The Exact Welfare Guarantee of Fixed-Price Bilateral Trade

Tingyi Lin
Tropic Research Institute and UIUC
tingyi3@illinois.edu

Yichen Shi
Zhongnan University of Economics and Law
shiyichen@stu.zuel.edu.cn

Ke Dong
Tsinghua University
dongk.23@pbcsf.tsinghua.edu.cn

Jiazhuo Li
University of Michigan, Ann Arbor
jiazhuo@umich.edu

Shawn Yu
Boston College
yulin@bc.edu

Huanxi Zhang
University of Wisconsin–Madison
hzhang2488@wisc.edu

Abstract

A seller and a buyer with independent private values can trade only at a posted price. We determine the worst case of this mechanism exactly: the best posted price always guarantees a $\beta_* = 0.738024\ldots$ fraction of first-best welfare, where β_* is given in closed form by the root of an explicit equation, the worst case is a limit of distributions that no pair attains, and its shape is unique up to scaling. This closes the gap $[0.7292, 0.73805]$ left by a line of work from SODA 2016 through two STOC 2023 papers and an AAAI 2026 follow-up. The proof is an explicit certificate of optimality: after one change of variables, the gap between the optimal value and the value of any buyer distribution is a sum of manifestly nonnegative integrals, as in a linear-programming dual, and equality identifies the worst-case shape.

The certificate also gives the complete tradeoff between gains from trade and the seller's initial welfare: when first-best gains from trade are a fraction κ of initial seller welfare, the exact worst-case fraction $\rho(\kappa)$ of gains obtained by the best price satisfies $\rho(\kappa) \sim 2/\log(1/\kappa)$ as $\kappa \to 0$. The worst-case buyer's survival function has two constant segments joined by an explicit nonexponential curve, and the worst case is approached through a vanishing atom whose value tends to infinity. The same constant is the exact guarantee of dominant-strategy mechanisms with individual rationality and strong budget balance in every realization.

For two units with increasing submodular valuations, an explicit finite instance has ratio below 0.7290804, so multi-unit trade is strictly harder than single-unit trade. Fixing the buyer, we characterize the least probability that a random common price needs to guarantee a given ratio against every seller; bounding this quantity over all buyers would determine the exact two-unit constant. Within an explicit family the worst ratio is $0.729080\ldots$, which we conjecture to be the two-unit constant.

Contents

1 Introduction **3**

2 Technical overview **6**

3 Model and the scalar curve **10**

4 The complete gains-from-trade frontier **11**

5 From controls to a local energy **12**

6 A reference curve at every endpoint **14**

7 Fixed-endpoint comparison and the endpoint envelope **16**

8 Global calibration and the extremal shape **19**

9 Prices and bounded hard pairs **20**

10 The exact welfare guarantee **22**

11 Proof of the gains-from-trade frontier **24**

12 The dominant-strategy welfare ceiling **25**

13 Extension to two units **26**

14 Related work **31**

15 Open questions **32**

AI methodology **33**

References **33**

A Notation **36**

B Omitted proofs **37**
B.1 Endpoint asymptotics 37
B.2 Proof of the price representation 37
B.3 Proof of the two-unit tail normalization 39
B.4 Proof of the two-unit pricing theorem 41
B.5 Finite-node evaluation 43

C Auxiliary results **43**
C.1 Quantitative finite approximation 43
C.2 Nonconcavity in the original control 44

D Optimization within a two-unit family **44**

E The two-unit instance **53**

F Routes we tried **53**

1 Introduction

A seller owns an item that a buyer may want. Myerson and Satterthwaite [MS83] showed that no mechanism can be at once efficient, incentive compatible, voluntary, and budget balanced. The simplest response is a fixed price: post one price and let the two parties take it or leave it. Welfare is the value of whoever ends up holding the item, and first-best welfare gives the item to whoever values it more. How much welfare does this simplicity cost, in the worst case over all pairs of independent value distributions? This paper gives the exact answer: the best price always keeps a fraction $\beta_* = 0.738024\ldots$ of first-best welfare. It also describes the distributions that come arbitrarily close to forcing this fraction. Two things are needed: for every pair, a price that secures this fraction; and pairs on which no price does noticeably better.

No pair of distributions attains β_*. It is approached by pairs in which the buyer's value has an explicit density on an interval, no mass below that interval, and a vanishing atom far above it, while the seller's value has an atom at zero, an explicit distribution on the same interval, and an atom above it (Figure 2b,c). The shape carries more than the number: the same family of pairs, indexed by one parameter, gives every sharp guarantee on gains from trade and the exact rate at which such guarantees fail as gains become small relative to the seller's initial welfare.

The upper bound came first, and it came close. Colini-Baldeschi, de Keijzer, Leonardi, and Turchetta [CBdKLT16] gave a pair of distributions on which no fixed price exceeds 0.7485, and Kang and Vondrák [KV19] lowered this to 0.7385. Cai and Wu [CW23] and Liu, Ren, and Wang [LRW23] developed complementary optimization approaches at STOC 2023, each with hard instances at 0.7381. Giambartolomei and de Keijzer [GdK26] subsequently narrowed the welfare ratio to $[0.7292, 0.73805]$.

We close the gap by proving a matching lower bound for every pair of distributions and identifying the worst-case pair exactly. Numerical programs bound the constant but cannot decide whether it is attained, what the worst case looks like, or how the guarantee changes with the instance; those are the questions answered here. The exact value $0.738024\ldots$ raises the previous lower guarantee by about 0.009; the reported upper bound was already within 3×10^{-5} of this value. Our results were obtained concurrently with and independently of the work of Tao Jiang, Minbo Gao, and Shaowei Cai [JGC26], who report the same exact single-unit welfare ratio; they do not study gains-from-trade guarantees, the dominant-strategy welfare ceiling, or multiple units. They reach the same variational problem, in other coordinates, by saturating a constraint on the seller, and prove optimality by showing that a maximizer exists and solving its first-order conditions; our certificate compares every competitor directly and assumes no maximizer. Their saturated seller is the worst-case seller of [LRW23, Theorem 3.5], and their threshold equivalence [JGC26, Proposition 1.2] corresponds to the criterion of [LRW23, Theorem 3.3]; they cite [LRW23] only for its numerical bounds. Table 1 summarizes what was known and what is new.

Our two main results are the exact welfare guarantee in Theorem A and the complete gains-from-trade frontier in Theorem B. The global comparison in Theorem C underlies both results.

Theorem A (Exact welfare guarantee, informal)**.** *Let the seller's and the buyer's values be independent and nonnegative, with positive and finite first-best welfare. The best posted price always achieves at least a $\beta_* = 0.738024\ldots$ fraction of first-best welfare, and no larger constant is guaranteed. Here β_* is given in closed form by the unique root of an explicit equation in one variable. Every pair of distributions achieves strictly more than β_*, and an explicit family of bounded pairs approaches it.*

The equation is derived in Section 2 and stated in full in Theorem A; the family is in Lemma 9.2.

In words: a posted price never loses more than 26.2 percent of the available welfare. The adversary who comes closest to this loss uses two ingredients. The buyer is ordinary with probability

Question	Previously known	This paper
Worst-case welfare ratio	The interval $[0.7292, 0.73805]$: hard instances at 0.7485, 0.7385, 0.7381, and 0.73805 [CBdKLT16, KV19, CW23, LRW23, GdK26], and lower bounds from finite programs [CW23, GdK26].	The exact value $\beta_* = 0.738024\ldots$, the root of an explicit equation (Theorem A).
Worst-case pair	Necessary conditions on an unknown interval; an exponential trial family for upper bounds [LRW23].	A unique explicit shape with a nonexponential middle arc, never attained and approached through an escaping atom (Theorem C, Proposition 8.1).
Method	Finite programs [CW23, GdK26]; first-order conditions with an explicit trial family [LRW23].	A global optimality certificate for a nonconcave control problem: a concavifying change of variables (the reciprocal clock) and a calibration (a sufficiency certificate from the calculus of variations) with obstacle multipliers (Lemmas 5.1 and 7.1, Theorem C).
Gains-from-trade guarantees	No constant fraction [BD21]; logarithmic bounds in the efficient-trade probability [CBGdK$^+$17].	The exact affine frontier and the exact conditional guarantee $\rho(\kappa)$ at each ratio $\kappa = G/M$, with $\rho(\kappa) \sim 2/\log(1/\kappa)$ (Theorem B, Corollary B$'$).
Dominant-strategy mechanisms	The randomized-price representation [HR87, ČP16].	The exact welfare ceiling β_* under realizationwise individual rationality and strong budget balance (Corollary A$'$).
Two units	A numerically computed upper bound of 0.7291 and a separation conjecture [GdK26].	A finite instance below 0.7290804 and the separation theorem; a one-sided reduction of the exact constant and a conjectured value $0.729080\ldots$, proved when the second buyer's body is constant (Theorems D, E, and 13.4, Conjecture 13.3).

Table 1: What was known and what is new.

close to one and a lottery ticket with the remaining sliver of probability; the seller's costs spread over the ordinary buyer's range. A low price serves the ordinary buyer but shuts out every seller whose cost exceeds it; a high price lets every seller trade with the lottery ticket but loses the ordinary buyer. No single price serves both, and β_* is the exact price of that dilemma.

Theorem B (Sharp gains-from-trade guarantees, informal)**.** *Let the values be independent and nonnegative with finite first-best welfare, and write M for the seller's initial welfare and G for first-best gains from trade (GFT). For each $\beta \in (0,1)$ there is a smallest penalty $\delta(\beta)$ such that some price always obtains GFT at least $\beta G - \delta(\beta)M$, and the curve δ is explicit and strictly convex. At each prescribed ratio $\kappa = G/M > 0$, this curve determines the exact fraction $\rho(\kappa)$ of G guaranteed by the best fixed price, with*

$$\rho(\kappa) \sim \frac{2}{\log(1/\kappa)} \qquad (\kappa \downarrow 0).$$

Bounded independent hard pairs approach the guarantee at every prescribed κ.

In words: when the potential gains from trade equal what the seller already has, the best price secures 54 percent of them; when they are a tenth of it, 36 percent; a hundredth, 26 percent. The welfare worst case of Theorem A sits at a ratio near 2.6, where the best price secures 64 percent of the gains. The logarithmic rate is reached only slowly; at a ratio of one hundredth, $2/\log(1/\kappa)$ is still 0.43.

The complete gains-from-trade frontier. Welfare counts the seller's initial value; gains from trade do not. A posted price guarantees a constant fraction of welfare but no constant fraction of

gains from trade [BD21], and the second main result says exactly how the guarantee on gains from trade decays. The curve δ of Theorem B meets $\delta = 1 - \beta$ exactly at the welfare constant, because a price with welfare at least $\beta(M + G)$ is exactly a price with GFT at least $\beta G - (1 - \beta)M$. For a fixed ratio $\kappa = G/M > 0$, let $\rho(\kappa)$ denote the exact worst-case fraction of GFT obtained by the best price. Theorem B gives

$$\kappa\rho(\kappa) = \sup_{0<\beta<1} \{\kappa\beta - \delta(\beta)\}.$$

Thus the known failure of a constant GFT approximation has an exact logarithmic scale when parameterized by G/M; the asymptotic in Theorem B is proved in Corollary B′. The coefficient 2 is sharp: for each prescribed G/M, bounded hard pairs approach the corresponding conditional guarantee. Both results are read off one explicit curve, indexed by a single parameter, which Section 3 computes in closed form.

Why is this hard? The worst case is the value of an optimization problem over buyer distributions that is not concave, so first-order conditions cannot certify a candidate, and its supremum is not attained. The limiting buyer atom has positive contribution to expected value even as its probability vanishes. This loss of tail mass under weak convergence prevents a genuine distribution pair from attaining the worst-case welfare ratio. Identifying the value exactly requires controlling this limiting tail and proving a matching bound for every distribution. The upper bound of [LRW23, Theorem 3.8] comes from a buyer whose probability of valuing the item above s is of the form $\lambda_0 + \lambda_1 e^{\lambda_2 s}$ between two constant segments. The exact worst-case buyer also has two constant segments, but its middle segment is not of this form on any subinterval (Proposition 8.1), so no member of that family reaches the exact constant. Theorem C supplies the global comparison and identifies the unique limiting shape.

Cai and Wu [CW23, Lemmas 3.1 and 3.4] characterize the exact ratio as the value of an infinite-dimensional min-max program and bound it through two discretizations that converge to it; Theorem A determines that value analytically. Liu, Ren, and Wang supply three of our ingredients: a price density that depends on the buyer, the criterion that reduces the worst case to maximizing a functional of the buyer, and a seller distribution that makes every price equally good [LRW23, Theorems 3.1, 3.3, and 3.5]. Their Theorem 3.4 gives the first-order equation that a maximizing buyer satisfies between the point where its survival function leaves one and the point where it reaches zero. We locate these two points, integrate the equation between them, and prove the global comparison and the uniqueness of the limiting shape.

Here is the key idea. In a linear program, a dual solution certifies that a candidate is optimal by writing the gap between the candidate and any feasible point as a sum of nonnegative terms. Our problem is not a finite-dimensional program and is not concave (Section 2 gives a two-point example), so the first-order conditions of [LRW23] are necessary but not sufficient, and no trial family can be certified from them. Our proof imports the sufficiency tool of the calculus of variations, calibration [ABD03, BF16], and obtains a certificate of this kind. One change of variables, which we call the reciprocal clock, measures the buyer by the reciprocal of its tail integral and measures time by how much of the second-moment penalty is still to come; in these coordinates the problem is concave once a single boundary value, the endpoint, is fixed. For every endpoint we write down a reference curve and a multiplier that is positive only where the constraints bind. Any competitor with the same endpoint is worse than the reference curve by a sum of three nonnegative integrals, and the value of the reference curve is an explicit function of the endpoint with a unique maximum. Equality forces the competitor to coincide with the reference curve at the best endpoint, which

is the worst-case shape. The certificate assumes neither existence nor regularity of a maximizer. We need this generality: no genuine distribution pair attains the worst-case ratio, and the relaxed control problem has a maximizer only as a limit of distributions.

Two consequences place the result in the mechanism-design literature. By Cai and Wu's Theorem 4.1 [CW23], β_* is also the exact universal ratio when only the buyer's complete prior or only the seller's complete prior is known and randomized prices are allowed. Among all mechanisms for risk-neutral traders that are truthful in dominant strategies, balance the budget in every realization, and never make a trader worse off in any realization,[1] none guarantees more than β_* (Corollary A′). Up to null sets every such mechanism is a random posted price [HR87, ČP16], so on smooth priors it does no better than the best fixed price, and smoothing carries this bound to the hard pairs. Simplicity costs nothing here: among mechanisms that are robust in this sense, a single posted price is worst-case optimal. It is not a ceiling for Bayesian incentive-compatible trade: Dobzinski and Shaulker [DS26] obtain 0.746 with a mechanism that is DSIC for the seller and Bayesian IC for the buyer.

Extension to two units. When the seller owns two units and both traders have decreasing marginal values, one posted price must serve both marginal trades. We give an explicit finite instance on which no posted price reaches ratio 0.7290804, so two units are strictly harder than one, and an instance below 0.83693 in which the buyer's and the seller's valuations are identically distributed, improving the bound 0.8372 of [GdK26] for that case (Theorem D). With one unit and identically distributed values the exact ratio is $(2+\sqrt{2})/4 \approx 0.854$ [KPV22], so the second unit is costly in that case too. The first instance, evaluated in exact rational arithmetic, settles the separation conjecture of [GdK26]. We then reduce the exact two-unit constant to a question about buyers alone: fixing the buyer, Theorem E characterizes how much probability a random price needs to guarantee a target ratio against every seller, and a bound on this quantity over all buyers would determine the constant. Section 2 explains the construction. Within an explicit family of instances the worst ratio is $0.729080\ldots$, the unique solution of an explicit system of equations, with rigorous bounds (Proposition 13.2). We conjecture that it is the two-unit constant (Conjecture 13.3); what remains is the bound (66) on this quantity over all buyers. A second unit therefore costs the designer something: the worst-case guarantee drops by at least 0.0089, and if the conjecture holds, by exactly $\beta_* - 0.729080\ldots$. We prove the conjecture when the buyer's value for the second unit takes a single value apart from a vanishing high tail, a class that contains the family's minimizer (Theorem 13.4).

Formal verification. Theorems A–E, Corollaries A′ and B′, and Proposition 13.2 have machine-checked proofs in Lean 4 with Mathlib, with the decimal value of β_* checked by interval arithmetic; the formal proofs of Theorem E and Proposition 13.2 assume scalar facts certified by exact, interval, and SMT arithmetic, and the supplementary material lists what is formalized and what is not.

2 Technical overview

Figure 1 maps the single-unit chain.

From a price distribution to a functional of the buyer. Fix the buyer's distribution and a target ratio β. A random price certifies the guarantee β against every seller distribution if, for each seller value s, the expected gains from trade conditional on s are at least β times the conditional

[1] That is, ex post: for every realized pair of values, not only in expectation.

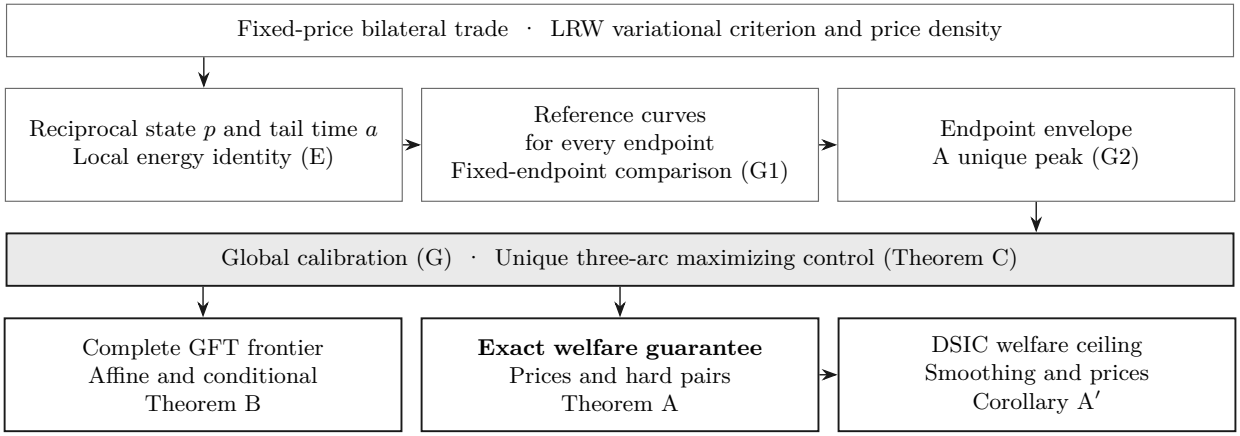


Figure 1: The single-unit proof of the guarantee runs along the central chain. The reference curve at the maximizing endpoint and bounded hard pairs give sharpness; varying the parameter gives the GFT frontier. The DSIC consequence additionally uses the price representation and smoothing.

first-best gains minus βd times s, for a constant $d > 0$. Averaging over the seller gives GFT at least $\beta G - \beta d \cdot M$, which is a welfare guarantee of β when $\beta d \le 1 - \beta$. Liu, Ren, and Wang [LRW23, Theorem 3.1] give a price density that meets this target with equality, and its total mass is β times a functional of the buyer alone. The density is a random price when this mass is at most one, with the remaining probability placed at price zero; because the target is met in expectation over this random price, some fixed price meets it. We rescale values so that the target vanishes at the cutoff $s = 1$; above the cutoff the target is negative and holds trivially, and at the cutoff the buyer's tail integral equals d. Writing h for the buyer's survival function in reversed and normalized time, g for its running integral started at d, and k for the running integral of h^2, that functional is

$$\mathcal{J}_d(h) = \int_0^1 \left[\frac{h}{g} + \frac{d^2 - k}{g^2} \right] \mathrm{d}t.$$

The guarantee holds whenever $\beta \mathcal{J}_d(h) \le 1$ for every buyer, and the equalizing seller of [LRW23] shows that nothing weaker suffices. The worst-case buyer therefore maximizes $\mathcal{J}_d$; we allow all measurable h with values in $[0, 1]$, a relaxation, since survival functions give only nondecreasing h. The seller has been eliminated, because the price density is the certificate against every seller (Lemma 9.1). Section 3 records the normalization.

The value curve. Everything then rests on the value $V(d) = \sup_h \mathcal{J}_d(h)$: the guarantee of GFT at least $\beta G - \beta d \cdot M$ holds for all pairs exactly when $\beta V(d) \le 1$. Theorem C computes this value along an explicit curve. For each $d > 0$ there is a unique $C_d \in (1/4, 1/2)$ with $\tau(C_d) = d$, and $V(d) = S(C_d)$, where τ and S are the closed forms of Section 3. The frontier of Theorem B is therefore the curve $(1/S(C), \tau(C)/S(C))$, and the welfare constant is its point on the line $\delta = 1 - \beta$, where $S(C_*) = 1 + \tau(C_*)$; numerically $C_* \approx 0.394$ and $\tau(C_*) \approx 0.355$.

The maximization is not concave. At $d = 1/2$, the controls equal to $(4/5, 3/5)$ and to $(1/5, 2/5)$ on the two halves of $[0, 1]$ have values $203/216$ and $103/96$. Their average is the constant control $1/2$, whose value is 1, yet the average of the two values exceeds 1 by $11/1728$; Appendix C.2 gives the general two-step formula. No first-order or Jensen argument can therefore certify a maximizer. This is the obstacle; everything below is built to get around it.

The reciprocal clock. Each contribution $h(u)^2$ to the accumulated second moment k is paid at every later time through the denominator g^2. Collecting these later payments into a single weight defines a new clock, the tail time $a = \int_t^1 g^{-2}$, under which the objective becomes local. We take the reciprocal $p = 1/g$ of the tail integral as the new state. In the new clock the slope of p is exactly h, so $h \le 1$ gives the upper obstacle $p(a) \le x + a$, and $h \ge 0$ gives $g \ge d$, hence the cap $p \le 1/d$; both obstacles are linear in p. The objective becomes $-\log(d \cdot x)$ plus an energy that is concave in $\log p$ once the endpoint $x = p(0)$, the reciprocal of the final tail integral, is fixed (Lemma 5.1). This is the only step that changes the structure of the problem; everything after it is a comparison.

A reference curve at every endpoint, and a certificate. Fixing x leaves a concave problem in $\log p$ with two upper obstacles. For every endpoint x we construct a reference curve P_x that follows the line $x + a$, then solves the Euler equation of the energy on a middle arc, then follows the cap $1/d$ (Lemma 6.1). Comparing any admissible p with the same endpoint to P_x, and writing $v = \log(p/P_x)$, the energy of P_x minus the energy of p is

$$\int \left[a(v')^2 + P_x^{-2}(e^{-2v} - 1 + 2v) - 2Q_x v \right] \mathrm{d}a.$$

The first two terms are nonnegative. The function Q_x is nonnegative and supported where the reference curve P_x meets an upper obstacle, so $v \le 0$ there and the last term is nonnegative as well. In the fixed-endpoint concave relaxation, $2Q_x$ is the obstacle multiplier. Its sign gives the weak-duality comparison familiar from LPs, and its contact-set support expresses complementary slackness. This is the fixed-endpoint gap, identity (G1) of Lemma 7.1. A second identity, (G2) of the same lemma, shows that the reference value as a function of the endpoint x has a unique maximum. Its derivative along the family of reference curves has the sign of a single explicit factor that changes sign exactly once; the nonconcavity of the original problem survives only in this one-dimensional comparison. Any competitor is therefore dominated in two steps, first by the reference curve at its own endpoint and then by the reference curve at the best endpoint. Equality in both steps forces it to coincide with the reference curve at the best endpoint, where the inverse coordinate change yields a control satisfying the original time constraint; this control is the unique maximizer (Theorem C). In the language of the calculus of variations, P_x is a calibration: an exact comparison identity that certifies global optimality, with $2Q_x$ as the obstacle multiplier.

What the relaxations cost. Three constraints were dropped on the way: monotonicity of h; the slope bound $0 \le p' \le 1$, kept only through its two integrated obstacles; and the normalization $\int p^{-2}\,\mathrm{d}a = 1$, which says that original time has length one. The maximizing control satisfies all three, so the relaxed value is the true value. It is still not a buyer: it vanishes near $t = 0$, so the survival function is zero just below $s = 1$, while the tail integral at $s = 1$ equals $d > 0$. Only a limit of buyers has both properties, through an atom of mass $\eta \to 0$ placed at $1 + d/\eta$, which is why the worst case is approached and never attained.

The worst-case pair. At the optimal endpoint the reference curve follows the line $x + a$, then a curved arc solving $P'' = -CP/a^2$, where C is the first integral of the Euler equation, then the cap $1/d$ (Figure 2a). Transforming back, the worst-case buyer's survival function equals 1 up to $s_1 = C^2/(1 - C)$, decreases along an explicit curve to 0 at $s_2 = 1 - C$, and the bounded hard pairs add an atom of mass η at value $1 + d/\eta$, whose contribution $\eta \cdot d/\eta = d$ to the tail integral at $s = 1$ survives as its probability vanishes (Figure 2b). The equalizing seller of [LRW23] has atoms at zero and at one, the second vanishing as $\eta \to 0$, and makes every price in $[0, 1)$ equally good (Figure 2c).

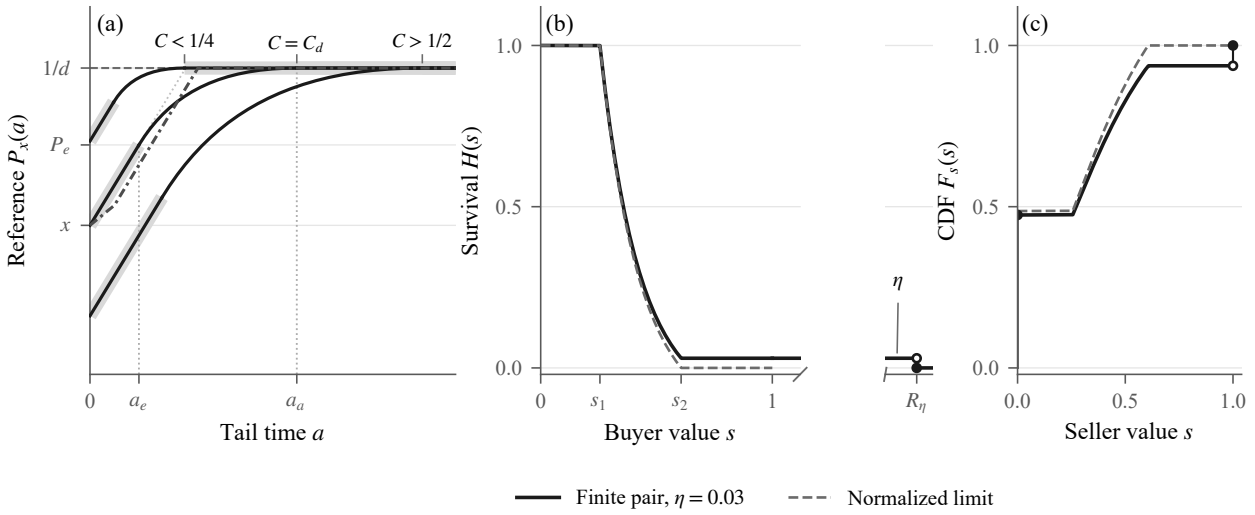


Figure 2: Reference curves, and the bounded hard pairs at $C = C_*$. (a) Reference curves at three endpoints, for $C < 1/4$, $C = C_d$, and $C > 1/2$, between the line obstacles $x + a$ and the cap $1/d$; contact segments are shaded, and x, P_e, a_e, a_a refer to the middle curve. The dash-dotted curve is a competitor in the obstacle class with the endpoint of the middle curve. (b,c) Solid curves give an actual independent pair with $\eta = 0.03$; dashed curves give the normalized limit. The buyer has an atom of mass η at $R_\eta = 1 + d/\eta$, shown after the axis break. The seller of the actual pair has atoms at zero and one; the atom at one vanishes in the limit. The buyer's free boundaries are $s_1 = C^2/(1 - C)$ and $s_2 = 1 - C$.

