Information Acquisition in Robust Trade

Hershdeep Chopra [*] Jeffrey C. Ely [†]

November 8, 2026

Abstract

We study timing of information acquisition by buyer and seller before bilateral trade. Traders flexibly and costlessly choose experiments about their independent private values. Observing experiments but not realizations, a designer selects a *robust* trading mechanism maximizing expected gains. We show that any equilibrium with simultaneous information acquisition is necessarily inefficient. In contrast, we construct an efficient sequential acquisition equilibrium in which the first mover partially learns and the second mover acquires full information. In this equilibrium, both traders strictly prefer moving first, regardless of initial ownership. Commitment through partial learning secures a larger share of the gains from trade.
Keywords: non-cooperative information design, robust mechanisms.

1 Introduction

This paper studies how the timing of information acquisition affects bilateral trade. A seller and a buyer with independent binary fundamentals choose experiments about their own values, either simultaneously or sequentially. Experiment choices are public, but their realizations are private. After the public experiment choices, a designer selects a *robust* (dominant strategy incentive-compatible, ex-post individually rational, and ex-post budget balanced) mechanism to maximize expected gains from trade, as in Hagerty and Rogerson (1987). The designer can be interpreted as a centralized clearinghouse

[*]London School of Economics. h.s.chopra@lse.ac.uk
[†]Department of Economics, Northwestern University. jeff@jeffely.com
This paper has benefited from discussion with Francesco Nava, Junrok Park, Philip Reny, Matthew Levy, and participants at the LSE Economic Theory Seminar and the Northwestern CET student theory lunch.

or regulated intermediary. The observable information conditions affect which rule the institution uses to facilitate trade, while incentive compatibility, budget balance, and participation are guaranteed type by type. This captures, for example, institutions that adapt trading protocols to observable market conditions without relying on distributional assumptions to sustain (truthful) participation.

Following the characterization of Hagerty and Rogerson (1987) and Čopič and Ponsatí (2016), robust trading mechanisms can be represented by lotteries over posted prices. We first show that every pure *simultaneous* or *sequential acquisition equilibrium* attains first-best total surplus. Moreover, on the equilibrium path there is no price dispersion, so it is implemented by a deterministic posted price. Additionally, the equilibrium features a "no-wasted" mass property. Given an equilibrium price, each trader's posterior distribution assigns no mass to outcomes where additional learning could increase gains from trade. Together, these properties rule out pure equilibrium in simultaneous acquisition.

Allowing mixed information strategies, we show there exist equilibria with simultaneous information acquisition, but every such equilibrium is necessarily inefficient. In contrast, we construct an efficient subgame-perfect equilibrium in pure strategies when the traders acquire information sequentially. In this equilibrium, the first mover partially learns her state and the second mover acquires full information. Fixing the acquisition order, all pure strategy sequential acquisition equilibria yield the same payoffs.

The sequential acquisition game allows us to compare how the order of information acquisition affects the bargaining power of the traders. Moving first commits the trader to an experiment, whereas moving second permits information acquisition tailored to the observed experiment. This contingent choice can strengthen the follower's bargaining position, so commitment is not a priori advantageous. Nevertheless, we show that in any pure strategy subgame-perfect equilibrium with sequential information acquisition, either trader is better off committing first.

For the seller-first acquisition game, we characterize the buyer's maximal payoff at a fixed price and derive his guaranteed continuation payoff. Additionally, we characterize the most informative buyer experiment that attains at least this guaranteed continuation payoff. Using this, we characterize the most informative seller experiment that minimizes the continuation payoff that the buyer can guarantee himself. We then construct a pure subgame-perfect equilibrium in which the seller partially learns, the buyer fully learns, and trade realizes first-best surplus. The buyer receives his minimum guaranteed continuation

payoff, leaving the seller the largest possible equilibrium payoff. Every pure seller-first equilibrium has these payoffs and realizes first-best surplus.

Reflecting the seller-first game's primitives along $1/2$ yields an isomorphism between the two acquisition orders, giving the corresponding buyer-first acquisition results. Finally, we show that both traders strictly prefer acquiring information first. The leader's experiment shapes both the designer's price and the follower's opportunities to influence that price through learning. Thus, commitment to partial learning secures a larger share of surplus while preserving efficient trade. All omitted proofs are in the appendix.

2 Related literature

Roesler and Szentes (2017) study how information choice shapes monopoly pricing, and Park (2025) extends this approach to buyers who cooperatively purchase a public good. Condorelli and Szentes (2020) study how a buyer shapes subsequent monopoly pricing by committing to a valuation distribution. Concurrent work by Kyriazis (2026) studies how the choice of buyer valuation distributions shapes optimal monopoly screening and the division of surplus. Ravid (2020) studies rational inattention in ultimatum bargaining, while Ravid et al. (2022) show that inefficiency can persist as private learning costs vanish. Schottmüller (2023) jointly optimizes information structures and the trading mechanism to maximize expected gains from trade.

Here both traders choose the private-information distributions fixed in Myerson and Satterthwaite (1983), before a designer selects a surplus-maximizing trading rule. Our trading institution follows Hagerty and Rogerson (1987); Čopič and Ponsatí (2016), within the broader literature on robust mechanisms and weaker assumptions about agents' beliefs (Wilson, 1987; Bergemann and Morris, 2005; Chung and Ely, 2007).

Our companion paper (Chopra and Ely, 2026a) studies simultaneous private learning before revenue-maximizing auction design. Here the designer maximizes gains from trade, and sequential acquisition lets the leader influence the follower's learning and the division of surplus. Relatedly, Terstiege and Wasser (2022) study bidders' disclosures before auction design, while Bergemann and Välimäki (2002) study efficient learning under a mechanism fixed before acquisition.

Other forms of noncooperative information design arise in competitive persuasion (Gentzkow

and Kamenica, 2017) and competition among proposals for a receiver's choice (Boleslavsky and Cotton, 2018; Au and Kawai, 2020). Yang (2015) studies flexible acquisition before a coordination game.

3 Model

3.1 Bilateral trade

A seller produces one indivisible object and a buyer demands it, with independent valuations. At the trading stage the seller has private cost $c \in [0,1]$ and the buyer has private value $v \in [0,1]$. If $x \in \{0,1\}$ denotes trade and t is the transfer from buyer to seller, the utilities of the seller and the buyer are

$$u_s(x,t,c) = t - xc \quad \text{and} \quad u_b(x,t,v) = xv - t$$

Conditional on (c,v), efficient trade occurs when $v > c$ and generates gains from trade $(v-c)^+$. As described below, these private valuations correspond to posterior means determined by endogenous information acquisition.

3.2 Information acquisition

At the outset, the seller and buyer are uninformed about their cost and value, respectively. Their payoff-relevant states are $\theta_s, \theta_b \in \{0,1\}$, with common prior means m_s and m_b. The states, and hence the seller cost and buyer valuation, are independent, with prior expectation

$$\Pr(\theta_s = 1) = m_s \quad \text{and} \quad \Pr(\theta_b = 1) = m_b$$

Throughout the main analysis, we maintain that there are ex-ante strict gains from trade

$$0 < m_s < m_b < 1$$

Thus trade is efficient at the prior. Traders publicly choose experiments about their own states and observe their realization privately. Conditional on the chosen experiments, the experiment realizations are independent. Because payoffs are linear in the state, an experiment is summarized by the distribution of its posterior mean: $c = \mathbb{E}[\theta_s \mid s_s]$ for the seller and $v = \mathbb{E}[\theta_b \mid s_b]$ for the buyer. Bayes-plausibility, see Kamenica and Gentzkow

(2011), therefore identifies the feasible posterior laws with the mean restrictions

$$\mathcal{G}(m_s) = \left\{ G \in \Delta([0,1]) : \int c \, dG(c) = m_s \right\}$$

$$\mathcal{F}(m_b) = \left\{ F \in \Delta([0,1]) : \int v \, dF(v) = m_b \right\}$$

In the *seller-first game*, the seller chooses $G \in \mathcal{G}(m_s)$ and the buyer chooses $F \in \mathcal{F}(m_b)$ after observing G. In the *buyer-first game* the order is reversed.[1]

In the simultaneous move game, the traders simultaneously choose their experiments.

Under full information, a strictly beneficial trade occurs when the buyer's state is high and the seller's is low. Expected first-best gains from trade are

$$W^{\mathrm{FB}}(m_s, m_b) = m_b(1 - m_s)$$

3.3 Robust continuation trade

A direct trading mechanism μ maps reported valuations into lotteries over the trade decision and transfers. A mechanism is robust if it satisfies *dominant-strategy incentive compatibility*, *exact budget balance*, and *ex-post individual rationality* as in Hagerty and Rogerson (1987). These conditions hold on the full valuation domain $[0,1]^2$, and expected payoffs average over the mechanism's randomization.

Following the posted-price characterization and the trading procedure of Čopič and Ponsatí (2016) and Hagerty and Rogerson (1987), we model robust continuation trade through lotteries over posted prices. We represent the price lottery of a continuation mechanism μ, in this implementation by λ in $\Delta([0,1])$.

After observing the experiments, the designer publicly announces the lottery. The price and private signals are then realized. Each trader observes his signal, the price, and the preceding public choices, and the traders simultaneously accept or reject. Trade occurs only if both accept, at a transfer of p from buyer to seller.

As in Hagerty and Rogerson (1987), we restrict participation to weakly dominant actions: the buyer accepts when $v > p$ and rejects when $v < p$, while the seller accepts when $c < p$

[1] We use "she" for the first mover and "he" for the second mover.

and rejects when $c > p$. At equality, either action or any mixture is permitted.[2] Write $\rho = (\rho_b, \rho_s)$ for the traders' participation choices at indifference. After each experiment pair and announced lottery, $\rho_b(p \mid F, G, \lambda)$ and $\rho_s(p \mid F, G, \lambda)$ give the acceptance probabilities of a buyer and seller whose posterior valuations equal p and who face a price p. These are components of the traders' equilibrium strategies, including after alternative announcements.

Given the choice of experiments (F, G), the posted price lottery λ and participation rule ρ, the payoffs conditional on a realized price p are[3]

$$\begin{aligned} U_b^\rho(p; F, G, \lambda) &= \left[G([0, p)) + \rho_s(p) G(\{p\})\right] \int (v - p)^+ \, dF(v) \\ U_s^\rho(p; F, G, \lambda) &= \left[F((p, 1]) + \rho_b(p) F(\{p\})\right] \int (p - c)^+ \, dG(c) \end{aligned} \tag{1}$$

Total surplus is $W^\rho = U_b^\rho + U_s^\rho$.

The designer anticipates participation rule ρ after each candidate announcement and maximizes expected total surplus over the robust continuation mechanisms

$$\sup_{\lambda \in \Delta([0,1])} \int W^\rho(p; F, G, \lambda) \, \lambda(dp) \tag{2}$$

A pricing strategy λ assigns a lottery $\lambda(\cdot \mid F, G)$ to every feasible experiment pair. A participation strategy $\rho = (\rho_b, \rho_s)$ specifies acceptance probabilities in $[0, 1]$ at indifference after every experiment pair, candidate lottery announcement, and realized price. Participation follows the weakly dominant actions specified above for types who are not indifferent. An *optimal continuation* is a pair (λ, ρ) for which the assigned lottery attains Equation 2 after every experiment pair, taking account of the traders' participation after each alternative announcement.

[2] This excludes joint rejection of strictly profitable offers, while leaving participation at indifference unrestricted.

[3] For a CDF G, $G(x-) = \lim_{y \uparrow x} G(y)$ and $G(x+) = \lim_{y \downarrow x} G(y)$ denote its left and right limits; right continuity gives $G(x+) = G(x)$. The same notation applies to F.

4 Acquisition equilibrium

An acquisition equilibrium is a subgame-perfect equilibrium with posted price implementation and participation as specified above. We use the term *pure* to refer only to the trader's information acquisition strategy, and allow the designer and indifferent traders to randomize.

Given a continuation (λ, ρ), trader i's expected payoff is

$$u_i^{\lambda,\rho}(F,G) = \int U_i^{\rho}\big(p; F, G, \lambda(\cdot \mid F, G)\big)\, \lambda(dp \mid F, G)$$

In simultaneous acquisition, each trader holds the other's experiment fixed. In sequential acquisition, the leader anticipates the follower's response, whereas the follower holds the leader's observed choice of experiment fixed.

Definition 1. *A* **pure simultaneous acquisition equilibrium** *consists of experiments* (F^*, G^*) *and an optimal continuation* (λ, ρ) *satisfying*

$$\begin{aligned} u_b^{\lambda,\rho}(F^*,G^*) &\geq u_b^{\lambda,\rho}(F,G^*) && \text{for every } F \in \mathcal{F}(m_b), \\ u_s^{\lambda,\rho}(F^*,G^*) &\geq u_s^{\lambda,\rho}(F^*,G) && \text{for every } G \in \mathcal{G}(m_s). \end{aligned}$$

In the seller-first game the buyer's pure information strategy is a map $\sigma_b : \mathcal{G}(m_s) \to \mathcal{F}(m_b)$.

Definition 2. *A* **pure seller-first acquisition equilibrium** *is a profile* (G^*, σ_b) *along with an optimal continuation* (λ, ρ) *satisfying*

$$\begin{aligned} u_b^{\lambda,\rho}(\sigma_b(G),G) &\geq u_b^{\lambda,\rho}(F,G) && \text{for every feasible } G, F, \\ u_s^{\lambda,\rho}(\sigma_b(G^*),G^*) &\geq u_s^{\lambda,\rho}(\sigma_b(G),G) && \text{for every } G. \end{aligned}$$

Pure buyer-first acquisition equilibrium is defined by interchanging the traders. Collectively, we refer to these as pure sequential acquisition equilibria.

4.1 Pure-strategy equilibrium

Recall that pure information acquisition strategies allow the designer to randomize over prices and indifferent traders to randomize over participation. In the following, we show

that any pure strategy acquisition equilibrium necessarily achieves the first-best payoff.

Theorem 1. *Suppose $0 < m_s < m_b < 1$. Every pure simultaneous or sequential acquisition equilibrium inducing the profile (F, G, λ, ρ) attains the first-best total surplus. In particular, $u_b^{\lambda,\rho}(F,G) + u_s^{\lambda,\rho}(F,G) = W^{\mathrm{FB}}(m_s, m_b)$.*

Because trade is ex ante efficient, each trader can guarantee a positive payoff. The traders choose information to influence the terms of trade, while the designer chooses those terms in response to the trading opportunities the choice of information creates. More favorable terms raise a trader's payoff conditional on trade, but can also leave mutually beneficial trades unrealized. Information choices change the likelihood and size of those forgone gains, and hence which terms of trade maximize total surplus. When expected total surplus is below first best, the seller can split high-cost posteriors into lower costs and cost one to create a profitable deviation. Prices serving the new posteriors give her a larger share of surplus. At lower maximizing prices, the buyer's payoff falls below his original payoff. In both cases, the proof of Theorem 1 in Appendix A.1 shows that the resulting total surplus is large enough to leave the seller better off, including when the buyer best responds. We next use efficiency and the traders' acquisition incentives to rule out price dispersion.

Corollary 1. *Every pure simultaneous or sequential acquisition equilibrium, as defined in Definition 1 and Definition 2, respectively, necessarily induces a deterministic posted price in $(0, 1)$.*

The deterministic price in Corollary 1 and the proof of Theorem 1 imply necessary restrictions on the supports of equilibrium value and cost distributions in pure strategy acquisition equilibria.

Corollary 2. *In every pure strategy sequential acquisition equilibrium, the follower acquires full information. At the posted price p of a pure strategy acquisition equilibrium, the buyer's posterior law assigns no mass to $(0, p)$ when the buyer acquires last, and the seller's posterior law assigns no mass to $(p, 1)$ when the seller acquires last. Both support restrictions hold for pure simultaneous acquisition equilibria.*

A pure simultaneous acquisition equilibrium necessarily satisfies both support restrictions in Corollary 2. At the posted price p given by Corollary 1, both traders earn positive payoffs. Efficiency therefore requires acceptance by any trader indifferent at p. The payoff

formulas in Equation 1 reduce to

$$u_b^{\lambda,\rho}(F,G) = G(p)\big[m_b - p(1-F(0))\big],$$
$$u_s^{\lambda,\rho}(F,G) = (1-F(0))\big[1-m_s-(1-p)G(p)\big].$$

By Theorem 1, their sum must equal first-best surplus. Independence then requires at least one trader to learn fully. Facing a fully informed opponent, the other trader can obtain arbitrarily close to the entire first-best surplus. This contradicts the positive payoff each trader can guarantee and gives the following impossibility result.

Corollary 3. *For $0 < m_s < m_b < 1$, no pure simultaneous acquisition equilibrium exists in the sense of Definition 1.*

Holding the price and the other trader's acceptance behavior fixed, full information weakly maximizes each trader's expected positive trading margin. But when both traders are fully informed, every interior price generates efficient trade, leaving the division of surplus sensitive to a change in information. A trader can then gain by acquiring less information and inducing a more favorable price.

