Efficient Learning under Competition

Brian C. Albrecht[*] Mark Whitmeyer[†]

April 21, 2025

Abstract

This paper examines the efficiency of markets where consumers privately acquire costly information about product fit before purchasing. In a departure from the inefficiency result found in the monopoly setting of Ravid, Roesler, and Szentes (2022), we show that duopolistic competition can restore *ex post* efficiency as information costs vanish. The core economic mechanism is that competition limits firms' ability to fully extract consumer surplus based on learned preferences (mitigating the hold-up problem), thereby preserving consumer incentives to acquire information that leads to efficient matching, even when information is arbitrarily cheap. We also find that the relationship between information costs and consumer welfare is non-monotonic: reducing frictions can harm consumers when information is already cheap, by intensifying the hold-up problem that competition only partially resolves.

[*]International Center for Law & Economics. Email: mail@briancalbrecht.com

[†]Arizona State University. Email: mark.whitmeyer@gmail.com

This paper was previously titled "Comparison Shopping: Learning Before Buying from Duopolists." We thank Gary Biglaiser, Andy Kleiner, Fei Li, Jiangtao Li, David Rahman, Doron Ravid, Tom Wiseman, Jingyi Xue, Kun Zhang, Jidong Zhou, and seminar and conference audiences at the 2021 SEA Conference, the 2023 IIOC Conference, and Singapore Management University for their comments. Cole Williams provided exceptionally helpful feedback in his role as a discussant at the 2023 IIOC Conference. Part of this paper was written while MW was visiting the department of economics at Duke University, and he thanks them for their generous hospitality.

1 Introduction

Consumers often face uncertainty about which product best suits their needs and can invest effort in learning before making a purchase. This investment in information, such as acquiring knowledge about product fit, is crucial in many markets, from electronics to professional services. How does this unique form of investment interact with market structure, particularly competition, to shape market outcomes?

A key benchmark is the monopoly model studied by Ravid, Roesler, and Szentes (2022) (henceforth RRS). They analyze a setting where a buyer can flexibly acquire any costly signal about her valuation before the monopolist seller makes a take-it-or-leave-it offer. Their striking result is that as the cost of information vanishes, the market outcome converges to an inefficient equilibrium. This inefficiency stems from an informational hold-up problem: the buyer anticipates that the seller will appropriate most of the surplus generated by her costly investment in information about product fit. Knowing that the value created by learning will be largely captured by the seller diminishes the incentive to undertake such learning, leading to underinvestment in knowledge acquisition even when information becomes nearly free. This highlights a specific form of hold-up where the relationship-specific investment is intangible–the acquisition and processing of information itself.

Does competition alter this pessimistic conclusion? This paper investigates whether introducing competition between sellers overturns the inefficiency identified by RRS. We study a model with two horizontally differentiated firms competing in prices, facing a representative consumer who can flexibly and privately acquire costly information about her match value with each firm before prices are set.

Our main finding provides a clear contrast to RRS: **competition restores efficiency in the limit of vanishing information costs by mitigating the informational hold-up**. Specifically, we construct an equilibrium that converges to one where the consumer always purchases the product that provides her higher utility *ex post*, as the cost of information becomes vanishingly small. Competition breaks the severe informational hold-up present in the monopoly case. Because firms set prices simultaneously without observing the consumer's private signal realization, they cannot fully appropriate the value created by the consumer's investment in learning. This preservation of returns to information investment maintains the incentive to learn about the relative merits of the products, ultimately leading to efficient choices when learning is sufficiently cheap.

This result highlights a novel benefit of competition specifically in the context of flexible, private pre-purchase investment in information. While prior work has explored the effects of information on competition (e.g., Moscarini and Ottaviani (2001), where information is exogenous), information acquisition under competition when prices are posted before learning (Matějka and McKay, 2012), or how competition affects hold-up problems more generally (e.g., Felli and Roberts (2016), in a different investment and market setup),[1] our analysis focuses on the specific interaction between price competition and the buyer's endogenous choice over information structures **before** observing prices. It is this interaction that resolves the limiting inefficiency driven by informational hold-up identified by RRS.