No genuine pair attains the ratio, and every pair is strictly above β_*. The middle arc is not an exponential on any subinterval (Proposition 8.1). Moving the parameter C moves the pair along the frontier of Theorem B; adding the finite buyer tail gives bounded hard pairs at every point of that frontier, and the price representation and smoothing then give the dominant-strategy welfare ceiling (Corollary A′).

New and adapted ingredients. The change of variables (Lemma 5.1) and the two gap identities (Lemma 7.1) carry the single-unit argument. Lemmas 9.1 and 9.2 adapt the price density and the equalizing seller of [LRW23] to arbitrary buyer atoms and to bounded supports, and we give their short proofs to fix the conventions. Appendix B.2 proves the price representation of [HR87, ČP16] in an almost-everywhere form for measurable mechanisms, which is the form that smoothed hard pairs require. The family optimization for two units (Appendix D) is computer-assisted; its computations and its interval and positivity certificates, with a replay script, are in the supplementary material.

Extension to two units. The same architecture applies to two units with one new difficulty. After fixing the buyer's two marginal values, a common price must satisfy one inequality for every ordered pair of seller costs. The order between the two costs couples the two units. A nondecreasing function, the compensation, moves slack between them: if the first unit's slack is at least minus the compensation and the second unit's slack is at least the compensation, their sum is nonnegative for every ordered pair of costs. For each compensation the least feasible price mass is the fixed point of a contraction (Theorem E). This least mass is convex in the compensation; its minimum is a functional of the buyer, and the supremum of that functional over buyers is the one quantity that remains open (Conjecture 13.3). Within an explicit family the worst ratio is attained at the unique admissible

solution of the family's first-order conditions and lies in (0.72908, 0.729081) (Proposition 13.2); the tail normalization that reduces general instances to bounded ones is Proposition 13.1.

3 Model and the scalar curve

Welfare is the seller's initial value plus the gains created by trade; keeping these two terms separate also exposes the GFT frontier.

The independent values V_s, V_b are nonnegative and satisfy

$$0 < \mathbb{E}\max\{V_s, V_b\} < \infty.$$

A deterministic price $z \geq 0$ trades when $V_s \leq z \leq V_b$. We use this inclusive convention for fixed prices. Define

$$M = \mathbb{E}V_s, \qquad G = \mathbb{E}(V_b - V_s)_+, \qquad \Gamma(z) = \mathbb{E}[(V_b - V_s)\mathbf{1}_{\{V_s \leq z \leq V_b\}}]. \tag{1}$$

The price welfare is $W(z) = M + \Gamma(z)$, and first-best welfare is $W_{\mathrm{FB}} = M + G$. We write $\Gamma_{\max} = \max_{z \geq 0} \Gamma(z)$ and

$$r_{\mathrm{FP}} = \inf_{F_s, F_b} \frac{M + \Gamma_{\max}}{M + G}. \tag{2}$$

The maximum exists: the inclusive rule makes Γ upper semicontinuous by reverse Fatou, and $\Gamma(z) \leq \mathbb{E}[V_b \mathbf{1}_{\{V_b \geq z\}}] \to 0$ as $z \to \infty$. If Γ is identically zero, every price is optimal. Randomization over prices cannot improve the optimum for a fixed pair of priors.

One parameter C indexes the whole construction. It is the first integral (18) of the Euler equation on the middle arc of the reference curve, and it sets the arc's curvature through $P'' = -CP/a^2$. As the overview explains, $\tau(C)$ is the initial state d at which the arc with parameter C is optimal, and $S(C) = \max \mathcal{J}_{\tau(C)}$ is the optimal value there (Theorem C). The quadratic D_C and its integral $\mathcal{I}_C$ arise when the Euler equation is integrated along the arc, whose running parameter is y; $y_e(C)$ is the value of y where the optimal arc meets the line. We record the closed forms now and use them in Theorems A and B. The scalar functions governing the optimum use $1/4 < C < 1/2$; the comparison at arbitrary endpoints will require the larger range specified in Section 6. For $1/4 < C < 1/2$, define

$$\begin{aligned} D_C(y) &= 1 - y + Cy^2, \quad y_e(C) = \frac{1 - 2C}{C(1 - C)}, \\ \mathcal{I}_C(y) &= \int_0^y \frac{\mathrm{d}r}{D_C(r)}, \qquad \mathcal{I}(C) = \mathcal{I}_C(y_e(C)), \end{aligned} \tag{3}$$

and the scalar curve

$$\tau(C) = \frac{C^2}{1 - 2C} e^{-\mathcal{I}(C)/2}, \qquad S(C) = C(2 + \mathcal{I}(C)). \tag{4}$$

With $\omega = \sqrt{C - 1/4}$ and principal arctangents, integration gives

$$\mathcal{I}(C) = \frac{1}{\omega}\left[\arctan\frac{1 - 3C}{2(1 - C)\omega} + \arctan\frac{1}{2\omega}\right]. \tag{5}$$

The normalized functional of [LRW23] uses a buyer survival function $H(s) = \Pr(V_b > s)$, its tail integral $L(s) = \mathbb{E}(V_b - s)_+$, and $K(s) = \int_s^1 H(u)^2\,\mathrm{d}u$. At a cutoff normalized by $L(1) = d$, reversing

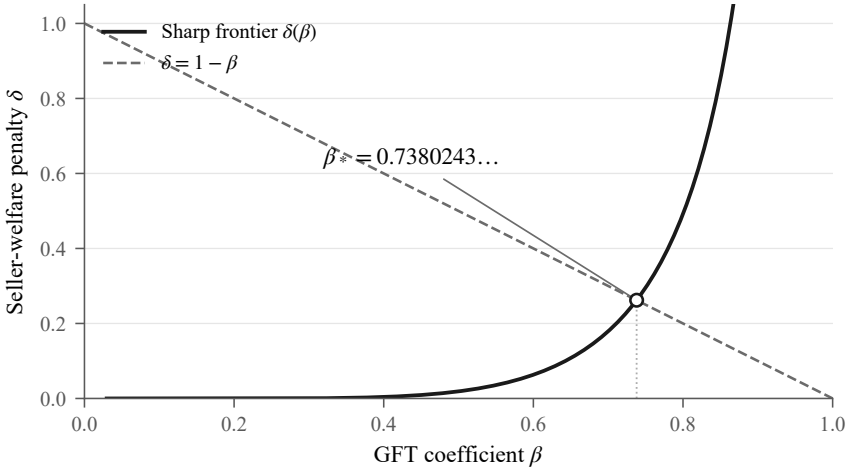


Figure 3: The sharp affine frontier. Its unique intersection with $\delta = 1 - \beta$ gives the welfare constant. Tangents with slope κ determine the exact conditional GFT guarantee through $\kappa\rho(\kappa) = \delta^*(\kappa)$.

time gives $(h, g, k)(t) = (H, L, K)(1 - t)$. For general $d > 0$ and measurable $h : [0, 1] \to [0, 1]$, define

$$\begin{aligned} g(t) &= d + \int_0^t h(u)\,\mathrm{d}u, \qquad k(t) = \int_0^t h(u)^2\,\mathrm{d}u, \\ \mathcal{J}_d(h) &= \int_0^1 \left[\frac{h(t)}{g(t)} + \frac{d^2 - k(t)}{g(t)^2}\right] \mathrm{d}t. \end{aligned} \tag{6}$$

Here $d > 0$ is the fixed initial state $g(0)$; the function $\tau(C)$ in (4) parameterizes the initial states of the maximizing family. Buyer survival functions correspond to nondecreasing controls after reversing time. We maximize $\mathcal{J}_d$ over the entire measurable box.

Differentiation of (3)–(5) gives

$$\mathcal{I}' = -\frac{2(C\mathcal{I} + 2)}{C(4C - 1)}, \qquad S' = \frac{(2C - 1)\mathcal{I} + 8C - 6}{4C - 1} < 0, \qquad \frac{\tau'}{\tau} = \frac{S'}{2C - 1} > 0. \tag{7}$$

The signs follow from $1/4 < C < 1/2$ and $\mathcal{I} > 0$. At the endpoints,

$$(S(C), \tau(C)) \longrightarrow (\infty, 0) \quad (C \downarrow 1/4), \qquad (S(C), \tau(C)) \longrightarrow (1, \infty) \quad (C \uparrow 1/2). \tag{8}$$

In particular, for every $d > 0$ there is a unique $C_d \in (1/4, 1/2)$ with $\tau(C_d) = d$. Table 2 in Appendix A lists every symbol by role with its definition site.

4 The complete gains-from-trade frontier

Define the sharp affine coefficient by

$$\delta(\beta) = \inf\left\{\delta \in \mathbb{R} : \Gamma_{\max} \geq \beta G - \delta M \text{ for every pair in Section 3}\right\}, \qquad 0 < \beta < 1. \tag{9}$$

For $\kappa > 0$, define

$$\rho(\kappa) = \inf_{\substack{F_s, F_b \text{ as in Section 3} \\ M > 0,\ G/M = \kappa}} \frac{\Gamma_{\max}}{G}. \tag{10}$$

Theorem B (Sharp affine and conditional GFT guarantees)**.** *The scalar curve* (4) *gives the entire frontier:*

$$\beta(C) = \frac{1}{S(C)}, \qquad \delta(\beta(C)) = \frac{\tau(C)}{S(C)}, \qquad \frac{1}{4} < C < \frac{1}{2}. \tag{11}$$

Every coefficient smaller than $\delta(\beta)$ is refuted by a bounded hard pair of independent distributions.

The function δ is strictly increasing and strictly convex on $(0,1)$. Its derivative ranges from zero to infinity.

With $u = 1 - 2C$, the map

$$\kappa(C) = \frac{\tau(C)(S(C)+u)}{u} \tag{12}$$

is a strictly increasing bijection from $(1/4, 1/2)$ onto $(0,\infty)$, and

$$\rho(\kappa(C)) = \frac{1}{S(C)+u} = \frac{1}{1+C\mathcal{I}(C)}. \tag{13}$$

Equivalently, extending δ by $+\infty$ outside $(0,1)$ for the purpose of conjugation,

$$\kappa\rho(\kappa) = \delta^*(\kappa) := \sup_{0<\beta<1} \{\kappa\beta - \delta(\beta)\}. \tag{14}$$

For each fixed κ, bounded hard pairs with exactly $G/M = \kappa$ approach (13)*.*

Corollary B′ (Endpoint asymptotics)**.** *The sharp curves satisfy*

$$\begin{aligned} &\delta(\beta) \sim \frac{e}{8}\,\beta e^{-2/\beta} \quad (\beta \downarrow 0), \qquad \delta(\beta) \sim \frac{1}{4(1-\beta)} \quad (\beta \uparrow 1), \\ &\rho(\kappa) \sim \frac{2}{\log(1/\kappa)} \quad (\kappa \downarrow 0), \quad 1-\rho(\kappa) \sim \frac{1}{\sqrt{\kappa}} \qquad (\kappa \to \infty). \end{aligned} \tag{15}$$

The frontier and its hard pairs are proved in Section 11; the endpoint expansions are proved in Appendix B.1.

In (12), $\kappa(C) = \delta'(\beta(C))$ is the slope of the frontier at $\beta(C)$, so the conditional guarantee at a ratio κ is read off where the frontier has slope κ (Figure 3). The welfare constant of Theorem A is the point of this frontier on the line $\delta = 1 - \beta$. There $\tau(C_*)/S(C_*) = 1 - 1/S(C_*)$, that is, $S(C_*) = 1 + \tau(C_*)$ and $\beta_* = 1/S(C_*)$.

5 From controls to a local energy

Each contribution $h(u)^2$ to the accumulated second moment affects every later time. The reciprocal clock collects those later contributions into a single weight, making the objective local in the reciprocal state p. The natural state for (6) is g; in that variable the objective is not concave and $0 \le h \le 1$ constrains a slope. The reciprocal state and the tail time are chosen so that the second-moment term becomes a Dirichlet energy with a positive weight and the two slope constraints integrate to two upper obstacles.

For a control in (6), introduce

$$A(t) = \int_t^1 g(u)^{-2}\,\mathrm{d}u, \qquad a_0 = A(0), \qquad p(A(t)) = \frac{1}{g(t)}, \qquad x = p(0) = \frac{1}{g(1)}. \tag{16}$$

Extend p by the constant $1/d$ for $a \ge a_0$.

Lemma 5.1 (energy representation)**.** *For every $d > 0$, a measurable $h : [0,1] \to [0,1]$ induces through* (16) *an absolutely continuous p with*

$$p'(A(t)) = h(t) \quad a.e., \qquad p(a) \le \min\{x + a, 1/d\}, \qquad \frac{1}{d+1} \le x \le \frac{1}{d}, \tag{17}$$

and satisfies

$$\mathcal{J}_d(h) - 1 = \mathcal{F}_d[p] := -\log(d \cdot x) + \int_0^\infty \left[d^2 - p^{-2} - a((\log p)')^2\right] \mathrm{d}a. \tag{E}$$

Proof. Since $d \le g \le d+1$, the map A is strictly decreasing and bi-Lipschitz. The chain rule gives $p'(A(t)) = h(t)$ and $\mathrm{d}t = -p^{-2}\,\mathrm{d}a$; the slope bound and constant extension give the two obstacles in (17).

The term $h(u)^2$ enters $k(t)$ for every $t \ge u$, so the objective has memory, and in the original variables no pointwise comparison of two controls is available. We reverse the order of integration to collect the total weight of $h(u)^2$:

$$\int_0^1 \frac{k(t)}{g(t)^2}\,\mathrm{d}t = \int_0^1 h(u)^2 \left(\int_u^1 g(t)^{-2}\,\mathrm{d}t\right)\mathrm{d}u = \int_0^1 A(u)h(u)^2\,\mathrm{d}u.$$

This is the source of the weight $a = A(u)$. Using $g' = h$ for the remaining first-order term and then changing variables gives

$$\mathcal{J}_d(h) = \log\frac{g(1)}{d} + d^2 a_0 - \int_0^1 A(t)h(t)^2\,\mathrm{d}t = -\log(d \cdot x) + \int_0^{a_0} \left[d^2 - a((\log p)')^2\right] \mathrm{d}a.$$

The original time interval has length one, so $1 = \int_0^{a_0} p(a)^{-2}\,\mathrm{d}a$. Subtracting this normalization yields the energy; its integrand vanishes on the constant extension:

$$\mathcal{J}_d(h) - 1 = -\log(d \cdot x) + \int_0^{a_0} \left[d^2 - p^{-2} - a((\log p)')^2\right] \mathrm{d}a.$$

□

At fixed x, the integral in (E) is concave in $\ell = \log p$, and the obstacles become

$$\ell(a) \le \min\{\log(x + a), -\log d\}.$$

The obstacle class relaxes the original problem, since we keep the two integrated consequences of $0 \le h \le 1$ and drop the slope constraint itself. This is what makes the fixed-endpoint problem concave. The price is that a maximizer of the relaxed problem need not come from an admissible control, and attaining the bound also requires the inverse coordinate change to satisfy the original slope and time constraints; this is settled only in Theorem C.

For this concave problem, a reference curve P_x must solve the Euler equation between its contacts with the two obstacles. Writing $\ell_x = \log P_x$, the equation on that interior interval is

$$(a\ell_x')' + e^{-2\ell_x} = 0.$$

Along a solution, differentiation shows that

$$C = ae^{-2\ell_x} + a\ell_x' - a^2(\ell_x')^2 \quad \text{is constant}, \qquad P_x'' = -\frac{C}{a^2}P_x. \tag{18}$$

We integrate this equation between tangency to the line $x + a$ and tangency to the cap $1/d$. The construction below gives a reference curve for every endpoint, including endpoints where that reference curve fails the original time constraint.

6 A reference curve at every endpoint

At a fixed initial state d, which sits at the far end of the tail-time axis where $p = 1/d$, each reference curve follows the line, solves the Euler equation on its middle arc, and then meets the constant cap. Covering every endpoint is necessary for a global comparison: a bound at the optimal endpoint alone says nothing about a competitor with a different endpoint. The family of reference curves therefore has to be continued in C below $1/4$, where D_C has real roots and the arctangent form (5) no longer applies, and above $1/2$ up to a terminal value $\bar{C}(d)$.

Fix $d > 0$. To compare all endpoints, we let the curvature parameter C vary over positive values for which the contact equations are defined. Lemma 6.1 identifies this range as $(0, \bar{C}(d))$. The endpoint envelope in the next section still uses this full range; its maximum selects $C_d \in (1/4, 1/2)$. Throughout this construction, D_C and $\mathcal{I}_C$ retain their formulas in (3) on intervals where $D_C > 0$. Define

$$\chi(C, y) = \frac{\sqrt{D_C(y)}}{y} e^{-\mathcal{I}_C(y)/2}, \qquad Cy < 1, \qquad D_C(r) > 0 \ (0 \le r \le y). \tag{19}$$

The running parameter is y, and $Y_d(C)$ denotes the contact parameter solving $\chi(C, Y_d(C)) = d$. This equation is a tangency condition: traced from the cap, where its slope is zero, the middle arc of the reference curve reaches slope one at the parameter value $Y_d(C)$, and there it touches the line $x + a$ (Lemma 6.1). Below we often write Y_d for $Y_d(C)$, and similarly Z_d, U_d, and $\mathcal{B}_d$ for the functions of C defined below. The corresponding reciprocal-state endpoint and contact data are

$$\begin{aligned} x(C) &= \frac{1 - D_C(Y_d)}{D_C(Y_d)}, \qquad a_e = \frac{CY_d^2}{D_C(Y_d)}, \\ a_a &= \frac{C}{d^2}, \qquad P_e = \frac{Y_d}{D_C(Y_d)}. \end{aligned} \tag{20}$$

For the endpoint derivative, define

$$Z_d(C) = \int_0^{Y_d} \frac{r^2}{D_C(r)^2}\,\mathrm{d}r, \qquad U_d(C) = \int_0^{Y_d} \frac{Y_d - r}{D_C(r)^2}\,\mathrm{d}r. \tag{21}$$

Lemma 6.1 (reference curves for every endpoint)**.** *For $d > 0$, there is $\bar{C}(d) > 1/4$ such that Y_d from* (19) *exists uniquely on $0 < C < \bar{C}(d)$. The map $C \mapsto x(C)$ is a continuous strictly decreasing bijection from this interval onto $(0, 1/d)$. For each such endpoint, using the contact data* (20)*, the function*

$$P_x(a) = \begin{cases} x + a, & 0 \le a \le a_e, \\ \textit{the middle arc below}, & a_e \le a \le a_a, \\ 1/d, & a \ge a_a, \end{cases} \tag{22}$$

$$a(r) = \frac{C}{d^2} e^{-\mathcal{I}_C(r)}, \quad P_x(a(r)) = \frac{e^{-\mathcal{I}_C(r)/2}}{d\sqrt{D_C(r)}} \quad (0 \le r \le Y_d(C))$$

is continuously differentiable and lies below both obstacles in (17)*. The function*

$$Q_x = P_x^{-2} + \left(a(\log P_x)'\right)' \tag{23}$$

has no atoms and is given almost everywhere by

$$Q_x(a) = \begin{cases} (1 + x)/(x + a)^2, & 0 < a < a_e, \\ 0, & a_e < a < a_a, \\ d^2, & a > a_a. \end{cases} \tag{24}$$

Proof. The contact equation determines an endpoint branch. Its monotone endpoint map gives coverage; the three pieces then determine the reference curve and its multiplier.

Existence of the endpoint branch. The partial derivatives of (19) are

$$\partial_y \log\chi = -\frac{1}{yD_C(y)} < 0, \qquad \partial_C \log\chi = \frac{y^2}{2D_C(y)} + \frac{1}{2}\int_0^y \frac{r^2}{D_C(r)^2}\,\mathrm{d}r > 0. \tag{25}$$

For $0 < C \le 1/4$, the permitted interval ends at the first positive root of D_C. As y ranges from zero to that root, χ decreases from infinity to zero. For $C > 1/4$, the permitted interval is $(0, 1/C)$, and its terminal value is

$$\chi_{\mathrm{end}}(C) = Ce^{-\mathcal{I}_C(1/C)/2}.$$

Along this endpoint, (25) gives

$$(\log\chi_{\mathrm{end}})' = \partial_C \log\chi(C, 1/C) + 1/C > 0.$$

This endpoint value tends to zero as $C \downarrow 1/4$ and to infinity as $C \to \infty$; for the latter limit, use $\mathcal{I}_C(1/C) \le [C(1 - 1/(4C))]^{-1}$. There is therefore a unique $\bar{C}(d) > 1/4$ with $\chi_{\mathrm{end}}(\bar{C}(d)) = d$, and the branch exists exactly on $(0, \bar{C}(d))$.

Coverage of every reciprocal-state endpoint. Along $\chi(C, Y_d) = d$, implicit differentiation using (25) gives

$$Y_d' = -\left.\frac{\partial_C \log\chi}{\partial_y \log\chi}\right|_{y=Y_d} = \frac{Y_d^3 + Y_d D_C(Y_d) Z_d}{2} > 0.$$

To express U_d through Z_d, we integrate two exact derivatives over $[0, Y_d]$:

$$\frac{\mathrm{d}}{\mathrm{d}r}\frac{r}{D_C(r)} = \frac{1 - Cr^2}{D_C(r)^2}, \qquad \frac{\mathrm{d}}{\mathrm{d}r}\frac{r^2}{D_C(r)} = \frac{2r - r^2}{D_C(r)^2}.$$

They give

$$\int_0^{Y_d} \frac{\mathrm{d}r}{D_C(r)^2} = \frac{Y_d}{D_C(Y_d)} + CZ_d, \qquad 2\int_0^{Y_d} \frac{r\,\mathrm{d}r}{D_C(r)^2} = \frac{Y_d^2}{D_C(Y_d)} + Z_d,$$

and together with Y_d' they give

$$\begin{aligned} 2U_d &= \frac{Y_d^2}{D_C(Y_d)} + (2CY_d - 1)Z_d, \\ \frac{\mathrm{d}}{\mathrm{d}C} D_C(Y_d) &= Y_d^2 + (2CY_d - 1)Y_d' \\ &= Y_d D_C(Y_d)(Y_d + U_d). \end{aligned} \tag{26}$$

The first line holds because U_d is Y_d times the first integral above minus the second. The second line substitutes Y_d' into the total derivative of $D_C(Y_d)$; after inserting the first line and $D_C(Y_d) = 1 - Y_d + CY_d^2$, both sides of its last equality expand to $Y_d^2 - \frac{1}{2}Y_d^3 + CY_d^4 + \frac{1}{2}(2CY_d - 1)Y_d D_C(Y_d) Z_d$. Since $U_d > 0$ by its integral definition,

$$x'(C) = -\frac{Y_d(Y_d + U_d)}{D_C(Y_d)} < 0. \tag{27}$$

At $C = 0$, the continuous extension of (19) is $\chi(0, y) = (1 - y)/y$. The implicit root tends to $1/(1 + d)$, so $x(C) \to 1/d$ as $C \downarrow 0$. At the other end, Y_d is increasing and bounded by $1/C$, so it has a limit $\bar{Y} \le 1/\bar{C}(d)$ as $C \uparrow \bar{C}(d)$, and continuity of χ gives $\chi(\bar{C}(d), \bar{Y}) = d$. The function

$\chi(\bar{C}(d),\cdot)$ is strictly decreasing and takes the value d at its terminal point $1/\bar{C}(d)$, so $\bar{Y} = 1/\bar{C}(d)$. Consequently $D_C(Y_d) \to 1 - 1/\bar{C}(d) + 1/\bar{C}(d) = 1$ and $x(C) \to 0$.

The reference curve and its multiplier. We differentiate the parametric middle arc:

$$\frac{\mathrm{d}a}{\mathrm{d}r} = -\frac{a}{D_C(r)}, \qquad P_x' = \frac{r}{P_x D_C(r)}, \qquad P_x'' = -\frac{C}{a^2}P_x < 0. \tag{28}$$

The identity $\chi(C, Y_d) = d$ gives $a(Y_d) = a_e$ and $P_x(a_e) = P_e = x + a_e$, with slope one. At $r = 0$, the curve reaches $(a_a, 1/d)$ with slope zero, so the pieces fit in C^1. Figure 2(a) shows the curves at three endpoints. Concavity puts the middle arc below both obstacles and keeps its slope between zero and one. On that arc $a(\log P_x)' = Cr$, so (23) vanishes. Since the joins are C^1, there are no jump terms, and on the contact sets

$$Q_x(a) = \frac{1+x}{(x+a)^2} \quad (0 < a < a_e), \qquad Q_x(a) = d^2 \quad (a > a_a).$$

□

At the limiting endpoint, set $Y_d(0) = 1/(1+d)$, $x(0) = 1/d$, and $P_{1/d} \equiv 1/d$. A reference curve need not meet the original time constraint; it bounds the energy of every reciprocal state arising from a control with the same endpoint.

7 Fixed-endpoint comparison and the endpoint envelope

At a fixed endpoint, the energy gap splits into a square, an exponential remainder, and a contact term. Once these terms are nonnegative, only the endpoint envelope remains to be maximized.