4.2 Mixed simultaneous acquisition

In this section we allow the traders to choose mixed strategies over experiments in the simultaneous acquisition game. The result below establishes the existence of such an equilibrium and shows that any such equilibrium is necessarily inefficient.

Proposition 1. *For $0 < m_s < m_b < 1$, a simultaneous acquisition equilibrium exists in mixed information acquisition strategies. Every simultaneous acquisition equilibrium has expected total surplus strictly below W^{FB}.*

Randomization over experiments therefore gives existence of equilibrium but leads to inefficient outcomes. The inefficiency follows as achieving first-best surplus requires that at least one trader is fully informed, but fully informative experiments lead to deviations by the opponent of the form used in the proof of Corollary 3 and thus can not be sustained in an equilibrium. In contrast, acquiring information first commits a trader to an experiment before the other trader chooses his. This commitment affects subsequent learning and the division of gains from trade, as the comparison in Section 8 shows. Each sequential acquisition game admits a pure strategy equilibrium, as shown in Corollary 5, and every

such equilibrium is efficient by Theorem 1.

5 Call and put representation

Represent a buyer experiment F and a seller experiment G by their expected positive margins at each price, the *call* and *put*:

$$\psi_F(p) := \mathbb{E}_F[(v-p)^+] \qquad \text{and} \qquad \varphi_G(p) := \mathbb{E}_G[(p-c)^+]$$

Each function determines its posterior law, and pointwise comparison corresponds to convex order.[4]

$$\Psi(m_b) = \left\{ \psi : [0,1] \to [0,1] : \begin{array}{l} \psi \text{ is convex, non-increasing, and one-Lipschitz,} \\ \psi(0) = m_b, \quad \psi(1) = 0 \end{array} \right\},$$

$$\Phi(m_s) = \left\{ \varphi : [0,1] \to [0,1] : \begin{array}{l} \varphi \text{ is convex, non-decreasing, and one-Lipschitz,} \\ \varphi(0) = 0, \quad \varphi(1) = 1 - m_s \end{array} \right\}.$$

For a given call and put we write F_ψ and G_φ for the associated distributions. The classes are compact in the uniform norm, and equivalently under weak convergence of the distributions

$$\|\varphi_n - \varphi\|_\infty \longrightarrow 0 \quad \Longleftrightarrow \quad G_{\varphi_n} \Rightarrow_w G_\varphi \tag{3}$$

with the same equivalence for calls.[5] Full buyer information has call $\psi^{\mathrm{FI}}(t) = m_b(1-t)$.

In this representation the payoff in Equation 1 becomes

$$U_b^\rho(p;\psi,\varphi,\lambda) = \Big[(\varphi)'_-(p) + \rho_s(p \mid F_\psi, G_\varphi, \lambda)\big((\varphi)'_+(p) - (\varphi)'_-(p)\big)\Big]\psi(p)$$
$$U_s^\rho(p;\psi,\varphi,\lambda) = \Big[-(\psi)'_+(p) + \rho_b(p \mid F_\psi, G_\varphi, \lambda)\big((\psi)'_+(p) - (\psi)'_-(p)\big)\Big]\varphi(p)$$

The derivative jumps are the corresponding point masses, of which the fractions $\rho_s(p)$ and $\rho_b(p)$ accept.[6] The buyer's strategy assigns a call function to each seller put: $\sigma_b : \Phi(m_s) \to$

[4] See Gentzkow and Kamenica (2016) for the integrated-distribution representation and Shaked and Shanthikumar (2007, Section 3.A) for convex order. The call-function method is developed in Chopra and Ely (2026c) to study selection markets.

[5] These classes are uniformly bounded, equicontinuous, and closed in the uniform norm, so compactness follows from Arzelà–Ascoli.

[6] At the endpoints we have $(\varphi)'_-(0) = 0$, $(\varphi)'_+(1) = 1$, $(\psi)'_-(0) = -1$, and $(\psi)'_+(1) = 0$.

$\Psi(m_b)$.

6 Seller-first acquisition

With seller-first acquisition, the buyer observes the seller's experiment before choosing his own. The designer then selects a price lottery that maximizes expected total surplus.

Independence and $(v-c)^+ \leq v(1-c)$ imply $W^\rho(p;\psi,\varphi,\lambda) \leq W^{\mathrm{FB}}(m_s,m_b)$

Given a seller put φ and an optimal continuation (λ,ρ), the buyer's *best-response value* is

$$B_{\lambda,\rho}(\varphi) = \sup_{\psi\in\Psi(m_b)} u_b^{\lambda,\rho}(\psi,\varphi)$$

A call attaining this supremum is a *buyer best response*.

We first fix the seller's experiment and characterize the buyer's highest payoff from implementing a given price in an auxiliary problem. We then calculate the buyer's guaranteed continuation payoff and minimize it over seller experiments. Finally, we construct an equilibrium with full buyer information in which the buyer receives this minimum guarantee and the seller receives the remaining first-best surplus.

6.1 Buyer's auxiliary problem

Fix a seller put φ and a target price $p \in (0,1)$. We first study an auxiliary problem that describes the buyer's best experiment that implements a given price.[7] The auxiliary participation rule ρ_p prescribes acceptance at indifference by the seller at every offered price and by the buyer only at or below the target.

$$\rho_{s,p}(p' \mid F,G,\lambda) = 1 \qquad \text{and} \qquad \rho_{b,p}(p' \mid F,G,\lambda) = \mathbb{1}_{\{p'\leq p\}}.$$

Thus an indifferent buyer rejects an offer $p' > p$, while a buyer with valuation strictly above p' still accepts. This rule is independent of the announced price lottery and is one admissible participation choice in the general model. This gives an upper bound on buyer payoff at prices maximizing total surplus under inclusive participation. Here, inclusive

[7] The approach of this section also generalizes buyer optimal learning in Roesler and Szentes (2017) to the case of uncertain cost with known distribution when the objective is seller-profit maximization as opposed to total-surplus maximization here.

participation means $\rho = (1, 1)$, so both traders accept with probability one when indifferent. We then compare attainable payoffs across all implementable target prices.

At the target price p,

$$U_b^{\rho_p}(p; \psi, \varphi) = (\varphi)'_+(p)\psi(p) \qquad \text{and} \qquad U_s^{\rho_p}(p; \psi, \varphi) = -\varphi(p)(\psi)'_-(p)$$

Write $W^{\rho_p} = U_b^{\rho_p} + U_s^{\rho_p}$. A target price p is *implementable in the auxiliary problem* if some feasible buyer experiment makes it a total surplus-maximizing price under ρ_p.

Note that it is without loss to restrict the buyer's support to $\{0\} \cup [p, 1]$. We can split every posterior value in $(0, p)$ into posterior values in $\{0, p\}$. This preserves the call at and above p, weakly raises total surplus at p, and leaves total surplus at every higher price unchanged. As such, we let the set of calls implementing price p after seller choice φ and under the participation rule ρ_p be given by

$$\mathcal{C}_p(\varphi) := \left\{ \psi \in \Psi(m_b) : \begin{array}{l} F_\psi((0, p)) = 0, \\ (\varphi)'_+(t)\psi(t) - \varphi(t)(\psi)'_+(t) \leq W^{\rho_p}(p; \psi, \varphi) \quad \text{for all } p < t < 1 \end{array} \right\}$$

By definition, every ψ in $\mathcal{C}_p(\varphi)$ makes p a total-surplus-maximizing price under ρ_p. Thus $\mathcal{C}_p(\varphi)$ is nonempty exactly when p is implementable in the auxiliary problem.

For every ψ in $\mathcal{C}_p(\varphi)$, $F_\psi((0, p)) = 0$ and the buyer accepts price p with probability $-(\psi)'_-(p)$. His expected margin conditional on accepting, the *conditional trading rent*, is $\eta_\psi(p) := \mathbb{E}[v - p \mid v \geq p] = -\psi(p)/(\psi)'_-(p)$. The mean restriction gives

$$m_b = -(\psi)'_-(p)[p + \eta_\psi(p)]$$

His payoff can therefore be written as

$$U_b^{\rho_p}(p; \psi, \varphi) = m_b(\varphi)'_+(p)\frac{\eta_\psi(p)}{p + \eta_\psi(p)}$$

At a fixed price, the buyer's payoff is increasing in $\eta_\psi(p)$ while the probability of trade decreases. The mean restriction and $p \leq v \leq 1$ conditional on accepting price p imply

$$\max\{0, m_b - p\} \leq \eta_\psi(p) \leq 1 - p$$

Equality in the upper bound holds exactly when all accepting values equal one, in which case the experiment is fully informative. Define the greatest buyer payoff at p in the auxiliary problem by

$$V(p;\varphi) = \sup_{\psi \in \mathcal{C}_p(\varphi)} U_b^{\rho_p}(p;\psi,\varphi)$$

The value is $-\infty$ if $\mathcal{C}_p(\varphi)$ is empty. Otherwise, define its convex *majorant call* as the pointwise supremum

$$\widehat{\psi}_\varphi^p(t) := \sup_{\psi \in \mathcal{C}_p(\varphi)} \psi(t)$$

Whenever the target price p is implementable in the auxiliary problem, the next theorem shows that this majorant call gives the most informative experiment in $\mathcal{C}_p(\varphi)$ and maximizes buyer payoff at price p under ρ_p.

Theorem 2. *If $\mathcal{C}_p(\varphi)$ is nonempty, its pointwise supremum $\widehat{\psi}_\varphi^p$ belongs to $\mathcal{C}_p(\varphi)$. Its distribution dominates every distribution associated with a call in $\mathcal{C}_p(\varphi)$ in convex order. Among calls in $\mathcal{C}_p(\varphi)$, $\widehat{\psi}_\varphi^p$ attains the greatest buyer payoff $V(p;\varphi)$ and also maximizes total surplus at p under ρ_p.*

Fix an implementable price p with $G_\varphi(p) > 0$. If a buyer support point in $(p,1)$ yielded strictly less total surplus under ρ_p than at p, a small mean-preserving spread would raise the call while keeping p optimal, contradicting Theorem 2. A similar local spread also rules out an atom at p. Both arguments are formalized in Appendix B.1. On intervals without buyer mass, surplus is nondecreasing as additional sellers accept. Since p maximizes total surplus, such intervals must also have no seller mass. Hence equality of total surplus under ρ_p holds throughout $[p,1)$. To describe the majorant call above the target price, let $c_0 = \inf\{t : \varphi(t) > 0\}$ and define[8]

$$H_\varphi(t) = \varphi(t) I_\varphi(t), \qquad I_\varphi(t) = \int_t^1 \frac{du}{\varphi(u)^2}, \qquad t > c_0. \tag{5}$$

Constant total surplus on the support of $\widehat{\psi}_\varphi^p$ under ρ_p implies that $(\varphi, \widehat{\psi}_\varphi^p)$ satisfies the following ordinary differential equation (ODE) for $p < t < 1$

$$(\varphi)'_+(t)\widehat{\psi}_\varphi^p(t) - \varphi(t)(\widehat{\psi}_\varphi^p)'_+(t) = W^{\rho_p}(p;\widehat{\psi}_\varphi^p,\varphi)$$

[8] If $G_\varphi(c_0) > 0$, define $H_\varphi(c_0)$ and $\eta_\varphi^*(c_0)$ by their right limits. In particular,

$$H_\varphi(c_0) := \lim_{t \downarrow c_0} H_\varphi(t) = \frac{1}{G_\varphi(c_0)}. \tag{4}$$

Solving the ODE above with the integration factor $1/\varphi$ and boundary condition $\widehat{\psi}^p_\varphi(1) = 0$ gives

$$\widehat{\psi}^p_\varphi(t) = W^{\rho_p}(p; \widehat{\psi}^p_\varphi, \varphi) H_\varphi(t)$$

With no buyer mass in $(0, p]$, the call is linear below p, and its right and left derivatives agree at p. Its conditional trading rent is therefore

$$\eta^*_\varphi(p) = -\frac{H_\varphi(p)}{(H_\varphi)'_+(p)}. \tag{6}$$

The function H_φ is decreasing and convex, with $H_\varphi(1) = 0$. In particular Equation 6 implies that $p + \eta^*_\varphi(p)$ is nondecreasing in p and $p + \eta^*_\varphi(p) \leq 1$. The mean restriction implies that the buyer's acceptance probability is $m_b/[p + \eta^*_\varphi(p)]$. Putting these observations together, we get the following shape of the majorant $\widehat{\psi}^p_\varphi$.

Corollary 4. *Fix a price p in $(0,1)$ with $G_\varphi(p) > 0$. The target price p is implementable in the auxiliary problem if and only if $p + \eta^*_\varphi(p) \geq m_b$. When this condition holds, its majorant call is*

$$\widehat{\psi}^p_\varphi(t) = \frac{m_b}{p + \eta^*_\varphi(p)} \times \begin{cases} p + \eta^*_\varphi(p) - t, & 0 \leq t \leq p, \\ \eta^*_\varphi(p) \dfrac{H_\varphi(t)}{H_\varphi(p)}, & p \leq t \leq 1 \end{cases} \tag{7}$$

Its conditional trading rent is $\eta^*_\varphi(p)$ and the corresponding buyer payoff under price p and participation ρ_p is

$$V(p; \varphi) = m_b G_\varphi(p) \frac{\eta^*_\varphi(p)}{p + \eta^*_\varphi(p)} \tag{8}$$

The law $F_{\widehat{\psi}^p_\varphi}$ assigns mass $1 - m_b/[p + \eta^*_\varphi(p)]$ to zero, $W^{\rho_p}(p; \widehat{\psi}^p_\varphi, \varphi)/(1 - m_s)$ to one, and none to $(0, p]$. Seller atoms in $(p, 1)$ induce buyer atoms at the same values. For $p < s < t < 1$,

$$F_{\widehat{\psi}^p_\varphi}(t) - F_{\widehat{\psi}^p_\varphi}(s) = W^{\rho_p}(p; \widehat{\psi}^p_\varphi, \varphi) \left[(H_\varphi)'_+(t) - (H_\varphi)'_+(s) \right]$$

6.2 Buyer's guaranteed continuation payoff

We now relax the auxiliary participation restrictions and give a sufficient condition for the buyer to obtain a payoff arbitrarily close to his fixed-price upper bound $V(p; \varphi)$. A

price p in $(0,1)$ with $G_\varphi(p) > 0$ is *strictly feasible* if $p + \eta^*_\varphi(p) > m_b$, and *binding feasible* if $p + \eta^*_\varphi(p) = m_b$. By Equation 7, under the majorant buyer call $\widehat{\psi}^p_\varphi$, all buyers accept a binding feasible target price, while a strictly feasible target price leaves positive mass at zero. This slack lets the buyer modify his experiment to secure nearly the optimal payoff $V(p;\varphi)$ from the auxiliary problem under every optimal continuation (λ, ρ).

Proposition 2. *Fix a seller put φ and a price $0 < p < 1$ with $G_\varphi(p) > 0$ and $p + \eta^*_\varphi(p) > m_b$. For every $\delta > 0$, there exists a feasible buyer experiment yielding a buyer payoff greater than $V(p;\varphi) - \delta$ under every optimal continuation (λ, ρ).*

Taking the supremum across strictly feasible target prices defines the buyer's guaranteed continuation payoff,

$$\Gamma(\varphi) = \sup_{\substack{0<p<1,\ G_\varphi(p)>0 \\ p+\eta^*_\varphi(p)>m_b}} V(p;\varphi)$$

By Proposition 2, for every fixed optimal continuation (λ, ρ), the buyer's best-response value satisfies

$$B_{\lambda,\rho}(\varphi) \geq \Gamma(\varphi) \tag{9}$$

The buyer's guaranteed payoff $\Gamma(\varphi)$ is positive because prices in $(m_b, 1)$ are strictly feasible and give positive buyer payoff. When binding feasible prices exist, they form an interval because $p + \eta^*_\varphi(p)$ is nondecreasing, and strictly feasible prices lie above its upper endpoint. On this interval the seller CDF is constant and $V(p;\varphi) = G_\varphi(p)(m_b - p)$ decreases with the price. Approaching the upper endpoint, p, by a decreasing sequence of strictly feasible prices gives $V(p;\varphi) \leq \Gamma(\varphi)$. Lower binding prices can yield greater buyer payoff under favorable continuations, so $\Gamma(\varphi)$ need not equal $B_{\lambda,\rho}(\varphi)$.

The supremum over strictly feasible prices, $\Gamma(\varphi)$, is a lower bound on the buyer's best-response value $B_{\lambda,\rho}(\varphi)$ under every optimal continuation. In what follows, we first minimize this guarantee over seller experiments. In Section 6.4, we construct optimal continuation strategies under which full information is a buyer best response, giving the buyer this minimum guarantee and the seller the remaining first-best surplus.