However, competition does not eliminate the underlying tension completely. Outside of the limiting case, we find an ambiguous effect of information frictions on consumer welfare. When information costs are high, reducing these costs benefits consumers as they can make better-informed decisions without drastically increasing firms' market power. But when information costs are already low, further reductions can actually harm consumers. Cheaper information induces consumers to learn more precisely, which, while improving match quality, also allows firms to extract more surplus through prices (softening price competition). Below a certain threshold, this price effect dominates the improved matching effect, making consumers worse off as information becomes nearly free. This highlights that even with competition, a residual hold-up problem persists and interacts subtly with the cost of information.

[1] Armstrong and Zhou (2022) construct firm- and consumer-optimal information structures under duopolistic competition.

1.1 Related Work

Our paper contributes primarily to the growing literature on costly information acquisition in markets and its implications for efficiency, particularly concerning the interaction between learning incentives and market structure.

The most closely related work is Ravid, Roesler, and Szentes (2022) (RRS), who study a bilateral trade (monopoly) setting where the buyer can flexibly acquire costly information about her valuation before the seller makes a take-it-or-leave-it offer. Their main result establishes that as information costs vanish, the equilibrium converges to an inefficient outcome. This inefficiency stems from a hold-up problem: the buyer anticipates the seller will extract most of the surplus revealed by any acquired information, which diminishes her incentive to become fully informed, even when doing so is nearly free. Our paper directly addresses this finding by asking whether competition can resolve this inefficiency. We show that under duopoly, there is a limiting outcome that *is* efficient, identifying the mitigation of this informational hold-up via price competition as the key mechanism.

We also connect to the broader literature on information acquisition by consumers. Many papers explore this in monopoly settings (e.g., Branco et al. (2012, 2016); Martin (2017); Pease (2023); Lang (2019)). A crucial distinction is often the timing of information acquisition relative to price observation. In Matějka and McKay (2012) the consumer learns both about her value and the terms of trade. In auction/mechanism design settings like Persico (2000) and Shi (2012), the terms of trade (price or mechanism) are known before the consumer decides what information to acquire. In Cusumano, Fabbri, and Pieroth (2024) and Ravid (2020), the consumer learns only about the terms of trade. These alternative specifications fundamentally alter the strategic problem compared to our "learning before trading" timing (adopted from RRS), where the hold-up problem related to unknown future prices is central and the only learning is about the consumer's valuation.

Our work complements contemporaneous research exploring related themes. Biglaiser, Gu, and Li (2024) study a duopoly model where firms set prices and consumers learn simultaneously. While related, their model differs in its single-dimensional state space (value 1 for one firm, 0 for the other) and focuses primarily on product differentiation and platform design rather than the limiting efficiency question we address as a contrast to RRS. The multi-dimensional nature of uncertainty in our model is key to our efficiency result. In his study of costly voter learning, Vaeth (2023) derives a similar collapse of dimensions in agents' learning as in our comparison-shopping theorem. Jain and Whitmeyer (2021) also consider flexible private learning before prices but in a large oligopoly with search frictions à la Wolinsky (1986).

Finally, from a technical perspective, solving our model requires analyzing equilibrium pricing involving mixed strategies and solving a multi-dimensional information design problem where the value of information is endogenous to the firms' anticipated pricing strategies. To accomplish this, we use recent technical results and insights from Dworczak and Kolotilin (2024), Yoder (2021), and Kleiner, Moldovanu, Strack, and Whitmeyer (2025). In turn, holding fixed the consumer's learning (when information is expensive), the pricing game faced by the firms in §4.1 is essentially that faced by the firms in Moscarini and Ottaviani (2001), who study the pricing-only game between firms facing an exogenously (and privately) informed consumer.

The rest of the paper is as follows: §1.1 covers related work; §2 sets up the model; §3 presents the binary-state version of RRS as another benchmark; §4 solves the optimal pricing game, conjecturing the optimal learning; §5 solves for optimal learning, given the equilibrium pricing, and proves the main results; and §6 discusses our assumptions, illustrates how our results extend when the consumer's prior has a density (§6.1) and sets up the benchmark where sellers observe what the buyers learn (§6.2).