For the reciprocal state p induced by a control h, let

$$v(a) = \log \frac{p(a)}{P_x(a)}, \qquad \Phi_d(x) = \mathcal{F}_d[P_x], \qquad \mathcal{B}_d(C) = 2C - 1 + C(1-C)Y_d. \tag{29}$$

Lemma 7.1 (nonnegative gaps)**.** *For p from* (16) *and the notation* (29)*, the fixed-endpoint gap is*

$$\Phi_d(x) - \mathcal{F}_d[p] = \int_0^\infty \left[a(v')^2 + \frac{e^{-2v} - 1 + 2v}{P_x^2} - 2Q_x v \right] \mathrm{d}a \ \geq 0. \tag{G1}$$

It vanishes exactly when $p = P_x$. The endpoint envelope has a unique maximum at $C = C_d$, and

$$\begin{aligned} \Phi_d(x(C)) &= C\mathcal{I}_C(Y_d(C)) - \frac{C^2 Y_d(C)}{1 - CY_d(C)}, \\ \frac{\mathrm{d}}{\mathrm{d}C}\Phi_d(x(C)) &= -\frac{(Y_d(C) + U_d(C))\mathcal{B}_d(C)}{(1 - CY_d(C))^2}, \\ \Phi_d(x(C_d)) - \Phi_d(x(C)) &= \int_{\min\{C,C_d\}}^{\max\{C,C_d\}} \frac{(Y_d(u) + U_d(u))|\mathcal{B}_d(u)|}{(1 - uY_d(u))^2}\,\mathrm{d}u. \end{aligned} \tag{G2}$$

Proof. The fixed-endpoint gap bounds the energy at a prescribed endpoint; the endpoint envelope then selects the maximizing endpoint.

The fixed-endpoint gap (G1). The reciprocal state p and reference curve P_x agree at zero and eventually equal $1/d$, so $v(0) = 0$ and v has compact support. Expanding the energy gives

$$\begin{aligned}\Phi_d(x) - \mathcal{F}_d[p] &= \int_0^\infty \left[\frac{e^{-2v}-1}{P_x^2} + a(2\ell_x' v' + (v')^2)\right] \mathrm{d}a \\ &= \int_0^\infty \left[a(v')^2 + \frac{e^{-2v}-1+2v}{P_x^2} + 2a\ell_x' v' - \frac{2v}{P_x^2}\right] \mathrm{d}a \\ &= \int_0^\infty \left[a(v')^2 + \frac{e^{-2v}-1+2v}{P_x^2} - 2Q_x v\right] \mathrm{d}a.\end{aligned}$$

Integration by parts contributes $2a\ell_x' v$, which vanishes at both ends; the C^1 joins produce no internal boundary terms.

The exponential remainder has the positive-kernel representation

$$e^{-2v} - 1 + 2v = 4v^2 \int_0^1 (1-\theta) e^{-2\theta v}\, \mathrm{d}\theta. \tag{30}$$

Where $Q_x > 0$, the reference curve P_x equals an upper obstacle. By (17), $v \le 0$ there; this sign is the only place where the constraint $0 \le h \le 1$ enters the comparison. Figure 2(a) shows a competitor with the endpoint of the middle curve: it lies below that curve on both contact segments and above it in between, where $Q_x = 0$. Every summand in (G1) is nonnegative, and (30) forces $v = 0$ almost everywhere in the equality case. Continuity gives $p = P_x$.

Value of the envelope. We first evaluate Φ_d by integrating the three pieces of the reference curve P_x. The contact data (20) and $\chi(C, Y_d) = d$ give three identities that are used throughout this step:

$$\begin{gathered}1 - D_C(Y_d) = Y_d(1 - CY_d), \qquad \frac{x}{P_e} = 1 - CY_d, \\ \log(d \cdot P_e) = -\tfrac{1}{2}\left[\log D_C(Y_d) + \mathcal{I}_C(Y_d)\right].\end{gathered}$$

On the line,

$$\int_0^{a_e} \frac{\mathrm{d}a}{(x+a)^2} = \frac{1}{x} - \frac{1}{P_e} = \frac{CD_C(Y_d)}{1 - CY_d}, \qquad \int_0^{a_e} \frac{a\,\mathrm{d}a}{(x+a)^2} = \log\frac{P_e}{x} + \frac{x}{P_e} - 1.$$

On the middle arc we integrate in the parameter r of (22). By (28), $\mathrm{d}a = -(a/D_C(r))\,\mathrm{d}r$, $P_x^{-2}\,\mathrm{d}a = -C\,\mathrm{d}r$, and $a(\log P_x)' = Cr$, so

$$\begin{aligned}\int_{a_e}^{a_a} P_x^{-2}\,\mathrm{d}a &= CY_d, \\ \int_{a_e}^{a_a} a\big((\log P_x)'\big)^2\,\mathrm{d}a &= C^2 \int_0^{Y_d} \frac{r^2\,\mathrm{d}r}{D_C(r)} \\ &= CY_d - C\mathcal{I}_C(Y_d) + \tfrac{1}{2}\left[\log D_C(Y_d) + \mathcal{I}_C(Y_d)\right].\end{aligned}$$

The last equality uses $Cr^2 = D_C(r) - 1 + r$ and $2C(1-r) = (2C-1) - D_C'(r)$. The cap contributes nothing, and $d^2 a_a = C$. The two integrals of P_x^{-2} give

$$\int_0^{a_a} P_x^{-2}\,\mathrm{d}a = CY_d + \frac{CD_C(Y_d)}{1 - CY_d}. \tag{31}$$

In the sum of all pieces the logarithms cancel against $-\log(d \cdot x)$ by the third contact identity, and $x/P_e = 1 - CY_d$ removes the remaining constants:

$$\Phi_d(x(C)) = C\big(1 + \mathcal{I}_C(Y_d)\big) - CY_d - \frac{CD_C(Y_d)}{1 - CY_d} = C\mathcal{I}_C(Y_d) - \frac{C^2 Y_d}{1 - CY_d},$$

which is the first identity in (G2).

Derivative along the branch. The derivative of Φ_d along the branch is the one computation in the single-unit proof that does not reduce to a single identity: it needs the branch derivative Y_d', the contact identities, and the following relation together:

$$(4C - 1)Z_d = 2\mathcal{I}_C(Y_d) + \frac{Y_d(Y_d - 2)}{D_C(Y_d)}. \tag{32}$$

To prove (32), we replace Y_d by a running value y. Both sides vanish at $y = 0$, and their y-derivatives agree because $(4C - 1)y^2 = 2yD_C(y) - y(y - 2)(2Cy - 1)$. No division by $4C - 1$ occurs, so the identity also holds at $C = 1/4$. Since $\partial_C \mathcal{I}_C(y) = -\int_0^y r^2 D_C(r)^{-2}\,\mathrm{d}r$, the first identity in (G2) gives

$$\frac{\mathrm{d}}{\mathrm{d}C}\Phi_d(x(C)) = \mathcal{I}_C(Y_d) - CZ_d + \frac{CY_d'}{D_C(Y_d)} - \frac{2CY_d - C^2Y_d^2 + C^2Y_d'}{(1 - CY_d)^2}.$$

We eliminate $\mathcal{I}_C(Y_d)$ by (32) and substitute $2Y_d' = Y_d^3 + Y_d D_C(Y_d) Z_d$. The result is affine in Z_d. Its Z_d-coefficient is

$$\frac{(2C - 1) + CY_d(3 - 5C) - 2C^2(1 - C)Y_d^2}{2(1 - CY_d)^2} = -\frac{(2CY_d - 1)\,\mathcal{B}_d}{2(1 - CY_d)^2},$$

and the part free of Z_d is

$$-\frac{Y_d\big(2D_C(Y_d) + Y_d\big)\mathcal{B}_d}{2D_C(Y_d)(1 - CY_d)^2};$$

both are polynomial identities in C and Y_d after clearing denominators. By the first line of (26), their sum is $-(Y_d + U_d)\mathcal{B}_d/(1 - CY_d)^2$, the derivative in (G2).

Location of the maximum. Concavity at a fixed endpoint does not give global optimality, because the nonconcavity of the original problem survives in the endpoint. A second local maximum of the envelope would break the comparison, and excluding it requires the sign of $\mathcal{B}_d$ on the whole branch $0 < C < \bar{C}(d)$, including the ranges $C \le 1/4$ and $C \ge 1/2$ where the closed forms of Section 3 are not available.

To locate the maximum, use the rational formula for $y_e(C)$ in (3) on all of $0 < C < 1/2$. Within the permitted branch, $\mathcal{B}_d = 0$ is equivalent to $Y_d = y_e(C)$. The contact parameter Y_d increases strictly with C, whereas

$$y_e'(C) = -\frac{1}{C^2} - \frac{1}{(1 - C)^2} < 0.$$

For $C \le 1/4$, the permitted branch lies before the first root of D_C, so $2CY_d < 1$ and $\mathcal{B}_d < 0$. For $1/2 \le C \le 1$, both terms in $2C - 1 + C(1 - C)Y_d$ are nonnegative and their sum is positive. For $C > 1$, the bound $CY_d < 1$ gives $\mathcal{B}_d > C > 0$. If $\bar{C}(d) < 1/2$, the right-end limit instead gives $\mathcal{B}_d \to \bar{C}(d) > 0$. Thus there is exactly one zero in $(1/4, 1/2)$. At that zero,

$$D_C(y_e(C)) = \frac{C^2}{(1 - C)^2}, \qquad \chi(C, y_e(C)) = \tau(C),$$

so this zero is C_d. The endpoint envelope increases before C_d and decreases after it. Integrating its derivative from C_d gives (G2). □

Dividing the derivative in (G2) by (27) gives

$$\frac{\mathrm{d}\Phi_d}{\mathrm{d}x} = \frac{D_C(Y_d)\,\mathcal{B}_d}{Y_d(1-CY_d)^2} = \frac{1}{x^2} - \frac{1+x}{P_e^2},$$

where the second form uses $x = (1-D_C(Y_d))/D_C(Y_d)$, $P_e = Y_d/D_C(Y_d)$, and $D_C(Y_d)-(1-CY_d)^2 = Y_d\mathcal{B}_d$. This is the value that the multiplier predicts. Moving the endpoint changes $-\log(d\cdot x)$ at rate $-1/x$ and relaxes the line obstacle $\log(x+a)$ at rate $1/(x+a)$, and by (24)

$$-\frac{1}{x} + \int_0^{a_e} \frac{2Q_x(a)}{x+a}\,\mathrm{d}a = \frac{1}{x^2} - \frac{1+x}{P_e^2}.$$

8 Global calibration and the extremal shape

Theorem C (Global bound and maximizing control)**.** *Let $d > 0$, let $h : [0,1] \to [0,1]$ be measurable, and let $C_d \in (1/4, 1/2)$ satisfy $\tau(C_d) = d$. Set $x = (d + \int_0^1 h(t)\,\mathrm{d}t)^{-1}$. For $x < 1/d$, let $C \in (0, \bar{C}(d))$ be the unique solution of $D_C(Y_d(C)) = (1+x)^{-1}$; for $x = 1/d$, set $C = 0$ and $P_x \equiv 1/d$. With the reciprocal state p, reference curve P_x, and logarithmic difference v from* (16)*,* (22)*, and* (29)*, respectively,*

$$\begin{aligned} S(C_d) - \mathcal{J}_d(h) = &\int_{\min\{C,C_d\}}^{\max\{C,C_d\}} \frac{(Y_d(u)+U_d(u))|\mathcal{B}_d(u)|}{(1-uY_d(u))^2}\,\mathrm{d}u \\ &+ \int_0^\infty \left[a(v')^2 + \frac{e^{-2v}-1+2v}{P_x^2} - 2Q_x v\right]\mathrm{d}a \;\geq 0. \end{aligned} \tag{G}$$

Consequently,

$$\max_{0\leq h\leq 1} \mathcal{J}_d(h) = S(C_d). \tag{33}$$

The maximizing control is unique up to equality almost everywhere. It has a zero arc of length C_d, a strictly increasing arc of length $(1-2C_d)/(1-C_d)$, and a one arc of length $C_d^2/(1-C_d)$. On the middle arc,

$$t = C_d(1+y), \qquad g(t) = d\sqrt{D_{C_d}(y)}\,e^{\mathcal{I}_{C_d}(y)/2}, \qquad h(t) = \frac{y\,g(t)}{D_{C_d}(y)}, \qquad 0 \leq y \leq y_e(C_d). \tag{34}$$

Proof. The two gaps give the global bound. Attainment requires reconstructing original time; the time constraint then gives uniqueness.

The global bound. At $C = C_d$, the first formula in (G2) reduces to

$$1 + \Phi_d(x(C_d)) = C_d(2 + \mathcal{I}(C_d)) = S(C_d).$$

Lemma 5.1 and Lemma 7.1 now give (G), with four nonnegative contributions. Equality requires $C = C_d$ and $p = P_{x(C_d)}$.

Attainment. So far a reference curve is only a bound: at a general endpoint it may need more than the available unit of original time. We now take $C = C_d$ and $x = x(C_d)$, and reconstruct original time from the reference curve by $\mathrm{d}t = -P_x^{-2}\,\mathrm{d}a$. On the middle arc, (22) gives $\mathrm{d}t = C\,\mathrm{d}y$, and (31) at $y = y_e(C)$ becomes

$$Cy_e(C) + \frac{CD_C(y_e(C))}{1-Cy_e(C)} = 1 - C. \tag{35}$$

Thus the middle and one arcs occupy time $1 - C$. We append a zero arc of length C at the initial state $g = d$. The resulting control is feasible on $[0, 1]$, and substituting (22) gives (34). Its one-arc length is $CD_C(y_e(C))/(1 - Cy_e(C)) = C^2/(1 - C)$. The derivative $P_x'' < 0$ in reversed time gives strict increase of the control on the middle arc.

Uniqueness. The reference curve P_x at $C = C_d$ first reaches $1/d$ at $a_a = C/d^2$. Any maximizing control must induce this reciprocal state up to some tail-time horizon $a_0 \geq a_a$; its time constraint gives

$$1 = \int_0^{a_0} P_x^{-2} \, \mathrm{d}a = 1 - C + d^2(a_0 - a_a), \qquad a_0 = \frac{2C}{d^2}. \tag{36}$$

The maximizing control is therefore fixed by

$$t(a) = \int_a^{a_0} P_x(u)^{-2} \, \mathrm{d}u, \qquad h(t(a)) = P_x'(a).$$

□

Proposition 8.1 (No exponential subarc)**.** *For $d > 0$, let h be the maximizing control of Theorem C, represented by* (34) *on its middle arc. There is no nonempty open subinterval of that arc on which $h(t) = \lambda_0 + \lambda_1 e^{\lambda_2 t}$ for real constants $\lambda_0, \lambda_1, \lambda_2$ with $h' > 0$.*

Proof. We take $C = C_d$, so that $d = \tau(C)$. An exponential subarc with $h' > 0$ would make h''/h' constant. From (34),

$$\frac{h''}{h'} = \frac{2 - 3Cy}{CD_C(y)}, \qquad \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{h''}{h'}\right) = \frac{2 - 3C - 4Cy + 3C^2y^2}{C^2 D_C(y)^2}. \tag{37}$$

The second numerator is a nonzero quadratic. Since y varies over every interior subinterval, h''/h' cannot be constant there. □

9 Prices and bounded hard pairs

Use the buyer survival function H and tail integral L from Section 3. The tail integral is absolutely continuous with derivative $-H$ almost everywhere. The following density is [LRW23, Theorem 3.1] in cutoff normalization. We give the direct argument to fix the treatment of arbitrary buyer atoms.

Lemma 9.1 (pricing from the variational bound)**.** *Let $0 < \beta < 1$ and $d > 0$. If*

$$\beta \sup_{0 \leq h \leq 1} \mathcal{J}_d(h) \leq 1, \tag{38}$$

then every independent pair of finite-mean distributions admits a price with

$$\Gamma(z) \geq \beta G - \beta d \cdot M. \tag{39}$$

Proof. The cutoff fixes the target conditional GFT; the price density realizes that target after averaging over the seller.

Cutoff normalization. We choose the cutoff where the target conditional GFT becomes zero. If $\mathbb{E}V_b = 0$, the right side of (39) is nonpositive. Otherwise the continuous function $d \cdot s - L(s)$ is strictly increasing from a negative value to infinity, so it has a unique positive zero. We rescale values by this zero. On $[0, 1]$, the normalized tail functions satisfy

$$L(s) = d + \int_s^1 H(u) \, \mathrm{d}u, \qquad K(s) = \int_s^1 H(u)^2 \, \mathrm{d}u.$$

The price density. We choose a random price with density

$$q(s) = \beta \left[\frac{H(s)}{L(s)} + \frac{d^2 - K(s)}{L(s)^2} \right], \qquad 0 \le s \le 1, \tag{40}$$

and put any remaining probability at zero. Since H is nonincreasing, $K(s) \le H(s)(L(s) - d)$; the density is nonnegative. Its integral is $\beta \mathcal{J}_d(H(1-\cdot)) \le 1$.

The absolutely continuous function $(d^2 - K)/L$ has derivative Hq/β, giving

$$L(s)q(s) + \int_s^1 H(z)q(z)\,\mathrm{d}z = \beta H(s) + \beta d \quad \text{a.e.} \tag{41}$$

The conditional guarantee. For a seller of value $s \le 1$, the conditional GFT is at least

$$\gamma(s) = \int_s^1 [L(z) + H(z)(z - s)]q(z)\,\mathrm{d}z.$$

The lower bound includes every buyer with value strictly above the price; buyers at an atom equal to the price can only add GFT. By (41), $\gamma'(s) = -\beta H(s) - \beta d$ almost everywhere, and $\gamma(1) = 0$. Hence $\gamma(s) = \beta L(s) - \beta d \cdot s$ for all $s \in [0, 1]$. For $s > 1$, this target is nonpositive because $L(s) \le d$. Independence allows us to average the conditional bound over the seller value:

$$\int_0^1 q(z)\Gamma(z)\,\mathrm{d}z \ge \beta \mathbb{E} L(V_s) - \beta d\, \mathbb{E} V_s = \beta G - \beta d \cdot M.$$

□

The seller CDF below is the equalizing construction of [LRW23, Theorem 3.5]. We use it with an explicit finite buyer tail and compute the endpoint atoms and residual gap needed for the sharp frontier. Figure 2 shows the reference curve and the associated hard pair.

Lemma 9.2 (bounded hard pairs)**.** *Let $H : [0, 1] \to [0, 1]$ be continuous and nonincreasing, with $H(1) = \eta > 0$, and let $d > 0$. Extend H to a buyer survival function by keeping it equal to η on $[1, R_\eta)$, where $R_\eta = 1 + d/\eta$, and setting it to zero at and above R_η. Include an atom of size $1 - H(0)$ at zero if needed. Using the tail integral L, define the auxiliary integral T and seller CDF F_s by*

$$L(s) = d + \int_s^1 H(u)\,\mathrm{d}u, \qquad T(s) = \int_0^s L(u)^{-2}\,\mathrm{d}u, \qquad F_s(s) = d\big(L(s)^{-1} - H(s)T(s)\big) \quad (0 \le s < 1), \tag{42}$$

and complete F_s with an atom at one. This is a valid seller CDF. With $T_\eta = T(1)$ and $J_\eta = \mathcal{J}_d(H(1-\cdot))$, the pair satisfies

$$M_\eta = 1 - d^2 T_\eta, \qquad \Gamma_{\max,\eta} = d + \eta d^2 T_\eta, \qquad G_\eta = d + d \cdot J_\eta - d^3 T_\eta. \tag{43}$$

In particular, for every β,

$$\Gamma_{\max,\eta} - \beta G_\eta + \beta d \cdot M_\eta = d(1 - \beta J_\eta) + \eta d^2 T_\eta. \tag{44}$$

Proof. The seller distribution. In the Stieltjes differential of F_s, the two terms containing $\mathrm{d}s$ cancel, leaving

$$\mathrm{d}F_s = -d \cdot T\,\mathrm{d}H \ge 0 \qquad (0 < s < 1).$$

At zero, the proposed seller law has mass $d/L(0)$. At the right endpoint, $F_s(1-) = 1 - d \cdot \eta T_\eta$, so the completion adds mass $d \cdot \eta T_\eta$. These are valid probabilities: $L(s) \geq d + \eta(1-s)$ gives $d \cdot \eta T_\eta \leq \eta/(d+\eta) < 1$.

GFT at a price. The identity $(d \cdot LT)' = F_s$ implies

$$\int_0^z F_s(s)\,\mathrm{d}s = d \cdot L(z)T(z), \qquad F_s(z)L(z) + H(z)\int_0^z F_s(s)\,\mathrm{d}s = d \quad (0 \leq z < 1). \tag{45}$$

The left side is the GFT at price z, by splitting $V_b - V_s = (V_b - z) + (z - V_s)$ on the trade event and using independence. For prices $1 \leq z \leq R_\eta$, all sellers trade with the remote buyer atom, and GFT equals $d + \eta(1 - M_\eta)$. The inclusive convention includes $z = R_\eta$. Prices above R_η give no trade. Integrating (45) gives the first two identities in (43).

First-best GFT. We integrate by parts:

$$G_\eta = \mathbb{E}L(V_s) = d + \int_0^1 H(s)F_s(s)\,\mathrm{d}s.$$

Fubini identifies

$$\int_0^1 H(s)^2 T(s)\,\mathrm{d}s = \int_0^1 \frac{K(s)}{L(s)^2}\,\mathrm{d}s.$$

Since the two iterated integrals agree, substitution in (42) gives the third identity in (43). Subtracting yields

$$\Gamma_{\max,\eta} - \beta G_\eta + \beta d \cdot M_\eta = d(1 - \beta J_\eta) + \eta d^2 T_\eta.$$

□

For the rest of the paper, let h_C be the maximizing control of Theorem C at $d = \tau(C)$, and put

$$H_C(s) = h_C(1-s), \qquad H_{\eta,C}(s) = \eta + (1-\eta)H_C(s) \quad (0 \leq s \leq 1). \tag{46}$$

We use Lemma 9.2 with this survival function. Every member has bounded supports; only the family has an escaping buyer atom.

10 The exact welfare guarantee

Theorem A (Exact welfare guarantee)**.** *Let V_s, V_b be independent nonnegative values with $0 < \mathbb{E}\max\{V_s, V_b\} < \infty$. A fixed price trades when $V_s \leq z \leq V_b$. Its optimal universal welfare guarantee is*

$$\inf_{F_s,F_b} \max_{z \geq 0} \frac{\mathbb{E}V_s + \mathbb{E}[(V_b - V_s)\mathbf{1}_{\{V_s \leq z \leq V_b\}}]}{\mathbb{E}\max\{V_s, V_b\}} = \beta_* \approx 0.73802.$$

The constant is $\beta_* = [C_*(2 + \mathcal{I}(C_*))]^{-1}$, where C_* is the unique root in $(1/4, 1/2)$ of

$$\begin{aligned} C(2 + \mathcal{I}(C)) &= 1 + \frac{C^2}{1-2C}e^{-\mathcal{I}(C)/2}, \\ \mathcal{I}(C) &= \frac{1}{\sqrt{C - 1/4}}\left[\arctan\frac{1-3C}{2(1-C)\sqrt{C-1/4}} + \arctan\frac{1}{2\sqrt{C-1/4}}\right]. \end{aligned} \tag{47}$$

Every pair has ratio strictly greater than β_*. For every $\varepsilon > 0$, the explicit bounded family in Lemma 9.2 contains a pair with ratio below $\beta_* + \varepsilon$.

The special functions in (47) come from integrating the reciprocal quadratic $D_C(y) = 1 - y + Cy^2$ along the middle arc and then reconstructing the original state: the integral gives arctangents, and reconstruction gives its exponential. The quadratic acquires a double root at $C = 1/4$; the middle arc shrinks to zero length as $C \uparrow 1/2$.

In the notation of Section 3, as noted after Theorem B, the root equation (47) reads

$$S(C_*) = 1 + \tau(C_*), \tag{48}$$

and the constant is

$$\beta_* = \frac{1}{S(C_*)} = \frac{1}{1 + \tau(C_*)}. \tag{49}$$

Numerical values. Interval arithmetic gives $0.73802433573 < \beta_* < 0.73802433574$ from (3)–(5) and the root equation (48). The corresponding parameter is $C_* \approx 0.3937571625$, with $7/20 < \tau(C_*) < 9/25$. At $\eta = 10^{-12}$, the independent pair (46) has optimal fixed-price welfare ratio $r_\eta < 0.738024335797$ and support contained in $[0, R_\eta]$, where $R_\eta < 360000000001$, by Proposition C.1.

Proof of Theorem A. The variational maximum gives the guarantee and the limiting hard pairs give sharpness. The zero arc then rules out attainment by a finite-mean pair.

The guarantee. By (7)–(8), S strictly decreases from infinity to one while τ strictly increases from zero to infinity, so (48) determines a unique C_*. Theorem C gives $\beta_* \sup_h \mathcal{J}_{\tau(C_*)}(h) = 1$. Lemma 9.1 and $\beta_* \tau(C_*) = 1 - \beta_*$ imply

$$M + \Gamma_{\max} \geq \beta_*(M + G).$$

Sharpness. We fix $C = C_*$ and let $\eta \downarrow 0$ in (46). Uniform convergence of the controls and the bound $g \geq d$ give $J_\eta \to S(C_*) = 1/\beta_*$ and convergence of T_η. Equation (44) becomes

$$(M_\eta + \Gamma_{\max,\eta}) - \beta_*(M_\eta + G_\eta) = d(1 - \beta_* J_\eta) + \eta d^2 T_\eta \longrightarrow 0.$$

The denominator $M_\eta + G_\eta$ is at least the buyer mean, which is at least $d > 0$. Consequently

$$r_{\mathrm{FP}} \leq \lim_{\eta \downarrow 0} \frac{M_\eta + \Gamma_{\max,\eta}}{M_\eta + G_\eta} = \beta_* \leq r_{\mathrm{FP}}.$$

Nonattainment. First suppose $G > 0$. In the cutoff normalization of Lemma 9.1, $L(1) = d > 0$, so $H(1) > 0$. The reversed control is bounded below by $H(1)$; it cannot equal the control in Theorem C, which has an initial zero arc. Thus the density in (40) has mass

$$\alpha = \beta_* \mathcal{J}_d(H(1 - \cdot)) < 1, \qquad \alpha \Gamma_{\max} \geq \beta_* G - (1 - \beta_*) M. \tag{50}$$

Since $G > 0$, the event $\{V_s < V_b\}$ has positive probability; it is the union over rational q of $\{V_s < q < V_b\}$, so some rational q has $\Gamma(q) > 0$, and $\Gamma_{\max} > 0$. We obtain

$$M + \Gamma_{\max} - \beta_*(M + G) \geq (1 - \alpha)\Gamma_{\max} > 0.$$

If $G = 0$, the welfare ratio is one. In all cases,

$$\frac{M + \Gamma_{\max}}{M + G} > \beta_*.$$

□

11 Proof of the gains-from-trade frontier

For the derivative calculations, set $\sigma = S'(C)$, with $C \in (1/4, 1/2)$ as in Theorem B.

Proof of Theorem B. The affine guarantee and its sharpness determine the frontier. Convexity then identifies the tangent giving each conditional guarantee.