6.3 Seller experiments and buyer payoff bounds

By Equation 9, a buyer best-response value of at most γ requires $\Gamma(\varphi) \leq \gamma$. The seller puts satisfying this necessary condition form the set

$$\mathcal{D}_\gamma = \{\varphi \in \Phi(m_s) : \Gamma(\varphi) \leq \gamma\}$$

As noted before, for a put in $\mathcal{D}_\gamma$ the buyer's continuation value $B_{\lambda,\rho}(\varphi)$ can exceed γ under some optimal continuations. When $\mathcal{D}_\gamma$ is nonempty, define its *convex majorant* as its pointwise supremum.

$$\varphi_\gamma(t) := \sup_{\varphi \in \mathcal{D}_\gamma} \varphi(t), \qquad \text{for } t \text{ in } [0,1] \tag{10}$$

Theorem 3. *Each set $\mathcal{D}_\gamma$ is compact in the uniform topology. Hence Γ is lower semicontinuous. Moreover, if $\mathcal{D}_\gamma$ is nonempty, its convex majorant φ_γ belongs to $\mathcal{D}_\gamma$.*

The next proposition identifies the shape of the convex majorant φ_γ and the buyer's continuation payoff along it.

Proposition 3. *Fix $0 < \gamma < W^{\text{FB}}$ with $\mathcal{D}_\gamma \neq \varnothing$. The seller distribution G_{φ_γ} has a positive atom at zero and is continuous on $(0,1)$ with interior support $[\underline{p}, \bar{p}]$, where $0 < \underline{p} < \bar{p} < 1$. Every price in $(0, \underline{p}]$ is binding feasible, and every price in $(\underline{p}, 1)$ is strictly feasible. Moreover, the buyer's continuation payoff is such that $\Gamma(\varphi_\gamma) = \gamma$ and $V(p; \varphi_\gamma) = \gamma$ for every price p in $[\underline{p}, \bar{p}]$.*

By Theorem 3, the convex majorant φ_γ belongs to $\mathcal{D}_\gamma$. Proposition 3 implies that for every price p in $[\underline{p}, \bar{p}]$, the interior support of G_{φ_γ}, the buyer experiment with majorant call $\widehat{\psi}^p_{\varphi_\gamma}$ gives continuation payoff γ at p whenever an optimal continuation selects that price. Define the minimum buyer guarantee β across all seller experiments by[9]

$$\beta(m_s, m_b) = \min_{\varphi \in \Phi(m_s)} \Gamma(\varphi)$$

After every seller experiment, every buyer best response under an optimal continuation gives the buyer at least β by Equation 9. Since total surplus cannot exceed W^{FB}, seller payoff cannot exceed $W^{\text{FB}} - \beta$. The minimum of Γ over feasible seller puts is attained at a seller put φ^*, by compactness of $\Phi(m_s)$ in the uniform topology and lower semicontinuity of Γ from Theorem 3, hence $\mathcal{D}_\beta$ is non-empty. For any price p in $(m_b, 1)$, Corollary 4 gives

[9] Whenever unambiguous, we drop the argument (m_s, m_b) from the $\beta(m_s, m_b)$ to optimize notation.

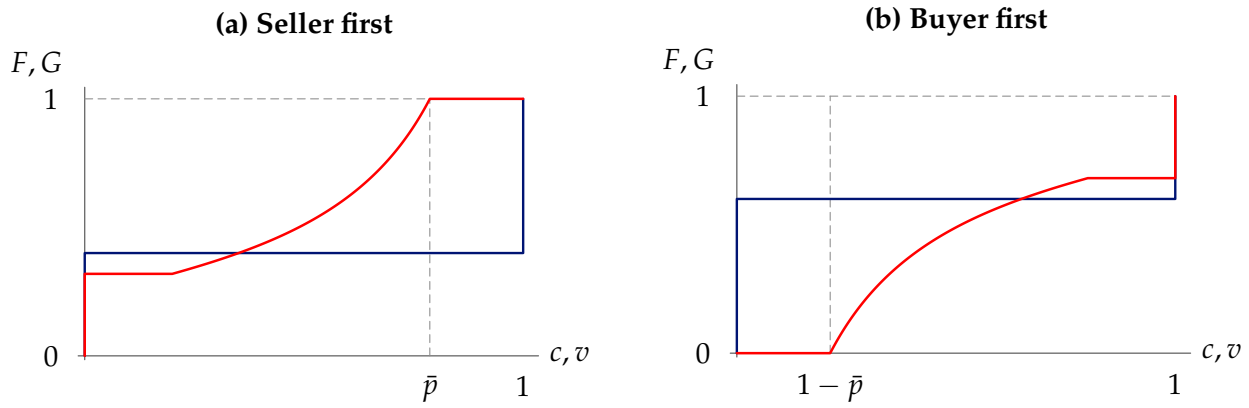


Figure 1: Equilibrium CDFs for $m_s = 0.4$, $m_b = 0.6$. The partially informed first mover (red) and fully informed follower (blue). Dashed lines mark the equilibrium price.

strict feasibility, and Equation 8 gives

$$\beta = \Gamma(\varphi^*) \geq V(p; \varphi^*) > 0$$

By Lemma A.1, some seller experiment guarantees the seller a strictly positive payoff. After that experiment the buyer receives a payoff strictly below W^{FB}, so Equation 9 gives $\beta < W^{\mathrm{FB}}$. Thus Proposition 3 applies at $\gamma = \beta$. The majorant put φ_β belongs to $\mathcal{D}_\beta$ by Theorem 3, and its buyer guarantee is $\Gamma(\varphi_\beta) = \beta$. Additionally, this put φ_β also excludes an atom at one.

Lemma 1. *If $\varphi \in \Phi(m_s)$ and $\Gamma(\varphi) = \beta$, then $G_\varphi(\{1\}) = 0$.*

6.4 Sequential acquisition equilibrium

The seller's payoff cannot exceed $W^{\mathrm{FB}} - \beta$, since the buyer can secure at least β after any seller experiment. Choosing the experiment corresponding to φ_β attains this bound under an optimal continuation in which full information is a buyer best response and the designer selects $\bar{p}$. Trade is efficient, and the buyer receives β, leaving $W^{\mathrm{FB}} - \beta$ to the seller. Thus the seller has no profitable deviation.

Theorem 4. *The seller-first game admits a pure strategy equilibrium in which the seller chooses the experiment corresponding to φ_β, the buyer acquires full information, and the designer selects $\bar{p}$. Every pure seller-first equilibrium gives buyer payoff $\beta(m_s, m_b)$ and seller payoff $W^{\mathrm{FB}} - \beta(m_s, m_b)$.*

The left panel of Figure 1 shows the equilibrium in Theorem 4, with partial seller infor-

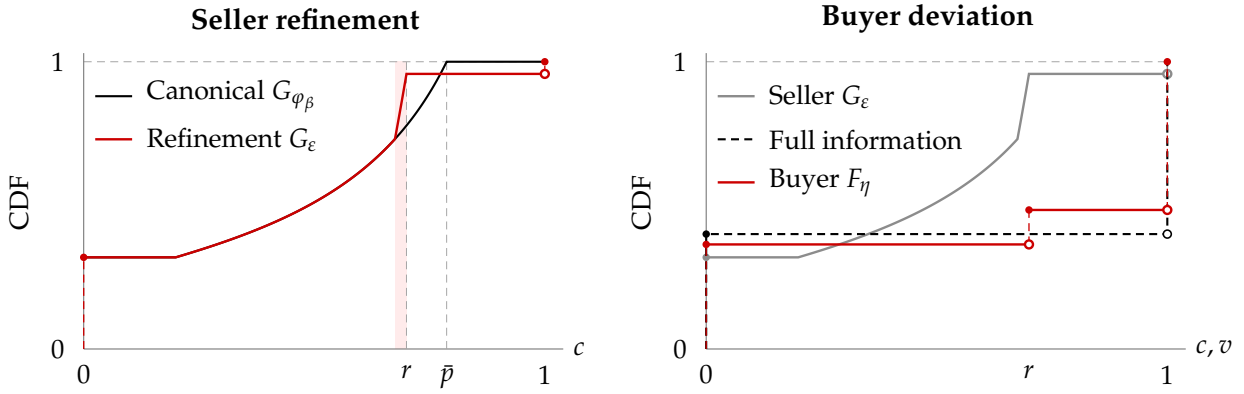


Figure 2: Seller deviation (left) and buyer deviation implementing price r (right). The buyer law F_η assigns mass η to r.

mation (red), full buyer information (blue), and price $\bar{p}$. The proof constructs an optimal continuation (λ, ρ) supporting a buyer best response after every seller experiment in Appendix E.

For payoff uniqueness, consider any pure seller-first equilibrium in which the seller receives less than $W^{\mathrm{FB}} - \beta$. For $\underline{p} < r < \bar{p}$ and small $0 < \varepsilon < r - \underline{p}$, the seller can refine the experiment corresponding to φ_β by splitting each cost above r between a cost in $[r - \varepsilon, r]$ and one, preserving its mean. The refined seller law G_ε has put φ_ε satisfying $\varphi_\varepsilon(p) = \varphi_\beta(p)$ for all $0 \leq p \leq r - \varepsilon$ and strictly exceeding φ_β on $(r, 1)$. Hence $H_{\varphi_\varepsilon}(\underline{p}) < H_{\varphi_\beta}(\underline{p})$. For every buyer experiment F, the designer's maximum total surplus satisfies

$$\max_{p \in [0,1]} W^{(1,1)}(p; F, G_{\varphi_\varepsilon}) \geq \frac{m_b - \underline{p}}{H_{\varphi_\varepsilon}(\underline{p})} > \frac{m_b - \underline{p}}{H_{\varphi_\beta}(\underline{p})} = m_b G_{\varphi_\beta}(0).$$

The equality above follows from binding feasibility at $\underline{p}$. Only zero-cost sellers trade at prices at or below $\underline{p}$, so total surplus there is at most $m_b G_{\varphi_\beta}(0)$. The strict inequality rules out these prices for every buyer experiment.

At an implemented price in $(\underline{p}, r - \varepsilon)$, the buyer's continuation payoff is at most β. For $p \geq r - \varepsilon$, the inequality $(v - p)^+ \leq v(1 - p)$ bounds this payoff by $m_b G_{\varphi_\varepsilon}(p)(1 - p)$. For small ε, this bound increases up to r and decreases above r, with value $m_b G_{\varphi_\varepsilon}(r)(1 - r) > \beta$ at r. Thus, no buyer experiment yields more than the full-information payoff at r with implemented price r.

After the seller chooses the experiment corresponding to φ_ε and the buyer acquires full information, the designer is indifferent between r and prices in $(r, 1)$. If an optimal con-

tinuation selects a price above r, the buyer can profitably deviate to an experiment on $\{0, r, 1\}$ with mean m_b and a sufficiently small positive mass at r. For this experiment, price r with acceptance by value-r buyers realizes all positive gains from trade. Lower prices exclude sellers just below r, and higher prices exclude value-r buyers. Thus every optimal continuation selects r and value-r buyers accept. As the buyer's mass at r vanishes, his payoff approaches $m_b G_{\varphi_\varepsilon}(r)(1-r)$. Subgame perfection therefore requires his best response after the seller deviation to attain this bound through full information and price r.

Figure 2 illustrates the two deviations. The seller then receives at least $m_b \varphi_\beta(r)$, which tends to $W^{\mathrm{FB}} - \beta$ as $r \uparrow \bar{p}$. Any smaller seller equilibrium payoff therefore permits a profitable deviation. Together with the buyer's guarantee β and the total-surplus bound W^{FB}, this establishes both equilibrium payoffs in Theorem 4.

7 Reflection symmetry

For a price $p \in [0,1]$, let $\mathsf{R}p := 1 - p$. For any call or put h, define its *reflection* by $(\mathsf{R}h)(p) = h(1-p)$. This reflection exchanges calls and puts. In particular,

$$\psi \in \Psi(m_b) \quad \Longrightarrow \quad \mathsf{R}\psi \in \Phi(1-m_b) \quad \text{and} \quad \varphi \in \Phi(m_s) \quad \Longrightarrow \quad \mathsf{R}\varphi \in \Psi(1-m_s)$$

Thus a seller-first problem with call ψ, put φ, and price p is mapped into a buyer-first problem with call $\mathsf{R}\varphi$, put $\mathsf{R}\psi$, and price $\mathsf{R}p$. The original seller becomes the reflected buyer, while the original buyer becomes the reflected seller. Seller-first with prior (m_s, m_b) is therefore mapped into buyer-first with prior $(1-m_b, 1-m_s)$.

Reflect each price lottery by $p \mapsto 1-p$ and exchange the traders' participation decisions after the corresponding experiments, announcement, and price. Write $\mathsf{R}\lambda$ and $\mathsf{R}\rho$ for the resulting lottery and participation rule. The payoffs then satisfy

$$\begin{aligned} U_s^{\rho}(p;\psi,\varphi,\lambda) &= U_b^{\mathsf{R}\rho}(1-p;\mathsf{R}\varphi,\mathsf{R}\psi,\mathsf{R}\lambda) \\ U_b^{\rho}(p;\psi,\varphi,\lambda) &= U_s^{\mathsf{R}\rho}(1-p;\mathsf{R}\varphi,\mathsf{R}\psi,\mathsf{R}\lambda) \end{aligned} \tag{11}$$

Proposition 4. *Reflection gives a one-to-one correspondence between pure seller-first acquisition equilibria with prior (m_s, m_b) and pure buyer-first acquisition equilibria with prior $(1-m_b, 1-m_s)$. It preserves the leader's and follower's payoffs.*

In Figure 1, $m_s + m_b = 1$, so $(1 - m_b, 1 - m_s) = (m_s, m_b)$. Reflection therefore maps the left panel's seller-first equilibrium to the right panel's buyer-first equilibrium with the same priors. The next result uses the reflection described above to characterize the equilibrium in the buyer-first game.

Corollary 5. *Both acquisition orders admit pure sequential acquisition equilibria. Every such equilibrium is efficient, and equilibrium payoffs are unique within each acquisition order.*

Proof. Existence and payoff uniqueness follow from Theorem 4 and reflection in Proposition 4. Efficiency follows from Theorem 1. □

Thus the follower's equilibrium payoff is $\beta(m_s, m_b)$ in the seller-first game and $\beta(1 - m_b, 1 - m_s)$ in the buyer-first game.

8 Timing of acquisition

By Theorem 4 and Proposition 4, all pure sequential acquisition equilibria in each order have the same payoff division. The divisions of first-best surplus are given in the following table.

	Seller	Buyer
Seller first	$W^{\text{FB}} - \beta(m_s, m_b)$	$\beta(m_s, m_b)$
Buyer first	$\beta(1 - m_b, 1 - m_s)$	$W^{\text{FB}} - \beta(1 - m_b, 1 - m_s)$

Fixing the prior, the next result compares each trader's payoff across these two games. Each trader earns strictly more when acquiring first.

Theorem 5. *For every $0 < m_s < m_b < 1$, each trader earns strictly more when acquiring first than when acquiring second, across all pure sequential acquisition equilibria in the two orders. Each trader's payoff gain from moving first relative to moving second is*

$$W^{\text{FB}} - \beta(m_s, m_b) - \beta(1 - m_b, 1 - m_s) > 0.$$

As noted earlier, holding the price and the other trader's acceptance probability fixed, full information weakly maximizes a trader's expected trading margin. Learning therefore helps the trader make an informed choice. Additionally, choice of experiments affects the designer's pricing incentives and the opponent's opportunities to influence prices through

his own information acquisition. These strategic effects explain why the incentive to learn as much as possible need not determine the equilibrium choice.

The two acquisition positions offer countervailing advantages. The first mover commits before observing the opponent's information choice, whereas the second mover can tailor his experiment to that commitment and seek more favorable trading terms. The leader must therefore anticipate both the price her experiment induces and how the follower can influence it through learning. Theorem 5 establishes that each trader earns strictly more by committing first despite the follower's ability to condition his information choice on the observed experiment. Fixing the prior and the trader, Theorem 5 compares equilibrium payoff across the two acquisition orders. In equilibrium the first mover chooses an experiment that minimizes the follower's guaranteed continuation payoff, and achieves efficiency. Together with Proposition 1, the results present two roles of timing. First, endogenous sequential acquisition leads to efficient trade in equilibrium. Second, acquisition order affects bargaining power, with each trader obtaining a larger equilibrium share when acquiring first.

9 Conclusion

Simultaneous acquisition admits mixed equilibrium, but every equilibrium is inefficient, whereas each sequential acquisition game admits efficient pure-strategy equilibria with a unique payoff division. Each trader earns strictly more by acquiring first, holding the priors fixed. The impossibility of Myerson and Satterthwaite (1983) takes private information as given. Relative to that benchmark, endogenous information acquisition restores first-best efficiency under robust mechanisms when acquisition is sequential, even though simultaneous acquisition does not. Additionally, in the sequential acquisition equilibria we construct, the leader commits to partial learning while the follower learns fully and this commitment to partial learning gives greater bargaining power than the additional flexibility from the ability to design experiments contingent on the first mover's actions.

A natural next step allows traders to begin with private information, as in Chopra and Ely (Forthcoming). Further learning must then refine that initial information. Our most-informative implementations offer a natural starting point for imposing this constraint. If an intermediary provides tailored experiments, eliciting the traders' initial information would also require incentive compatibility, as in Chopra and Ely (2026b).