2 Model

We analyze a market with two horizontally differentiated firms ($i = 1, 2$) and a representative consumer. The consumer's value for firm i's product, Z_i, is initially unknown. To simplify the analysis and isolate the effects of competition on learning incentives, we assume Z_i is a binary random variable with support $\{0, 1\}$ and $\mathbb{P}(Z_i = 1) = \frac{1}{2}$. The values Z_1

and Z_2 are symmetrically distributed: $\mathbb{P}(Z_1 = 1|Z_2 = v) = \mathbb{P}(Z_2 = 1|Z_1 = v)$ for $v \in \{0,1\}$. Let $\omega := \mathbb{P}(Z_1 = 0, Z_2 = 1)$. We assume a **Full-Support Prior**: $0 < \omega \leq \frac{2}{5}$. In § 6.1, we allow for a continuous (mean $\frac{1}{2}$) prior.

The consumer privately learns about the state of the world at a cost, formalized as follows. Let $\mathcal{F}$ denote the set of all distributions supported on $[0,1]^2$ that are fusions (mean-preserving contractions or MPCs) of the prior, i.e., that can be obtained by observing some signal. The consumer may acquire any fusion $F \in \mathcal{F}$ at cost $C\colon \mathcal{F} \to \mathbb{R}$, where C satisfies the following assumptions:

Assumptions on the Cost Functional: We assume for any $F \in \mathcal{F}$,

$$C(F) = \kappa \int c dF ,$$

where $\kappa > 0$ is a scalar and c takes the form

$$c(x,y) = \varphi(x) + \varphi(1-x) + \varphi(y) + \varphi(1-y),$$

where φ is strictly convex and thrice differentiable and also satisfies

(i) $\varphi(z) < \infty$ for all $z \in (0,1)$,

(ii) $\varphi'''(1-z) - \varphi'''(z) \geq 0$ for all $z \in \left(0, \frac{1}{2}\right)$, and

(iii) $\lim_{z' \downarrow 0} |\varphi'(z')| = \infty$.

The following function (which we use for our figures) fits our specification:

$$c(x,y) = x \log x + (1-x)\log(1-x) + y \log y + (1-y)\log(1-y) - 2\log\frac{1}{2}. \tag{1}$$

We also specify that the consumer's utility is additively separable in her value for the good she purchases, its price, and her cost of acquiring information: if she purchases a product with expected value x at price p and at posterior (x,y), her utility is $x - p - \kappa c(x,y)$. We specify **Full Market Coverage**: the consumer has a negligible (or nonexistent) outside option and so will always purchase from one of the firms. For simplicity, we set the marginal costs of production for the two firms to 0.

The timing of the game is straightforward:

(i) **Private Learning:** The consumer acquires information about the two products. Neither firm observes this learning.

(ii) **Simultaneous Price Setting:** The firms simultaneously post prices.

(iii) **Purchase Decision** Given posterior value (x, y) and prices (p_1, p_2), the consumer purchases from Firm 1 (2) if $x - p_1 > (<) \; y - p_2$.

In the first (information acquisition) stage, the consumer solves

$$\max_{F \in \mathcal{F}} \int (u - \kappa c) \, dF,$$

where $u\colon [0,1]^2 \to \mathbb{R}$ is the consumer's reduced form utility from posterior (x, y).

3 Benchmark: Informational Hold-up under Monopoly

To establish the benchmark inefficiency result that competition potentially overcomes, we first analyze the monopoly scenario studied by RRS, adapted to our binary state space $\{0,1\}$. This clarifies the informational hold-up mechanism in the absence of competition. When only one seller exists, it possesses significant power to extract surplus based on the information the buyer acquires.

