C Price Selection and Refinement

C.1 Local passivity

Assumption 2 (Locally passive beliefs)**.** For price selection, let $q > 0$ be an active offer and let $\mu_B(\cdot \mid q)$ be the seller's on-path posterior over the buyer's match value. After unused prices p approaching q from below, the posterior is locally passive:

$$\mu_B(\cdot \mid p) \Rightarrow \mu_B(\cdot \mid q) \qquad \text{as } p \uparrow q,$$

where $\Rightarrow$ denotes weak convergence.

Assumption 2 implies that every active price leaves the seller state at the acceptance margin indifferent. To see this, suppose that state instead earned strictly positive expected rent at an active price q. For all sufficiently close undercuts $p < q$, local passivity would keep that state's expected minimum early price below p, so it would still accept. A probing undercut would retain acceptance by seller L and rejection by seller H; a knockout undercut would retain acceptance by both states. Any buyer using q could then induce the same seller responses at the lower price p, a profitable deviation. Hence the probing price satisfies $q_L = \mathcal{C}_L(q_L)$, while the knockout price satisfies $q_H = \mathcal{C}_H(q_H)$.

Local passivity selects prices that leave the marginal accepting seller indifferent. Criterion D1 instead restricts what the seller may infer about a buyer who offers an unused price: it compares which types would gain from that offer under different seller responses (Cho and Kreps, 1987; Banks and Sobel, 1987). To make this comparison when the seller has private information, describe a response by its acceptance probability in each state. A D1-compatible completion of the reported outcome need not retain the locally passive beliefs used to construct it.

C.2 The D1 test with state-contingent seller responses

Lemma C.1 (Contingent-plan representation for criterion D1)**.** *In the auction described in Section 2, each buyer optimally bids its value regardless of competing bids or the seller's threshold. Replace that stage by its expected payoffs $u_s(b)$ and $R_s(b)$ when examining early-offer deviations. In this reduced game, the buyer chooses whether to offer and at what price, does not observe s, has a type independent of s, and has no further decision after the seller responds. After any price p, the one-offer game is outcome- and payoff-equivalent to a signaling game in which the seller chooses a contingent acceptance plan $a = (a_L, a_H) \in [0,1]^2$. Criterion D1 in this representation is the strict-gain set-inclusion test used below.*

Proof. A contingent acceptance plan specifies the probability $a_s \in [0,1]$ that the seller accepts p in state s. Thus $(1,0)$ means acceptance in L and rejection in H, while $(1,1)$ means acceptance in both states. The plan records both possible responses; the seller still observes its own state before responding. Let μ be a common posterior over buyer type after p, and let $\lambda_s > 0$ be the prior probability of state s. After subtracting the constant $\sum_s \lambda_s d_{S,s}$, the seller's expected payoff from plan a is

$$\sum_{s\in\{L,H\}} \lambda_s \left\{ a_s p + (1-a_s) \int C_s(b)\, d\mu(b,d) \right\}.$$

Each acceptance probability affects only the corresponding state's payoff. Thus $a_s = 1$ is optimal when $p > \int C_s(b)\, d\mu(b,d)$, $a_s = 0$ is optimal when the inequality is reversed, and any acceptance probability is optimal at equality. A plan is a best response exactly when each state's response is optimal at the original seller information set (p,s). These acceptance probabilities also give the buyer the same expected payoff as in the original game.

Because the buyer's type is independent of s and its offer cannot depend on that unobserved state, Bayes' rule gives $\mu(\cdot \mid p,s) = \mu(\cdot \mid p)$ at on-path prices. We maintain this common-posterior restriction after unused prices. It also holds for posteriors obtained by allowing each buyer type a small probability of making the unused offer. The seller's two states may respond differently because their continuation values differ, while drawing the same inference about the buyer.