References

Au, Pak Hung and Keiichi Kawai, "Competitive Information Disclosure by Multiple Senders," *Games and Economic Behavior*, 2020, *119*, 56–78.

Bergemann, Dirk and Juuso Välimäki, "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 2002, *70* (3), 1007–1033.

_ **and Stephen Morris**, "Robust Mechanism Design," *Econometrica*, 2005, *73* (6), 1771–1813.

Boleslavsky, Raphael and Christopher Cotton, "Limited Capacity in Project Selection: Competition through Evidence Production," *Economic Theory*, 2018, *65* (2), 385–421.

Chopra, Hershdeep and Jeffrey C. Ely, "Equilibrium Information Acquisition in Optimal Auctions," 2026. Working Paper.

_ **and** _ , "Incentive-Compatible Information Design," 2026. Working paper.

_ **and** _ , "Optimal Learning in Selection Markets," 2026. Working paper.

_ **and** _ , "Interim Information Design and Buyer-Optimal Pricing," *American Economic Review: Insights*, Forthcoming.

Chung, Kim-Sau and Jeffrey C. Ely, "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 2007, *74* (2), 447–476.

Condorelli, Daniele and Balázs Szentes, "Information Design in the Holdup Problem," *Journal of Political Economy*, 2020, *128* (2), 681–709.

Čopič, Jernej and Clara Ponsatí, "Optimal Robust Bilateral Trade: Risk Neutrality," *Journal of Economic Theory*, 2016, *163*, 276–287.

Gentzkow, Matthew and Emir Kamenica, "A Rothschild–Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, 2016, *106* (5), 597–601.

_ **and** _ , "Competition in Persuasion," *Review of Economic Studies*, 2017, *84* (1), 300–322.

Hagerty, Kathleen M. and William P. Rogerson, "Robust Trading Mechanisms," *Journal of Economic Theory*, 1987, *42* (1), 94–107.

Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Kyriazis, Panagiotis, "Market Composition and the Consumer Surplus-Profit Frontier in Monopoly Screening," 2026. Working paper, arXiv:2604.09340.

Myerson, Roger B. and Mark A. Satterthwaite, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 1983, *29* (2), 265–281.

Park, Junrok, "Buyer-optimal Learning in Collective Purchase Problems," June 2025. Working paper.

Ravid, Doron, "Ultimatum Bargaining with Rational Inattention," *American Economic Review*, 2020, *110* (9), 2948–2963.

_ , Anne-Katrin Roesler, and Balázs Szentes, "Learning before Trading: On the Inefficiency of Ignoring Free Information," *Journal of Political Economy*, 2022, *130* (2), 346–387.

Roesler, Anne-Katrin and Balázs Szentes, "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 2017, *107* (7), 2072–2080.

Schottmüller, Christoph, "Optimal Information Structures in Bilateral Trade," *Theoretical Economics*, 2023, *18* (1), 421–461.

Shaked, Moshe and J. George Shanthikumar, *Stochastic Orders* Springer Series in Statistics, New York: Springer, 2007.

Simon, Leo K. and William R. Zame, "Discontinuous Games and Endogenous Sharing Rules," *Econometrica*, 1990, *58* (4), 861–872.

Terstiege, Stefan and Cédric Wasser, "Competitive Information Disclosure to an Auctioneer," *American Economic Journal: Microeconomics*, 2022, *14* (3), 622–664.

Wilson, Robert, "Game-Theoretic Analyses of Trading Processes," in Truman F. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*, Cambridge: Cambridge University Press, 1987, chapter 2, pp. 33–70.

Yang, Ming, "Coordination with Flexible Information Acquisition," *Journal of Economic Theory*, 2015, *158*, 721–738.

Appendix

A Efficient surplus and equilibrium prices

Let F and G denote the distributions of the buyer's posterior value and the seller's posterior cost. For U_i^ρ and W^ρ, we use superscript $(1,1)$ when both traders accept at indifference, and $(0,1)$ when the indifferent buyer rejects and the indifferent seller accepts. For any participation rule, each trader's payoff and total surplus are bounded from above by $U_i^{(1,1)}$ and $W^{(1,1)}$, respectively. Moreover, $W^{(1,1)}$ is upper semicontinuous in price.

Lemma A.1. *Suppose $0 < m_s < m_b < 1$. Each trader has a feasible experiment guaranteeing a strictly positive expected payoff uniformly over the opponent's experiments and optimal continuations. Consequently, each trader's payoff is strictly positive in every pure strategy acquisition equilibrium.*

Proof. For the seller, fix $0 < \varepsilon < m_s(m_b - m_s)/4$, set $p_0 = m_s + \varepsilon$, and choose $\widetilde{G} = (\varepsilon/p_0)\delta_0 + (m_s/p_0)\delta_{p_0}$. Fix any buyer experiment F and optimal continuation (λ, ρ). Since F has at most countably many atoms, choose $p \in (p_0, p_0 + \varepsilon)$ with $F(\{p\}) = 0$. At this price all sellers strictly prefer to accept, and no buyer is indifferent with positive probability. Buyer payoff is at least $m_b - p > m_b - m_s - 2\varepsilon > (m_b - m_s)/2$, so optimal expected total surplus exceeds $(m_b - m_s)/2$. At any price below p_0, only zero-cost sellers can trade, yielding a total surplus below $m_b\varepsilon/p_0 < (m_b - m_s)/4$. At every price $p \geq p_0$, the zero-cost sellers receive at least ε times total surplus. Integrating over prices therefore gives seller payoff greater than $\varepsilon(m_b - m_s)/4 > 0$, uniformly over F and optimal continuations.

For the buyer, choose $p_1 < m_b$ sufficiently close to m_b and $\widetilde{F} = \frac{1-m_b}{1-p_1}\delta_{p_1} + \frac{m_b - p_1}{1-p_1}\delta_1$. A similar argument gives a uniform positive buyer payoff. □

A.1 Efficient surplus

Proof of Theorem 1. Fix a pure simultaneous or seller-first equilibrium with experiments (F, G) and optimal continuation (λ, ρ). Write $u_i = u_i^{\lambda,\rho}(F, G)$ for trader i's on-path payoff, $i \in \{b, s\}$. The payoffs (u_b, u_s) and continuation (λ, ρ) remain fixed throughout the proof. The continuation is optimal after every experiment pair, including deviations, with participation as in Section 3.3. Lemma A.1 gives positive payoffs $u_i > 0$. Independence and $(v - c)^+ \leq v(1 - c)$ bound total surplus by W^{FB}.

Suppose for contradiction that $u_b + u_s < W^{\mathrm{FB}}$. Choose $p_0 \in (0,1)$ so that

$$\frac{u_s}{u_b + u_s} < \frac{\varphi_G(p_0)}{1 - m_s} < 1 - \frac{u_b}{W^{\mathrm{FB}}}. \tag{12}$$

Since $u_b > 0$ and $u_b + u_s < W^{\mathrm{FB}}$, we have $u_s/(u_b + u_s) < 1 - u_b/W^{\mathrm{FB}}$. Continuity of φ_G, together with $\varphi_G(0) = 0$ and $\varphi_G(1) = 1 - m_s$, therefore implies such a p_0 exists. We will show that the seller now has a profitable deviation. We first move costs in $(p_0, 1)$ to p_0 or one, preserving their means, then smooth the interior costs.

Claim 1. *For any $p \in [0,1)$, if a feasible seller law $\widehat{G}$ is supported on $[0,p] \cup \{1\}$, the buyer can guarantee payoffs arbitrarily close to $m_b \widehat{G}(p)(1-p)$ under every optimal continuation $(\widehat{\lambda}, \widehat{\rho})$.*

Proof. Fix $q \in (p,1)$ and let $\widehat{F}$ be uniform on $(q, (1+q)/2)$, with mean $\bar{v}$. For sufficiently small $\varepsilon > 0$, the law

$$\widetilde{F} = [1 - m_b - \varepsilon(1 - \bar{v})]\delta_0 + \varepsilon \widehat{F} + (m_b - \varepsilon \bar{v})\delta_1$$

is feasible. At every price in (p,q), all positive-value buyers and all sellers with cost at most p strictly accept. These are exactly the pairs with positive gains from trade, so maximal total surplus is attained independently of participation at indifference. Every price above q excludes some of the continuous buyer mass and gives strictly less surplus. Hence every optimal continuation selects prices at most q and implements all positive-gain trades. Value-one buyers therefore contribute at least $(m_b - \varepsilon\bar{v})\widehat{G}(p)(1-q)$ to buyer payoff. Letting $q \downarrow p$ and $\varepsilon \downarrow 0$ proves the claim. □

If $G((p_0,1)) = 0$, Claim 1 gives buyer payoffs arbitrarily close to $m_b G(p_0)(1-p_0) = m_b[1 - m_s - \varphi_G(p_0)] > u_b$, where the strict inequality follows from Equation 12. This contradicts that the original buyer choice F was optimal, so $G((p_0,1)) > 0$.

Construction 1 (seller deviation). Perturb G by replacing each cost $c > p_0$ by p_0 with probability $(1-c)/(1-p_0)$ and by one otherwise. The resulting distribution G_0 is supported on $[0,p_0] \cup \{1\}$ and has the same mean as G and the same put on $[0,p_0]$. For sufficiently small $\varepsilon > 0$ and for each $0 < c \le p_0$ under G_0, draw $\widetilde{c}$ uniformly from $[(1-\varepsilon)c, c]$, then replace c by $\widetilde{c}$ with probability $(1-c)/(1-\widetilde{c})$ and by one otherwise. This gives a feasible distribution G_ε supported on $[0,p_0] \cup \{1\}$ and continuous on $(0,1)$.[10]

[10] For any x in $(0,1)$, $G_\varepsilon(\{x\}) \le \int_{(0,p_0]} \Pr(\widetilde{c} = x \mid c)\, dG_0(c) = 0$, because each conditional draw is uniform

Perturbing G to G_ε is a mean-preserving spread that shifts mass from a posterior c to some $\widetilde{c} \le c$ and to one. At a fixed price s, under inclusive participation $(1,1)$, a pair $c \le s \le v$ that traded before under G contributes expected gains from trade of at least $v - c$ under either G_0 or G_ε. By construction of G_ε from G_0, $c - \varepsilon \le \widetilde{c} \le c$, giving the following inequality for s in $[0,1]$,

$$0 \le \frac{1-c}{1-\widetilde{c}}(s-\widetilde{c})^+ - (s-c)^+ \le (s-\widetilde{c})^+ - (s-c)^+ \le \varepsilon.$$

Also, by construction of G_ε, $\widetilde{c} \le s$ requires $c \le s + \varepsilon$. Taking expectations over the uniform draw of $\widetilde{c}$ conditional on c, then integrating over c with respect to G_0, gives

$$\begin{aligned} 0 \le \varphi_{G_\varepsilon}(s) - \varphi_{G_0}(s) \le \varepsilon \quad &\text{for } s \text{ in } [0,1], \\ G_\varepsilon(s) \le G_0(s+\varepsilon) = G(s+\varepsilon) \quad &\text{for } s \text{ in } [0, p_0 - \varepsilon). \end{aligned} \tag{13}$$

Thus each perturbation weakly raises $W^{(1,1)}$ at every price, for every buyer law $\widehat{F}$,

$$W^{(1,1)}(s;\widehat{F},G) \le W^{(1,1)}(s;\widehat{F},G_0) \le W^{(1,1)}(s;\widehat{F},G_\varepsilon). \tag{14}$$

Claim 2. *Fix a feasible buyer law $\widehat{F}$ and any optimal continuation $(\widehat{\lambda},\widehat{\rho})$. The lottery $\widehat{\lambda}(\cdot \mid \widehat{F}, G_\varepsilon)$ assigns probability one to maximizers of $W^{(1,1)}(\cdot;\widehat{F},G_\varepsilon)$. Moreover, for every $\gamma > 0$, trader i's expected payoff satisfies*

$$\left| u_i^{\widehat{\lambda},\widehat{\rho}}(\widehat{F},G_\varepsilon) - \int U_i^{(1,1)}(s;\widehat{F},G_\varepsilon)\,\widehat{\lambda}(ds \mid \widehat{F},G_\varepsilon) \right| < \gamma.$$

Proof. Choose a maximizer p of $W^{(1,1)}(\cdot;\widehat{F},G_\varepsilon)$, which exists by upper semicontinuity. For $p > 0$, take $q_n \uparrow p$ in $(0,1)$ avoiding buyer atoms. Buyers with value $v \ge p$ strictly accept at q_n. The pairs that traded at p under $(1,1)$ but are excluded at q_n have $v \ge p$ and $c \in (q_n, p]$. For $p < 1$, their expected gains are at most $G_\varepsilon(p) - G_\varepsilon(q_n)$, which tends to zero by continuity. For $p = 1$, these gains are at most $1 - q_n$. For $p = 0$, take $q_n \downarrow 0$ avoiding buyer atoms. Zero-cost sellers strictly accept, and the excluded buyers have $v < q_n$, so their expected gains are at most q_n. In every case, $W^{\widehat{\rho}}(q_n;\widehat{F},G_\varepsilon,\delta_{q_n}) \le W^{(1,1)}(p;\widehat{F},G_\varepsilon)$ by maximality of p, and these bounds give $W^{\widehat{\rho}}(q_n;\widehat{F},G_\varepsilon,\delta_{q_n}) \longrightarrow W^{(1,1)}(p;\widehat{F},G_\varepsilon)$. For every $\gamma > 0$, convergence and optimality give, for sufficiently large n,

$$\int W^{\widehat{\rho}}\big(s;\widehat{F},G_\varepsilon,\widehat{\lambda}(\cdot \mid \widehat{F},G_\varepsilon)\big)\,\widehat{\lambda}(ds \mid \widehat{F},G_\varepsilon) \ge W^{\widehat{\rho}}(q_n;\widehat{F},G_\varepsilon,\delta_{q_n}) > \max_{t\in[0,1]} W^{(1,1)}(t;\widehat{F},G_\varepsilon) - \gamma.$$

on an interval of positive length.

Thus $W^{\widehat{\rho}}(s;\widehat{F},G_\varepsilon,\widehat{\lambda}(\cdot \mid \widehat{F},G_\varepsilon)) = \max_{t\in[0,1]} W^{(1,1)}(t;\widehat{F},G_\varepsilon)$ for $\widehat{\lambda}(\cdot \mid \widehat{F},G_\varepsilon)$-almost every s. As the seller distribution G_ε is continuous on the interior, her payoff under an arbitrary continuation $(\widehat{\lambda},\widehat{\rho})$ approaches her payoff under inclusive trade $\rho = (1,1)$ along an increasing sequence of implemented prices. Thus for any $\gamma > 0$ we have

$$u_s^{\widehat{\lambda},\widehat{\rho}}(\widehat{F},G_\varepsilon) > \int U_s^{(1,1)}(s;\widehat{F},G_\varepsilon)\,\widehat{\lambda}(ds \mid \widehat{F},G_\varepsilon) - \frac{\gamma}{2}$$

$$\int U_b^{(1,1)}(s;\widehat{F},G_\varepsilon)\,\widehat{\lambda}(ds \mid \widehat{F},G_\varepsilon) > u_b^{\widehat{\lambda},\widehat{\rho}}(\widehat{F},G_\varepsilon) - \frac{\gamma}{2}$$

Under the selected price lottery $\widehat{\lambda}$, the preceding equality gives the same expected total surplus under $\widehat{\rho}$ and under inclusive participation $(1,1)$. Any additional trade under $(1,1)$ gives zero surplus to the trader who was indifferent and nonnegative surplus to the other trader. Thus neither trader can lose from inclusive participation, and a strict gain for either would raise expected total surplus. Each trader therefore earns the same expected payoff under both rules. □

At every price s, $W^{(1,1)}(s;\widehat{F},G_\varepsilon) \le [1-\widehat{F}(s-)](1-m_s)$, since $v-c \le 1-c$ and the traders' values are independent. For $s \ge p_0-\varepsilon$, Equation 13 and the Lipschitz bound for puts give $\varphi_{G_\varepsilon}(s) \ge \varphi_{G_0}(s) \ge \varphi_G(p_0)-\varepsilon$. Under acceptance at indifference, $\rho=(1,1)$, for $s \ge p_0-\varepsilon$,

$$\begin{aligned} U_s^{(1,1)}(s;\widehat{F},G_\varepsilon) &= [1-\widehat{F}(s-)]\varphi_{G_\varepsilon}(s) \\ &\ge [1-\widehat{F}(s-)](\varphi_G(p_0)-\varepsilon) \\ &\ge \frac{\varphi_G(p_0)-\varepsilon}{1-m_s} W^{(1,1)}(s;\widehat{F},G_\varepsilon). \end{aligned} \tag{15}$$

By Equation 12, the fraction $(\varphi_G(p_0)-\varepsilon)/(1-m_s)$ exceeds $u_s/(u_b+u_s)$ for small ε. Thus every price $s \ge p_0-\varepsilon$ with $W^{(1,1)}(s;\widehat{F},G_\varepsilon) \ge u_b+u_s$ gives $U_s^{(1,1)}(s;\widehat{F},G_\varepsilon) > u_s$. We use the following claim to complete the proof in both timings, and then prove the claim.