The buyer has value 1 for the seller's good with probability $\mu \in (0,1)$. Otherwise, her value is 0. She can privately invest in information about this value–a distribution over posteriors F–at cost

$$C(F) := \kappa \int (c(x) - c(\mu)) dF(x),$$

where $\kappa > 0$ is a scaling parameter and c is a strictly convex function that is continuously differentiable on $(0,1)$. Moreover, further translating our assumptions on the cost functional to the single-firm environment in the natural manner, we specify $\lim_{x \downarrow 0} |c(x)| = \lim_{x \uparrow 1} |c(x)| = \infty$.

Following RRS, we assume that the consumer has an outside option of 0. Since the seller makes a take-it-or-leave-it offer p *after* learning occurs (though without observing the specific realization x), it can strategically set the price anticipating the buyer's maximum willingness to pay conditional on purchase, thereby capturing much of the value the information investment created.

To characterize the equilibrium, we define the following distribution of posteriors

$$F(x) = 1 - \frac{\underline{x}}{x}, \text{ on } [\underline{x}, \bar{x}], \tag{2}$$

and distribution of prices

$$G(p) = \kappa\left(c'(p) - c'(\underline{x})\right), \text{ on } [\underline{x}, \bar{x}], \tag{3}$$

where $\underline{x}$ and $\bar{x}$ are pinned down by the two equations

$$1 = \kappa\left(c'(\bar{x}) - c'(\underline{x})\right), \quad \text{and} \quad 1 + \log\frac{\bar{x}}{\underline{x}} - \frac{\mu}{\underline{x}} = 0,$$

which correspond simply to the requirements that G is a cumulative distribution function (CDF) and F is Bayes-plausible.[2]

Proposition 3.1 (Monopoly Equilibrium)**.** *If $\kappa > 0$, the unique equilibrium involves the consumer randomizing over posteriors according to distribution F and the firm randomizing over prices according to distribution G defined by Expressions 2 and 3.*

The proof is adapted from RRS; see Appendix §A.1.

Our primary interest is the limit as information becomes cheap ($\kappa \downarrow 0$), corresponding to the main result in RRS:

Proposition 3.2 (Monopoly Limiting Inefficiency)**.** *As $\kappa \downarrow 0$, the monopolist's price converges to $p = 1$. The consumer's posterior distribution F converges to the truncated Pareto distribution on $[a, 1]$, where $a \in (0, \mu)$ uniquely solves $1 - \log(a) - \mu/a = 0$.*

See Appendix §A.2.

Crucially, this limiting outcome is *ex post* inefficient. Even though information is virtually free, the consumer does not fully learn the state. With positive probability (specifically, when her posterior $x \in [a, 1)$), her true value is 1, yet she does not purchase because the price is $p = 1 > x$.

[2] Bayes-plausibility is equivalent to the feasibility of obtaining the specified distribution over posteriors under some learning strategy, i.e., that the consumer's belief is a martingale.

This failure to learn perfectly, even when information is free, is the essence of the informational hold-up problem: the buyer's return on investing in information is expropriated *ex post* by the seller's pricing strategy, destroying the incentive for efficient investment in knowledge. We now turn to whether duopoly competition alters this inefficient outcome.

4 Pricing Subgame Under Duopoly Competition

Having established the monopoly benchmark where informational hold-up leads to inefficiency, we now analyze the duopoly case. Eventually, we will characterize the symmetric equilibria of the game with flexible learning by the consumer and price-setting by the firms. To do that, we solve the game by backward induction, and so we start by examining two games of pure price setting. In this "second stage," we are holding fixed the consumer's learning; *viz.*, fixing some conjectured distribution over consumer valuations for the two firms' products. With this conjecture fixed, we solve for the equilibrium pricing strategy. In §5, we then show these two potential distributions of beliefs are correct equilibrium distributions of valuations, given the consumer's correct conjectures of the firms' pricing strategies.

To elaborate, as we will later verify, when information costs are sufficiently *high*, there is an equilibrium in which the consumer's learning has symmetric binary support. Accordingly, the analysis in §4.1 characterizes the equilibrium pricing by the firms when the consumer learns in this manner–naturally, the consumer's distribution over posteriors is an equilibrium object, correctly anticipated by the firms; likewise, the firms' pricing will be conjectured by the consumer in her information-acquisition problem. When information costs are sufficiently *low*, there is an equilibrium in which the consumer's learning has symmetric ternary support. In §4.2, we characterize the equilibrium pricing by the firms who anticipate this variety of learning by the consumer.