Affine guarantee and sharpness. We fix β and choose the unique C with $S(C) = 1/\beta$. Put $d = \tau(C)$. Since Theorem C makes (38) an equality, Lemma 9.1 gives the guarantee with $\delta = \beta d$.

For sharpness, apply (44) to (46). Its right side tends to zero. The limiting seller mean is strictly positive; more precisely, (36) gives

$$T_* = 2C/d^2, \qquad M_* = 1 - 2C > 0. \tag{51}$$

If $\tilde{\delta} < \beta d$, the expression with coefficient $\tilde{\delta}$ has limit

$$\Gamma_{\max,\eta} - \beta G_\eta + \tilde{\delta} M_\eta \longrightarrow -(\beta d - \tilde{\delta})(1 - 2C) < 0.$$

Convexity. With u as in (12), formula (7) gives $\sigma < 0$ and $\tau'(C) = -d \cdot \sigma/u$. Differentiating the parametric curve gives

$$\frac{\mathrm{d}\beta}{\mathrm{d}C} = -\frac{\sigma}{S^2} > 0, \qquad \delta'(\beta) = d\left(1 + \frac{S}{u}\right) > 0, \qquad \frac{\mathrm{d}}{\mathrm{d}C}\delta'(\beta(C)) = \frac{dS(2-\sigma)}{u^2} > 0. \tag{52}$$

As $C \downarrow 1/4$, d decays exponentially in $\mathcal{I}$ while S grows linearly, so the middle expression tends to zero. As $C \uparrow 1/2$, both d and S/u tend to infinity. Thus δ' increases from zero to infinity.

Conditional guarantee. Optimizing the affine bound gives

$$\frac{\Gamma_{\max}}{G} \geq \sup_{0<\beta<1} \left\{ \beta - \frac{\delta(\beta)}{\kappa} \right\}.$$

By (52), the maximizing point is unique and satisfies $\delta'(\beta) = \kappa$. Equations (52) and (11) identify this point with (12); substitution gives $1/(S+u)$. These derivatives also show that $\kappa(C)$ is a bijection.

For the matching sequence, (43) and (51) give the limits

$$M_* = u, \qquad \Gamma_{\max,*} = d, \qquad G_* = d(S+u). \tag{53}$$

For the pairs in (46), the functions $\kappa_\eta(C) = G_{\eta,C}/M_{\eta,C}$ are continuous and converge locally uniformly to $\kappa(C)$ as $\eta \downarrow 0$: the controls in (34) vary continuously across their moving joins, and $\tau(C)$ and the limiting seller mean stay bounded away from zero on compact parameter intervals, so (43) gives uniform convergence of the ratios. Since κ is strictly increasing, we apply the intermediate value theorem on shrinking brackets around C to obtain $C_\eta \to C$ with

$$\frac{G_{\eta,C_\eta}}{M_{\eta,C_\eta}} = \kappa(C), \qquad \frac{\Gamma_{\max,\eta,C_\eta}}{G_{\eta,C_\eta}} \longrightarrow \frac{1}{S(C)+u}.$$

□

12 The dominant-strategy welfare ceiling

A mechanism specifies a measurable lottery over trade and transfers for every pair of reports. Truthfulness is dominant in expected utility. We require individual rationality relative to retaining the initial endowment and equal buyer-to-seller transfers in every realization. These requirements imply that there is no transfer without trade and that every realized trade price lies between the reported values at a truthful profile.

The structural fact is the randomized-price representation of Hagerty and Rogerson [HR87] and Čopič and Ponsatí [ČP16]. We use it in an almost-everywhere form, which accommodates weak-indifference conventions at price atoms: Proposition 12.2 shows that, for measurable mechanisms under an absolutely continuous joint prior on a common bounded report interval, mechanism welfare is at most best-price welfare. Smoothing (Lemma 12.1) makes this sufficient for the hard pairs.

Lemma 12.1 (smoothing bounded hard pairs)**.** *An independent pair supported on* $[0,\Lambda]$ *has independent absolutely continuous approximations with full support on* $(0,\Lambda)$ *for which first-best welfare converges and the limit superior of best fixed-price welfare is at most the original optimum.*

Proof. For an original value X, we take independent uniforms $\zeta_0,\zeta_1\in(0,1)$ and $\zeta_2\in[-1/2,1/2]$ and set

$$X_\varepsilon=\begin{cases}\Lambda\zeta_1, & \zeta_0<\varepsilon,\\ (1-2\varepsilon)X+\varepsilon\Lambda+\varepsilon\Lambda\zeta_2, & \zeta_0\geq\varepsilon.\end{cases} \tag{54}$$

For $0<\varepsilon<1/2$, this variable has a full-support density on $(0,\Lambda)$ and converges almost surely to X. Using independent couplings for the two traders preserves independence, and bounded convergence gives convergence of first-best welfare.

The realized price welfare

$$w(s,b,z)=s+(b-s)_+\mathbf{1}_{\{s\leq z\leq b\}}$$

is bounded and upper semicontinuous on the compact cube. We choose approximately optimal prices along a sequence realizing the limit superior and pass to a convergent subsequence $z_\varepsilon\to z$. Reverse Fatou gives

$$\limsup_{\varepsilon\downarrow 0}\max_{z'}\mathbb{E}w(V_{s,\varepsilon},V_{b,\varepsilon},z')\leq\mathbb{E}w(V_s,V_b,z)\leq\max_{z'}\mathbb{E}w(V_s,V_b,z').$$

□

Proposition 12.2 (Price representation under absolutely continuous priors)**.** *Fix a common report interval* $(0,\Lambda)$*, where* $\Lambda>0$*. For every measurable, risk-neutral DSIC mechanism with realizationwise individual rationality and strong budget balance, there is a positive Borel measure* ν *on* $(0,\Lambda)$ *of mass* $m\leq 1$ *such that its allocation probability* φ *and expected transfer* Π *satisfy*

$$\varphi(s,b)=\nu((s,b]),\qquad \Pi(s,b)=\int_{(s,b]}z\,\nu(\mathrm{d}z)\quad \textit{for Lebesgue-a.e. } 0<s<b<\Lambda. \tag{55}$$

Under any absolutely continuous joint prior on this report square, its expected welfare is at most that of the best deterministic price.

The proof is given in Appendix B.2.

Corollary A′ (exact dominant-strategy guarantee)**.** *For the independent-prior model of Section 3 and measurable mechanisms on the common nonnegative real report domain, suppose traders are risk neutral, truthfulness is dominant in expected lottery payoffs, and individual rationality and strong budget balance hold in every realization. The optimal universal welfare guarantee, allowing full prior information, is* β_**.*

Proof. Fixed prices belong to this class, so Theorem A gives a guarantee of β_*. For the reverse bound, we choose a bounded hard pair from Theorem A and apply Lemma 12.1. Every mechanism restricted to the open support interval satisfies Proposition 12.2, so on each smooth pair its welfare is at most the best price welfare. We first take the smoothing parameter to zero and then the hard-pair parameter to zero. These inequalities hold even for a mechanism selected separately for each prior pair, giving

$$\beta_* \le r_{\mathrm{DSIC}} \le r_{\mathrm{FP}} = \beta_*.$$

□

13 Extension to two units

13.1 Model and separation

We use the increasing submodular model of [GGdK$^+$19, GdK26].[2] The seller owns two units. The buyer's marginal values and the seller's marginal costs of relinquishing successive units satisfy $B_1 \ge B_2 > 0$ and $0 < S_1 \le S_2$. The two vectors are independent, with finite first moments; their coordinates may be dependent. A common price z trades successive units while both agents accept, using the inclusive convention of Section 3. Since the accepting units form an initial segment,

$$\begin{aligned} M_2 &= \mathbb{E}(S_1 + S_2), & O_2 &= \sum_{i=1}^{2} \mathbb{E}\max\{B_i, S_i\}, \\ \Gamma_2(z) &= \sum_{i=1}^{2} \mathbb{E}[(B_i - S_i)\mathbf{1}_{\{S_i \le z \le B_i\}}], & r_2^* &= \inf \frac{M_2 + \sup_{z \ge 0} \Gamma_2(z)}{O_2}. \end{aligned} \tag{56}$$

The symmetric subclass requires (B_1, B_2) and (S_2, S_1) to have the same law, so the valuation functions are identically distributed. We denote its infimum by $r^*_{2,\mathrm{sym}}$.

Theorem D (Two units are strictly harder)**.** *There are independent, finitely supported laws of strictly positive ordered buyer and seller vectors for which every common price has welfare ratio less than* 0.7290804*. There is also a symmetric such instance with ratio less than* 0.83693*. Consequently*

$$r_2^* < 0.7290804 < 0.738024 < \beta_*, \qquad r^*_{2,\mathrm{sym}} < 0.83693.$$

The general instance has 131 buyer types and 131 seller types. It has a common high buyer atom, a single second-buyer body atom, and two second-seller atoms. Its exact integer evaluation at all 264 valuation events proves the bound for every real price. Appendix E gives the construction parameters and identifies the rational data in the supplement. Together with Theorem A, this resolves the separation conjecture of [GdK26, Section 1] already for two units.

[2] Increasing and submodular over two units means $B_1 \ge B_2$ and $S_1 \le S_2$: marginal values decrease and marginal costs increase.

13.2 Tail normalization

Proposition 13.1 below shows that every instance can be replaced, without raising its ratio, by a limit of instances in which the buyer has a common atom of vanishing probability far above all other values. We normalize the atom's limiting first moment to one per unit and describe the rest of each side by a bounded vector, its body: Z for the buyer and Y for the seller. Let $Z = (Z_1, Z_2)$ and $Y = (Y_1, Y_2)$ be independent bounded nonnegative vectors with $Z_1 \geq Z_2$ and $Y_1 \leq Y_2$. Define

$$\begin{aligned} \Gamma_0(z) &= \sum_{i=1}^{2} \mathbb{E}\Big[\mathbf{1}_{\{Y_i<z\}}\left(1 + (Z_i - Y_i)\mathbf{1}_{\{z\leq Z_i\}}\right)\Big], \\ \mathcal{R}_2(Z,Y) &= \frac{2 + \mathbb{E}(Y_1 + Y_2)}{2 + \sum_{i=1}^{2} \mathbb{E}\max\{Z_i, Y_i\}}. \end{aligned} \tag{57}$$

The constant 1 in each gain term is the contribution of that atom. The normalized model uses strict seller acceptance and inclusive buyer acceptance, as induced by a positive seller displacement before taking the limit. We call (Z, Y) *admissible bodies* if $\Gamma_0(z) \leq 2$ for every price. Above the bodies, $\Gamma_0(z) = 2$.

Proposition 13.1 (A common escaping tail)**.** *For the admissible bodies and normalized ratio defined in* (57)*,*

$$r_2^* = \inf_{(Z,Y)\ \text{admissible}} \mathcal{R}_2(Z,Y).$$

Every admissible body pair is approached by independent, strictly positive, finite-mean ordered instances whose best common-price welfare ratios tend to $\mathcal{R}_2(Z,Y)$*.*

Appendix B.3 proves both directions, including the price constraint at the clipping atom.

13.3 The pricing functional

As in the single-unit argument, we first fix the buyers and look for the least total price mass that meets a target against every seller. With two units the target must hold for every ordered pair of seller costs $s_1 \leq s_2$. We describe a price by its cumulative mass $\varpi(s)$, the mass of prices above s, and measure the slack of unit i at seller cost s by a potential $\psi_i^{\varpi}(s)$. Feasibility is the joint condition $\psi_1^{\varpi}(s_1) + \psi_2^{\varpi}(s_2) \geq 0$ on $s_1 \leq s_2$. A nondecreasing compensation ϑ splits it into one inequality per unit, and for fixed ϑ the least feasible mass is the fixed point of a contraction.

We fix buyer bodies $Z_1 \geq_{\text{st}} Z_2$ supported on $[0, \bar{b}]$ and write $H_i(s) = \Pr(Z_i > s)$, so that the buyer Stieltjes measures are $-\mathrm{d}H_i$; for seller bodies we write $F_i(s) = \Pr(Y_i \leq s)$. For $s, z \geq 0$, define

$$\begin{aligned} \bar{L}_i(s) &= 1 + \mathbb{E}(Z_i - s)_+, \\ g_i(s,z) &= \mathbf{1}_{\{s<z\}}\left[1 + \mathbb{E}\big((Z_i - s)\mathbf{1}_{\{z\leq Z_i\}}\big)\right]. \end{aligned} \tag{58}$$

Thus $s + \bar{L}_i(s) = 1 + \mathbb{E}\max\{Z_i, s\}$ is normalized efficient welfare. A nonnegative price measure π gives conditional gain $\int g_i(s,z)\,\pi(\mathrm{d}z)$, and we minimize its total mass; any unused probability is placed at zero.

Fix $d > 0$; as with one unit, the target charges βd per unit of seller cost, and $d = (1-\beta)/\beta$ is the welfare case (Theorem E(iii)). A normalized price consists of a body measure $\widehat{\pi}_0$ on $(0, \bar{b}]$ and an ideal tail mass $\varpi(\bar{b})$, which contributes one unit of gain per unit of mass to every seller in $[0, \bar{b}]$. Its cumulative mass

$$\varpi(s) = \varpi(\bar{b}) + \widehat{\pi}_0((s, \bar{b}])$$

is nonnegative, nonincreasing, and right-continuous. The gain and the seller potential of unit i are

$$\begin{aligned} G_i^{\varpi}(s) &= \bar{L}_i(s)\varpi(s) - \int_{(s,\bar{b}]} (t-s)\varpi(t)\,(-\mathrm{d}H_i(t)), \\ \psi_i^{\varpi}(s) &= d\cdot s - \bar{L}_i(s) + G_i^{\varpi}(s). \end{aligned} \tag{59}$$

The atom at s is excluded from $\varpi(s)$; with this trace convention a buyer atom at the posted price still trades.

The order between the two seller costs enters through a compensation. Let $\mathcal{V}$ consist of bounded, nondecreasing, right-continuous compensations $\vartheta : [0,\bar{b}] \to \mathbb{R}$. If $\psi_1^{\varpi} \geq -\vartheta$ and $\psi_2^{\varpi} \geq \vartheta$, then $\psi_1^{\varpi}(s_1) + \psi_2^{\varpi}(s_2) \geq \vartheta(s_2) - \vartheta(s_1) \geq 0$ whenever $s_1 \leq s_2$, because ϑ is nondecreasing. Conversely, a nonnegative potential sum permits the canonical compensation $\Theta_{\varpi}(s) = -\inf_{t\leq s} \psi_1^{\varpi}(t)$. For fixed ϑ the two inequalities become a fixed-point problem for ϖ: with $\epsilon_1 = -1$ and $\epsilon_2 = 1$, dividing $\psi_i^{\varpi} \geq \epsilon_i\vartheta$ by $\bar{L}_i$ and rearranging gives $\varpi \geq \mathcal{A}_{i,d,\vartheta}\varpi$. The operator below also takes a supremum over future seller costs, which keeps the cumulative mass nonincreasing. On bounded Borel functions ϖ, define

$$\begin{aligned} (\mathcal{T}_{d,\vartheta}\varpi)(s) &= \sup_{s\leq t\leq \bar{b}} \max\{0, \mathcal{A}_{1,d,\vartheta}\varpi(t), \mathcal{A}_{2,d,\vartheta}\varpi(t)\}, \\ \mathcal{A}_{i,d,\vartheta}\varpi(t) &= 1 - \frac{d\cdot t}{\bar{L}_i(t)} + \frac{\epsilon_i\vartheta(t)}{\bar{L}_i(t)} + \frac{\int_{(t,\bar{b}]}(v-t)\varpi(v)\,(-\mathrm{d}H_i(v))}{\bar{L}_i(t)}. \end{aligned} \tag{60}$$

Theorem E (Pricing for ordered buyer pairs)**.** *For every ordered buyer pair supported on* $[0,\bar{b}]$ *and every* $d > 0$*:*

(i) For each $\vartheta \in \mathcal{V}$*, the operator* $\mathcal{T}_{d,\vartheta}$ *has a unique bounded fixed point* $\varpi_{d,\vartheta}$*, the pointwise least feasible cumulative mass. Iteration from zero converges uniformly with contraction factor* $\bar{b}/(1+\bar{b})$*.*

(ii) The value $\varpi_{d,\vartheta}(0)$ *is jointly convex in* (d,ϑ)*, and the minimum*

$$\mathcal{J}_d^{(2)}(H_1,H_2) = \min_{\vartheta\in\mathcal{V}} \varpi_{d,\vartheta}(0) \tag{61}$$

is attained.

(iii) For $0 < \beta < 1$*, there is a price probability with bounded support satisfying*

$$\sum_{i=1}^{2} \int g_i(s_i,z)\,\pi(\mathrm{d}z) \geq \beta \sum_{i=1}^{2} \bar{L}_i(s_i) - \beta d(s_1+s_2) \quad (0\leq s_1\leq s_2) \tag{62}$$

if and only if $\beta\mathcal{J}_d^{(2)}(H_1,H_2) \leq 1$*. At* $d = (1-\beta)/\beta$*, this is a welfare guarantee of* β *against every ordered seller distribution in the normalized model.*

(iv) If H_i^{δ} *is the survival function of* $\delta\lfloor Z_i/\delta \rfloor$*, then*

$$\mathcal{J}_d^{(2)}(H_1,H_2) \leq \mathcal{J}_d^{(2)}(H_1^{\delta},H_2^{\delta}) + \delta.$$

The supremum over bounded ordered buyer bodies consequently equals the supremum over finite-support ordered buyer bodies.

The tail is realized by density βd on $(\bar{b}, \bar{b} + \varpi_{d,\vartheta}(\bar{b})/d)$, an empty interval when the tail mass is zero. This measure need not be finitely atomic, and it is not the only such price: the same tail mass may instead be placed as a single atom at $\bar{b} + \varpi_{d,\vartheta}(\bar{b})/d$. Appendix B.4 proves the theorem and gives the finite-node recurrence. This extends the buyer-dependent pricing step of [LRW23, Theorem 3.1]; the new ingredient is the prefix compensation.

13.4 The structured family

All quantities carrying the subscript f belong to the family defined here. The family describes limiting two-unit instances, constructed in Appendix D: both sellers have the same mass at zero, the second buyer's body is a single value, the second seller has two atoms, and the first buyer is described by one concave curve $\mathcal{P}$. We use t_{f} and p_{f} for its scalar parameters, ξ for its curve coordinate, and $\mathcal{P}$ for its reciprocal state. They are distinct from the single-unit coordinates in the main text. For $0 < t_{\mathrm{f}} < 1$, $t_{\mathrm{f}} \le p_{\mathrm{f}} \le 1$, and $\xi_0 > 0$, define

$$\begin{aligned} m_{\mathrm{f}} &= \frac{2t_{\mathrm{f}}}{1+t_{\mathrm{f}}}, \qquad N_{\mathrm{f}} = \frac{2}{1+t_{\mathrm{f}}}, \qquad c_{\mathrm{f}} = \frac{p_{\mathrm{f}}-t_{\mathrm{f}}}{2}, \qquad \delta_{\mathrm{f}} = \frac{p_{\mathrm{f}}+t_{\mathrm{f}}}{2}, \\ e_{\mathrm{f}} &= \frac{1+t_{\mathrm{f}}}{2}, \qquad v_{\mathrm{f}} = 1-e_{\mathrm{f}}, \qquad a_{\mathrm{f}} = \frac{c_{\mathrm{f}}}{t_{\mathrm{f}}p_{\mathrm{f}}}, \qquad h_{\mathrm{f}} = \frac{v_{\mathrm{f}}}{\xi_0}. \end{aligned} \tag{63}$$

We require $\xi_0 \ge c_{\mathrm{f}}$. The closed curve class $\mathfrak{F}$ consists of positive, concave Lipschitz functions on $[c_{\mathrm{f}}, \xi_0]$ satisfying

$$\mathcal{P}(c_{\mathrm{f}}) = p_{\mathrm{f}}, \qquad \mathcal{P}(\xi_0) = 1, \qquad \mathcal{P}(\xi) \le U_{\mathrm{f}}(\xi) := \min\{\xi + \delta_{\mathrm{f}}, h_{\mathrm{f}}\xi + e_{\mathrm{f}}\}. \tag{64}$$

A parameter triple $(t_{\mathrm{f}}, p_{\mathrm{f}}, \xi_0)$ is feasible when this class is nonempty. For a nondegenerate interval, concavity and the two endpoint obstacles imply $h_{\mathrm{f}} \le \mathcal{P}' \le 1$ almost everywhere. The case $p_{\mathrm{f}} = 1$ is included as a degenerate interval by continuity. We distinguish the subclass whose endpoint slopes are 1 and h_{f}; the closed class also admits endpoint atoms and corners. Its functionals are

$$\begin{aligned} T_{\mathrm{f}} &= \int_{c_{\mathrm{f}}}^{\xi_0} \mathcal{P}^{-2}\,\mathrm{d}\xi, & Q_{\mathrm{f}} &= \int_{c_{\mathrm{f}}}^{\xi_0} \xi(\mathcal{P}'/\mathcal{P})^2\,\mathrm{d}\xi, \\ M_{\mathrm{f}} &= N_{\mathrm{f}}(a_{\mathrm{f}} + T_{\mathrm{f}} - \xi_0), & G_{\mathrm{f}} &= 2 + 2m_{\mathrm{f}}a_{\mathrm{f}} - N_{\mathrm{f}}\log p_{\mathrm{f}} - N_{\mathrm{f}}Q_{\mathrm{f}}, \\ R_{\mathrm{f}} &= \frac{M_{\mathrm{f}}+2}{M_{\mathrm{f}}+G_{\mathrm{f}}}. \end{aligned} \tag{65}$$

Both sellers have mass m_{f} at zero, so t_{f} fixes this common mass, and the second buyer's body is the single value a_{f}. The curve $\mathcal{P}$ is the reciprocal of the first buyer's normalized tail integral as a function of the coordinate ξ, and its right derivative is that buyer's survival function. The first buyer's body starts at a_{f}, where $\mathcal{P} = p_{\mathrm{f}}$, ends at the coordinate value ξ_0, and has length T_{f} in value units. The constant N_{f} scales the first seller's CDF on the body, and $c_{\mathrm{f}}, \delta_{\mathrm{f}}, e_{\mathrm{f}}, v_{\mathrm{f}}, h_{\mathrm{f}}$ locate the coordinate interval and the two affine obstacles in (64); the choice $h_{\mathrm{f}} = v_{\mathrm{f}}/\xi_0$ puts the second obstacle through $(\xi_0, 1)$. In these instances M_{f} and G_{f} are the limiting seller welfare and efficient gains, and R_{f} is the welfare ratio (Lemma D.1).

Proposition 13.2 (The minimum within the two-unit curve family)**.** *Let $\mathfrak{F}$ and R_{f} be defined by (63)–(65). The infimum $r_{\mathrm{fam}} = \inf_{\mathfrak{F}} R_{\mathrm{f}}$ is attained at a unique curve and parameter triple $(t_{\mathrm{f}}, p_{\mathrm{f}}, \xi_0)$. Its curve consists of a positive initial affine contact, one strictly concave interval satisfying $\mathcal{P}'' + C_{\mathrm{f}}\mathcal{P}/\xi^2 = 0$ for a constant $C_{\mathrm{f}} > 0$, and a positive final affine contact. The same infimum holds in the subclass with endpoint slopes $1, h_{\mathrm{f}}$, and*

$$\frac{18227}{25000} < r_{\mathrm{fam}} < \frac{729081}{10^6}.$$

The minimizing body laws are approached by strictly positive finite-mean ordered instances with a common escaping buyer atom.

The family value is approximately 0.729080. Appendix D proves the proposition and characterizes its value by the unique self-consistent physical stationary solution; the computations are in the supplementary material.

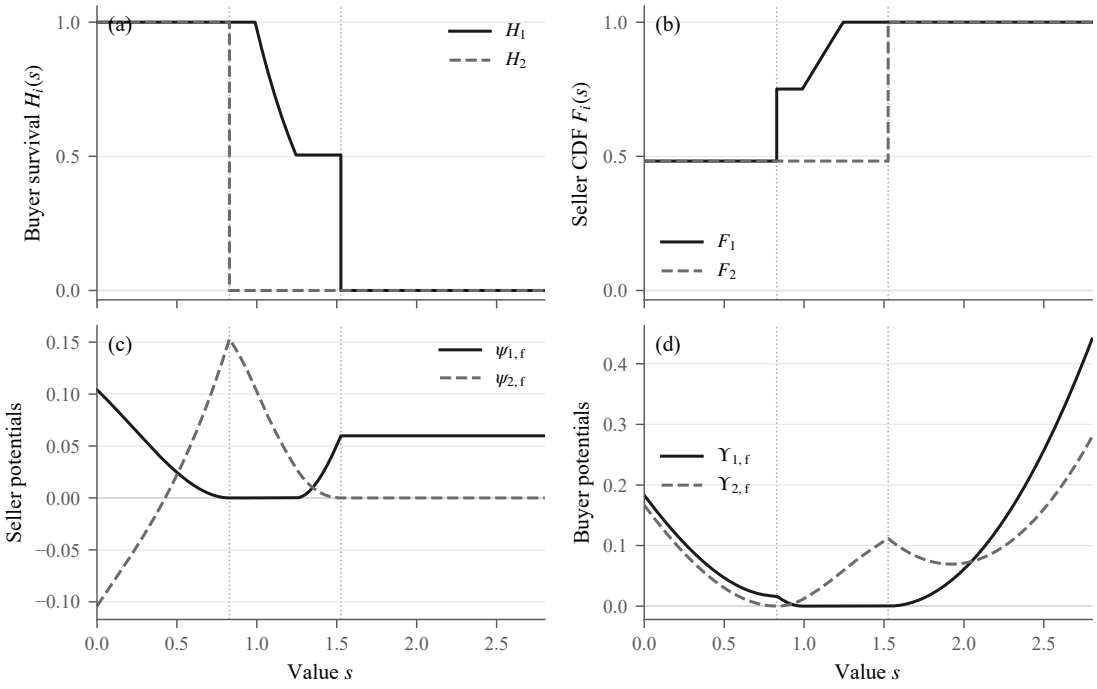


Figure 4: The two-unit family minimizer. The top panels show the buyer survival functions and seller CDFs. The lower panels show the seller and buyer measure potentials for the reference price of the one-sided certificate in the supplementary material. Solid and dashed curves refer to the first and second units. Their ordered sums are nonnegative. Vertical lines mark the second buyer's body atom and the common terminal body value.

13.5 The remaining comparison

Conjecture 13.3. *The optimal two-unit common-price welfare guarantee equals* $r_2^* = r_{\mathrm{fam}}$.