Claim 3. *For all sufficiently small $\varepsilon > 0$ and every buyer law $\widehat{F}$ that is either F or a best response to G_ε under the fixed equilibrium continuation (λ,ρ), every price $p < p_0-\varepsilon$ maximizing total surplus $W^{\rho}(\cdot;\widehat{F},G_\varepsilon,\lambda(\cdot \mid \widehat{F},G_\varepsilon))$ satisfies $U_b^{(1,1)}(p;\widehat{F},G_\varepsilon) < u_b$.*

In simultaneous acquisition, the buyer chooses F. By Equation 14, $\max_s W^{(1,1)}(s;F,G_\varepsilon) \ge u_b+u_s$, the original equilibrium's expected total surplus. By Claim 2, every maximizer of total surplus $W^{\rho}(\cdot;F,G_\varepsilon,\lambda(\cdot \mid F,G_\varepsilon))$ also maximizes $W^{(1,1)}(\cdot;F,G_\varepsilon)$. At such a price

below $p_0 - \varepsilon$, Claim 3 gives $U_b^{(1,1)} < u_b$, leaving $U_s^{(1,1)} > u_s$. At such a price $s \geq p_0 - \varepsilon$, Equation 15 gives

$$\begin{aligned} U_s^{(1,1)}(s; F, G_\varepsilon) &\geq \frac{\varphi_G(p_0) - \varepsilon}{1 - m_s} W^{(1,1)}(s; F, G_\varepsilon) \\ &\geq \frac{\varphi_G(p_0) - \varepsilon}{1 - m_s}(u_b + u_s) > u_s. \qquad \text{(by Equation 12)} \end{aligned}$$

Applying Claim 2 with $(\widehat{\lambda}, \widehat{\rho}) = (\lambda, \rho)$, total surplus equals $\max_t W^{(1,1)}(t; F, G_\varepsilon)$ almost everywhere under $\lambda(\cdot \mid F, G_\varepsilon)$, so $\int U_s^{(1,1)}(s; F, G_\varepsilon)\, \lambda(ds \mid F, G_\varepsilon) > u_s$. Choose γ smaller than this difference. The same claim then gives $u_s^{\lambda,\rho}(F, G_\varepsilon) > u_s$, contradicting the assumption $u_b + u_s < W^{\mathrm{FB}}$ in simultaneous acquisition equilibrium.

In seller-first acquisition, the buyer may change his experiment in response to the seller's deviation. Let $\widehat{F}$ be the buyer best response prescribed by the equilibrium after G_ε, with subsequent price lottery $\lambda(\cdot \mid \widehat{F}, G_\varepsilon)$. By Claim 1, the buyer's payoff after the seller's deviation satisfies

$$\begin{aligned} u_b^{\lambda,\rho}(\widehat{F}, G_\varepsilon) &\geq m_b(1 - p_0)G_\varepsilon(p_0) \\ &\geq m_b[1 - m_s - \varphi_G(p_0) - \varepsilon] \\ &> u_b. \end{aligned} \tag{16}$$

Here $(1 - p_0)G_\varepsilon(p_0) = 1 - m_s - \varphi_{G_\varepsilon}(p_0)$ by the support of G_ε, and Equation 13 bounds its put by $\varphi_G(p_0) + \varepsilon$. The strict inequality follows from Equation 12 for small ε.

We first show that $\max_s W^{(1,1)}(s; \widehat{F}, G_\varepsilon) > u_b + u_s$. Suppose instead that this maximum is at most $u_b + u_s$. At a price below $p_0 - \varepsilon$ maximizing total surplus $W^\rho(\cdot; \widehat{F}, G_\varepsilon, \lambda(\cdot \mid \widehat{F}, G_\varepsilon))$, Claim 3 gives $U_b^{(1,1)} < u_b$. At any other price s maximizing this function, subtracting Equation 15 from $W^{(1,1)}(s; \widehat{F}, G_\varepsilon)$ gives

$$\begin{aligned} U_b^{(1,1)}(s; \widehat{F}, G_\varepsilon) &\leq \left(1 - \frac{\varphi_G(p_0) - \varepsilon}{1 - m_s}\right) W^{(1,1)}(s; \widehat{F}, G_\varepsilon) \\ &\leq \left(1 - \frac{\varphi_G(p_0) - \varepsilon}{1 - m_s}\right)(u_b + u_s) \\ &< u_b. \qquad \text{(by Equation 12)} \end{aligned}$$

By Claim 2 applied to (λ, ρ), the lottery $\lambda(\cdot \mid \widehat{F}, G_\varepsilon)$ assigns probability one to prices that maximize $W^\rho(\cdot; \widehat{F}, G_\varepsilon, \lambda(\cdot \mid \widehat{F}, G_\varepsilon))$. Every such price falls into one of the two cases above, $s < p_0 - \varepsilon$ or $s \geq p_0 - \varepsilon$. Integrating these bounds gives $u_b^{\lambda,\rho}(\widehat{F}, G_\varepsilon) \leq u_b$, contradicting

Equation 16. Thus $\max_s W^{(1,1)}(s;\widehat{F},G_\varepsilon) > u_b + u_s$, where $u_b + u_s$ is the original equilibrium's fixed expected total surplus. The same two price cases and choice of γ used for the simultaneous case give $u_s^{\lambda,\rho}(\widehat{F},G_\varepsilon) > u_s$. This contradicts the assumption $u_b + u_s < W^{\mathrm{FB}}$ in seller-first equilibrium. The buyer-first case follows by symmetry. □

Proof of Claim 3. By Claim 2, it suffices to prove the assertion for maximizers of $W^{(1,1)}$. It therefore suffices to prove the assertion for maximizers of $W^{(1,1)}$. Assume for contradiction that there exist $\varepsilon_N \downarrow 0$, laws F_N equal to F or best responses to G_{ε_N} under (λ,ρ), and maximizers $p_N < p_0 - \varepsilon_N$ of $W^{(1,1)}(\cdot;F_N,G_{\varepsilon_N})$ with $U_b^{(1,1)}(p_N;F_N,G_{\varepsilon_N}) \geq u_b$. By compactness of $\mathcal{F}(m_b)$ under weak convergence and of $[0,p_0]$, and by Construction 1, we can choose a subsequence with $F_N \Rightarrow \widehat{F}$, $G_{\varepsilon_N} \Rightarrow G_0$, and $p_N \to p \leq p_0$.

Then $\psi_{F_N} \to \psi_{\widehat{F}}$ and $\varphi_{G_{\varepsilon_N}} \to \varphi_{G_0}$ pointwise, hence uniformly by their common 1-Lipschitz bound. Weak convergence also implies $\limsup_N F_N([p_N,1]) \leq \widehat{F}([p,1])$, while Equation 13 and right continuity of G give $\limsup_N G_{\varepsilon_N}(p_N) \leq G(p)$. These bounds, together with $U_b^{(1,1)}(p_N;F_N,G_{\varepsilon_N}) \geq u_b$ and $\varphi_{G_0}(p) = \varphi_G(p)$, yield

$$G(p)\psi_{\widehat{F}}(p) \geq u_b \quad \text{and} \quad W^{(1,1)}(p;\widehat{F},G) \geq \limsup_N W^{(1,1)}(p_N;F_N,G_{\varepsilon_N}) > u_b. \tag{17}$$

If $F_N = F$, Equation 14 gives $W^{(1,1)}(p_N;F_N,G_{\varepsilon_N}) \geq u_b + u_s > u_b$. Otherwise, the expected buyer payoff, and hence maximal surplus, is bounded below by $m_b[1 - m_s - \varphi_G(p_0) - \varepsilon_N] \to m_b[1 - m_s - \varphi_G(p_0)] > u_b$ by Equation 16 and Equation 12. The strict inequality in Equation 17 follows from these two cases.

Write $I_\varphi(s) = \int_s^1 \varphi(t)^{-2}\,dt$. For $\varphi_{G_0}(s) > 0$, uniform convergence of the puts gives $I_{\varphi_{G_{\varepsilon_N}}}(s) \to I_{\varphi_{G_0}}(s) < \infty$. Integrating $-(\psi_{F_N}/\varphi_{G_{\varepsilon_N}})' = W^{(0,1)}(\cdot;F_N,G_{\varepsilon_N})/\varphi^2_{G_{\varepsilon_N}}$ from s to 1 and using maximality of p_N and Equation 17 yields

$$\psi_{\widehat{F}}(s) \leq W^{(1,1)}(p;\widehat{F},G)\varphi_{G_0}(s)I_{\varphi_{G_0}}(s).$$

Define $\psi^*(s) = W^{(1,1)}(p;\widehat{F},G)\varphi_G(s)I_{\varphi_G}(s)$ for $s > p$, extending continuously to p. This function is convex, decreasing, and vanishes at one. By Construction 1 and the preceding bound, it strictly exceeds $\psi_{\widehat{F}}$ on $(p,p_0]$, and at p if $\varphi_G(p) > 0$. It also solves $G(s)\psi^*(s) - \varphi_G(s)\psi^{*\prime}(s+) = W^{(1,1)}(p;\widehat{F},G)$. If $\varphi_G(p) > 0$, this equation and $pG(p) \geq \varphi_G(p)$ give

$$\psi^*(p) - p\psi^{*\prime}(p+) = \int_{[p,1]} v\,d\widehat{F}(v) - \left(\frac{pG(p)}{\varphi_G(p)} - 1\right)[\psi^*(p) - \psi_{\widehat{F}}(p)] \leq m_b.$$

The integral term above is at most m_b, and the subtracted term is nonnegative.

If $\varphi_G(p) = 0$, then $p < p_0$ and Equation 17 gives $G(\{p\}) = G(p) > 0$. Right continuity gives $\varphi_G(s) = \int_p^s G(t)\,dt \sim G(p)(s-p)$ as $s \downarrow p$, so $\psi^*(p+) = W^{(1,1)}(p; \widehat{F}, G)/G(p) = \psi_{\widehat{F}}(p)$, since $\varphi_G(p) = 0$ implies that seller payoff at p is zero. Equality of ψ^* and $\psi_{\widehat{F}}$ at p and the strict comparison to its right imply $\psi^{*\prime}(p+) \geq \psi'_{\widehat{F}}(p+)$. Hence $\psi^*(p) - p\psi^{*\prime}(p+) \leq \int_{(p,1]} v\,d\widehat{F}(v) \leq m_b$.

The call comparison and $(v-p)^+ \geq v - p$ give $\psi^*(p) \geq \psi_{\widehat{F}}(p) \geq m_b - p$. For $p > 0$, extend ψ^* linearly from $(0, m_b)$ to $(p, \psi^*(p))$. Its slope satisfies $-1 \leq (\psi^*(p) - m_b)/p \leq \psi^{*\prime}(p+) \leq 0$. Thus ψ^* is the call function of a feasible law F^*.

By definition of ψ^*, we have $W^{(0,1)}(s; F^*, G) = W^{(1,1)}(p; \widehat{F}, G)$ for $p < s < 1$. If $\varphi_G(p) > 0$, the strict call comparison gives $G(p)\psi^*(p) > G(p)\psi_{\widehat{F}}(p) \geq u_b$ by Equation 17. If $\varphi_G(p) = 0$, seller surplus at p is zero, so $G(p)\psi^*(p) = W^{(1,1)}(p; \widehat{F}, G) > u_b$ by the same equation. Right continuity gives $G(q)\psi^*(q) > u_b$ for $q > p$ close to p, while the strict inequality $\psi^* > \psi_{\widehat{F}}$ to the right of p implies that $\psi^*(q) > \psi_{\widehat{F}}(q) \geq m_b - q$. Fix such a q with $G(\{q\}) = 0$. We perturb F^* so that prices above q give less total surplus than prices approaching q from below.

Construct F^*_δ from F^* by drawing v_δ uniformly from $[v - 2\delta(v-q), v]$ for each $v > q$. By construction, F^*_δ is continuous above q. This leftward shift reduces the mean. Restore the mean to m_b by moving the mass of $[q - \delta', q)$ to q, with $0 \leq q - \delta' \leq p$. If the interval reaches zero, split the zero atom if needed, but move it all when $p = 0$. Since $\psi^*(q) > m_b - q$, there is mass below q, so the perturbation is well defined for small enough $\delta > 0$ and q close enough to p. For q close enough to p, choose $0 < \delta < 1/2$ and δ' to preserve the mean and satisfy $(1-\delta)G(q)\psi^*(q) > u_b$. By construction, ψ^* is linear on $[0, p]$, so $F^*((0,p)) = 0$. The perturbation then gives, for q close enough to p, $F^*_\delta((0,q)) = 0$ and $F^*_\delta(\{q\}) > 0$.

For $s \leq q$, every trade with positive gains also occurs at q under inclusive participation $(1,1)$. Choose $s_n \uparrow q$ in $(0,q)$ with $F^*_\delta(\{s_n\}) = G(\{s_n\}) = 0$. Neither trader is indifferent with positive probability at these prices, so announcing the deterministic price s_n yields $W^{(1,1)}(s_n; F^*_\delta, G)$ under every participation rule. These surpluses converge to $W^{(1,1)}(q; F^*_\delta, G)$ because all positive-value buyers accept and $G(\{q\}) = 0$.

At a price $s > q$, a buyer who accepts under F^*_δ had value $v > s$ under F^*. Thus, every such trade at price s also occurs under F^* with participation $(0,1)$. For a pair (v,c) drawn from $F^* \times G$ with $c \leq s < v$, the expected gains from trade after perturbing F^* to F^*_δ are at most

$(1-\delta)(v-c)+\delta(q-c)^+$. Integrating the $v-c$ terms gives $W^{(0,1)}(q;F^*,G)$ for $s<1$, and zero at $s=1$. Including all $v>q$ bounds the remaining terms by $F^*((q,1])\varphi_G(q)$. At q, the payoff formula gives $W^{(0,1)}(q;F^*,G)=G(q)\psi^*(q)+F^*((q,1])\varphi_G(q)$. The perturbation gives $\psi_{F^*_\delta}(q)=(1-\delta)\psi^*(q)$ and preserves the mass strictly above q. Substituting these identities yields

$$\begin{aligned}W^{(1,1)}(s;F^*_\delta,G)&\leq(1-\delta)W^{(0,1)}(q;F^*,G)+\delta F^*((q,1])\varphi_G(q)\\&=G(q)\psi_{F^*_\delta}(q)+F^*_\delta((q,1])\varphi_G(q)\\&<W^{(1,1)}(q;F^*_\delta,G).\end{aligned}$$

The strict inequality holds because acceptance by buyers with value q adds seller payoff $F^*_\delta(\{q\})\varphi_G(q)>0$, where $\varphi_G(q)\geq(q-p)G(p)>0$. No price can yield more than $W^{(1,1)}(q;F^*_\delta,G)$, whereas the deterministic prices s_n approach this bound. Optimality therefore forces expected total surplus to equal $W^{(1,1)}(q;F^*_\delta,G)$ under every optimal continuation $(\widehat{\lambda},\widehat{\rho})$. Since every price above q yields strictly less, $\widehat{\lambda}(\cdot\mid F^*_\delta,G)$ assigns no probability above q. Seller payoff at any selected price is at most $F^*_\delta([q,1])\varphi_G(q)$. Subtracting this bound from expected total surplus gives buyer payoff at least

$$\begin{aligned}W^{(1,1)}(q;F^*_\delta,G)-F^*_\delta([q,1])\varphi_G(q)&=G(q)\psi_{F^*_\delta}(q)\\&=(1-\delta)G(q)\psi^*(q).\end{aligned}$$

This is greater than u_b by $G(p)\psi^*(p)>u_b$, right continuity, and the choices of q and δ. Thus the deviation is profitable under the continuation (λ,ρ), contradicting buyer optimality. □

A.2 No price dispersion

Proof of Corollary 1. In a seller-first equilibrium with experiments (F,G) and continuation (λ,ρ), Lemma A.1 and Theorem 1 give positive payoffs and first-best total surplus. Since $(v-c)^+\leq v(1-c)$ is strict when both values are interior, independence forces at least one trader to learn fully. Full seller learning would give $G(0)=1-m_s$, so by construction in Claim 1 with cutoff close to zero, the buyer can obtain payoff arbitrarily close to W^{FB}, contradicting $u_s^{\lambda,\rho}(F,G)>0$. Thus the buyer learns fully and $G((0,1))>0$.

Set $p_1=\sup(\operatorname{supp}G\cap[0,1))$, by above $p_1>0$. Any price below p_1 excludes positive-gain trades with value-one buyers. Efficiency therefore requires $\operatorname{supp}\lambda(\cdot\mid F,G)\subseteq[p_1,1]$. If $p_1=1$, this forces price one and $u_b^{\lambda,\rho}(F,G)=0$. Hence $p_1<1$.