Call a plan *rationalizable* after p if it is optimal under at least one common posterior over buyer types. Let $\mathcal{A}(p)$ be the set of these plans, including random acceptance when a seller state is indifferent. These are the responses compared by D1. For buyer type $\theta = (b,d)$, let $U^*(\theta)$ be its equilibrium payoff and $u_\theta(p,a)$ its payoff from the unused offer p under plan a. The strict-gain set $\mathcal{D}(\theta,p)$ contains the plans making that deviation profitable; the indifference set $\mathcal{D}^0(\theta,p)$ contains those giving the equilibrium payoff:

$$\mathcal{D}(\theta,p) = \{a \in \mathcal{A}(p) : u_\theta(p,a) > U^*(\theta)\},$$
$$\mathcal{D}^0(\theta,p) = \{a \in \mathcal{A}(p) : u_\theta(p,a) = U^*(\theta)\}.$$

The D1 rule excludes θ as an explanation for the unused offer if another type θ' satisfies

$$\mathcal{D}(\theta,p) \cup \mathcal{D}^0(\theta,p) \subsetneq \mathcal{D}(\theta',p).$$

Every rationalizable response that makes θ weakly willing to deviate then makes θ' strictly willing, and at least one additional response makes θ' strictly willing. The unused offer is attributed only to types surviving this comparison. Our D1 statements use this test

in the stated contingent-plan representation; payoff equivalence does not imply that a different extensive-form representation yields the same refinement. □

C.3 D1 compatibility of the probing equilibrium

For the sufficient D1 statement in Proposition 3.1, impose

$$\bar{D}_L \geq 0. \tag{C.1}$$

To verify D1 for $\kappa > 0$, first identify the possible seller responses. Since $C_L(b) < C_H(b)$ for every valuation and both states use the same posterior, seller H can accept with positive probability only if seller L accepts with certainty. Consequently, the rationalizable plans satisfy

$$\mathcal{A}(p) \subseteq \{(x, 0) : x \in [0, 1]\} \cup \{(1, y) : y \in [0, 1]\}, \tag{C.2}$$

The first branch allows L to accept with probability x while H rejects; the second allows H to accept with probability y while L accepts for certain. For an undercut $p < q_L$, the wait–probe marginal type $(b, D_L(b; q_L))$ requires, on the first branch, the acceptance probability

$$x^*(p) = \frac{\kappa}{\kappa + (1 - \pi)(q_L - p)}, \tag{C.3}$$

independently of b; every nonmarginal type requires strictly more. Along the marginal curve, the second branch weakly favors higher b because $w_H(b) - w_L(b)$ is weakly increasing. Hence the highest feasible marginal value

$$b_L^0 = \sup\{b \in [\underline{b}, \bar{b}] : D_L(b; q_L) \geq 0\}$$

is D1-permitted and has the largest minimum early price among permitted marginal types. A belief concentrated on it deters every undercut when $q_L \leq C_L(b_L^0)$.

To verify this inequality under (C.1), define

$$\Xi_L(b) = C_L(b) - w_L(b) + \frac{\kappa}{1 - \pi} = \Gamma_L(b) - d_{S,L} + \frac{\kappa}{1 - \pi}.$$

The competition wedge makes Ξ_L weakly decreasing. If $b_L^0 < \bar{b}$, then $D_L(b_L^0; q_L) = 0$, so

$$C_L(b_L^0) - q_L = \Xi_L(b_L^0) \geq \Xi_L(\bar{b}) = \bar{D}_L \geq 0.$$

If $b_L^0 = \bar{b}$, interiority of the fixed point directly gives $q_L < C_L(\bar{b})$. For an overcut $p > q_L$, no type gains on the first branch of (C.2), since it gives no more L-acceptance than the cheaper equilibrium probe. On the second branch, the required H-acceptance probability is minimized by the highest-value, probing type with the largest resolution benefit. A

belief concentrated on $\bar{b}$ makes H reject below $C_H(\bar{b})$; at or above that price, the top-type inequality in (3.7) rules out a gain even with acceptance in both states. These beliefs cover every off-path price and give the claimed D1 completion.

C.4 D1 compatibility of the full-menu equilibrium

When $\kappa > 0$, the full-menu PBE in Proposition 3.2 likewise admits a D1-compatible completion in the contingent-plan representation. Boundary pricing equates each on-path price to the corresponding seller's posterior minimum early price. Global interiority gives positive probing and knockout mass at values arbitrarily close to $\bar{b}$; because C_L and C_H are strictly increasing on a top interval, their conditional means are strictly below their endpoint values. Thus $q_L < C_L(\bar{b})$ and $q_H < C_H(\bar{b})$. For $p < q_L$, the argument leading to (C.3) makes the highest-value wait–probe marginal type $(\bar{b}, D_L(\bar{b}))$ D1-permitted. A belief concentrated on that type makes seller L reject because $p < q_L < C_L(\bar{b})$, and pointwise ordering makes seller H reject as well.