Theorem 13.4 (The conjectured value for a constant second-buyer body)**.** *Let* (Z, Y) *be admissible bodies whose second buyer's body* Z_2 *is almost surely constant. Then* $\mathcal{R}_2(Z, Y) \geq r_{\mathrm{fam}}$. *The family minimizer of Proposition 13.2 has a constant second buyer's body, so the bound is attained.*

Appendix D.5 proves the theorem.

Proposition 13.1 and Theorem E give the following sufficient condition for the matching lower bound:

$$\sup_{H_1 \geq H_2 \text{ bounded buyer bodies}} \mathcal{J}^{(2)}_{d_{\mathrm{fam}}}(H_1, H_2) \leq \frac{1}{r_{\mathrm{fam}}}, \qquad d_{\mathrm{fam}} = \frac{1 - r_{\mathrm{fam}}}{r_{\mathrm{fam}}}. \tag{66}$$

Finite buyer support suffices in this supremum by part (iv) of Theorem E. If this budget bound holds, averaging $\Gamma_0 \leq 2$ under the price in Theorem E gives $2 + \mathbb{E}(Y_1 + Y_2) \geq r_{\mathrm{fam}}[2 + \sum_i \mathbb{E} \max\{Z_i, Y_i\}]$ for every admissible pair. At the family minimizer, a certificate in the supplementary material checks this comparison from each side: in the limiting model with escaping buyer first moment one per unit, one explicit price distribution guarantees the ratio r_{fam} against every ordered seller pair when the buyers are held at the minimizer, and against every ordered buyer pair when the sellers are held there. The budget bound (66) asks for the first property at every bounded ordered buyer pair, with a price that may depend on the buyers. The identity $1 + d_{\mathrm{fam}} = 1/r_{\mathrm{fam}}$ is exactly the budget condition under which Appendix B.5 extends the finite-node constraints to every ordered seller pair.

Numerical evidence. A floating-point search over 20,000 finite-support ordered buyer pairs, broad random samples together with random perturbations of the discretized family minimizer, found no pair above the budget. The largest value, 1.3715902, came from a perturbed minimizer, against $1/r_{\mathrm{fam}} \approx 1.3715909$. This is search evidence for the conjecture and does not establish the supremum in (66); the sampling design is in Appendix B.5.

14 Related work

Welfare and fixed prices. Myerson and Satterthwaite [MS83] establish the general incompatibility of full efficiency, Bayesian incentive compatibility, voluntary participation, and no outside subsidy. Fixed-price approximation keeps strong budget balance and dominant strategies while relaxing efficiency. Blumrosen and Dobzinski [BD21] obtain a $1 - 1/e$ welfare guarantee. Kang, Pernice, and Vondrák [KPV22] improve that guarantee and determine the exact ratio $(2 + \sqrt{2})/4$ for identically distributed values. On the impossibility side, Colini-Baldeschi, de Keijzer, Leonardi, and Turchetta [CBdKLT16] constructed a pair on which no fixed price exceeds 0.7485; Kang and Vondrák [KV19] lowered this to 0.7385, the two STOC 2023 papers to 0.7381, and Giambartolomei and de Keijzer to 0.73805. The asymmetric problem leads to the programs of Cai and Wu [CW23], the variational approach of Liu, Ren, and Wang [LRW23], and the tighter numerical bounds of Giambartolomei and de Keijzer [GdK26].

The variational ingredients from [LRW23] and their role in our proof are described in Section 1. Cai and Wu's Lemma 3.3 already produces genuine hard instances from finite upper-program solutions. Our hard pairs approach the exact constant and every point of the frontier, including at an exactly prescribed moment ratio.

Gains from trade. Welfare includes the seller's initial value; GFT measures only the improvement over keeping the item. This distinction is essential for approximation. Fixed prices admit no uniform positive GFT ratio [BD21]. Colini-Baldeschi, Goldberg, de Keijzer, Leonardi, and Turchetta [CBGdK$^+$17] give asymptotically tight logarithmic guarantees in terms of efficient-trade probability. Theorem B instead determines an exact curve in the moment ratio G/M, including its leading logarithmic coefficient.

Allowing Bayesian incentives yields a different GFT theory. Blumrosen and Mizrahi [BM16] analyze offering mechanisms under distributional assumptions. Brustle, Cai, Wu, and Zhao [BCWZ17] obtain a constant approximation to the second-best GFT benchmark in two-sided markets. Deng, Mao, Sivan, and Wang [DMSW22] establish a constant approximation to first-best GFT for arbitrary independent bilateral priors, improved by Fei [Fei22]; Jo [Jo26] further sharpens the bounds for the random-offerer mechanism. Liu, Qin, Ren, and Wang [LQRW26] determine the exact second-best-to-first-best GFT ratio as 1/2, using Bayesian incentive compatibility and interim individual rationality. Their benchmark optimizes over feasible Bayesian mechanisms. Our frontier optimizes fixed prices and conditions on initial welfare, so these are different exact-value questions.

Information and larger markets. Babaioff, Goldner, and Gonczarowski [BGG20] study additional competition and sample pricing; Dütting, Fusco, Lazos, Leonardi, and Reiffenhäuser [DFL$^+$21] establish welfare guarantees from a single seller sample in broader two-sided markets. Cai and Wu [CW23] compare full and one-sided priors and analyze fixed numbers of samples. Deng, Mao, Sivan, Wang, and Wu [DMS$^+$25] study sample-based offering mechanisms for GFT. The one-sided-prior consequence of Theorem A uses the existing minimax equivalence; it does not supply a new sample bound.

McAfee's double auction [McA92] obtains efficiency as markets grow while preserving dominant strategies and avoiding a deficit. Babaioff, Cai, Gonczarowski, and Zhao [BCGZ18] combine asymptotic efficiency with a finite-market GFT guarantee. Gerstgrasser, Goldberg, de Keijzer, Lazos, and Skopalik [GGdK$^+$19] characterize multi-unit bilateral mechanisms and analyze their welfare; [GdK26] further improves the fixed-price bounds in that setting. With unrestricted correlated values, Dobzinski and Shaulker [DS24] show that DSIC mechanisms have no positive universal welfare ratio, while a mechanism with one side DSIC and the other Bayesian IC does. Their later reserve mechanism [DS26] exceeds the fixed-price welfare ceiling even for independent values.

Giambartolomei and de Keijzer [GdK26] report a two-unit upper bound of 0.7291, just below their single-unit lower bound of 0.7292, and conjecture strict separation from single-unit trade. The margin between these two numerically computed bounds is 10^{-4}; Theorem D widens it to 0.0089, with an exact rational instance below 0.7290804 on one side and the exact constant β_* on the other, and thereby proves the separation. Their marginal formulation retains the buyer and seller order constraints; Theorem E uses the seller order through its prefix compensation. Their symmetric two-unit bound is 0.8372.

Representation and variational comparison. The interpretation through dominant strategies uses the randomized-price representations of Hagerty and Rogerson [HR87] and Čopič and Ponsatí [ČP16]. The almost-everywhere representation and smoothing argument in Section 12 handle the measure-theoretic conventions needed by the hard pairs. The representation relies on strong budget balance. When only no deficit is required, dominant-strategy mechanisms need not be randomized prices. The generalized double auctions of Zhang [Zha26] trade when the bid–ask difference exceeds a random spread, and they maximize the worst-case gains from trade when values in $[0, 1]$ may be correlated and only the first-best GFT is known.

Calibration is a classical way to prove global variational optimality by an exact comparison, including in nonconvex problems [ABD03, BF16]. Contact multipliers and complementary slackness belong to the theory of obstacle problems and variational inequalities [KS00]. Here the specific construction is the reciprocal clock, a reference curve for every endpoint, and the endpoint envelope controlling the remaining nonconcavity. The conditional frontier uses ordinary convex conjugacy [Roc70]; related minimax methods [Sio58] enter the information equivalence of [CW23] and the second-best analysis of [LQRW26]. We evaluate the frontier and construct hard pairs at each slope.

15 Open questions

For randomized DSIC mechanisms, what is the exact welfare guarantee when realizationwise strong budget balance is weakened to no deficit? Beyond dominant strategies, the optimal Bayesian welfare guarantee under ex-post individual rationality remains to be determined [DS26]. Unrestricted correlation destroys a positive DSIC ratio [DS24]; a quantitative frontier under restricted dependence could describe how the independent guarantee deteriorates. For two units, Conjecture 13.3 asks for the candidate constant, with (66) giving a sufficient buyer-functional comparison; by Theorem 13.4, a counterexample needs a random second-buyer body. For $k \geq 3$, exact common-price constants and extremal structures remain open [GGdK$^+$19, GdK26]. Direct compression of the second buyer to its mean and Jensen comparisons for the optimized first-unit value fail on explicit instances, which are given in the supplementary material. Which other single-parameter mechanism families admit a state transformation that makes the adversary's objective concave after fixing finitely many boundary values?

AI methodology

Generative AI tools, including Codex and Claude, assisted with mathematical exploration, literature searches, and prose polishing. In particular, they were used to explore the reciprocal-state representation and calibration in Lemmas 5.1 and 7.1, the common-tail normalization in Proposition 13.1, the structured-family optimization in Proposition 13.2, and the rational comparison construction in Lemma D.7. They were also used to search for counterexamples to intermediate claims and to implement linear programming, symbolic algebra, exact rational arithmetic, interval arithmetic, and SMT checks. With AI assistance, the authors conceived the paper's core ideas and developed its methods. Appendix F describes the routes they tried for the exact single-unit constant and why each was kept or set aside. The authors also improved the paper's writing, readability, and economic interpretation. One of the authors verified the analytic derivations by hand. The computer-assisted parts are too large to be checked by hand: the interval and positivity certificates for the two-unit family optimization (Proposition 13.2 and Appendix D), the exact rational evaluation of the two-unit instances in Theorem D, including the 131-type instance (Appendix E), and the numerical enclosures in Section 10. These checks are carried out by deterministic software (exact rational arithmetic, interval arithmetic, and SMT solvers), not by language models. Although parts of the code were written with AI assistance, the checks themselves do not depend on any language model and can be rerun independently; the code, certificates, and a replay script are available at `https://github.com/lintingyee-del/fixed-price-bilateral-trade`.

Another author used Archon, an agent for Lean 4 proof development, to assist in formalizing the core results of the paper in Lean 4 with Mathlib. The Lean code is available at `https://github.com/lintingyee-del/fixed-price-bilateral-trade`; the supplementary material describes the formalization. The authors checked the correctness and originality of all content, including the references, and take full responsibility for the paper.

References

[ABD03] Giovanni Alberti, Guy Bouchitté, and Gianni Dal Maso. The calibration method for the Mumford–Shah functional and free-discontinuity problems. *Calculus of Variations and Partial Differential Equations*, 16(3):299–333, 2003. `https://cvgmt.sns.it/paper/1428/`.

[BCGZ18] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, page 373, 2018. doi:`10.1145/3219166.3219203`. Extended version: arXiv:1802.08023.

[BCWZ17] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, pages 589–590, 2017. arXiv:1706.04637.

[BD21] Liad Blumrosen and Shahar Dobzinski. (Almost) efficient mechanisms for bilateral trading. *Games and Economic Behavior*, 130:369–383, 2021. doi:`10.1016/j.geb.2021.08.011`.

[BF16] Guy Bouchitté and Ilaria Fragalà. A duality theory for non-convex problems in the calculus of variations. arXiv:1607.02878, 2016.

[BGG20] Moshe Babaioff, Kira Goldner, and Yannai A. Gonczarowski. Bulow–Klemperer-style results for welfare maximization in two-sided markets. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2452–2471, 2020. arXiv:1903.06696.

[BM16] Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In *Web and Internet Economics (WINE)*, pages 400–413, 2016. doi:10.1007/978-3-662-54110-4_28.

[CBdKLT16] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1424–1443, 2016.

[CBGdK+17] Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Fixed price approximability of the optimal gain from trade. In *Web and Internet Economics (WINE)*, pages 146–160, 2017. Extended version: arXiv:1710.08394.

[ČP16] Jernej Čopič and Clara Ponsatí. Optimal robust bilateral trade: risk neutrality. *Journal of Economic Theory*, 163:276–287, 2016. doi:10.1016/j.jet.2015.11.008.

[CW23] Yang Cai and Jinzhao Wu. On the optimal fixed-price mechanism in bilateral trade. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 737–750, 2023. doi:10.1145/3564246.3585171. Extended version: arXiv:2301.05167v2.

[DFL+21] Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 1452–1465, 2021. doi:10.1145/3406325.3451076.

[DMS+25] Yuan Deng, Jieming Mao, Balasubramanian Sivan, Kangning Wang, and Jinzhao Wu. Approximately efficient bilateral trade with samples. arXiv:2502.13122, 2025.

[DMSW22] Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Approximately efficient bilateral trade. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 718–721, 2022. doi:10.1145/3519935.3520054.

[DS24] Shahar Dobzinski and Ariel Shaulker. Bilateral trade with correlated values. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 237–246, 2024. doi:10.1145/3618260.3649659.

[DS26] Shahar Dobzinski and Ariel Shaulker. Welfare maximization in bilateral trade: improved approximation guarantees beyond the fixed price barrier. arXiv:2606.04890v1, 2026.

[Fei22] Yumou Fei. Improved approximation to first-best gains-from-trade. In *Web and Internet Economics (WINE)*, volume 13778 of *Lecture Notes in Computer Science*, pages 204–218, 2022. doi:10.1007/978-3-031-22832-2_12. Extended version: arXiv:2205.00140.

[GdK26] Giordano Giambartolomei and Bart de Keijzer. On the approximation ratio of optimal fixed-price mechanisms for single and multi-unit bilateral trade. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(20):16946–16953, 2026. doi:10.1609/aaai.v40i20.38741.

[GGdK+19] Matthias Gerstgrasser, Paul W. Goldberg, Bart de Keijzer, Philip Lazos, and Alexander Skopalik. Multi-unit bilateral trade. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):1973–1980, 2019. doi:10.1609/aaai.v33i01.33011973.

[HR87] Kathleen M. Hagerty and William P. Rogerson. Robust trading mechanisms. *Journal of Economic Theory*, 42(1):94–107, 1987. doi:10.1016/0022-0531(87)90104-9.

[JGC26] Tao Jiang, Minbo Gao, and Shaowei Cai. The exact approximation ratio of the optimal fixed-price mechanism in bilateral trade. arXiv:2609.27878v1, 2026.

[Jo26] Sunghyeon Jo. Tighter bounds for the random-offerer mechanism in bilateral trade. arXiv:2607.13959, 2026.

[KPV22] Zi Yang Kang, Francisco Pernice, and Jan Vondrák. Fixed-price approximations in bilateral trade. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2964–2985, 2022. arXiv:2107.14327.

[KS00] David Kinderlehrer and Guido Stampacchia. *An Introduction to Variational Inequalities and Their Applications*. SIAM, 2000. Reprint of the 1980 edition. doi:10.1137/1.9780898719451.

[KV19] Zi Yang Kang and Jan Vondrák. Fixed-price approximations to optimal efficiency in bilateral trade. SSRN working paper, ssrn.com/abstract=3460336, 2019. Earlier version: Strategy-proof approximations of optimal efficiency in bilateral trade, 2018.

[LQRW26] Zhengyang Liu, Ying Qin, Zeyu Ren, and Zihe Wang. Second-best bilateral trade is 1/2 efficient. arXiv:2606.03849, 2026.

[LRW23] Zhengyang Liu, Zeyu Ren, and Zihe Wang. Improved approximation ratios of fixed-price mechanisms in bilateral trades. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 751–760, 2023. doi:10.1145/3564246.3585160. Extended version: arXiv:2303.15711v1.

[McA92] R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, 1992. doi:10.1016/0022-0531(92)90091-U.

[MS83] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983. doi:10.1016/0022-0531(83)90048-0.

[Roc70] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970. doi:10.1515/9781400873173.

[Sio58] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958. https://msp.org/pjm/1958/8-1/pjm-v8-n1-p14-p.pdf.

[Zha26] Wanchang Zhang. Generalized double auctions: maximizing welfare guarantees in bilateral trade. *American Economic Review*, 2026. Forthcoming.

A Notation

Table 2 lists the symbols of the single-unit analysis by role; its continuation lists the two-unit symbols.

Symbol	Meaning	Definition
Model quantities		
V_s, V_b	Seller and buyer values.	Section 3
M, G	Initial seller welfare and first-best GFT.	(1)
$\Gamma(z), \Gamma_{\max}$	GFT at price z and its maximum over prices.	(1)–(2)
r_{FP}	Optimal universal fixed-price welfare ratio.	(2)
$W(z), W_{\mathrm{FB}}$	Price welfare $M + \Gamma(z)$ and first-best welfare $M + G$.	Section 3
Scalar curve and parameter roles		
C	Curvature parameter: $(1/4, 1/2)$ for the scalar curve, $(0, \bar{C}(d))$ for reference curves; the auxiliary endpoint-value test uses $C > 1/4$.	Sections 3, 6–7
$D_C(y), y_e(C)$	Quadratic denominator and middle-arc endpoint; y_e also extends to $0 < C < 1/2$ for comparison.	(3); Section 7
$\mathcal{I}_C(y), \mathcal{I}(C)$	Integral to a running endpoint y, and its value at $y_e(C)$.	(3)
$\tau(C), S(C)$	Initial-state and objective-value coordinates of the scalar curve.	(4)
d, C_d	Fixed input to $\mathcal{J}_d$, and the selected parameter satisfying $\tau(C_d) = d$.	(6)–(8)
C_*, β_*	Welfare parameter solving $S(C_*) = 1 + \tau(C_*)$, and $1/S(C_*)$.	Theorem A
Control and original states		
$h, g, k, \mathcal{J}_d$	Control, accumulated first state, accumulated square, and objective.	(6)
H, L, K	Buyer survival function, tail integral, and truncated squared-survival integral.	Before (6)
s, t	Normalized value and reversed time, with $t = 1 - s$.	Before (6)
Transformed coordinates		
$A(t), a, a_0$	Tail-time map, a point in tail time, and the horizon $A(0)$.	(16)
p, x	Reciprocal state $p(A(t)) = 1/g(t)$ and its endpoint $x = p(0)$.	(16)
ℓ, ℓ_x	Logarithms of p and the reference curve P_x.	Section 5
$\mathcal{F}_d[p]$	Energy equal to $\mathcal{J}_d(h) - 1$.	(E)
Reference curves at a fixed initial state		
$\chi(C, y), Y_d(C), \bar{C}(d)$	Contact function, contact parameter solving $\chi(C, Y_d(C)) = d$, and upper limit of the curvature range.	(19); Lemma 6.1
y, r	Running interior parameters; the contact parameter is $Y_d(C)$, often written Y_d.	(3), (22)
$Z_d(C), U_d(C)$	Auxiliary integrals for the endpoint derivative.	(21)
$x(C)$	Endpoint of the reference curve with curvature C.	(20)
a_e, a_a, P_e	Line/middle-arc contact, middle-arc/cap contact, and first contact height.	(20)
P_x, Q_x	Reference curve at endpoint x, and half its obstacle multiplier.	(22)–(24)
$\Phi_d(x), \mathcal{B}_d(C)$	Endpoint envelope and the factor in its endpoint derivative.	(29)
v	Logarithmic difference $\log(p/P_x)$.	(29)
Bounded hard pairs		
η, R_η	Remote buyer-atom mass and its value $1 + d/\eta$.	Lemma 9.2
T, T_η, J_η	Integral of L^{-2}, its value at one, and $\mathcal{J}_d(H(1-\cdot))$.	(42)–(43)
$h_C, H_C, H_{\eta,C}$	Maximizing control, its reversed survival function, and the survival function with finite tail mass.	(46)
GFT frontier		
$\beta, \delta(\beta)$	Coefficient of first-best GFT and the sharp seller-welfare penalty.	(9)
$\kappa, \rho(\kappa), \kappa(C)$	Ratio G/M, its conditional GFT guarantee, and its parameterization.	(10)–(12)
u, σ	Limiting seller mean $u = 1 - 2C$ and scalar derivative $\sigma = S'(C)$.	(12); Section 11

Table 2: Notation by role. Definition references locate the formulas; h is the control, p the reciprocal state, and P_x the reference curve.

Symbol	Meaning	Definition
Two units		
$B_i, S_i; Z_i, Y_i$	Ordered marginal values and costs; their normalized bodies.	Section 13.1–13.2
$M_2, O_2, \Gamma_2, r_2^*$	Initial and efficient welfare, common-price gain, and worst-case ratio.	(56)
$H_i, F_i, \bar{L}_i, \bar{b}$	Buyer survival, seller CDF, normalized tail integral, and body bound.	Section 13.3; (58)
$\pi, \varpi, \psi_i^{\varpi}$	Price measure, normalized cumulative mass, and seller potential.	(59)
$\vartheta, \Theta_{\varpi}, \mathcal{T}_{d,\vartheta}$	Compensation, canonical prefix compensation, and obstacle operator.	(60)
$\mathcal{J}_d^{(2)}$	Least normalized common-price mass for a fixed buyer pair.	(61)
$\mathfrak{F}, r_{\mathrm{fam}}, d_{\mathrm{fam}}$	Structured family, its minimum ratio, and associated penalty.	Section 13.4; (66)

Table 2: Notation by role (continued): two units.

B Omitted proofs

This appendix proves the statements whose proofs were deferred, in the order in which they appear: Corollary B′, Proposition 12.2, Proposition 13.1, and Theorem E.

B.1 Endpoint asymptotics

Proof of Corollary B′. Along the scalar curve, we write $d = \tau(C)$.

Small gains. As $C \downarrow 1/4$, formula (5) gives $\mathcal{I} = O((C-1/4)^{-1/2})$. Since $1/\beta = S = C(2+\mathcal{I})$, we have $\mathcal{I}/2 = 1/(2\beta C) - 1$ and $1/(2\beta C) = 2/\beta - (4C-1)/(2C\beta)$, so that $e^{-\mathcal{I}/2} = e\, e^{-2/\beta} \exp((4C-1)/(2C\beta))$. Because $C^2/(1-2C) \to 1/8$,

$$\frac{\tau(C)}{(e/8)e^{-2/\beta}} = \frac{8C^2}{1-2C} \exp\left(\frac{4C-1}{2C\beta}\right) \longrightarrow 1.$$

The exponent tends to zero because $(C-1/4)/\beta = (C-1/4)S(C) = O((C-1/4)^{1/2})$. Multiplication by β gives the first asymptotic. In the same limit, $u \to 1/2$, $\kappa = d(S+u)/u \sim (e/4)Se^{-2S}$, and $\rho = 1/(S+u) \sim 1/S$. Thus $\log(1/\kappa) \sim 2S$.

Large gains. We write $C = 1/2 - \varepsilon$. Then $y_e \sim 8\varepsilon$ and $D_C(y) = 1 + O(\varepsilon)$ uniformly on the shrinking integration interval. Hence

$$\mathcal{I} \sim 8\varepsilon, \qquad S = 1 + 2\varepsilon + o(\varepsilon), \qquad d \sim \frac{1}{8\varepsilon}, \qquad u = 2\varepsilon.$$

Equations (11), (12), and (13) give

$$\delta \sim \frac{1}{8\varepsilon} \sim \frac{1}{4(1-\beta)}, \qquad \kappa \sim \frac{1}{16\varepsilon^2}, \qquad 1 - \rho \sim 4\varepsilon \sim \frac{1}{\sqrt{\kappa}}.$$

□

B.2 Proof of the price representation

The representation is due to [HR87, ČP16]. We prove the almost-everywhere form stated in Proposition 12.2, because the welfare comparison on smoothed hard pairs uses exactly this form and pointwise statements depend on tie conventions at price atoms.

On the report triangle $\mathcal{T} = \{(s,b) : 0 < s < b < \Lambda\}$, partial derivatives below are understood in the sense of distributions.

Proof of Proposition 12.2. The envelope identity forces the mixed derivative of the allocation probability to vanish, so the probability is locally a sum of one function of each report; gluing the local forms gives the price measure.

The envelope identity. Realizationwise IR and equal transfers force zero transfer without trade and a price in $[s, b]$ with trade, so $\varphi = \Pi = 0$ for $s > b$ and both truthful utilities vanish on the diagonal. The buyer's DSIC inequalities make $\varphi(s, \cdot)$ nondecreasing. By the envelope theorem her truthful utility $b\varphi(s, b) - \Pi(s, b)$ is the integral of $\varphi(s, \cdot)$ from the diagonal, which is the first identity below; the seller's inequalities make $\varphi(\cdot, b)$ nonincreasing and give the second in the same way:

$$b\varphi(s, b) - \Pi(s, b) = \int_s^b \varphi(s, t)\,\mathrm{d}t, \tag{67}$$

$$\Pi(s, b) - s\varphi(s, b) = \int_s^b \varphi(t, b)\,\mathrm{d}t. \tag{68}$$

Adding gives the Hagerty–Rogerson envelope identity

$$(b - s)\varphi(s, b) = \int_s^b [\varphi(s, t) + \varphi(t, b)]\,\mathrm{d}t. \tag{69}$$

Vanishing mixed derivative. We apply $\partial_s\partial_b$ to (69) against tests compactly supported away from the diagonal. The two integrals on the right contribute $\partial_s\varphi - \partial_b\varphi$, which cancels the same terms on the left and leaves $(b - s)\partial_s\partial_b\varphi = 0$. Since $1/(b - s)$ is smooth on $\mathcal{T}$,

$$\partial_s\partial_b\varphi = 0 \quad \text{in } \mathcal{D}'(\mathcal{T}). \tag{70}$$

Local separation. On an open rectangle $\mathcal{U} \times \mathcal{V} \subset \mathcal{T}$, every zero-integral test has a compactly supported smooth primitive, so (70) annihilates products of two such tests. A distribution with this property is a sum of one function of each variable, obtained by averaging against fixed tests of integral one. There are therefore functions $F_{\mathcal{U},\mathcal{V}}$ and $H_{\mathcal{U},\mathcal{V}}$ with

$$\varphi(s, b) = F_{\mathcal{U},\mathcal{V}}(b) + H_{\mathcal{U},\mathcal{V}}(s) \quad \text{a.e. on } \mathcal{U} \times \mathcal{V}. \tag{71}$$

Gluing. Apply (71) on the countable rectangles $(0, q) \times (q, \Lambda)$, with rational $q \in (0, \Lambda)$ (Figure 5, left). On each rectangle, Fubini lifts the four exceptional evaluation sets to four-dimensional null sets. These rectangles cover all ordered quadruples, giving

$$\varphi(s, b) = \varphi(\underline{s}, b) + \varphi(s, \overline{b}) - \varphi(\underline{s}, \overline{b}) \tag{72}$$

for almost every $0 < \underline{s} < s < b < \overline{b} < \Lambda$. Fubini supplies a full-measure set of anchors $(\underline{s}, \overline{b})$ for which (72) holds for almost every pair between them. We choose nested anchors with $\underline{s}_n \downarrow 0$ and $\overline{b}_n \uparrow \Lambda$ (Figure 5, right), each from a positive-area rectangle near the endpoints. Outside one two-dimensional null set, (72) then holds for every sufficiently large n.