If $\lambda(\cdot \mid F, G) \neq \delta_{p_1}$, efficiency and full buyer information give

$$u_b^{\lambda,\rho}(F,G) = m_b G(p_1) \int_{[p_1,1]} (1-p)\, \lambda(dp \mid F, G) < m_b G(p_1)(1-p_1).$$

By Claim 1 at cutoff p_1, the buyer then has a profitable deviation, contradicting optimality of full learning. Thus $\lambda(\cdot \mid F, G) = \delta_{p_1}$. Buyer-first follows by symmetry. The same proof applies to simultaneous acquisition, where F and G are mutual best responses. □

A.3 Equilibrium support

Proof of Corollary 2. The proof of Corollary 1 shows that the buyer learns fully in seller-first acquisition, and, by symmetry, the seller learns fully in buyer-first acquisition. The stated support restrictions follow, since the equilibrium price is interior. The same arguments establish these support restrictions in simultaneous acquisition. □

A.4 Simultaneous acquisition

Proof of Corollary 3. Suppose a pure simultaneous acquisition equilibrium exists with experiments (F, G) and optimal continuation (λ, ρ). By Theorem 1 and Lemma A.1,

$$\begin{aligned} u_b^{\lambda,\rho}(F,G) + u_s^{\lambda,\rho}(F,G) &= W^{\mathrm{FB}}, \\ 0 < u_i^{\lambda,\rho}(F,G) &< W^{\mathrm{FB}}, \qquad i \in \{b,s\}. \end{aligned}$$

Since $(v-c)^+ \leq v(1-c)$ is strict when both posterior values are interior, efficiency and independence require at least one trader to learn fully. If the seller learns fully, $G(0) = 1 - m_s$, and Claim 1 at cutoff zero gives the buyer payoffs arbitrarily close to W^{FB} under every optimal continuation. A sufficiently close approximation exceeds $u_b^{\lambda,\rho}(F,G)$, contradicting buyer optimality. If the buyer learns fully, a similar argument gives the seller a profitable deviation. □

A.5 Mixed simultaneous acquisition

Proof of Proposition 1. We fist establish the existence of a mixed equilibrium using the endogenous sharing rules studied in Simon and Zame (1990).

1. The Simon–Zame theorem. We apply the existence theorem in Section 2 of Simon and

Zame (1990). A game with an endogenous sharing rule consists of a player's strategy space and a payoff correspondence. The payoff correspondence maps a strategy profile (an element of the cartesian product of players' strategy spaces) to all conceivable payoff pairs. A sharing rule is a Borel measurable selection from the payoff correspondence. A solution for the game is a sharing rule and a mixed strategy profile with the property that, given the sharing rule, each player's action is a best response to the actions of the other player. If the strategy spaces are compact metric and the payoff correspondence is bounded and upper hemicontinuous with nonempty, compact, convex values, the theorem implies that every game with an endogenous sharing rule has a solution. Here the pure strategies are feasible experiments F and G, and a sharing rule is a choice $q(F, G)$ of payoffs attainable under optimal continuations. In the following, we verify these properties for the payoff correspondence $\mathcal{U}$ defined below and implement the selected sharing rule through an optimal continuation.

2. Sharing rules and optimal continuations.

Let $\mathcal{A}(F, G)$ be the payoff pairs attainable from arbitrary price lotteries λ and admissible participation strategies ρ. In particular, these lotteries need not be total surplus maximizing. For (F, G), a triple $(p, \alpha_b, \alpha_s) \in [0,1]^3$ is *feasible* if $1 - F(p) \leq \alpha_b \leq 1 - F(p-)$ and $G(p-) \leq \alpha_s \leq G(p)$, with $F(0-) = G(0-) = 0$. The aggregate acceptance probabilities α_b, α_s then give payoff pair $(\alpha_s \psi_F(p), \alpha_b \varphi_G(p))$ by Equation 1. Equivalently, $\mathcal{A}(F, G)$ consists of expected payoffs from probability distributions over these triples.

Let $\underline{W}(F, G) = \sup_p W^{(0,0)}(p; F, G)$ be the supremum of total surplus when indifferent traders reject. For each fixed p, the nonnegative integrand $(v - c)\mathbb{1}\{c < p < v\}$ is lower semicontinuous, so $W^{(0,0)}(p; F, G)$ is lower semicontinuous under weak convergence of (F, G). Taking the supremum over p preserves lower semicontinuity.

Consider the following payoff correspondence

$$\mathcal{U}(F, G) = \{x \in \mathcal{A}(F, G) \mid x_b + x_s \geq \underline{W}(F, G)\}. \tag{18}$$

$\mathcal{U}(F, G)$ consists precisely of the payoff pairs attainable under continuations optimal after (F, G). Indeed, the designer can obtain at least $W^{(0,0)}(p; F, G)$ at every deterministic price, so any optimal continuation yielding x satisfies $x_b + x_s \geq \underline{W}(F, G)$. Conversely, take $x \in \mathcal{U}(F, G)$ and a price lottery $\lambda(\cdot \mid F, G)$ with participation decisions generating x. Prescribe those decisions after this announcement and rejection at indifference after every

other lottery λ'. Deviating to λ' then gives the designer

$$\int W^{(0,0)}(p; F, G)\, d\lambda'(p) \leq \underline{W}(F, G) \leq x_b + x_s.$$

Thus the prescribed lottery is optimal and implements x.

3. Verification and equilibrium existence.

Let $F_n \Rightarrow F$, $G_n \Rightarrow G$, and let feasible triples $(p_n, \alpha_{b,n}, \alpha_{s,n})$ for (F_n, G_n) converge to (p, α_b, α_s). Since $F_n \otimes \delta_{p_n} \Rightarrow F \otimes \delta_p$, the weak-convergence bounds for the open set $\{v > t\}$ and the closed set $\{v \geq t\}$ in $[0,1]^2$ give

$$1 - F(p) \leq \liminf_n [1 - F_n(p_n)] \leq \alpha_b \leq \limsup_n [1 - F_n(p_n-)] \leq 1 - F(p-).$$

Similarly, applying these bounds to $G_n \otimes \delta_{p_n} \Rightarrow G \otimes \delta_p$ and the sets $\{c < t\}$ and $\{c \leq t\}$ gives

$$G(p-) \leq \liminf_n G_n(p_n-) \leq \alpha_s \leq \limsup_n G_n(p_n) \leq G(p).$$

Thus (p, α_b, α_s) is feasible, including at $p = 0, 1$.

Take $x^n \in \mathcal{A}(F_n, G_n)$ and let ν_n be a law on feasible triples generating x^n. Compactness of $\mathcal{F}(m_b)$, $\mathcal{G}(m_s)$, and $\Delta([0,1]^3)$ gives a subsequence along which (F_n, G_n, ν_n) converges weakly to (F, G, ν). The preceding inequalities show that the set of $(F, G, p, \alpha_b, \alpha_s)$ with (p, α_b, α_s) feasible for (F, G) is closed. The joint laws $\delta_{(F_n, G_n)} \otimes \nu_n$ assign this set probability one, so their weak limit does too. Thus ν is concentrated on triples feasible for (F, G). Uniform convergence of calls and puts gives convergence of x^n to the payoff generated by ν. Hence the graph of $\mathcal{A}$, the set of all (F, G, x) with $x \in \mathcal{A}(F, G)$, is compact. Each $\mathcal{A}(F, G)$ is convex because the laws generating two payoff pairs can be mixed.

By Equation 18, the graph of $\mathcal{U}$ is the subset of this compact graph satisfying $x_b + x_s \geq \underline{W}(F, G)$. Lower semicontinuity of $\underline{W}$ makes this subset closed. For $(F_n, G_n) \to (F, G)$, any sequence $x^n \in \mathcal{U}(F_n, G_n)$ has a convergent subsequence in $[0,1]^2$, and the closed graph puts its limit in $\mathcal{U}(F, G)$. This gives upper hemicontinuity. Each value is compact and convex as the intersection of $\mathcal{A}(F, G)$ with a closed half-space. Each value is nonempty because the maximum of $x_b + x_s$ over $\mathcal{A}(F, G)$ is at least $W^{(0,0)}(p; F, G)$ for every p, and hence at least $\underline{W}(F, G)$. By Simon and Zame (1990), there is a Borel measurable sharing rule $q(F, G) \in \mathcal{U}(F, G)$ and a mixed Nash equilibrium (π_b, π_s) for q. For each experiment

pair, select measurably a law on feasible triples generating $q(F, G)$. The construction following Equation 18 implements these payoffs through an optimal continuation (λ, ρ), so $(\pi_b, \pi_s, \lambda, \rho)$ is a mixed simultaneous acquisition equilibrium.

Strict inefficiency. Suppose a mixed equilibrium attains expected total surplus W^{FB}. Both traders have positive expected payoffs by Lemma A.1. If each chose experiments with interior posterior mass with positive probability, independence would give positive probability to $v, c \in (0, 1)$. On this event, $(v - c)^+ < v(1 - c)$, so expected gains from trade would be strictly below $m_b(1 - m_s) = W^{\mathrm{FB}}$. Thus one trader chooses full information almost surely.

The deviation argument in Appendix A.4 now gives the other trader payoffs arbitrarily close to W^{FB} under every optimal continuation. The fully informed trader's strictly positive expected payoff leaves the other trader with expected payoff strictly below W^{FB}. Choosing a sufficiently close approximation gives a profitable pure deviation from that trader's mixed strategy, a contradiction. □

B The buyer's fixed-price problem

B.1 The buyer's majorant and payoff

If at least one of F and G has no atoms in $(0, 1)$, every optimal continuation (λ, ρ) satisfies, for $\lambda(\cdot \mid F, G)$-almost every p,

$$\begin{aligned} U_i^{\rho}\big(p; F, G, \lambda(\cdot \mid F, G)\big) &= U_i^{(1,1)}(p; F, G), \qquad i \in \{b, s\}, \\ W^{(1,1)}(p; F, G) &= \max_{t \in [0,1]} W^{(1,1)}(t; F, G). \end{aligned} \tag{19}$$

The designer can obtain surplus arbitrarily close to $\max_t W^{(1,1)}(t; F, G)$ using deterministic prices, by the argument of Claim 2. Since $W^{\rho} \le W^{(1,1)} \le \max_t W^{(1,1)}(t; F, G)$, optimality of the designer's lottery makes both inequalities equalities. Each payoff satisfies $U_i^{\rho} \le U_i^{(1,1)}$, so equality of their sum implies equality of both payoffs.

Proof of Theorem 2. The pointwise supremum $\widehat{\psi}_{\varphi}^{p}$ is convex and nonincreasing, has slopes in $[-1, 0]$, and satisfies $\widehat{\psi}_{\varphi}^{p}(0) = m_b$ and $\widehat{\psi}_{\varphi}^{p}(1) = 0$. Thus it is a feasible call and linear on $[0, p]$.

For any feasible call ψ linear on $[0, p]$, total surplus at p under ρ_p is

$$W^{\rho_p}(p; \psi, \varphi) = G_\varphi(p)\psi(p) + \frac{\varphi(p)}{p}\big[m_b - \psi(p)\big].$$

This expression is continuous and nondecreasing in $\psi(p)$ because $G_\varphi(p) \geq \varphi(p)/p$.

To prove compactness of $\mathcal{C}_p(\varphi)$, let ψ_n in this set converge uniformly to ψ in the compact set $\Psi(m_b)$. Weak convergence gives $F_\psi((0, p)) = 0$. For $p < t < 1$, convexity gives $(\psi_n)'_+(t) \leq [\psi_n(t+\varepsilon) - \psi_n(t)]/\varepsilon$. Taking $n \to \infty$ and then $\varepsilon \downarrow 0$ gives $\limsup_n (\psi_n)'_+(t) \leq (\psi)'_+(t)$. Hence

$$\begin{aligned} W^{\rho_p}(t; \psi, \varphi) &\leq \liminf_n W^{\rho_p}(t; \psi_n, \varphi) \\ &\leq \liminf_n W^{\rho_p}(p; \psi_n, \varphi) \\ &= W^{\rho_p}(p; \psi, \varphi). \end{aligned}$$

The first inequality uses $W^{\rho_p}(t; \psi, \varphi) = G_\varphi(t)\psi(t) - \varphi(t)(\psi)'_+(t)$. The second holds because each ψ_n implements p, and the equality follows from continuity of surplus at p. Thus $\psi \in \mathcal{C}_p(\varphi)$, proving that $\mathcal{C}_p$ is closed and compact.

For each t in $(p, 1)$, compactness of $\mathcal{C}_p(\varphi)$ gives a call ψ_t in this set with $\psi_t(t) = \widehat{\psi}^p_\varphi(t)$. Since $\widehat{\psi}^p_\varphi \geq \psi_t$ and the calls agree at t, comparison of their right derivatives gives $(\widehat{\psi}^p_\varphi)'_+(t) \geq (\psi_t)'_+(t)$. Therefore

$$\begin{aligned} W^{\rho_p}(t; \widehat{\psi}^p_\varphi, \varphi) &\leq W^{\rho_p}(t; \psi_t, \varphi) \\ &\leq W^{\rho_p}(p; \psi_t, \varphi) \\ &\leq W^{\rho_p}(p; \widehat{\psi}^p_\varphi, \varphi). \end{aligned}$$

The first inequality uses equality of the calls at t and the derivative comparison, the second uses $\psi_t \in \mathcal{C}_p(\varphi)$, and the third uses monotonicity of surplus at p in the call's value there. Thus $\widehat{\psi}^p_\varphi \in \mathcal{C}_p(\varphi)$.

Since $\widehat{\psi}^p_\varphi \geq \psi$ for every $\psi \in \mathcal{C}_p(\varphi)$, it maximizes buyer payoff $G_\varphi(p)\psi(p)$ and total surplus at p, attaining $V(p; \varphi)$. By construction, $F_{\widehat{\psi}^p_\varphi}$ is the most informative law with a call in $\mathcal{C}_p(\varphi)$. □

Proof of Corollary 4. Let p be implementable and write $W^* = W^{\rho_p}(p; \widehat{\psi}^p_\varphi, \varphi) > 0$.

Suppose for contradiction that $F_{\widehat{\psi}^p_\varphi}(\{p\}) > 0$, and let ψ_δ denote the buyer call after either perturbation below. If $\varphi(p) > 0$, acceptance at p and rejection at indifference above p under

ρ_p imply $\lim_{t\downarrow p} W^{\rho_p}(t;\widehat{\psi}^p_\varphi,\varphi) = W^* - \varphi(p)F_{\widehat{\psi}^p_\varphi}(\{p\}) < W^*$. This limit implies that we can choose $0 < \varepsilon < 1 - p$ and $0 < \delta \leq F_{\widehat{\psi}^p_\varphi}(\{p\})$ small enough that $W^{\rho_p}(t;\widehat{\psi}^p_\varphi,\varphi) + 2\delta < W^*$ for every $p < t < p + \varepsilon$. Move mass δ from p to zero and $p + \varepsilon$ with respective probabilities $\varepsilon/(p+\varepsilon)$ and $p/(p+\varepsilon)$, preserving its mean. Total surplus at p weakly rises because $pG_\varphi(p) \geq \varphi(p)$. Above p, total surplus changes only on $(p, p+\varepsilon)$, by at most 2δ.

When $\varphi(p) = 0$, the right limit equals W^*, so the preceding inequality cannot hold throughout $(p, p+\varepsilon)$ even though $\varphi(t) > 0$ for $t > p$. Here $\widehat{\psi}^p_\varphi(p) > 0$ since $W^* > 0$. Multiply buyer mass on $(p, 1]$ by $1 + \delta$ and add mass $\delta\widehat{\psi}^p_\varphi(p)/p$ at zero, subtracting the total added mass $\frac{\delta}{p}\int_{(p,1]} v\, dF_{\widehat{\psi}^p_\varphi}(v)$ from the atom at p. Choose $\delta > 0$ so that this removed mass does not exceed the atom's mass. This preserves the mean and gives $\psi_\delta(t) = (1+\delta)\widehat{\psi}^p_\varphi(t)$ for $t \geq p$. Moreover, $W^{\rho_p}(t;\psi_\delta,\varphi) = (1+\delta)W^{\rho_p}(t;\widehat{\psi}^p_\varphi,\varphi)$ for $p \leq t < 1$, and total surplus is zero for $t < p$ and $t = 1$.

In both cases p remains optimal under ρ_p, so $\psi_\delta \in \mathcal{C}_p(\varphi)$, but $\psi_\delta(p) > \widehat{\psi}^p_\varphi(p)$ contradicts the definition of $\widehat{\psi}^p_\varphi$. Thus $F_{\widehat{\psi}^p_\varphi}(\{p\}) = 0$.

Suppose $W^{\rho_p}(\cdot;\widehat{\psi}^p_\varphi,\varphi) < W^*$ at a buyer support point in $(p,1)$. The function $W^{\rho_p}(\cdot;\widehat{\psi}^p_\varphi,\varphi)$ is right-continuous on $(p,1)$, with no downward jumps away from buyer atoms. Choose two nearby $p < \underline{v} < \bar{v} < 1$ and a cutoff $x \in (\underline{v}, \bar{v})$ such that $[x, \bar{v})$ contains positive buyer mass and surplus is at most $W^* - \delta$ there for some $\delta > 0$. Since φ is positive above p, choose them close enough that $\bar{v} - \underline{v} < \varphi(\underline{v})$. For each value $v \in [x, \bar{v})$, perturb the buyer's distribution by replacing a fraction $0 < \varepsilon < \min\{1, \delta/2\}$ of its mass with a lottery over $\underline{v}, \bar{v}$ with mean v.