For an intermediate price $q_L < p < q_H$, no type gains on the first branch of (C.2). On the second branch, the probe–knockout marginal type $(b, D_H(b))$ requires

$$y^*(p) = \frac{(1-\pi)(p-q_L)}{\pi(r_0 - p)}, \qquad r_0 = w_H(b) + D_H(b) = \frac{q_H - (1-\pi)q_L}{\pi}, \tag{C.4}$$

which is independent of b; types on either side of the boundary require weakly more H-acceptance. The highest-value marginal type is therefore D1-permitted. A belief concentrated on $\bar{b}$ makes H reject because $p < q_H < C_H(\bar{b})$, while any L-acceptance cannot generate a gain at a price above q_L. At $p \geq q_H$, even acceptance in both states is weakly worse than the equilibrium knockout. These cases exhaust the nested response set. D1 therefore supports the equilibrium without uniquely selecting it or making implementation belief-independent.

D Welfare Accounting and Numerical Examples

D.1 Gains and losses from a ban

In the boundary-price full-menu equilibrium, the seller is indifferent between accepting and continuing after a probe in L and after knockout in H. Its gain after knockout in L is $\int [C_H(b) - C_L(b)] h_K(b)\, db$. Hence

$$U_S^D - U_S^B = (1-\pi) M_K \int [C_H(b) - C_L(b)] h_K(b)\, db > 0. \tag{D.1}$$

The early buyer's expected gain relative to a ban is the sum of the gains from probing and knockout among the types choosing those actions:

$$U_E^D - U_E^B = \int f(b) \left\{ \int_{D_L(b)}^{D_H(b)} V_L(b,d)\, dG(d) + \int_{D_H(b)}^{\bar{d}} V_H(b,d)\, dG(d) \right\} db \geq 0. \tag{D.2}$$

Discretion gives the early buyer the option to offer in addition to its auction payoff U_E^B. Its gain is therefore nonnegative, and strictly positive when a positive mass offers at a strict gain. A later buyer loses only when an early offer is accepted, so late buyers' aggregate gain from the ban is

$$U_{\text{late}}^B - U_{\text{late}}^D = \mathbb{E}[\mathbf{1}_{A_c} \Lambda_s(b)] \geq 0. \tag{D.3}$$

Combining these equations with (4.3) also gives $W^B - W^D = \mathbb{E}[\mathbf{1}_{A_c} \Lambda_s(b)] - (U_E^D - U_E^B) - (U_S^D - U_S^B)$. The buyer and seller gain in expectation at their own information sets. This need not hold for every realized (b, s): a seller accepting 80 against equally likely minimum early prices of 60 and 100 breaks even in expectation but gives up 20 in the latter case.

D.2 The paired welfare examples

Set $(n_L, n_H) = (2, 6)$, $(v_L, v_H) = (.2, .4)$, $\pi = .5$, $d_{S,L} = d_{S,H} = 0$, $\kappa = .02$, and $\bar{d} = 1$. On $[0, 1]$, the buyer's resolution benefit has distribution

$$G(d) = 10^{-4} d + (1 - 10^{-4})[1 - (1 - d)^{30}].$$

This mixes a beta distribution concentrated near zero with a small uniform component: $10^{-4} U[0,1] + (1 - 10^{-4})\,\text{Beta}(1, 30)$. The uniform component ensures a positive density throughout the support. Holding this distribution fixed, compare the common early- and late-buyer valuation distributions

$$\begin{aligned} F^S &= .01 U[0,1] + .99\,\text{Beta}(40, 40), \\ F^T &= .01 U[0,1] + .99 \left[\frac{16}{17}\,\text{Beta}(38, 42) + \frac{1}{17}\,\text{Beta}(9, 1) \right]. \end{aligned}$$

Both means equal .5. The uniform components ensure positive densities. The selected price pairs are approximately $(.51098059, .55437879)$ and $(.51561331, .64062633)$, respectively.