Coordinate monotonicity gives the limits

$$F_0(b) = \lim_{\underline{s}\downarrow 0} \varphi(\underline{s}, b), \qquad H_0(s) = \lim_{\overline{b}\uparrow\Lambda} \varphi(s, \overline{b}), \qquad m = \lim_n \varphi(\underline{s}_n, \overline{b}_n) \le 1. \tag{73}$$

The anchors eventually enclose every point, so $m = \sup_{\mathcal{T}} \varphi = F_0(\Lambda-) = H_0(0+)$. We set $F_1 = m - H_0$ and pass to the limit in (72):

$$\varphi(s, b) = F_0(b) - F_1(s) \quad \text{a.e. on } \mathcal{T}, \qquad F_1(0+) = 0. \tag{74}$$

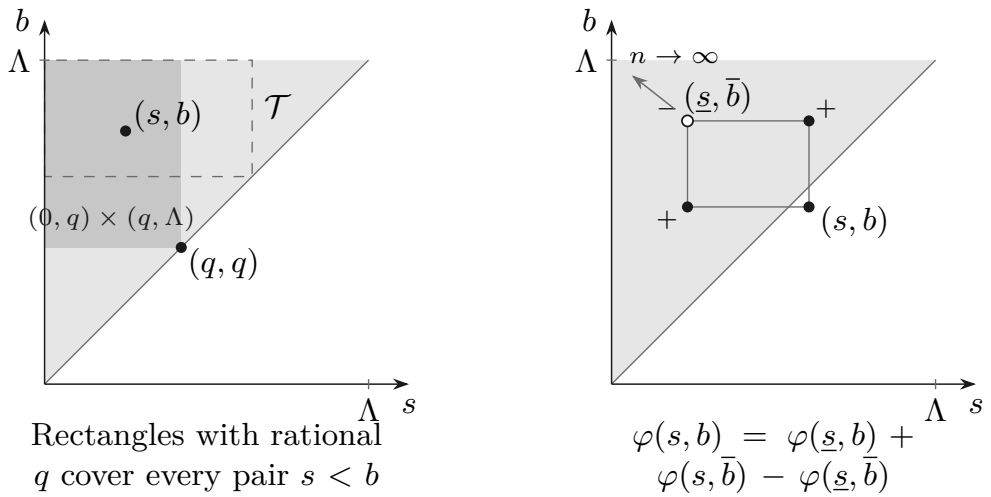


Figure 5: Gluing in the proof of Proposition 12.2. Left: the rectangles $(0, q) \times (q, \Lambda)$, one for each rational q, cover the report triangle $\mathcal{T}$. Right: identity (72) at a point (s, b) with anchor $(\underline{s}, \bar{b})$; the anchors move to $(0, \Lambda)$ along the sequence.

Both functions are nondecreasing and lie in $[0, m]$.

The price measure and transfer. Substituting (74) into both integrals of (69) gives

$$\int_s^b (F_0(t) - F_1(t))\,\mathrm{d}t = 0 \quad \text{for a.e. } (s, b) \in \mathcal{T}. \tag{75}$$

The primitive of $F_0 - F_1$ is continuous and takes equal values at almost every ordered pair, so it is constant and $F_0 = F_1$ almost everywhere. Their common right-continuous representative $\widehat{F}$ satisfies $\widehat{F}(0+) = 0$ and $\widehat{F}(\Lambda-) = m$. Its Stieltjes measure ν is positive, has mass $m \le 1$, and gives $\varphi(s, b) = \nu((s, b])$ almost everywhere. Substituting in (67) gives

$$\Pi(s, b) = b\nu((s, b]) - \int_{(s,b]} (b - z)\,\nu(\mathrm{d}z) = \int_{(s,b]} z\,\nu(\mathrm{d}z). \tag{76}$$

The welfare bound. Changing $(s, b]$ to $[s, b]$ affects only the seller-type lines through the countably many atoms of ν. These lines and the diagonal have zero probability under the absolutely continuous prior, and the lower triangle has no trade. With $W_\varnothing = M$,

$$W_{\mathrm{mech}} = (1 - m)W_\varnothing + \int_{(0,\Lambda)} W(z)\,\nu(\mathrm{d}z). \tag{77}$$

Bounded convergence and the absence of marginal atoms make W continuous on $[0, \Lambda]$, where it attains a maximum. Every price has welfare at least $W_\varnothing$, so

$$W_{\mathrm{mech}} \le \max_{0 \le z \le \Lambda} W(z). \tag{78}$$

□

B.3 Proof of the two-unit tail normalization

Proof of Proposition 13.1. We cut both buyers at a level c_2 at which the removed tails have total first moment equal to the best price gain, and return this moment as a common remote atom, split

equally between the two units. The orders $B_1 \geq B_2$ and $S_1 \leq S_2$ make the equal split no better for any price.

Every instance reduces to admissible bodies. We write $\Gamma_{2,\max} = \sup_z \Gamma_2(z)$ for the original maximum price gain. For an instance with $R_2 = (M_2 + \Gamma_{2,\max})/O_2 < 1$, we have $\Gamma_{2,\max} > 0$. Indeed, if every rational price had zero gain, the union of $\{S_i < z < B_i\}$ over rational z would give $O_2 = M_2$. Also $\Gamma_{2,\max} < O_2 - M_2 \leq \mathbb{E}(B_1 + B_2)$. The function $c \mapsto \sum_i \mathbb{E}(B_i - c)_+$ is continuous, starts above $\Gamma_{2,\max}$, and tends to zero. We therefore choose $c_2 > 0$ so that

$$e_i = \mathbb{E}(B_i - c_2)_+, \qquad e_1 + e_2 = \Gamma_{2,\max}. \tag{79}$$

Let $Z_i = B_i \wedge c_2$ and $Y_i = S_i \wedge c_2$. We take the new buyer to equal Z with probability $1 - \eta$ and $(\Gamma_{2,\max}/(2\eta), \Gamma_{2,\max}/(2\eta))$ with probability η. We choose η small enough that the tail exceeds $c_2 + \varepsilon$, and take the new seller to be $Y + (\varepsilon, \varepsilon)$, independently of the buyer.

For $z \leq c_2$, define

$$\begin{aligned} f_i(z) &= \Pr(S_i \leq z - \varepsilon), \\ t_i(z) &= \mathbb{E}[(S_i + \varepsilon)\mathbf{1}_{\{S_i \leq z - \varepsilon\}}], \\ \mathscr{K}_i(z) &= \mathbb{E}[(Z_i - S_i - \varepsilon)\mathbf{1}_{\{S_i + \varepsilon \leq z \leq Z_i\}}], \\ \mathscr{L}_i(z) &= \mathbb{E}[(B_i - S_i)\mathbf{1}_{\{S_i \leq z - \varepsilon,\ z \leq B_i\}}]. \end{aligned}$$

The displacement excludes sellers clipped to c_2. Independence and $Z_i = B_i \wedge c_2$ give

$$\mathscr{L}_i(z) = \mathscr{K}_i(z) + e_i f_i(z) + \varepsilon f_i(z) \Pr(B_i \geq z).$$

Writing $\Gamma_{\eta,\varepsilon}$ for the new gain, we obtain

$$\begin{aligned} \sum_i \mathscr{L}_i(z) - \Gamma_{\eta,\varepsilon}(z) = \eta \sum_i \mathscr{K}_i(z) + \frac{(e_1 - e_2)(f_1(z) - f_2(z))}{2} \\ + \varepsilon \sum_i f_i(z) \Pr(B_i \geq z) + \eta \sum_i t_i(z) \geq 0. \end{aligned} \tag{80}$$

Every term is nonnegative: in the second, $e_1 \geq e_2$ because $B_1 \geq B_2$, and $f_1 \geq f_2$ because $S_1 \leq S_2$. The original price includes all trades counted by $\mathscr{L}_i(z)$, so $\Gamma_{\eta,\varepsilon}(z) \leq \Gamma_2(z) \leq \Gamma_{2,\max}$. For $z > c_2$, only the common tail can buy, and its gain is at most $\Gamma_{2,\max}$. Every price strictly between $c_2 + \varepsilon$ and the tail value attains $\Gamma_{2,\max} - \eta M_\varepsilon$, where $M_\varepsilon = \sum_i \mathbb{E}(Y_i + \varepsilon)$. Consequently

$$\Gamma_{2,\max} - \eta M_\varepsilon \leq \sup_z \Gamma_{\eta,\varepsilon}(z) \leq \Gamma_{2,\max}. \tag{81}$$

Define the removed welfare quantities

$$\Delta_S = \sum_i \mathbb{E}(S_i - c_2)_+, \qquad \Delta_O = \sum_i \mathbb{E}[S_i - \max\{B_i, c_2\}]_+.$$

They satisfy $0 \leq \Delta_O \leq \Delta_S$ and $M_\varepsilon = M_2 - \Delta_S + 2\varepsilon$. The pointwise identity

$$\max\{b, s\} - \max\{b \wedge c, s \wedge c\} - (b - c)_+ = [s - \max\{b, c\}]_+$$

gives $\Gamma_{2,\max} + \sum_i \mathbb{E} \max\{Z_i, Y_i\} = O_2 - \Delta_O$. The new efficient welfare is $\Gamma_{2,\max} + (1 - \eta) \sum_i \mathbb{E} \max\{Z_i, Y_i + \varepsilon\}$. By (81), its ratio tends to $(M_2 - \Delta_S + \Gamma_{2,\max})/(O_2 - \Delta_O)$. The comparison with R_2 follows from

$$(M_2 + \Gamma_{2,\max})(O_2 - \Delta_O) - (M_2 - \Delta_S + \Gamma_{2,\max})O_2 = O_2(\Delta_S - \Delta_O) + (O_2 - M_2 - \Gamma_{2,\max})\Delta_O \geq 0.$$

At each fixed price, letting $\eta, \varepsilon \downarrow 0$ in the gain formula gives (57) with $\Gamma_{2,\max}/2$ in place of 1. Scaling all values by $2/\Gamma_{2,\max}$ therefore gives admissible bodies with no larger ratio.

Every admissible pair is approached. We fix admissible (Z, Y) bounded by $\bar{b}$. We take the buyer to equal Z with probability $1-\eta$ and $(1/\eta, 1/\eta)$ with probability η, and use the seller $Y + (\varepsilon, \varepsilon)$. We choose $1/\eta > \bar{b} + \varepsilon$. For $z \le 1/\eta$, the gain is

$$\begin{aligned}(1-\eta) \sum_i \mathbb{E}[(Z_i - Y_i - \varepsilon)\mathbf{1}_{\{Y_i+\varepsilon \le z \le Z_i\}}] \\ + \sum_i \Pr(Y_i \le z - \varepsilon) - \eta \sum_i \mathbb{E}[(Y_i + \varepsilon)\mathbf{1}_{\{Y_i \le z-\varepsilon\}}] \le \Gamma_0(z) \le 2.\end{aligned}$$

Prices above $1/\eta$ gain zero; prices between $\bar{b} + \varepsilon$ and $1/\eta$ have gain tending to 2. The initial and efficient welfare tend to $\mathbb{E}(Y_1 + Y_2)$ and $2 + \sum_i \mathbb{E}\max\{Z_i, Y_i\}$. We add a common positive displacement to all values to make them strictly positive. It translates prices and adds the same amount to both welfare quantities. Letting this displacement vanish gives

$$r_2^* \le \inf_{\substack{(Z,Y) \\ \text{admissible}}} \mathcal{R}_2(Z, Y) \le r_2^*.$$

□

B.4 Proof of the two-unit pricing theorem

Proof of Theorem E. For a fixed compensation the two unit constraints say $\varpi \ge \mathcal{T}_{d,\vartheta}\varpi$, and a buyer body bounded by $\bar{b}$ leaves the kernel of $\mathcal{T}_{d,\vartheta}$ with mass at most $\bar{b}/(1+\bar{b})$. The operator is therefore an order-preserving contraction, its fixed point is the least feasible mass, and leastness carries the other parts.

The potentials. A unit of price mass at $z > s$ contributes to the gain of unit i

$$\bar{L}_i(s) - \int_{s<t<z} (t-s)\,(-\mathrm{d}H_i(t)) = 1 + \mathbb{E}[(Z_i - s)\mathbf{1}_{\{Z_i \ge z\}}] = g_i(s, z).$$

A unit of ideal tail mass contributes 1, the same expression with the integral taken over $(s, \bar{b}]$; exchanging the order of integration gives (59). For nonnegative nonincreasing ϖ, the potential constraints are equivalent to $\varpi \ge \mathcal{T}_{d,\vartheta}\varpi$. If the potential sum is nonnegative on every ordered pair, the canonical prefix Θ_ϖ satisfies both constraints, and it is bounded and right-continuous because the potentials are.

Part (i). The positive kernel in (60) has mass

$$\frac{\int_{(t,\bar{b}]}(v-t)\,(-\mathrm{d}H_i(v))}{\bar{L}_i(t)} = \frac{\bar{L}_i(t) - 1}{\bar{L}_i(t)} \le \frac{\bar{b}}{1+\bar{b}} < 1.$$

Maxima and suprema do not increase uniform differences, so $\mathcal{T}_{d,\vartheta}$ is a contraction on bounded Borel functions, and its iterates from $\varpi^{(0)} = 0$ converge uniformly to its unique fixed point. Its output is nonnegative, nonincreasing, and right-continuous, because the obstacles are right-continuous in t and so is their future supremum; the uniform limit inherits these properties. The positive kernel also makes the operator order-preserving, so for any feasible nonnegative cumulative mass $\widetilde{\varpi}$, induction gives

$$\varpi^{(0)} \le \widetilde{\varpi}, \qquad \varpi^{(n+1)} \le \mathcal{T}_{d,\vartheta}\widetilde{\varpi} \le \widetilde{\varpi}.$$

The limit is therefore pointwise least.

Part (ii). The potential inequalities are affine in (ϖ, d, ϑ), so convex combinations of feasible triples are feasible, and leastness gives the convexity. The constant $\varpi = 1 + \bar{b}$ with $\vartheta = 0$ is feasible. Along a minimizing sequence with $\varpi_n(0) \le \bar{\varpi}$, the masses lie in $[0, \bar{\varpi}]$. The canonical prefixes of the compensations remain feasible and are bounded by $d \cdot \bar{b} + (1+\bar{b})(1+\bar{\varpi})$, because $0 \le G_i^{\varpi_n} \le (1+\bar{b})\bar{\varpi}$. By Helly's selection theorem a subsequence converges pointwise to monotone limits $\varpi_{\mathrm{raw}}, \vartheta_{\mathrm{raw}}$, and bounded convergence passes the integral inequalities to the limit. We take right limits at $s < \bar{b}$, keeping the values at $\bar{b}$, to obtain $\varpi^+ \le \varpi_{\mathrm{raw}}$ and ϑ^+. The regularized pair is feasible: passing the raw inequalities along $t \downarrow s$ and then replacing ϖ_{raw} by ϖ^+ in the positive kernel only decreases the right side. Leastness and the definition of the infimum give

$$\mathcal{J}_d^{(2)} \le \varpi_{d,\vartheta^+}(0) \le \varpi^+(0) \le \varpi_{\mathrm{raw}}(0) = \mathcal{J}_d^{(2)}.$$

The choice of $\bar{b}$. The minimum in (61) is unchanged when the operator is defined on $[0, \bar{b}']$ with $\bar{b}' > \bar{b}$. A feasible pair on $[0, \bar{b}']$ restricts to a feasible pair on $[0, \bar{b}]$ with the same value at 0, because the buyers put no mass above $\bar{b}$ and a future supremum over a shorter range is no larger. A feasible pair on $[0, \bar{b}]$ extends by the constants $\varpi(\bar{b})$ and $\vartheta(\bar{b})$. For $\bar{b} < t \le \bar{b}'$ the kernel integral vanishes and $\bar{L}_i = 1$, so $\mathcal{A}_{i,d,\vartheta}\varpi(t) = 1 - d \cdot t + \epsilon_i \vartheta(\bar{b})$, which decreases in t because $d > 0$ and stays below $\mathcal{A}_{i,d,\vartheta}\varpi(\bar{b}) \le \varpi(\bar{b})$. The two values are equal by leastness.

Part (iii). We choose an optimal compensation, write $\varpi = \varpi_{d,\vartheta}$, and put $-\beta \, \mathrm{d}\varpi$ on $(0, \bar{b}]$, density βd on $(\bar{b}, \bar{b} + \varpi(\bar{b})/d)$, and the remaining mass $1 - \beta\varpi(0) \ge 0$ at zero. For $s \le \bar{b}$, the tail contributes $\beta\varpi(\bar{b})$, exactly as the ideal tail does. For $s = \bar{b} + w$, the realized potential, with gain divided by β, differs from its value at $\bar{b}$ by

$$d \cdot w - \min\{d \cdot w, \varpi(\bar{b})\} = d(w - \varpi(\bar{b})/d)_+ \ge 0.$$

A single atom of the same mass at $\bar{b} + \varpi(\bar{b})/d$ gives the difference $d \cdot w - \varpi(\bar{b}) \mathbf{1}_{\{d \cdot w \ge \varpi(\bar{b})\}} \ge 0$ instead. Clipping seller values at $\bar{b}$ preserves their order and weakly decreases their potential sum. On the clipped domain,

$$\psi_1^{\varpi}(s_1) + \psi_2^{\varpi}(s_2) \ge -\vartheta(s_1) + \vartheta(s_2) \ge 0 \qquad (s_1 \le s_2).$$

Multiplication by β proves (62), including $\varpi(\bar{b}) = 0$.

Conversely, we take a probability satisfying (62), discard its zero-price mass, and divide by β. For sellers in $[0, \bar{b}]$, prices above $\bar{b}$ can be collected into the ideal tail without changing gain. The resulting cumulative mass $\widetilde{\varpi}$ has a nonnegative potential sum and a feasible canonical compensation, and leastness gives

$$\beta \mathcal{J}_d^{(2)} \le \beta \varpi_{d, \Theta_{\widetilde{\varpi}}}(0) \le \beta \widetilde{\varpi}(0) \le 1.$$

For the welfare guarantee, let $d = (1-\beta)/\beta$, so that $\beta(1+d) = 1$, and let a price π satisfy (62). For bounded seller bodies $Y_1 \le Y_2$, independent of the buyers and of the price, we put $(s_1, s_2) = (Y_1, Y_2)$ and take expectations. By independence, the left side becomes the expected gain $\int \Gamma_0 \, \mathrm{d}\pi$ from (57); since $s + \bar{L}_i(s) = 1 + \mathbb{E} \max\{Z_i, s\}$, the right side becomes $\beta \sum_i \mathbb{E}(1 + \max\{Z_i, Y_i\}) - \mathbb{E}(Y_1 + Y_2)$. Adding the initial welfare to both sides gives

$$\mathbb{E}(Y_1 + Y_2) + \int \Gamma_0 \, \mathrm{d}\pi \ge \beta \Big[2 + \sum_{i=1}^{2} \mathbb{E} \max\{Z_i, Y_i\} \Big].$$

Part (iv). Since $Z_i^\delta \le Z_i < Z_i^\delta + \delta$, we have $0 \le \bar{L}_i(s) - \bar{L}_i^\delta(s) \le \delta$, and $g_i(s,z) \ge g_i^\delta(s,z)$ because the function $(Z_i - s)\mathbf{1}_{\{z \le Z_i\}}$ is nondecreasing in Z_i when $s < z$, including its jump at z. We take an optimal rounded price in its ideal-tail representation and add tail mass δ. Each potential then

increases by at least $-(\bar{L}_i - \bar{L}_i^{\delta}) + \delta \geq 0$, so the same compensation is feasible at the additional mass δ, and leastness gives the inequality of part (iv). Rounding down is nondecreasing and does not raise a value, so the rounded bodies are again ordered, take finitely many values, and lie in $[0, \bar{b}]$, on which both values are computed because $\mathcal{J}_d^{(2)}$ does not depend on $\bar{b}$. Taking suprema and then letting $\delta \downarrow 0$ gives the last claim. □

B.5 Finite-node evaluation

Suppose the buyers are supported on nodes $0 = v_0 < \cdots < v_n = \bar{b}$. Write $p_{i\ell} = \Pr(Z_i = v_\ell)$, $\bar{L}_{ij} = \bar{L}_i(v_j)$, and $\vartheta_j = \vartheta(v_j)$. For prescribed nondecreasing node compensations, the least solution of the node inequalities is obtained backwards:

$$\begin{aligned} \varpi_n &= \max\{0, 1 - d \cdot \bar{b} - \vartheta_n, 1 - d \cdot \bar{b} + \vartheta_n\}, \\ \varpi_j &= \max\left\{\varpi_{j+1},\ \max_{i=1,2}\left[1 - \frac{d \cdot v_j}{\bar{L}_{ij}} + \frac{\epsilon_i \vartheta_j}{\bar{L}_{ij}} + \sum_{\ell > j} \frac{(v_\ell - v_j) p_{i\ell}}{\bar{L}_{ij}} \varpi_\ell\right]\right\}. \end{aligned} \tag{82}$$

The value is piecewise affine and convex in $(d, \vartheta_0, \ldots, \vartheta_n)$, and the price at v_j has normalized mass $\varpi_{j-1} - \varpi_j$.

The node solution certifies a price only if it also covers the seller values between nodes. At the relevant budget $\varpi_0 \leq 1 + d$ it does. Between consecutive nodes,

$$(\psi_i^{\varpi})'(s) = d + H_i(s) - \int_{(s,\bar{b}]} \Pr(Z_i \geq z)\, \widehat{\pi}_0(\mathrm{d}z) \geq d - H_i(s)(\varpi_0 - 1) \geq 0.$$

The first inequality holds because $\Pr(Z_i \geq z) \leq H_i(s)$ for $z > s$ and $\widehat{\pi}_0((s, \bar{b}]) \leq \varpi(s) \leq \varpi_0$, and the second because $0 \leq H_i \leq 1$ and $\varpi_0 \leq 1 + d$. At a node, the potential drops by the price atom's mass times $\bar{L}_i(s)$. Its minimum on each interval $[v_j, v_{j+1})$ is therefore its value at v_j, and ordered node constraints cover all ordered real seller values. At arbitrary budgets, (82) is only a node optimization; the full interval problem is (60).

At the numerical value $d_{\mathrm{fam}} \approx 0.371591$, the search of Section 13.5 evaluated 20,000 finite-support ordered buyer pairs: 16,000 broad random samples with 2–32 common support nodes and value scales between 0.01 and 100, and 4,000 random perturbations of discretized reference buyers with up to 65 atoms in each marginal. The broad samples had maximum 1.3645392. The largest values agreed under cumulative-mass and price-mass formulations of the inner linear program.

C Auxiliary results

C.1 Quantitative finite approximation

The hard-pair approximation error is linear in the buyer tail mass.

Proposition C.1 (finite-tail error)**.** *Let $d = \tau(C_*)$, and let r_η be the optimal fixed-price welfare ratio of* (46) *at C_*. With*

$$L_{\mathrm{err}}(d) = \frac{3}{d} + \frac{3}{d^2} + \frac{2}{d^3},$$

one has

$$0 \leq r_\eta - \beta_* \leq \eta\left[\beta_* L_{\mathrm{err}}(d) + \frac{1}{d}\right]. \tag{83}$$

Proof. For $0 \le h \le 1$, the perturbation $h_\eta = \eta + (1-\eta)h$ and its states satisfy

$$0 \le h_\eta - h \le \eta, \quad 0 \le g_\eta - g \le \eta, \quad 0 \le k_\eta - k \le 2\eta, \quad g_\eta, g \ge d.$$

Hence $|g_\eta^{-1} - g^{-1}| \le \eta d^{-2}$ and $|g_\eta^{-2} - g^{-2}| \le 2\eta d^{-3}$, and since $|d^2 - k| \le d^2 + 1$, integration over the unit interval gives

$$|\mathcal{J}_d(h_\eta) - \mathcal{J}_d(h)| \le \eta \left[\frac{1}{d} + \frac{1}{d^2} + \frac{2}{d^2} + \frac{2(d^2+1)}{d^3} \right] = \eta L_{\mathrm{err}}(d). \tag{84}$$

For $h = h_{C_*}$, equations (44) and (49) imply

$$r_\eta - \beta_* = \frac{\beta_* d(\mathcal{J}_d(h) - \mathcal{J}_d(h_\eta)) + \eta d^2 T_\eta}{M_\eta + G_\eta}.$$

The first numerator term is nonnegative by Theorem C. We bound the other term by $d^2 T_\eta \le 1$ and the denominator by $M_\eta + G_\eta \ge d$. This gives (83). □

C.2 Nonconcavity in the original control

The objective in (6) is not concave in h. If $h = u_0$ on $[0, \theta]$ and $h = u_1$ on $(\theta, 1]$, direct integration gives

$$\mathcal{J}_d(h) = 1 + \frac{\theta(1-\theta)(d+u_0)(u_1-u_0)}{(d+\theta u_0)(d+\theta u_0 + (1-\theta)u_1)}. \tag{85}$$

At $d = \theta = 1/2$, the controls $(4/5, 3/5)$ and $(1/5, 2/5)$ have values $203/216$ and $103/96$. Their midpoint is the constant $1/2$ control, whose value is 1. The Jensen gap is

$$\frac{1}{2}\left(\frac{203}{216} + \frac{103}{96}\right) - 1 = \frac{11}{1728} > 0. \tag{86}$$

D Optimization within a two-unit family

This appendix proves Proposition 13.2, for the family and functionals defined before it. The proof identifies the fixed-endpoint energy minimizer (Lemma D.3), excludes the outer boundaries (Lemma D.4), and isolates the stationary branch (Lemmas D.5 and D.6). The family value is a limiting upper bound for ordered two-unit trade; the strict bound of Theorem D comes from a separate finite instance and does not use the optimization argument. The computations behind Lemmas D.2, D.4, D.5, and D.6 are in the supplementary material, which also holds the certificates that replay them.

D.1 Trading interpretation of the curve family

The trading interpretation uses the increasing map and body endpoint

$$s_{\mathrm{f}}(\xi) = a_{\mathrm{f}} + \int_{c_{\mathrm{f}}}^{\xi} \mathcal{P}(r)^{-2}\,\mathrm{d}r, \qquad b_{\mathrm{f}} = s_{\mathrm{f}}(\xi_0). \tag{87}$$

The first buyer body has survival 1 below a_{f} and

$$H_{1,\mathrm{f}}(s_{\mathrm{f}}(\xi)) = \mathcal{P}'_+(\xi), \qquad F_{1,\mathrm{f}}(s_{\mathrm{f}}(\xi)) = N_{\mathrm{f}}\{\mathcal{P}(\xi) - \xi\mathcal{P}'_+(\xi)\} \quad (c_{\mathrm{f}} \le \xi < \xi_0). \tag{88}$$

The first seller has mass m_{f} at zero, has constant CDF m_{f} between zero and a_{f}, follows (88) on the body interval, and is completed to mass one at b_{f}. If $a_{\mathrm{f}} = 0$, coincident initial masses are combined. The second buyer body is $Z_2 = a_{\mathrm{f}}$, and the second seller has law $m_{\mathrm{f}}\delta_0 + (1 - m_{\mathrm{f}})\delta_{b_{\mathrm{f}}}$. Endpoint atoms complete the first buyer survival to a probability law. Coupling the bodies on each side by common quantiles gives the body vector laws. A sequence of instances converges to $\mathcal{P}$ when its buyer and seller vector laws converge weakly to them.