For prices t in $[\underline{v}, x)$, the mass moved to $\underline{v}$ reduces mass of buyers accepting the trade. This perturbation preserves the mean and raises the call by $t - \underline{v}$ per unit of the mass moved to $\underline{v}$, so the change in total surplus per unit is $(t - \underline{v})G_\varphi(t) - \varphi(t) < 0$. Indeed, $G_\varphi(t) \leq 1$ and $t - \underline{v} < \bar{v} - \underline{v} < \varphi(\underline{v}) \leq \varphi(t)$.

For prices t in $[x, \bar{v})$, the call and acceptance probability each change by at most ε, so surplus increases by at most $2\varepsilon < \delta$. Elsewhere, including at p, surplus is unchanged.

For the fixed slack $\delta > 0$, we can take $\varepsilon > 0$ arbitrarily small with $2\varepsilon < \delta$. No price then yields surplus above W^*, so the resulting call belongs to $\mathcal{C}_p(\varphi)$ and exceeds its pointwise supremum $\widehat{\psi}^p_\varphi$, a contradiction.

We have shown that $W^{\rho_p} = W^*$ at every buyer support point in $(p,1)$. On an interval

without buyer mass between support points or the endpoints $p, 1$, buyer acceptance is constant and newly accepting sellers add nonnegative gains. Total surplus is therefore nondecreasing from W^* at its left endpoint. Since p maximizes W^{ρ_p}, total surplus cannot exceed W^*. Thus $W^{\rho_p} = W^*$ throughout $[p, 1)$.

As in the proof of Claim 3, constant total surplus under ρ_p yields the ordinary differential equation (ODE)

$$\varphi'_+ \widehat{\psi}^p_\varphi - \varphi(\widehat{\psi}^p_\varphi)'_+ = W^*$$

The integrating factor $1/\varphi$ and boundary condition $\widehat{\psi}^p_\varphi(1) = 0$ yield $\widehat{\psi}^p_\varphi(t) = W^* H_\varphi(t)$ for $t > p$. When $\varphi(p) = 0$, $p = c_0$ and Equation 4 extends this solution to $\widehat{\psi}^p_\varphi(p) = W^* H_\varphi(p)$. As $F_{\widehat{\psi}^p_\varphi}$ has no atom at p, the left and right slopes of $\widehat{\psi}^p_\varphi$ at p agree. Hence $[W^* H_\varphi(p) - m_b]/p = W^*(H_\varphi)'_+(p)$. This, along with the ODE formulated above, implies Equation 7.

For $\widehat{\psi}^p_\varphi$ to be a feasible buyer call, its slope $-m_b/[p + \eta^*_\varphi(p)]$ below p must be at least -1, requiring $p + \eta^*_\varphi(p) \geq m_b$. Conversely, if $p + \eta^*_\varphi(p) \geq m_b$, the convex call in Equation 7 has slopes in $[-1, 0]$, value m_b at zero, and value zero at one. It is feasible and makes p maximize W^{ρ_p}, hence belongs to $\mathcal{C}_p(\varphi)$. □

C The buyer's guaranteed continuation payoff

Proof of Proposition 2. Fix $\delta > 0$. For every feasible $p < q < 1$, Equation 8 and the monotonicity of G_φ and $p + \eta^*_\varphi(p)$ imply

$$V(q; \varphi) = m_b G_\varphi(q) \left(1 - \frac{q}{q + \eta^*_\varphi(q)}\right) \geq m_b G_\varphi(p) \left(1 - \frac{q}{p + \eta^*_\varphi(p)}\right)$$

Using Equation 8 again we get

$$\begin{aligned} V(q; \varphi) &\geq V(p; \varphi) - \frac{m_b G_\varphi(p)}{p + \eta^*_\varphi(p)}(q - p) \\ &\geq V(p; \varphi) - (q - p). \end{aligned} \tag{20}$$

The second inequality follows from $p + \eta^*_\varphi(p) \geq m_b$ and $G_\varphi(p) \leq 1$. Choose q in $(p, 1)$ with $q - p < \delta/2$ and $G_\varphi(\{q\}) = 0$. Then $\varphi(q) \geq (q - p) G_\varphi(p) > 0$, and monotonicity of $p + \eta^*_\varphi(p)$ implies that q remains strictly feasible. Hence $V(q; \varphi) > V(p; \varphi) - \delta/2$.

Apply the perturbation in the proof of Claim 3 to the law of $\widehat{\psi}_\varphi^q$. Draw each value $v > q$ uniformly from $[v - 2\varepsilon(v-q), v]$ and restore the mean by moving mass $\varepsilon\widehat{\psi}_\varphi^q(q)/q$ from zero to q. Strict feasibility implies positive mass at zero, so this perturbation is feasible for sufficiently small $\varepsilon > 0$. Write F_ε for the resulting law. The estimate in the proof of Claim 3 gives

$$W^{(1,1)}(q; F_\varepsilon, G_\varphi) - W^{(1,1)}(t; F_\varepsilon, G_\varphi) \geq \frac{\varepsilon\widehat{\psi}_\varphi^q(q)}{q}\varphi(q) > 0 \qquad (t > q).$$

Thus, price q yields strictly more total surplus than any higher price. When ε is small enough we get $\varepsilon V(q;\varphi) < \delta/2$, so the perturbation reduces buyer payoff at q by less than $\delta/2$.

Since $F_\varepsilon((0,q)) = 0$, total surplus under inclusive participation at every positive price below q is at most $W^{(1,1)}(q; F_\varepsilon, G_\varphi)$. Deterministic prices approaching q from below and avoiding seller atoms yield total surplus converging to this value, independently of participation, because $G_\varphi(\{q\}) = 0$. An optimal price lottery thus attains total surplus arbitrarily close to $W^{(1,1)}(q; F_\varepsilon, G_\varphi)$ and excludes prices above q. Seller payoff at every selected price is at most $F_\varepsilon([q,1])\varphi(q)$. Subtracting gives

$$\begin{aligned} u_b^{\lambda,\rho}(F_\varepsilon, G_\varphi) &\geq W^{(1,1)}(q; F_\varepsilon, G_\varphi) - F_\varepsilon([q,1])\varphi(q) \\ &= (1-\varepsilon)V(q;\varphi) \\ &> V(p;\varphi) - \delta. \end{aligned}$$

□

D Seller experiments and payoff bounds

D.1 The convex majorant

Proof of Theorem 3. We first show that $\mathcal{D}_\gamma$ is closed in the compact set $\Phi(m_s)$ under the uniform norm. Let $\varphi_n \in \mathcal{D}_\gamma$ converge uniformly to $\varphi \in \Phi(m_s)$. At a strictly feasible price q with $\varphi(q) > 0$ and $G_\varphi(\{q\}) = 0$, as q is a point of continuity Equation 3 gives $G_{\varphi_n}(q) \to G_\varphi(q)$. Since $\varphi_n \geq \varphi(q)/2$ on $[q,1]$ when n is large enough, uniform convergence also gives $I_{\varphi_n}(q) \to I_\varphi(q)$. By Equation 6, $\eta^*_{\varphi_n}(q) \to \eta^*_\varphi(q)$, so q remains strictly feasible for φ_n when n is large enough. Equation 8 implies

$$V(q;\varphi) = \lim_{n\to\infty} V(q;\varphi_n) \leq \gamma.$$

For any strictly feasible p, choose such prices q decreasing to p and avoiding atoms of G_φ. These prices are strictly feasible because $p + \eta^*_\varphi(p)$ is nondecreasing, and Equation 20 gives $V(p;\varphi) \le \gamma + (q-p)$. Letting $q \downarrow p$ proves $\varphi \in \mathcal{D}_\gamma$. Thus $\mathcal{D}_\gamma$ is compact. Every sublevel set of Γ is therefore closed, which implies lower semicontinuity of Γ.

Now suppose $\mathcal{D}_\gamma$ is nonempty. For $\varphi_1, \varphi_2 \in \mathcal{D}_\gamma$, their pointwise maximum $\widetilde{\varphi}$ is feasible. Fix any strictly feasible p with $\widetilde{\varphi}(p) > 0$. Such prices exist, since $p > m_b$ implies $\widetilde{\varphi}(p) \ge p - m_s > 0$ and $p + \eta^*_{\widetilde{\varphi}}(p) > m_b$. Choose $i \in \{1,2\}$ with the larger value at p, breaking a tie by the larger right derivative. Then $\widetilde{\varphi}(p) = \varphi_i(p)$ and $(\widetilde{\varphi})'_+(p) = (\varphi_i)'_+(p)$. Since $\widetilde{\varphi} \ge \varphi_i$ we get,

$$I_{\widetilde{\varphi}}(p) = \int_p^1 \frac{dt}{\widetilde{\varphi}(t)^2} \le \int_p^1 \frac{dt}{\varphi_i(t)^2} = I_{\varphi_i}(p).$$

Equation 6 gives $p + \eta^*_{\varphi_i}(p) \ge p + \eta^*_{\widetilde{\varphi}}(p) > m_b$, so p is strictly feasible for φ_i. Equation 8 then implies

$$V(p;\widetilde{\varphi}) \le V(p;\varphi_i) \le \gamma.$$

If $\widetilde{\varphi}(p) = 0$ at a strictly feasible target, prices $q \downarrow p$ satisfy $\widetilde{\varphi}(q) \ge (q-p) G_{\widetilde{\varphi}}(p) > 0$ and remain strictly feasible. Hence Equation 20 gives $V(p;\widetilde{\varphi}) \le \gamma + (q-p)$, and taking the limit proves $\widetilde{\varphi} \in \mathcal{D}_\gamma$.

Compactness of $\mathcal{D}_\gamma$ and continuity of $\varphi \mapsto \int_0^1 \varphi(t)\,dt$ imply that some $\varphi^* \in \mathcal{D}_\gamma$ maximizes this integral. If $\widehat{\varphi} \in \mathcal{D}_\gamma$ exceeded φ^* at a point, continuity of both functions would imply

$$\int_0^1 \max\{\varphi^*(t), \widehat{\varphi}(t)\}\,dt > \int_0^1 \varphi^*(t)\,dt.$$

The preceding argument implies $\max\{\varphi^*, \widehat{\varphi}\} \in \mathcal{D}_\gamma$, contradicting maximality. Hence φ^* dominates every member and equals φ_γ from Equation 10. □

D.2 The seller majorant

Proof of Proposition 3. Let $[\underline{p}, \bar{p}]$ be the smallest closed interval containing $\operatorname{supp} G_{\varphi_\gamma} \cap (0,1)$. If G_{φ_γ} were supported on $\{0,1\}$, then $\eta^*_{\varphi_\gamma}(p) = 1 - p$ for every interior price. Each interior price would be strictly feasible with $V(p;\varphi_\gamma) = W^{\mathrm{FB}}(1-p)$ by Equation 8. Taking the supremum gives $\Gamma(\varphi_\gamma) = W^{\mathrm{FB}} > \gamma$, contradicting $\varphi_\gamma \in \mathcal{D}_\gamma$.

First suppose, for a contradiction, that G_{φ_γ} has an atom at a cost c in $\operatorname{supp} G_{\varphi_\gamma} \cap (0,1)$.

If $\varphi_\gamma(c) > 0$, take $0 < \varepsilon < \min\{c, 1-c\}$ and replace this atom by equal masses at $c-\varepsilon$ and $c+\varepsilon$, preserving its mean c. The resulting put $\widetilde{\varphi}$ exceeds φ_γ at c and agrees with it outside $(c-\varepsilon, c+\varepsilon)$. Below $c-\varepsilon$, the put is unchanged and $I_{\widetilde{\varphi}} \leq I_{\varphi_\gamma}$. Thus every price below $c-\varepsilon$ that is strictly feasible under $\widetilde{\varphi}$ is also strictly feasible under φ_γ, with no higher buyer payoff, by Equation 6 and Equation 8.

For $c+\varepsilon \leq p < 1$, implementability of price p and the buyer's auxiliary payoff $V(p;\cdot)$ are unchanged. For sufficiently small ε and every $p \in [c-\varepsilon, c+\varepsilon)$, $G_{\widetilde{\varphi}}(p) \leq G_{\varphi_\gamma}(c) - \frac{1}{4}G_{\varphi_\gamma}(\{c\})$, whereas $|\widetilde{\varphi}(p) - \varphi_\gamma(c)| \leq 2\varepsilon$ and $|I_{\widetilde{\varphi}}(p) - I_{\varphi_\gamma}(c)| \leq 8\varepsilon/\varphi_\gamma(c)^2$. For fixed p, $\varphi(p)$, and $I_\varphi(p)$, the expressions for $p+\eta^*_\varphi(p)$ and $V(p;\varphi)$ in Equation 6 and Equation 8 strictly increase with $G_\varphi(p)$. If $c+\eta^*_{\varphi_\gamma}(c) \leq m_b$, then for small enough ε every price p in $[c-\varepsilon, c+\varepsilon)$ satisfies $p+\eta^*_{\widetilde{\varphi}}(p) < m_b$ and is therefore infeasible under $\widetilde{\varphi}$. If instead c is strictly feasible under φ_γ, then $V(c;\varphi_\gamma) \leq \gamma$ because $\varphi_\gamma \in \mathcal{D}_\gamma$. In this case, a price in the interval may remain feasible after the split, but for small enough ε every such price gives $V(p;\widetilde{\varphi}) < V(c;\varphi_\gamma) \leq \gamma$. Together with the comparison outside the interval, this gives $\widetilde{\varphi} \in \mathcal{D}_\gamma$. Yet $\widetilde{\varphi}(c) > \varphi_\gamma(c)$, contradicting the definition of φ_γ as the pointwise supremum over $\mathcal{D}_\gamma$.

If $\varphi_\gamma(c) = 0$, then $c = \underline{p}$ and there is no seller mass below c. Split the atom at c between 0 and $c+\varepsilon$ with probabilities $\varepsilon/(c+\varepsilon)$ and $c/(c+\varepsilon)$. Choose ε small enough such that $\varepsilon/(c+\varepsilon) + G_{\varphi_\gamma}((c, c+\varepsilon)) < \gamma/m_b$. Both terms can be made arbitrarily small, since $(c, c+\varepsilon)$ excludes the atom at c. Every price $p < c+\varepsilon$ strictly feasible under $\widetilde{\varphi}$ therefore gives buyer payoff $V(p;\widetilde{\varphi}) \leq m_b G_{\widetilde{\varphi}}(p) < \gamma$. For $c+\varepsilon \leq p < 1$, implementability of price p and the buyer's auxiliary payoff $V(p;\cdot)$ are unchanged. Thus $\widetilde{\varphi} \in \mathcal{D}_\gamma$ dominates φ_γ and is strictly larger at c, a contradiction. Hence G_{φ_γ} is continuous on $(0,1)$ and, by definition of $\underline{p}$, $G_{\varphi_\gamma}(p) > G_{\varphi_\gamma}(\underline{p})$ for every $p > \underline{p}$.

To establish feasibility at $\underline{p}$, we first show that $\lim_{p \downarrow \underline{p}}[p+\eta^*_{\varphi_\gamma}(p)] \geq m_b$. Otherwise, choose c in $\operatorname{supp} G_{\varphi_\gamma} \cap (\underline{p}, 1)$ sufficiently near $\underline{p}$ that $c+\eta^*_{\varphi_\gamma}(c) < m_b$. As before, split costs in $(c-\varepsilon, c+\varepsilon)$ between its endpoints, preserving each cost's mean. Continuity of G_{φ_γ} and Equation 6 imply that, for small enough ε, $p+\eta^*_{\widetilde{\varphi}}(p) < m_b$ throughout $[c-\varepsilon, c+\varepsilon)$. The comparison outside this interval is as above, so $\widetilde{\varphi}$ dominates φ_γ, is strictly larger at c, and belongs to $\mathcal{D}_\gamma$, a contradiction. Equation 5 and Equation 6 imply that $p+\eta^*_{\varphi_\gamma}(p)$ strictly increases whenever G_{φ_γ} increases. Thus every price p in $(\underline{p}, 1)$ is strictly feasible.

We next show that $V(c;\varphi_\gamma) = \gamma$ for every $c \in \operatorname{supp} G_{\varphi_\gamma} \cap (\underline{p}, 1)$. If $V(c;\varphi_\gamma) < \gamma$, continuity and the same construction of $\widetilde{\varphi}$ give $V(p;\widetilde{\varphi}) < \gamma$ throughout $[c-\varepsilon, c+\varepsilon)$ for small enough

ε, yielding a contradiction as above. On a gap in the support, G_{φ_γ} and $p + \eta^*_{\varphi_\gamma}(p)$ are constant, so Equation 8 makes V strictly decrease, whereas its value at either endpoint of the gap is γ, a contradiction. Hence $\operatorname{supp} G_{\varphi_\gamma} \cap [\underline{p}, \bar{p}]$ has no gaps and $V(p; \varphi_\gamma) = \gamma$ throughout $(\underline{p}, \bar{p})$.