Integration and root verification. The calculations below verify that each reported price pair is close to an exact equilibrium and that numerical error cannot reverse the welfare comparison in Table 4.1. For an integer beta distribution, its CDF is a finite polynomial,

$$B_{a,b}(x) = \sum_{j=a}^{a+b-1} \binom{a+b-1}{j} x^j (1-x)^{a+b-1-j}.$$

The mixture CDFs can therefore be enclosed at rational endpoints using outward-rounded arithmetic. For a fixed count n, let

$$A_n(t) = 1 - F(t)^n, \qquad B_n(t) = 1 - F(t)^n - n[1 - F(t)]F(t)^{n-1}.$$

These are the probabilities that at least one and at least two late values exceed t. They are decreasing in t. The seller's expected auction payoff and the expected allocation gain from continuing to auction are

$$R_{n,v}(b) = v + \int_v^{\bar{b}} B_n(t)\,dt + \mathbf{1}\{b > v\} \int_v^b n[1 - F(t)]F(t)^{n-1}\,dt,$$
$$L_{n,v}(b) = \max\{v - b, 0\} + \int_{\max\{v,b\}}^{\bar{b}} A_n(t)\,dt.$$

The remaining revenue integrand $n(1 - F)F^{n-1}$ is unimodal in F, so its endpoint values and its possible interior maximum bound its range on each cell. The calculation uses outward-rounded ball arithmetic, implemented with Arb through `python-flint`, to enclose numerical error. It uses 24,000 rational cells and 60 decimal digits. Outer value integrals use Lebesgue–Stieltjes sums: each integrand's interval range is multiplied by the exact enclosed probability mass of that cell. The same procedure bounds the price residuals, their derivatives, the action probabilities, and welfare integrals.

For root certification, use the unnormalized residuals

$$H(q) = \begin{pmatrix} \int [C_L(b) - q_L] m_P(b; q)\,dF(b) \\ \int [C_H(b) - q_H] m_K(b; q)\,dF(b) \end{pmatrix}.$$

Positive offer masses make $H(q) = 0$ equivalent to $T(q) = q$. For a candidate box X with center x, a numerical inverse Y of the central Jacobian, and an interval enclosure $J(X)$ of DH, the Krawczyk operator is

$$K(X) = x - YH(x) + [I - YJ(X)](X - x).$$

Strict inclusion $K(X) \subset \operatorname{int} X$ certifies a root; the nonsingular interval Jacobian also verifies local uniqueness within the box. Initial boxes have centers $(.510980593207608, .554378791664074)$ and $(.515613314493683, .640626332049219)$, and coordinate radii .0004 and .0008, respectively. Table D.1 reports outward enlargements of the resulting root and welfare bounds. This is a local certificate, not a global uniqueness result.

The strict cutoff margins establish $0 < D_L < D_H < 1$ over the whole value support and the reported root boxes. Positive seller-response margins and $C_s(1) - q_s > 0$ verify

Table D.1: Certified Bounds for the Welfare Examples

	F^S	F^T
q_L	$[0.51087978, 0.51108141]$	$[0.51550714, 0.51571949]$
q_H	$[0.55430282, 0.55445477]$	$[0.64035841, 0.64089426]$
$W^B - W^D$	$[-0.00366207, -0.00339311]$	$[0.00026167, 0.00034508]$
$\det J(X)$	$[0.01303747, 0.01911430]$	$[0.00017958, 0.00033740]$
$\min D_L$, lower bound	0.017569	0.002409
$\min(D_H - D_L)$, lower bound	0.003932	0.102852
$\max D_H$, upper bound	0.598030	0.766282

on-path acceptance and the endpoint inequalities used in Appendix C. The certified root and strict incentive inequalities establish the equilibrium conditions used in Appendix B.3.

The Calculation Workbook provides explanations, Python code, and recorded outputs for the participation example, welfare calculations, interval certificates, appendix exercises, and numerical figures. Readers can download a single Jupyter notebook and run it from top to bottom. The workbook lists the required dependencies; the certificates use Python 3.12 and `python-flint==0.9.0`.