Lemma D.1 (Realization by limiting instances)**.** *Every member of $\mathfrak{F}$ is the limit of strictly positive, finite-mean ordered two-unit instances with independent buyer and seller vectors, whose optimal common-price welfare ratios converge to R_{f}. Their seller welfare and efficient gains converge to M_{f} and G_{f}, and their best price gains converge to 2.*

Proof. Concavity and the obstacles make $H_{1,\mathrm{f}}$ a survival function and $F_{1,\mathrm{f}}$ a CDF with $m_{\mathrm{f}} \le F_{1,\mathrm{f}} \le 1$, both right-continuous because $\mathcal{P}'_+$ is. The common-quantile couplings give $Z_1 \ge Z_2$ and $Y_1 \le Y_2$, and we draw the sides independently. The limiting first-unit tail integral and integrated seller CDF are

$$\bar{L}_{1,\mathrm{f}} = \mathcal{P}^{-1}, \qquad A_{1,\mathrm{f}} = N_{\mathrm{f}}\xi/\mathcal{P}, \qquad F_{1,\mathrm{f}}\bar{L}_{1,\mathrm{f}} + H_{1,\mathrm{f}}A_{1,\mathrm{f}} = N_{\mathrm{f}}.$$

That is, the limiting first-unit gain is N_{f} at every body price.

As in the proof of Proposition 13.1, we give both buyer bodies weight $1 - \eta_{\mathrm{f}}$ and add a common buyer atom of mass η_{f} at $1/\eta_{\mathrm{f}}$, so that each escaping tail has first moment one. We shift both sellers up by $\epsilon_{\mathrm{f}} > 0$, and when $a_{\mathrm{f}} = 0$ we also translate all values by a common vanishing positive amount. These operations preserve the orders. At a price s, the shifted first-unit gain is bounded by the unshifted right-limit gain at $s - \epsilon_{\mathrm{f}}$, because the buyer stop-loss and its acceptance probability decrease as the price moves up. It is therefore at most N_{f} when $s - \epsilon_{\mathrm{f}} \in [a_{\mathrm{f}}, b_{\mathrm{f}})$, and at most $m_{\mathrm{f}}(a_{\mathrm{f}} + 1/p_{\mathrm{f}}) \le N_{\mathrm{f}}$ for $a_{\mathrm{f}} < s < a_{\mathrm{f}} + \epsilon_{\mathrm{f}}$; the second unit contributes at most m_{f} in these regions, and $N_{\mathrm{f}} + m_{\mathrm{f}} = 2$. For prices $s \le a_{\mathrm{f}}$, the total gain is at most $m_{\mathrm{f}}(2 + \bar{z}_{\mathrm{f}} + a_{\mathrm{f}}) = 2$, where $\bar{z}_{\mathrm{f}} = a_{\mathrm{f}} + 1/p_{\mathrm{f}} - 1$ is the first body mean. For $s \ge b_{\mathrm{f}} + \epsilon_{\mathrm{f}}$, only the escaping buyer tails can trade, and their combined first moment is two. Prices above the shifted sellers and below the remote buyer atom have gains tending to two. Integrating the body states gives (65), so

$$\lim_{\eta_{\mathrm{f}}, \epsilon_{\mathrm{f}} \downarrow 0} \frac{\text{seller welfare} + \text{best price gains}}{\text{seller welfare} + \text{efficient gains}} = \frac{M_{\mathrm{f}} + 2}{M_{\mathrm{f}} + G_{\mathrm{f}}}.$$

□

For a trial ratio $0 < \beta_{\mathrm{f}} < 1$, define $d_{\mathrm{f}} = (1 - \beta_{\mathrm{f}})/\beta_{\mathrm{f}}$ and

$$\begin{aligned} E_{\mathrm{f}}(\mathcal{P}) &= \int_{c_{\mathrm{f}}}^{\xi_0} \{\xi(\log \mathcal{P})'^2 + d_{\mathrm{f}}\mathcal{P}^{-2}\}\, \mathrm{d}\xi, \\ K_{\mathrm{f}}(\mathcal{P}) &= \frac{\beta_{\mathrm{f}}G_{\mathrm{f}} - (1 - \beta_{\mathrm{f}})M_{\mathrm{f}} - 2}{\beta_{\mathrm{f}}N_{\mathrm{f}}} \\ &= 1 - \frac{t_{\mathrm{f}}}{p_{\mathrm{f}}} - \log p_{\mathrm{f}} - d_{\mathrm{f}}a_{\mathrm{f}} - d_{\mathrm{f}}(1 + t_{\mathrm{f}}) + d_{\mathrm{f}}\xi_0 - E_{\mathrm{f}}(\mathcal{P}). \end{aligned} \tag{89}$$

Thus $R_{\mathrm{f}} < \beta_{\mathrm{f}}$ exactly when $K_{\mathrm{f}} > 0$. The normalization by N_{f} depends on t_{f}; all outer derivatives below are derivatives of K_{f} along outer variations, in which one of $t_{\mathrm{f}}, p_{\mathrm{f}}, h_{\mathrm{f}}$ moves and the other two stay fixed. A derivative is one-sided where only one direction of the variation is feasible.

D.2 Compactness and the fixed-endpoint minimizer

Lemma D.2 (Compactness below three quarters)**.** *Let $0 < \beta_{\mathrm{f}} \le 3/4$. Every curve with $R_{\mathrm{f}} < \beta_{\mathrm{f}}$ satisfies $1/32 < m_{\mathrm{f}} < 3/4$, $\xi_0 < 600$, and $p_{\mathrm{f}} < 1$. If K_{f} takes a positive value, its positive supremum is attained in this parameter region.*

Proof. Body means bound the functionals. The identity $\bar{z}_{\mathrm{f}} + a_{\mathrm{f}} = 2(1-m_{\mathrm{f}})/m_{\mathrm{f}}$ gives $G_{\mathrm{f}} \le 2/m_{\mathrm{f}}$ and $R_{\mathrm{f}} \ge m_{\mathrm{f}}$, hence $m_{\mathrm{f}} < 3/4$. For $m_{\mathrm{f}} \le 1/32$, the estimates $G_{\mathrm{f}} \le 4+2\log(2/m_{\mathrm{f}})$ and $M_{\mathrm{f}} \ge (1-m_{\mathrm{f}})^2/m_{\mathrm{f}}$ make $\beta_{\mathrm{f}} G_{\mathrm{f}} - (1-\beta_{\mathrm{f}}) M_{\mathrm{f}} - 2$ negative. When $K_{\mathrm{f}} > 0$, the bound $G_{\mathrm{f}} \le 2/m_{\mathrm{f}}$ then bounds M_{f}, and $M_{\mathrm{f}} > (5/16)\xi_0$ gives $\xi_0 < 600$. For attainment, we extend each curve to a concave curve with slopes in $[0,1]$ on $[0,600]$. Such curves are compact for uniform convergence, under which the obstacles and endpoint conditions persist and T_{f} and Q_{f} are continuous. At $p_{\mathrm{f}} = 1$ the limiting ratio is at least $3/4$. The estimates are in the supplementary material. □

Lemma D.3 (Three-arc energy minimizer)**.** *At every feasible fixed $(t_{\mathrm{f}}, p_{\mathrm{f}}, \xi_0)$, the energy E_{f} has a unique minimizer in* (64)*. It is entirely affine, or consists of an initial affine contact, one strictly concave free interval, and a final affine contact. Either contact may have zero length. On the free interval there is a constant $C_{\mathrm{f}} > 0$ such that*

$$\mathcal{P}'' + \frac{C_{\mathrm{f}}}{\xi^2}\mathcal{P} = 0, \qquad d_{\mathrm{f}} = \xi\mathcal{P}'^2 - \mathcal{P}\mathcal{P}' + \frac{C_{\mathrm{f}}\mathcal{P}^2}{\xi}. \tag{90}$$

If $c_{\mathrm{f}} = 0$, every non-affine minimizer has a positive initial contact interval.

Proof. We minimize over logarithms $z_{\mathrm{f}} = \log \mathcal{P}$ below the obstacle $\log U_{\mathrm{f}}$. This class is convex and does not impose concavity; we show that its minimizer is concave.

In this class the energy $\int(\xi z_{\mathrm{f}}'^2 + d_{\mathrm{f}} e^{-2z_{\mathrm{f}}})$ is strictly convex, and truncation below $\log p_{\mathrm{f}}$ decreases both terms. The direct method gives a unique minimizer: in H^1 when $c_{\mathrm{f}} > 0$, and through weak limits on compact intervals away from zero when $c_{\mathrm{f}} = 0$, where the bounds $p_{\mathrm{f}} \le \mathcal{P} \le p_{\mathrm{f}} + \xi$ preserve the trace at zero. Its variational inequality is

$$\mu_{\mathrm{f}} = (\xi z_{\mathrm{f}}')' + d_{\mathrm{f}} e^{-2z_{\mathrm{f}}} \ge 0, \qquad \operatorname{supp}\mu_{\mathrm{f}} \subseteq \{\mathcal{P} = U_{\mathrm{f}}\}. \tag{91}$$

Away from zero, $\xi z_{\mathrm{f}}'$ has bounded variation. At a contact with an affine part of the obstacle, the left derivative of $\mathcal{P}$ is at least the obstacle slope and the right derivative at most that slope, while (91) requires a nonnegative jump. Both derivatives therefore equal the obstacle slope. The same signs exclude contact at a genuine downward corner of the obstacle.

On a free interval, we put $\omega_{\mathrm{f}} = \xi\mathcal{P}'/\mathcal{P}$ to integrate the Euler equation. It gives

$$\omega_{\mathrm{f}}' = -d_{\mathrm{f}}/\mathcal{P}^2, \qquad C_{\mathrm{f}} = \omega_{\mathrm{f}} - \omega_{\mathrm{f}}^2 + d_{\mathrm{f}}\xi/\mathcal{P}^2, \qquad C_{\mathrm{f}}' = 0.$$

These relations rearrange to (90). The right endpoint slope is positive, while the left value of ω_{f} is less than one, by the positive intercept at a contact or by the slope bound at $c_{\mathrm{f}} > 0$. Since ω_{f} decreases, we obtain $0 < \omega_{\mathrm{f}} < 1$, $C_{\mathrm{f}} > 0$, and $\mathcal{P}'' < 0$ on the free interval. A strictly concave interval lies above its chord, so it cannot leave and return to the same affine obstacle. The first obstacle is active before their intersection and the second afterwards, so at most one free interval separates the initial and final contacts. If a free interval started at zero, bounded positive $\mathcal{P}$ and the equation for ω_{f}' would force $\omega_{\mathrm{f}}(0) = 0$ and then $\omega_{\mathrm{f}} < 0$, contradicting $\mathcal{P} \ge p_{\mathrm{f}}$.

For the resulting minimizer $\mathcal{P}_*$, write $\upsilon = \log(\mathcal{P}/\mathcal{P}_*)$. As for (G1), integration by parts in (91) gives

$$E_{\mathrm{f}}(\mathcal{P}) - E_{\mathrm{f}}(\mathcal{P}_*) = \int_{c_{\mathrm{f}}}^{\xi_0} \{\xi\upsilon'^2 + d_{\mathrm{f}}\mathcal{P}_*^{-2}(e^{-2\upsilon} - 1 + 2\upsilon)\}\,\mathrm{d}\xi - 2\int \upsilon\,\mathrm{d}\mu_{\mathrm{f}} \ge 0. \tag{92}$$

□

D.3 Boundary exclusion and the stationary system

In the two entirely affine cases, M_{f} and G_{f} are rational functions of $(t_{\mathrm{f}}, p_{\mathrm{f}})$ on the respective domains $t_{\mathrm{f}} \le p_{\mathrm{f}} < 1$ and $(1+t_{\mathrm{f}})/2 < p_{\mathrm{f}} < 1$. On both domains $R_{\mathrm{f}} \ge 73/100$: after multiplication by the positive denominators this is a polynomial inequality, proved by positive Bernstein expansions on rational subrectangles. The functions and the coefficients are in the supplementary material.

Lemma D.4 (Interior outer parameters)**.** *Suppose $\beta_{\mathrm{f}} \in [18227/25000, 729081/10^6]$. At a positive global maximum of K_{f}, or at a global minimum of R_{f} whose value is β_{f}, one has $t_{\mathrm{f}} < p_{\mathrm{f}} < 1$ and both affine contacts have positive length. The derivatives of K_{f} with respect to $(t_{\mathrm{f}}, p_{\mathrm{f}}, h_{\mathrm{f}})$ vanish there.*

Proof. Since

$$K_{\mathrm{f}} = \frac{M_{\mathrm{f}} + G_{\mathrm{f}}}{N_{\mathrm{f}}} \left(1 - \frac{R_{\mathrm{f}}}{\beta_{\mathrm{f}}}\right),$$

a ratio minimizer of value β_{f} maximizes K_{f} at zero, so in both cases we consider a maximizer of K_{f}. We replace its curve by the fixed-endpoint energy minimizer, which maximizes K_{f} at those endpoints by (89). The bound $R_{\mathrm{f}} \ge 73/100$ on the affine cases forces a nonempty free interval. Each boundary is excluded by an outer variation that is affine on the contact intervals and blends to zero inside a compact part of the strictly concave free interval, where strict curvature and the obstacle gap preserve feasibility for small changes. If the initial contact had zero length with $c_{\mathrm{f}} > 0$, lowering p_{f} would increase K_{f}; for $c_{\mathrm{f}} = 0$ the initial contact is positive by Lemma D.3. If the final contact had zero length, increasing ξ_0 would increase K_{f}, and at $p_{\mathrm{f}} = t_{\mathrm{f}}$ increasing p_{f} would. Lemma D.2 excludes $p_{\mathrm{f}} = 1$ and the noncompact boundaries. At all remaining parameters the variations work in both directions, so the three derivatives vanish. The first-order formulas and the constructions are in the supplementary material. □

Write $\xi_\ell < \xi_r$ for the contact times, $\mathcal{P}_\ell = \xi_\ell + \delta_{\mathrm{f}}$, and $\mathcal{P}_r = h_{\mathrm{f}}\xi_r + e_{\mathrm{f}}$. Define the two dimensionless contact slopes $q_{\ell,\mathrm{f}} = \xi_\ell/\mathcal{P}_\ell$ and $q_{r,\mathrm{f}} = h_{\mathrm{f}}\xi_r/\mathcal{P}_r$. Integrating (90) along the free arc expresses T_{f}, Q_{f}, and G_{f} through these contact data, and the three vanishing derivatives of Lemma D.4 become three explicit equations, which we call the outer equations. Both computations are in the supplementary material.

D.4 The stationary branch and the family minimum

We eliminate the contact parameters from the outer equations. From here on, $\mathcal{P}_\ell$, ξ_ℓ, and C_{f} denote the functions of $(t_{\mathrm{f}}, p_{\mathrm{f}})$ defined in the first line below; at a stationary point they agree with the contact data above, by (90) and the outer equation for p_{f}. Define

$$\mathcal{P}_\ell = p_{\mathrm{f}}\sqrt{\frac{\delta_{\mathrm{f}} + d_{\mathrm{f}}}{t_{\mathrm{f}} + d_{\mathrm{f}}}}, \qquad \xi_\ell = \mathcal{P}_\ell - \delta_{\mathrm{f}}, \qquad C_{\mathrm{f}} = \frac{(\delta_{\mathrm{f}} + d_{\mathrm{f}})\xi_\ell}{\mathcal{P}_\ell^2},$$
$$v_{\mathrm{f}}^2(q_{r,\mathrm{f}}^2 - q_{r,\mathrm{f}} + C_{\mathrm{f}}) = C_{\mathrm{f}} q_{r,\mathrm{f}}^2, \qquad 0 < q_{r,\mathrm{f}} < v_{\mathrm{f}},$$
$$\gamma_{\mathrm{f}} = -d_{\mathrm{f}}/t_{\mathrm{f}}^2 + 2d_{\mathrm{f}} + 1/p_{\mathrm{f}} + (t_{\mathrm{f}} + d_{\mathrm{f}})/p_{\mathrm{f}}^2, \qquad \mathcal{A}_{\mathrm{f}}(t_{\mathrm{f}}, p_{\mathrm{f}}) = -\gamma_{\mathrm{f}} + \frac{v_{\mathrm{f}}(v_{\mathrm{f}} - q_{r,\mathrm{f}})}{e_{\mathrm{f}}(v_{\mathrm{f}} + q_{r,\mathrm{f}})} = 0. \tag{93}$$

For every $d_{\mathrm{f}} > 0$ we call the set $0 < t_{\mathrm{f}} < p_{\mathrm{f}} < 1$ the algebraic branch. There $\mathcal{P}_\ell > p_{\mathrm{f}} > \delta_{\mathrm{f}}$, so $\xi_\ell > 0$ and $C_{\mathrm{f}} > 0$. The equation for $q_{r,\mathrm{f}}$ has one root in $(0, v_{\mathrm{f}})$, because $v_{\mathrm{f}}^2 q_{r,\mathrm{f}}(1 - q_{r,\mathrm{f}})/(v_{\mathrm{f}}^2 - q_{r,\mathrm{f}}^2)$ increases

from zero to infinity. With $\mathcal{D}_{\mathrm{f}}(u) = u^2 - u + C_{\mathrm{f}}$, define

$$\mathcal{I}_{\mathrm{f}} = \int_{q_{r,\mathrm{f}}}^{q_{\ell,\mathrm{f}}} \frac{\mathrm{d}u}{\mathcal{D}_{\mathrm{f}}(u)},$$

$$\mathcal{F}_{\mathrm{f}}(t_{\mathrm{f}}, p_{\mathrm{f}}) = \log \frac{1 - q_{r,\mathrm{f}}}{1 - q_{\ell,\mathrm{f}}} + \frac{1}{2} \log \frac{\mathcal{D}_{\mathrm{f}}(q_{\ell,\mathrm{f}})}{\mathcal{D}_{\mathrm{f}}(q_{r,\mathrm{f}})} + \frac{\mathcal{I}_{\mathrm{f}}}{2} - \log \frac{e_{\mathrm{f}}}{\delta_{\mathrm{f}}} = 0. \tag{94}$$

The recovered parameters are

$$h_{\mathrm{f}} = \frac{d_{\mathrm{f}} q_{r,\mathrm{f}}(1 - q_{r,\mathrm{f}})}{e_{\mathrm{f}} \mathcal{D}_{\mathrm{f}}(q_{r,\mathrm{f}})}, \qquad \mathcal{P}_r = \frac{e_{\mathrm{f}}}{1 - q_{r,\mathrm{f}}}, \qquad \xi_r = \frac{q_{r,\mathrm{f}} \mathcal{P}_r}{h_{\mathrm{f}}}, \qquad \xi_0 = \frac{v_{\mathrm{f}}}{h_{\mathrm{f}}}. \tag{95}$$

A physical stationary point means that these formulas produce a curve in (64), with $c_{\mathrm{f}} < \xi_\ell < \xi_r < \xi_0$, $0 < q_{r,\mathrm{f}} < q_{\ell,\mathrm{f}} < 1$, and positive $\mathcal{D}_{\mathrm{f}}$ on the integration interval.

Lemma D.5 (Stationary domain and monotonicity)**.** *For $37/100 \le d_{\mathrm{f}} \le 3/8$, every physical stationary point satisfies $1/5 < t_{\mathrm{f}} < 421/1000$. On the algebraic branch, $\partial_{p_{\mathrm{f}}} \mathcal{A}_{\mathrm{f}} > 0$.*

Proof. On the algebraic branch, $\partial_{p_{\mathrm{f}}} \mathcal{A}_{\mathrm{f}}$ is the sum of $p_{\mathrm{f}}^{-2} + 2(d_{\mathrm{f}} + t_{\mathrm{f}}) p_{\mathrm{f}}^{-3}$ and the response of the last term of $\mathcal{A}_{\mathrm{f}}$ to C_{f}. That term has slope between $-2/e_{\mathrm{f}}$ and 0 in C_{f}, and $\partial_{p_{\mathrm{f}}} C_{\mathrm{f}} < t_{\mathrm{f}}(d_{\mathrm{f}} + t_{\mathrm{f}})/p_{\mathrm{f}}^3$. The response is therefore positive when $\partial_{p_{\mathrm{f}}} C_{\mathrm{f}} < 0$ and otherwise exceeds $-2(d_{\mathrm{f}} + t_{\mathrm{f}}) p_{\mathrm{f}}^{-3}$, because $e_{\mathrm{f}} > t_{\mathrm{f}}$. We exclude $t_{\mathrm{f}} \le 1/5$ using $\gamma_{\mathrm{f}} > 0$, which holds at a physical stationary point. For such t_{f} it forces $p_{\mathrm{f}} < 7t_{\mathrm{f}}/4$ and makes $q_{\ell,\mathrm{f}}$ too small for the order $q_{r,\mathrm{f}} < q_{\ell,\mathrm{f}}$. The upper bound follows from $\gamma_{\mathrm{f}} < v_{\mathrm{f}}/e_{\mathrm{f}}$ and $p_{\mathrm{f}} < 1$. The estimates are in the supplementary material. □

Lemma D.6 (Validated stationary branch)**.** *For every $\beta_{\mathrm{f}} \in [18227/25000, 729081/10^6]$, the system (93)–(94) has exactly one physical stationary point. At $\beta_{\mathrm{f}} = 18227/25000$ it satisfies $\beta_{\mathrm{f}} G_{\mathrm{f}} - (1 - \beta_{\mathrm{f}}) M_{\mathrm{f}} - 2 < 0$.*

Proof. Lemma D.5 reduces the covering problem to one dimension: over each rational t_{f} interval, strict monotonicity of $\mathcal{A}_{\mathrm{f}}$ in p_{f} encloses its unique root. We cover the t_{f} range by finitely many rational intervals and evaluate them in interval arithmetic. An interval is discarded when it has no algebraic root, violates a physical phase condition, or has a nonzero enclosure of $\mathcal{F}_{\mathrm{f}}$. The retained intervals are contiguous. On them $C_{\mathrm{f}} > 1/4$, the function $\mathcal{F}_{\mathrm{f}}$ is strictly increasing along the root curve, and it has opposite signs at the two outer endpoints. For every β_{f} in the trial interval, exactly one solution of (93)–(94) therefore satisfies the phase conditions.

This solution is a physical stationary point. Interval evaluation of (95) over the retained root enclosure gives the contact and slope orders, and integrating $\mathrm{d}\xi/\xi = -\mathrm{d}u/\mathcal{D}_{\mathrm{f}}(u)$ from $q_{\ell,\mathrm{f}}$ to $q_{r,\mathrm{f}}$ reconstructs a curve of (64) that meets both obstacles tangentially and satisfies the outer equations. The covering, its interval bounds, and the reconstruction are in the supplementary material. At the lower endpoint the covering gives

$$\beta_{\mathrm{f}} G_{\mathrm{f}} - (1 - \beta_{\mathrm{f}}) M_{\mathrm{f}} - 2 \in [-2.116 \cdot 10^{-6}, -1.157 \cdot 10^{-6}] < 0.$$

□

Lemma D.7 (A rational comparison curve)**.** *The class $\mathfrak{F}$ contains a rational concave polygon with $R_{\mathrm{f}} < 729081/10^6$.*

Proof. We use the rational polygon specified in Appendix E. Exact segment integration with its rational vertices gives

$$R_{\mathrm{f}} < 729081/10^6.$$

□

Proof of Proposition 13.2. One implication is used three times; we call it $(*)$. If β_{f} lies in the trial interval $[18227/25000, 729081/10^6]$, then every curve at which K_{f} attains a positive maximum is a physical stationary point, and so is every curve that minimizes R_{f} with value β_{f}.

Proof of $(*)$. By Lemmas D.3 and D.4, such a curve has three arcs with positive contacts and satisfies the outer equations. On its free arc, $\omega_{\mathrm{f}} = \xi\mathcal{P}'/\mathcal{P}$ decreases from $q_{\ell,\mathrm{f}}$ to $q_{r,\mathrm{f}}$ inside $(0,1)$, and the first integral in (90) gives $\mathcal{D}_{\mathrm{f}}(\omega_{\mathrm{f}}) = d_{\mathrm{f}}\xi/\mathcal{P}^2 > 0$. Integrating $\mathrm{d}\xi/\xi = -\mathrm{d}\omega_{\mathrm{f}}/\mathcal{D}_{\mathrm{f}}(\omega_{\mathrm{f}})$ between the contacts gives (94), and the outer equations then reduce to (93) and (95) with the phase conditions. The reconstruction of Lemma D.6 returns the curve. The substitutions are in the supplementary material.

The lower bound. Set $\beta_{\mathrm{f}} = 18227/25000$. If K_{f} had a positive value, Lemma D.2 would give an attained positive maximum, which by $(*)$ would be a physical stationary point with $\beta_{\mathrm{f}}G_{\mathrm{f}} - (1-\beta_{\mathrm{f}})M_{\mathrm{f}} - 2 > 0$, contradicting Lemma D.6. A curve with $R_{\mathrm{f}} = \beta_{\mathrm{f}}$ would then minimize R_{f} and contradict the same lemma with the value zero. Hence $R_{\mathrm{f}} > 18227/25000$ on $\mathfrak{F}$.

Attainment and uniqueness. Lemma D.7 provides a curve below $729081/10^6$. A minimizing sequence below this curve's ratio has ratios below $3/4$, so Lemma D.2 places its parameters in $1/32 < m_{\mathrm{f}} < 3/4$, $\xi_0 < 600$, $p_{\mathrm{f}} < 1$, and the extraction in the proof of that lemma gives a subsequence converging to a curve in $\mathfrak{F}$ along which R_{f} converges. Its ratio $r_{\mathrm{fam}} < 3/4$ excludes $p_{\mathrm{f}} = 1$, where $R_{\mathrm{f}} \geq 3/4$. So the minimum r_{fam} is attained. At the trial value $\beta_{\mathrm{f}} = r_{\mathrm{fam}}$, every minimizing curve is a physical stationary point by $(*)$. Lemma D.6 makes this point unique, and (95) fixes the parameter triple. By (89), a maximizer of K_{f} minimizes E_{f} at its own parameters, so the uniqueness in Lemma D.3 fixes the curve.