Since $\gamma = V(p; \varphi_\gamma) \leq \min\{m_b G_{\varphi_\gamma}(p), m_b(1-p)\}$ on $(\underline{p}, \bar{p})$, the limits at $\underline{p}$ and $\bar{p}$ give $G_{\varphi_\gamma}(0) > 0$ and $\bar{p} < 1$. If $\underline{p} = 0$, then $V(p; \varphi_\gamma) = \gamma$ and feasibility imply $G_{\varphi_\gamma}(0) = \gamma/m_b$, $H_{\varphi_\gamma}(0) = m_b/\gamma$, and $-\infty < H_{\varphi_\gamma}{}'_+(0) < 0$. But Equation 5 and Equation 8 give $H_{\varphi_\gamma}(p) - H_{\varphi_\gamma}(0) = [p + (m_b/\gamma)\varphi_\gamma(p)]H_{\varphi_\gamma}{}'(p)$. Dividing by p and letting $p \downarrow 0$ yields $H_{\varphi_\gamma}{}'_+(0) = 2H_{\varphi_\gamma}{}'_+(0)$ because $\varphi_\gamma(p)/p \to \gamma/m_b$, a contradiction. Thus $\underline{p} > 0$, and continuity extends $V(p; \varphi_\gamma) = \gamma$ to $[\underline{p}, \bar{p}]$.

Finally, note that $\underline{p}$ is binding feasible. There is no seller mass in $(0, \underline{p}]$, so $G_{\varphi_\gamma}(p) = G_{\varphi_\gamma}(0)$ and $p + \eta^*_{\varphi_\gamma}(p) = \underline{p} + \eta^*_{\varphi_\gamma}(\underline{p})$ for $0 < p \leq \underline{p}$. If $\underline{p}$ were strictly feasible, Equation 8 would therefore give $V(p; \varphi_\gamma) > V(\underline{p}; \varphi_\gamma) = \gamma$ at every $0 < p < \underline{p}$, contradicting $\varphi_\gamma \in \mathcal{D}_\gamma$. Thus $\underline{p} + \eta^*_{\varphi_\gamma}(\underline{p}) = m_b$, and the same equality holds throughout $(0, \underline{p}]$. Since the strictly feasible prices in $(\underline{p}, \bar{p}]$ give payoff γ, $\Gamma(\varphi_\gamma) = \gamma$. □

D.3 No atom at one

Proof of Lemma 1. Suppose $(\varphi)'_-(1) < 1$. Choose $q \in (1 - \beta/m_b, 1)$ and define, for $c < 1$,

$$\widetilde{G}(c) = (1-\varepsilon)G_\varphi(c) + \varepsilon \frac{1-m_s}{1-q}\mathbf{1}_{\{c \geq q\}} \quad \text{and} \quad \widetilde{G}(1) = 1.$$

This moves an ε fraction of the mass below one to one, then moves mass $\varepsilon(1-m_s)/(1-q)$ from one to q, preserving the mean. The atom at one makes this feasible for small $\varepsilon > 0$. The corresponding put satisfies

$$\varphi_{\widetilde{G}}(p) = (1-\varepsilon)\varphi(p) + \varepsilon \frac{1-m_s}{1-q}(p-q)^+.$$

Recall that $\Gamma(\varphi_{\widetilde{G}})$ is the supremum of the buyer payoffs $V(p; \varphi_{\widetilde{G}})$ over strictly feasible prices. By Equation 20, it suffices to bound them where the put is positive. Below q, the put and CDF are multiplied by $1-\varepsilon$, while $I_{\varphi_{\widetilde{G}}}(p) \leq I_\varphi(p)/(1-\varepsilon)^2$. Hence Equation 6 gives $\eta^*_{\varphi_{\widetilde{G}}}(p) \leq \eta^*_\varphi(p)$. Every such price below q that is strictly feasible for $\varphi_{\widetilde{G}}$ is therefore strictly feasible for φ, and Equation 8 gives the upper bound $V(p; \varphi_{\widetilde{G}}) \leq (1-\varepsilon)\beta$. For $q \leq p < 1$, including the atom at q, Equation 8 gives $V(p; \varphi_{\widetilde{G}}) \leq m_b(1-p) \leq m_b(1-q)$.

Consequently,

$$\Gamma(\varphi_{\widetilde{G}}) \leq \max\{(1-\varepsilon)\beta, m_b(1-q)\} < \beta,$$

contradicting the definition of β. □

E Sequential acquisition equilibrium

For any seller put φ, recall that $c_0 = \inf\{s : \varphi(s) > 0\}$. For a feasible buyer call ψ, integrating $(\psi/\varphi)' = -W^{(0,1)}/\varphi^2$ from $t > c_0$ to 1 as in the proof of Claim 3 and using $H_\varphi = \varphi I_\varphi$ from Equation 5 gives

$$\psi(t) \leq H_\varphi(t) \max_{s\in[0,1]} W^{(1,1)}(s;\psi,\varphi), \qquad t > c_0. \tag{21}$$

Proof of Theorem 4. With seller law G_{φ_β} and full buyer information, price $\bar{p}$ maximizes total surplus and gives buyer payoff $m_b(1-\bar{p}) = \beta$ by Proposition 3 and Lemma 1. Neither law has an atom at $\bar{p}$, so the participation rule does not affect these payoffs.

After a buyer deviation, the designer chooses the highest maximizer p of $W^{(1,1)}$. Under G_{φ_β}, only zero-cost sellers trade at prices at or below $\underline{p}$, so total surplus at any price $p \leq \underline{p}$ is at most $m_b G_{\varphi_\beta}(0)$. If all buyer values are at least $\underline{p}$, the same buyers and sellers trade at every such price under inclusive participation, so $\underline{p}$ yields as much surplus as any lower price. If the buyer assigns positive mass below $\underline{p}$, then $\psi(\underline{p}) > m_b - \underline{p} = m_b G_{\varphi_\beta}(0) H_{\varphi_\beta}(\underline{p})$, where the equality follows from binding feasibility of $\underline{p}$. By Equation 21, maximum surplus then strictly exceeds $m_b G_{\varphi_\beta}(0)$. In either case, the highest surplus-maximizing price is at least $\underline{p}$. Both traders accept at indifference after the selected announcement, with the buyer rejecting any higher prices when indifferent and the seller accepting any higher prices when indifferent after any other choice of experiments. The buyer majorant in Theorem 2 and Corollary 4 bounds his payoff by $V(p;\varphi_\beta) \leq \beta$ for $p \geq \underline{p}$, and his payoff is zero at one. Full information is therefore a best response.

To complete the equilibrium construction, after seller deviations, choose, for each seller law, the buyer experiment and optimal continuation jointly giving the highest buyer payoff. Compactness of the graph of $\mathcal{U}$ ensures attainment. The construction following Equation 18 implements the maximizing payoff pair and specifies an optimal continuation after each alternative buyer experiment. The buyer then has a best response, since no alternative experiment can yield more under any optimal continuation. The seller attains her payoff

upper bound $W^{\mathrm{FB}} - \beta$ from choosing the experiment corresponding to φ_β.

For payoff uniqueness, consider any pure seller-first equilibrium. Fix $\underline{p} < r < \bar{p}$ and small $0 < \varepsilon < r - \underline{p}$. We construct a seller deviation after which subgame perfection requires the buyer to fully learn and yields the price r, allowing the seller to approach $W^{\mathrm{FB}} - \beta$ as $r \uparrow \bar{p}$. Refine G_{φ_β} by splitting each cost $c > r$ between a uniform draw s from $[r - \varepsilon, r]$ and one, assigning probability $(1 - c)/(1 - s)$ to s to preserve its mean. Leave lower costs unchanged. The resulting law G_ε has no interior atoms, and its put φ_ε agrees with φ_β below $r - \varepsilon$ and strictly exceeds it on $(r, 1)$. Thus $H_{\varphi_\varepsilon}(\underline{p}) < H_{\varphi_\beta}(\underline{p})$, so Equation 21 implies

$$\max_{t \in [0,1]} W^{(1,1)}(t; F, G_\varepsilon) \geq \frac{m_b - \underline{p}}{H_{\varphi_\varepsilon}(\underline{p})} > m_b G_{\varphi_\beta}(0)$$

for every buyer law F. Prices at or below $\underline{p}$ yield at most $m_b G_{\varphi_\beta}(0)$, so Equation 19 excludes them under every optimal continuation.

For $\underline{p} < p < r - \varepsilon$, the put and CDF are unchanged, while $I_{\varphi_\varepsilon}(p) \leq I_{\varphi_\beta}(p)$. Equation 6 and Equation 8 therefore bound buyer payoff at an implemented price by $V(p; \varphi_\varepsilon) \leq \beta$.

For $p \geq r - \varepsilon$, buyer payoff is at most $m_b G_\varepsilon(p)(1 - p)$, uniquely maximized at r for small ε.

For $r - \varepsilon \leq s < t \leq r$, integration by parts, the density added by the refinement, and $G_\varepsilon \leq 1$ give

$$\begin{aligned}(1 - t)G_\varepsilon(t) - (1 - s)G_\varepsilon(s) &\geq (t - s)\left[\frac{1}{\varepsilon}\int_{(r,1]}(1 - c)\, dG_{\varphi_\beta}(c) - 1\right] \\ &> 0\end{aligned}$$

for small ε, since $r < \bar{p}$ makes the integral positive.

On $(r, 1)$, G_ε is constant, and $G_\varepsilon(p)(1 - p)$ strictly decreases. Moreover,

$$m_b G_\varepsilon(r)(1 - r) \longrightarrow m_b[1 - m_s - \varphi_\beta(r)] > \beta \qquad (\varepsilon \downarrow 0).$$

Thus the buyer's payoff is bounded above by $m_b G_\varepsilon(r)(1 - r)$, with equality requiring selection of r and full information, since $(v - r)^+ \leq v(1 - r)$ is strict for $0 < v < 1$.

Under full information, any price $p > r$ also maximizes $W^{(1,1)}$ but gives buyer payoff $m_b G_\varepsilon(r)(1 - p) < m_b G_\varepsilon(r)(1 - r)$. The buyer can obtain payoff arbitrarily close to $m_b G_\varepsilon(r)(1 - r)$ under every optimal continuation by choosing the experiments on $\{0, r, 1\}$

in Section 6.4. These make r the unique surplus-maximizing price, and the buyer's payoff approaches this value as his mass at r vanishes. Subgame perfection requires the buyer's response after G_ε to attain this bound, hence to use full information and induce price r. The seller receives $m_b\varphi_\varepsilon(r) \geq m_b\varphi_\beta(r)$, and $m_b\varphi_\beta(r) \to W^{\mathrm{FB}} - \beta$ as $r \uparrow \bar{p}$. Any lower seller equilibrium payoff therefore admits a profitable deviation. The buyer's guarantee and the total-surplus bound rule out a higher seller payoff and force buyer payoff β. □

F Reflection symmetry

Proof of Proposition 4. For any signal z, $\mathbb{E}[1-\theta \mid z] = 1 - \mathbb{E}[\theta \mid z]$, so complementing the fundamental reflects its posterior mean. Reflection also preserves independence. Write $\mathsf{R}\lambda$ for the pushforward of λ under $p \mapsto 1-p$, and exchange the traders' participation decisions after the corresponding reflected experiments, announcement, and realized price. This defines $\mathsf{R}\rho$ after every history, including designer deviations. Under $(v, c, p) \mapsto (1-c, 1-v, 1-p)$ we obtain Equation 11 after every experiment announcement.

For a seller-first profile $(\varphi^*, \sigma_b, \lambda, \rho)$, let $\widetilde{\psi}$ be any buyer call in $\Psi(1-m_s)$ for the reflected game. Its reflection is an original-game seller put, so define $\widetilde{\sigma}_s(\widetilde{\psi}) = \mathsf{R}[\sigma_b(\mathsf{R}\widetilde{\psi})] \in \Phi(1-m_b)$. Reflect pricing after every experiment pair to obtain $\widetilde{\lambda}$. Equation 11 exchanges trader payoffs and preserves total surplus, hence preserves acquisition incentives and the designer's best responses. The rule $\mathsf{R}\rho$ preserves weakly dominant participation after every announcement. Thus, if the original profile is a subgame-perfect equilibrium of the seller-first game with priors (m_s, m_b), then $(\mathsf{R}\varphi^*, \widetilde{\sigma}_s, \widetilde{\lambda}, \mathsf{R}\rho)$ is a subgame-perfect equilibrium of the buyer-first game with priors $(1-m_b, 1-m_s)$. Reflection is its own inverse, proving the correspondence. □

G Timing of acquisition

Proof of Theorem 5. By Theorem 4 and Proposition 4, it suffices to bound the sum of follower payoffs below W^{FB}. Set $\alpha = \sqrt{1-m_b}/(\sqrt{m_s} + \sqrt{1-m_b})$. Then $\alpha\sqrt{m_s} = (1-\alpha)\sqrt{1-m_b}$, so $\alpha^2 m_s = (1-m_b)(1-\alpha)^2$. For $s_0 = \alpha^2(1-m_s)$ and $s_1 = 1 - m_b + m_b s_0$, we have $s_1 - s_0 = (1-m_b)(1-s_0) > 0$ and $1 - s_1 = m_b(1-s_0) > 0$. Thus $0 < s_0 < s_1 < 1$. Define $\gamma > 0$ by

$$\frac{W^{\mathrm{FB}}}{\gamma} = \frac{\log(s_1/s_0)}{1-s_0} + \frac{m_b}{\sqrt{s_1}}.$$

Since $\sqrt{s_1} = 1 - m_b + \alpha m_b$ and $2\alpha\sqrt{s_1} - s_0 - s_1 = \alpha^2 m_s m_b > 0$, the inequality $\log t > 2(t-1)/(t+1)$ for $t > 1$ gives $\gamma < \alpha W^{\mathrm{FB}}$. For $s \in [s_0, s_1]$, define $J(s) = W^{\mathrm{FB}}/\gamma - \log(s/s_0)/(1-s_0)$ and

$$p(s) = \frac{1-s}{1-s_0} - sJ(s), \qquad \varphi(p(s)) = \frac{\gamma}{m_b}\left[J(s) - \frac{1-s}{1-s_0}\right].$$

Here $p'(s) = -J(s) < 0$ since $J(s) \geq J(s_1) = m_b/\sqrt{s_1} > 0$, so price decreases as s increases. The endpoints are $\underline{p} = p(s_1) = m_b(1-\sqrt{s_1}) > 0$ and $\bar{p} = p(s_0) = 1 - s_0 W^{\mathrm{FB}}/\gamma < 1$. Differentiation gives an increasing, strictly convex put. Extend it linearly outside $[\underline{p}, \bar{p}]$, matching its endpoint values and slopes. The identities $\varphi(\underline{p}) = \underline{p}\varphi'(\underline{p})$ and $\varphi(\bar{p}) + (1-\bar{p})\varphi'(\bar{p}) = 1 - m_s$, with $\varphi'(\bar{p}) = (\gamma/(\alpha W^{\mathrm{FB}}))^2 < 1$, verify feasibility.

In Equation 5, substitute $\varphi(p) = \varphi'(\underline{p})p$ below $\underline{p}$ and $\varphi(p) = 1 - m_s - \varphi'(\bar{p})(1-p)$ above $\bar{p}$. On the interior, $I_\varphi(p(s))$ and $m_b s/(\gamma\varphi(p(s)))$ both have derivative $J(s)/\varphi(p(s))^2$ and agree at s_0. Hence $H_\varphi(p) = \sqrt{s_1}(m_b - p)/\gamma$ for $0 < p \leq \underline{p}$, $H_\varphi(p(s)) = m_b s/\gamma$ for $s \in [s_0, s_1]$, and $H_\varphi(p) = (1-p)/(1-m_s)$ for $\bar{p} \leq p < 1$. Equation 6 and Equation 8 give $\eta^*_\varphi(p(s)) = sJ(s)$, $p(s) + \eta^*_\varphi(p(s)) = (1-s)/(1-s_0)$, and $V(p(s);\varphi) = \gamma$. For $0 < p \leq \underline{p}$, $\eta^*_\varphi(p) = m_b - p$. Above $\bar{p}$, $\eta^*_\varphi(p) = 1 - p$ and $V(p;\varphi) = m_b\varphi'(\bar{p})(1-p)$ decreases from γ. Thus strictly feasible prices are exactly $(\underline{p}, 1)$ and $\beta(m_s, m_b) \leq \Gamma(\varphi) = \gamma < \alpha W^{\mathrm{FB}}$. Reflection replaces α by $1-\alpha$ and leaves W^{FB} unchanged, giving $\beta(1-m_b, 1-m_s) < (1-\alpha)W^{\mathrm{FB}}$. Adding the strict bounds yields

$$\beta(m_s, m_b) + \beta(1-m_b, 1-m_s) < \alpha W^{\mathrm{FB}} + (1-\alpha)W^{\mathrm{FB}} = W^{\mathrm{FB}}. \qquad \square$$