Self-consistency. Suppose a physical stationary solution in the trial interval has $K_{\mathrm{f}} = 0$. If a different curve had $R_{\mathrm{f}} < \beta_{\mathrm{f}}$, the attained positive maximum of K_{f} would be another physical stationary point by $(*)$, which Lemma D.6 excludes. Thus every self-consistent solution is the family minimum. Its positive contacts put it in the endpoint-slope subclass, and

$$\inf_{\mathfrak{F}} R_{\mathrm{f}} = \inf_{\mathfrak{F} \text{ with endpoint slopes } 1, h_{\mathrm{f}}} R_{\mathrm{f}} = r_{\mathrm{fam}} \in \left(\frac{18227}{25000}, \frac{729081}{10^6}\right).$$

□

D.5 Constant second-buyer bodies

This subsection proves Theorem 13.4. For admissible bodies and $d > 0$ put

$$\mathcal{Q}_d = \sum_{i=1}^{2} \mathbb{E}(Z_i - Y_i)_+ - d\,\mathbb{E}(Y_1 + Y_2). \tag{96}$$

Since $\max\{Z_i, Y_i\} = Y_i + (Z_i - Y_i)_+$, the bound $\mathcal{R}_2(Z,Y) \geq 1/(1+d)$ is equivalent to $\mathcal{Q}_d \leq 2d$, and $1/(1+d_{\mathrm{fam}}) = r_{\mathrm{fam}}$.

The reduction ends in a family that differs from $\mathfrak{F}$ only in that the first seller may have more mass at zero than the second. With the quantities of (63), let the first seller have mass $N_{\mathrm{f}}w$ at zero, where $t_{\mathrm{f}} \leq w < \min\{e_{\mathrm{f}}, p_{\mathrm{f}}\}$, and replace $a_{\mathrm{f}}, c_{\mathrm{f}}, \delta_{\mathrm{f}}$ in (64)–(65) by

$$a_w = \frac{p_{\mathrm{f}} - w}{p_{\mathrm{f}}(w + t_{\mathrm{f}})}, \qquad c_w = \frac{w(p_{\mathrm{f}} - w)}{w + t_{\mathrm{f}}}, \qquad \delta_w = p_{\mathrm{f}} - c_w, \tag{97}$$

and the term $2m_{\mathrm{f}}a_{\mathrm{f}}$ of G_{f} by $N_{\mathrm{f}}(w + t_{\mathrm{f}})a_w$. We write $\mathfrak{F}_{\geq}$, M_w, G_w, R_w for the resulting class and functionals; at $w = t_{\mathrm{f}} < p_{\mathrm{f}}$ they are $\mathfrak{F}$, M_{f}, G_{f}, R_{f}. The trading interpretation of Lemma D.1 carries over, as the supplementary material checks, and a_w is the value of the second body at which the prices below the body also have gain two.

Lemma D.8 (Unequal masses at zero)**.** $\inf_{\mathfrak{F}_{\geq}} R_w = r_{\mathrm{fam}}$.

Proof. The minimizer of Proposition 13.2 has $t_{\mathrm{f}} < p_{\mathrm{f}}$, so it lies in $\mathfrak{F}_{\geq}$, and we must show $R_w \geq r_{\mathrm{fam}}$. We take $\beta_{\mathrm{f}} = r_{\mathrm{fam}}$, so that $d_{\mathrm{f}} = d_{\mathrm{fam}} < 3/8$, and the comparison

$$K_w = \frac{\beta_{\mathrm{f}} G_w - (1-\beta_{\mathrm{f}}) M_w - 2}{\beta_{\mathrm{f}} N_{\mathrm{f}}} = 1 - \frac{w}{p_{\mathrm{f}}} - \log p_{\mathrm{f}} - d_{\mathrm{f}} a_w - d_{\mathrm{f}}(1+t_{\mathrm{f}}) + d_{\mathrm{f}}\xi_0 - E_{\mathrm{f}}(\mathcal{P}),$$

which is (89) at $w = t_{\mathrm{f}}$. Suppose $K_w > 0$ somewhere. The estimates of Lemma D.2, with the first seller's zero mass $N_{\mathrm{f}} w$ in the bounds for the first unit and $R_w \geq m_{\mathrm{f}}$ for the second, give an attained positive maximum on the closed parameter set, which adds $t_{\mathrm{f}} = 0$, $N_{\mathrm{f}} w = 1$ and $p_{\mathrm{f}} = w$. Lemma D.3 uses only the two affine obstacles, so the maximizing curve has the three-arc form, and it is not entirely affine, because $R_w \geq 73/100$ on the two affine curves. Its boundaries are excluded as in Lemma D.4. A zero initial contact with $c_w > 0$ gives $\partial_{p_{\mathrm{f}}} K_w = -c_w w\{\mathcal{P}'_+(c_w) - (w+t_{\mathrm{f}})/w\}^2/\{p_{\mathrm{f}}^2(w+t_{\mathrm{f}})\}$, which is negative unless $t_{\mathrm{f}} = 0$ and $\mathcal{P}'_+(c_w) = 1$; then extending the curve along the first obstacle to a smaller p_{f} keeps K_w and gives a positive contact. A zero final contact gives $\partial_{\xi_0} K_w = \xi_0 \mathcal{P}'_-(\xi_0)^2 > 0$. The faces $p_{\mathrm{f}} = 1$ and $N_{\mathrm{f}} w = 1$ have ratio at least $3/4$, and the face $\delta_w = e_{\mathrm{f}}$ has ratio at least $73/100$.

So both contacts are positive and the derivatives in p_{f} and h_{f} vanish; the latter gives the equation for $q_{r,\mathrm{f}}$ in (93), with C_{f} the constant of the free arc. Put $\Upsilon = (w + d_{\mathrm{f}}) a_w^2$. If $w > t_{\mathrm{f}} > 0$, the price derivative is positive at $p_{\mathrm{f}} = w$, so $a_w > 0$, and the two stationarity conditions turn the mass derivatives into

$$\partial_w K_w = -\frac{\Pi}{w}, \qquad \partial_{t_{\mathrm{f}}} K_w = \Upsilon - d_{\mathrm{f}} + \frac{v_{\mathrm{f}}(v_{\mathrm{f}} - q_{r,\mathrm{f}})}{2e_{\mathrm{f}}(v_{\mathrm{f}} + q_{r,\mathrm{f}})}, \tag{98}$$
$$\Pi = (2t_{\mathrm{f}} + w)\Upsilon - a_w(2d_{\mathrm{f}} + t_{\mathrm{f}} + 2w) + 1.$$

For $0 \leq t_{\mathrm{f}} \leq w$, $2w \leq 1 + t_{\mathrm{f}}$ and $d_{\mathrm{f}} \leq 3/8$, we have $\Pi > 0$ whenever $\Upsilon \leq d_{\mathrm{f}}$. Both masses are free, so $\partial_w K_w = 0$ forces $\Upsilon > d_{\mathrm{f}}$, and then $\partial_{t_{\mathrm{f}}} K_w > 0$ contradicts maximality. If $t_{\mathrm{f}} = 0$, the second seller has no mass at zero, so the second body can sit anywhere below the first body. We use this freedom at the stationary w, where $(w + d_{\mathrm{f}})/\mathcal{P}_\ell^2 = d_{\mathrm{f}}/w^2$; the one-sided derivative in t_{f} is then at least $(\sqrt{w + d_{\mathrm{f}}} + \sqrt{d_{\mathrm{f}}})^{-2} - d_{\mathrm{f}} > 256/625 - 3/8$, since $w < 1/2$. Hence the maximum has $w = t_{\mathrm{f}}$ and lies in $\mathfrak{F}$, where $K_w \leq 0$ by Proposition 13.2, a contradiction. The estimates, the two affine curves, the derivatives, the boundary ratios and the positivity of Π are in the supplementary material. □

Proof of Theorem 13.4. Let $Z_2 \equiv \zeta$ and $d = d_{\mathrm{fam}}$, and suppose, for contradiction, that $\mathcal{Q}_d > 2d$. We modify the bodies in steps that keep them admissible and the second body constant, and that do not lower $\mathcal{Q}_d$ except by an arbitrarily small amount in the first step. They end at the limiting instance of a member of $\mathfrak{F}_{\geq}$, which contradicts Lemma D.8. The main device is to drop the monotonicity of the second seller's CDF above ζ. The first unit's gain can then be raised pointwise along a fixed reciprocal curve, and the one region where the relaxed CDF can fall is removed by a translation.

Finite positive data. Clipping both sellers at the top b of the bodies does not lower $\mathcal{Q}_d$. Rounding the first buyer down and both sellers up to a mesh of width ϵ that contains 0, ζ and b keeps both orders and every price constraint and lowers $\mathcal{Q}_d$ by at most $(3 + 2d)\epsilon$; translating all bodies up by ϵ lowers it by $2d\epsilon$. If rounding lowers the top, we clip again. Next we send each seller value $y \in (0, \zeta)$ to 0 with probability $1 - y/\zeta$ and to ζ with probability y/ζ, using one uniform variable for both sellers. This keeps means and orders. Since $Z_i \geq \zeta$, each gain and surplus term is affine in y below ζ, so $\mathcal{Q}_d$ and the gains at prices above ζ stay, and the new gain at a price $z \leq \zeta$ is at most the old gain at ζ. With $\eta = \min\{1/2, 1/(2 + \mathbb{E}Z_1 + \zeta)\}$, the law $\eta\delta_0 + (1-\eta)\delta_b$ has total gain at most one at body prices, so mixing a small weight of it into both sellers keeps admissibility. We may therefore assume finitely supported bodies with $\zeta > 0$, both sellers supported on $\{0\} \cup [\zeta, b]$, and

masses $0 < m_2 \le m_1$ at zero. The prices $z \le \zeta$ then have gain

$$m_1(1 + \mathbb{E}Z_1) + m_2(1 + \zeta) \le 2. \tag{99}$$

If $b = \zeta$, both bodies equal ζ, and (99) gives $\mathcal{Q}_d \le 2/3 < 2d$; so $b > \zeta$.

A relaxed problem in the reciprocal clock. Write $A_1(s) = \mathbb{E}(s - Y_1)_+$ and $N_1 = \bar{L}_1 F_1 + H_1 A_1$ for the first unit's gain. We use right traces; actual prices use left traces, and each step below keeps their constraints by the same computation. On the body (ζ, b) the second unit's gain is F_2, admissibility reads $N_1 + F_2 \le 2$, seller order reads $F_2 \le F_1$, and

$$\mathcal{Q}_d = (1 + d)\zeta(m_1 + m_2) - 2d\zeta + \int_\zeta^b \{H_1 F_1 + d(F_1 + F_2 - 2)\}\,\mathrm{d}s. \tag{100}$$

We drop the monotonicity of F_2 on the body and keep $m_2 \le F_2 \le F_1$ and the price constraints. For given first-unit laws the best relaxed profile is $\hat{F}_2 = \min\{F_1, 2 - N_1\}$. With $N = 2 - m_2$, put

$$\mathcal{P} = \frac{1}{\bar{L}_1}, \qquad \xi = \frac{A_1}{N\bar{L}_1}, \qquad n = \frac{N_1}{N} \le 1, \qquad \mathcal{V} = \mathcal{P} - \xi\mathcal{P}',$$

with derivatives in ξ. Then $\mathrm{d}\xi/\mathrm{d}s = n\mathcal{P}^2$, $H_1 = n\mathcal{P}'$ and $F_1 = nN\mathcal{V}$, and the polygon $\mathcal{P}$ is concave because $\mathcal{P}'/\mathcal{V} = NH_1/F_1$ does not increase, while at fixed $(\xi, \mathcal{P})$ the ratio $\mathcal{P}'/(\mathcal{P} - \xi\mathcal{P}')$ increases with $\mathcal{P}'$. With $F_2 = \hat{F}_2$, the integral in (100) becomes

$$\int \frac{nN\mathcal{P}'\mathcal{V} + d\{N\mathcal{V} - 2/n + \min(N\mathcal{V}, 2/n - N)\}}{\mathcal{P}^2}\,\mathrm{d}\xi, \tag{101}$$

which does not decrease when n rises at a fixed curve. Since $F_2 \ge m_2$, $H_1 \le 1$ and $F_1 \le 1$, we have $n \le \bar{n} = \min\{1, 1/\mathcal{P}', 1/(N\mathcal{V})\}$. We raise n to $\bar{n}$ and recover prices from $\mathrm{d}s = \mathrm{d}\xi/(\bar{n}\mathcal{P}^2)$, starting at ζ. The curve, the states at ζ and the terminal value $\mathcal{P} = 1$ stay, so (99) and the first terms of (100) stay, and $\mathcal{Q}_d$ does not fall. Because $\mathcal{P}'$ decreases and $\mathcal{V}$ increases, the new laws $H_1 = \bar{n}\mathcal{P}'$ and $F_1 = \bar{n}N\mathcal{V}$ are monotone, and the body splits into three consecutive regimes, each possibly empty, on which respectively $H_1 = 1$, $N_1 = N$ and $F_1 = 1$. In the middle regime $\hat{F}_2 = m_2$. In the last one $\hat{F}_2 = 2 - N_1$ does not decrease, because $\mathrm{d}N_1 = A_1\,\mathrm{d}H_1$ there. So $\hat{F}_2$ can fall only in the first regime.

Removing the first regime. Let $[\zeta, c]$ be the first regime, $c > \zeta$, so that $\bar{L}_1(s) = \bar{L}_1(0) - s$ there. We choose the exit traces

$$\phi_1 = \min\Big\{1, \frac{N - A_1(c)}{\bar{L}_1(c)}\Big\}, \qquad \phi_2 = 2 - \bar{L}_1(c)\phi_1 - A_1(c).$$

The first lies between the one-sided traces of F_1 at c; the relaxed profile after c is at least $\phi_2 \ge m_2$, and $\phi_2 = m_2$ unless $\phi_1 = 1$. Let μ be the first seller's law on $[\zeta, c]$ with its atom at c counted up to ϕ_1, so that μ has mass $\phi_1 - m_1$, and let $I \le (\phi_1 - m_1)(c - \zeta)$ be its first moment about ζ. Then $A_1(c) = \phi_1 c - \zeta(\phi_1 - m_1) - I$, hence $\bar{L}_1(0)\phi_1 - \zeta(\phi_1 - m_1) - I = 2 - \phi_2$. If $\phi_1 > m_1$, we lower the first buyer and both sellers above c by $\Delta = I/\phi_1$, so that $c - \Delta - \zeta \ge m_1(c - \zeta)/\phi_1 > 0$. Below $c - \Delta$ we give the first buyer survival one, the first seller masses m_1 at zero and $\phi_1 - m_1$ at ζ, and the relaxed profile the value ϕ_2 on $(\zeta, c - \Delta)$. The first-unit gain there is $(\bar{L}_1(0) - \Delta - s)\phi_1 + \phi_1 s - \zeta(\phi_1 - m_1) = 2 - \phi_2$, the states at $c - \Delta$ are the old states at c, and the left side of (99) falls by $m_1\Delta$. On the old first regime, $\hat{F}_2 - \phi_2 \le 2 - \phi_2 - N_1(s) = \int_{(s,c]}(\bar{L}_1(0) - u)\,\mu(\mathrm{d}u)$, and $\int (u - \zeta)^2\,\mu(\mathrm{d}u) \ge I^2/(\phi_1 - m_1)$. So

$$\Delta\mathcal{Q}_d = d(2 - \phi_2)\Delta - d\int_\zeta^c (\hat{F}_2 - \phi_2)\,\mathrm{d}s \ge \frac{dIm_1\{\zeta(\phi_1 - m_1) + I\}}{\phi_1(\phi_1 - m_1)} \ge 0.$$

If $\phi_1 = m_1$, nothing moves, and raising $\hat{F}_2$ to ϕ_2 on the first regime is feasible and raises $\mathcal{Q}_d$.

The two branches. The relaxed profile is now a CDF: it is constant through the middle regime and does not decrease after it. If $F_1 = 1$ on the whole body, we are in the second branch below. Otherwise the profile starts at m_2 above ζ, so the second seller has no atom at ζ, and the total gain on the body is two.

In the first branch we freeze $n = 1$ at the fixed curve: $H_1 = \mathcal{P}'$, $F_1 = N\mathcal{V}$, and the second seller becomes $m_2\delta_0 + (1 - m_2)\delta_{\tilde{b}}$ at the new top $\tilde{b}$. In the clock, the body's part of $\mathbb{E}(Y_1 + Y_2)$ has density $N(1 - \mathcal{V})/\mathcal{P}^2$, which does not depend on n, and the efficient gain has density $nN\mathcal{P}'\mathcal{V}/\mathcal{P}^2$, which does not decrease in n. Where $N\mathcal{V}$ would exceed one, we replace $\mathcal{P}$ after its first crossing by the supporting line on which $N\mathcal{V} = 1$, up to the point where it reaches one. This line lies above $\mathcal{P}$, so the seller mean does not rise. It saturates the first seller, so the remaining efficient gain at the crossing is $1/\mathcal{P} - 1$, which bounds the old one. The first-unit gain is now N on the body, the first buyer's mean is unchanged, and $\mathcal{Q}_d$ has not fallen.

In the second branch $\hat{F}_2 = 2 - N_1$ on the body. Let $h_1 = H_1(\zeta)$ be the first buyer's mass above ζ, which is positive because $b > \zeta$. Replacing this part of its law by an atom of mass h_1 at its mean keeps the states at ζ and the price constraints and raises $\mathcal{Q}_d$ by d times its variance mass. With b now this mean,

$$Z_1 \sim (1 - h_1)\delta_\zeta + h_1\delta_b, \quad Y_1 \sim m_1\delta_0 + (1 - m_1)\delta_\zeta, \quad Y_2 \sim m_2\delta_0 + (f_2 - m_2)\delta_\zeta + (1 - f_2)\delta_b,$$

where $f_2 = 1 - h_1(b - \zeta + m_1\zeta) \geq m_2$ and $\mathcal{Q}_d = (1+d)\zeta(m_1+m_2) - 2d\zeta + h_1(b-\zeta)\{1 - d(b - \zeta + m_1\zeta)\}$. Keeping $h_1(b - \zeta)$ fixed and lowering b raises $\mathcal{Q}_d$ by $dh_1(b - \zeta)$ per unit, so we lower b until $h_1 = 1$ or $f_2 = m_2$. At $h_1 = 1$, $\mathcal{Q}_d$ is affine in the masses on a compact polygon, with slopes $(1 + d)\zeta$ in m_2 and $\{1 + d(1 - b + \zeta)\}\zeta > 0$ in m_1, since $f_2 \geq 0$ gives $b - \zeta \leq 1$. So a point with $f_2 > m_2$ is improved by raising m_2, by trading m_1 for m_2 along (99) at the rate $(1 + \zeta)/(1 + b) < 1$, or by raising both, and the corner $m_1 = m_2$ with equality in (99) has $f_2 < m_2$. These moves all raise m_2, so a maximum over the masses with m_2 at least its current value has $f_2 = m_2 > 0$. Both branches end with $N_1 = N$ and $F_2 = m_2$ on the body.

Lower-price gain two. If (99) is strict, we raise the lower-price gain without lowering $\mathcal{Q}_d$. When $m_2 \leq 2d/(1 + d)$, we fix the curve and $A_1(\zeta) = m_1\zeta$ and lower ζ, translating the body. The first two terms of (100) equal $(1 + d)A_1(\zeta) + \zeta\{(1 + d)m_2 - 2d\}$ and do not fall, and once $m_1 = A_1(\zeta)/\zeta$ reaches $F_1(\zeta+)$, the left side of (99) exceeds two by $A_1(\zeta)\{1 - H_1(\zeta)\} + m_2A_1(\zeta)/F_1(\zeta+) \geq 0$. When $m_2 > 2d/(1 + d) \geq 1/2$, we keep both masses and start the body later on the same curve, at $\zeta = N\xi/\{m_1\mathcal{P}(\xi)\}$. Then

$$\frac{\mathrm{d}\mathcal{Q}_d}{\mathrm{d}\xi} = \frac{F_1\{1 - H_1 + ((1 + d)m_2 - 2d)/m_1\} + dN}{\mathcal{P}^2} > 0,$$

the lower-price gain also increases in ξ, and at the end of the curve it is at least $2 + (m_1 - m_2)(m_1 + m_2 - 1)/m_1 \geq 2$. In both cases we stop when (99) holds with equality.

The bodies are now the limiting instance of a member of $\mathfrak{F}_{\geq}$, with $N_\mathrm{f} = N$, $m_2 = N_\mathrm{f}t_\mathrm{f}$, $m_1 = N_\mathrm{f}w$ and $p_\mathrm{f} = 1/\bar{L}_1(\zeta)$. Their laws on the body are those of (88), and equality in (99) is the formula for a_w. The supplementary material checks that this is a member with $\mathcal{R}_2(Z, Y) = R_w$. Since $\mathcal{Q}_d > 2d$, the ratio of this member is below r_fam, contradicting Lemma D.8. The same computation at $w = t_\mathrm{f}$ shows that the bodies of the family minimizer are admissible with $\mathcal{R}_2 = r_\mathrm{fam}$, so the bound is attained. □

E The two-unit instance

We approximate the three-piece curve of Proposition 13.2 by a rational concave polygon with 128 interior segments. Its initial parameters, rounded to six decimal places, are

$$p_{\mathrm{w}} \approx 0.671601, \quad \xi_{\ell,\mathrm{w}} \approx 0.258131,$$
$$\xi_{r,\mathrm{w}} \approx 0.429836, \quad \xi_{0,\mathrm{w}} \approx 0.675567.$$

For consecutive polygon vertices (X_j, P_j), set $h_j = (P_{j+1} - P_j)/(X_{j+1} - X_j)$. The physical-value increment is $(X_{j+1} - X_j)/(P_j P_{j+1})$; the first buyer survival is h_j, and the first seller CDF is $(2 - m_{\mathrm{w}})(P_j - X_j h_j)$. The finite marginal laws are obtained by rounding values and cumulative probabilities to fixed rational grids, specified with the full vertices and weights in the supplementary material. The table rounds body parameters to six decimal places.

Quantity	Value
Second buyer body	$a_{\mathrm{w}} \approx 0.828691$, mass $1 - \eta_{\mathrm{w}}$
Common buyer tail	value 10^{10}, mass $\eta_{\mathrm{w}} = 10^{-10}$
Second seller, lower atom	10^{-12}, mass $m_{\mathrm{w}} \approx 0.482351$
Second seller, upper atom	≈ 1.525881, mass $1 - m_{\mathrm{w}}$
Marginal support sizes	$131, 2$ for the buyer; $131, 2$ for the seller

Common-quantile coupling within each side gives 131 ordered vector types on each side. We draw the buyer and seller vectors independently. The seller displacement of 10^{-12} separates coincident body atoms. Every value is strictly positive.

Proof of Theorem D. We sum welfare exactly at every valuation event of the rational instance. Each type pair contributes its marginal gain on the closed price interval $[S_i, B_i]$. The inclusive event value contains all trades available in an adjacent open interval; welfare is constant between events and equals initial welfare outside the support. Thus the 264 event values cover every real price. Their maximum divided by efficient welfare is a rational number strictly below 0.7290804.

For the symmetric instance, we use the supplementary rational law for the seller vector and its coordinate reversal for the independent buyer vector. Exact evaluation of its 30 valuation events gives a ratio below 0.83693. Together with Theorem A, the two finite instances give

$$r_2^* < 0.7290804 < 0.738024 < \beta_*, \qquad r_{2,\mathrm{sym}}^* < 0.83693.$$

□

F Routes we tried

Our aim was the exact constant together with the shape of the worst case. A sharper bound would settle neither question, which ruled out the first two routes below.

Finer finite programs. The lower bounds 0.72 and 0.7292 come from finite programs [CW23, GdK26]. Refining their grids enlarges a nonconvex program, and the refined program still certifies only a lower bound, without deciding whether the worst ratio is attained or what the worst case looks like. We did not pursue this route.

A numerical verification function. The criterion of [LRW23, Theorem 3.3] reduces a guarantee to an upper bound on $\sup \mathcal{J}_d$, with $\mathcal{J}_d$ as in (6). Our first plan was to treat this maximization as a control problem with state (h, g, k), fit a value function to a dynamic-programming solution, and check its Hamilton–Jacobi–Bellman inequalities by interval arithmetic, which certifies a lower bound. Numerical solutions had already located the candidate value. We set the certificate aside before running it, because it could prove only a bound below β_* and would say nothing about the worst case.

The exponential trial family. The upper bound of [LRW23, Theorem 3.8] uses a buyer whose survival function has the form $\lambda_0 + \lambda_1 e^{\lambda_2 s}$ between two constant segments, and tuning its parameters brings the ratio very close to β_*. The family was therefore a natural guess for the worst case. The exact middle arc is not of this form on any subinterval (Proposition 8.1), so the family gives upper bounds only.

Concavity in the original control. If $\mathcal{J}_d$ were concave, the first-order equation of [LRW23, Theorem 3.4] would be sufficient for a maximum. It is not: Appendix C.2 gives two controls whose midpoint has value below the average of their values by $11/1728$. This led us to look for coordinates in which the problem becomes concave after a few boundary values are fixed. The reciprocal clock (Section 5) gives a relaxation that is concave once the endpoint is fixed, and the endpoint envelope (Section 7) compares the endpoints.

The proof of Theorem C. We worked out two arguments for the value and the shape. The first proves that $\mathcal{J}_d$ attains its maximum over the measurable box, shows that a maximizer is monotone, and uses necessary conditions of Pontryagin type to exclude every shape except the one in Theorem C. The second is the calibration of Sections 5–8, which writes the deficit of an arbitrary measurable control as a sum of nonnegative terms. We kept the second, which compares every competitor directly and so needs neither the existence nor the regularity of a maximizer. Its equality case also gives the shape and its uniqueness in the same step.

The price representation. For Corollary A′ we first planned to cite the randomized-price representation of [HR87, ČP16] in its pointwise form. Pointwise statements depend on tie conventions at price atoms, and the smoothed hard pairs need only a welfare comparison under absolutely continuous priors, so we proved the almost-everywhere form of Proposition 12.2 directly.

Other mechanisms. We also tried the comparison on neighboring constants. For the gains-from-trade guarantee of the random-offerer mechanism [DMSW22, Jo26], we derived the singular-arc equation for the seller when the buyer is restricted to equal-revenue distributions. An explicit finite example shows that the worst buyer for a fixed seller need not be equal-revenue, and the reduction that changes both sides together remained open. For buyer-offering with a reserve [DS26], the kernel that bounds the guarantee contains a tail integral of an unknown threshold. We found neither a calibration in the same coordinates nor a reduction of the seller, and the lower and upper bounds we computed did not meet. The version in which only one prior is known gives no second constant, since its ratio is again β_* by [CW23, Theorem 4.1].