4.1 Pricing with Symmetric Two-Point Support

Suppose the consumer has a symmetric distribution over valuations for the firms' products with support on $(0, \lambda)$ and $(\lambda, 0)$, each with probability $\frac{1}{2}$, and where $\lambda > 0$. Getting slightly ahead of ourselves, these beliefs arise from comparison shopping where the consumer learns only that she likes one of the goods more than the other by λ. This game is a modified, full-market-coverage, version of Moscarini and Ottaviani (2001). By standard, undercutting arguments with finite valuation types, firms must randomize in any equilibrium.

Lemma 4.1. *There exist no symmetric equilibria in which firms do not randomize over prices.*

The proof is in §A.3. As a result, we search for an equilibrium in which firms randomize over prices. We define the following piecewise distribution of prices, which is the equilibrium distribution:

$$\Gamma(p) = \begin{cases} \Gamma_L := \frac{p-\sqrt{2}\lambda}{\lambda+p}, & \sqrt{2}\lambda \leq p \leq \left(1+\sqrt{2}\right)\lambda \\ \Gamma_H := \frac{\left(3+\sqrt{2}\right)\lambda-2p}{\lambda-p}, & \left(1+\sqrt{2}\right)\lambda \leq p \leq \left(2+\sqrt{2}\right)\lambda \end{cases}. \tag{4}$$

Proposition 4.2. *In the price-setting game of this subsection, the unique symmetric equilibrium is for each firm to choose the distribution over prices* Γ *specified in Expression 4.*

The proof is in §A.4. This equilibrium distribution arises as the distribution that generates unit-elastic demand for the other firm (given the conjectured uncertainty), rendering it willing to randomize.

4.2 Pricing with Symmetric Three-Point Support

Now suppose the consumer has a symmetric distribution over valuations for the firms' products with support on 3 points as follows. For 2 points, her valuation for one of the firms is $\lambda > 0$ greater than that of the other firm, just as in the previous subsection. At the 3rd point, the consumer is indifferent between each of the firms.

After normalization, we specify that with probability $q \leq \frac{2}{5}$ the consumer's vector of valuations for the 2 firms is $(0, \lambda)$ and with probability $1 - 2q$ the consumer's vector of valuations is $(0,0)$. As in the case with two-point support, any equilibrium must involve randomization over prices.

Lemma 4.3. *There exist no symmetric equilibria in which firms do not randomize over prices. Moreover, firms' distributions over prices cannot have atoms.*

The proof is in §A.5. Again, we search for an equilibrium in which firms randomize over prices. Define the distribution $\Phi(p)$

$$\Phi(p) := \frac{(1-q)(p(1-2q)-\lambda q)}{p(1-2q)^2} \quad \text{on} \quad \left[\frac{q}{1-2q}\lambda, \frac{q}{1-2q}\lambda + \lambda\right]. \tag{5}$$

Proposition 4.4. *In the price-setting game of this subsection, it is an equilibrium for each firm to choose the distribution over prices Φ specified in Expression 5.*

The proof is in §A.6.

5 Equilibrium Information Investment and Efficiency

We now solve the consumer's problem: choosing an optimal information structure $F \in \mathcal{F}$ to maximize expected utility, anticipating the competitive pricing behavior ($\Gamma(p)$ or $\Phi(p)$) derived in §4. This section shows how the consumer's value function, shaped by the mitigated hold-up under competition, leads to learning outcomes that contrast sharply with the monopoly benchmark, culminating in our main result: the restoration of limiting efficiency.

First, we establish that as long as there exist frictions, no matter how small ($\kappa > 0$), the consumer only learns along the **Comparison Shopping** line $y = 1 - x$. That is, the consumer's learning exclusively focuses on the *relative* merits of each firm's product. Defining the set ℓ^* as

$$\ell^* := \left\{(x,y) \in [0,1]^2 \ : \ y = 1 - x\right\},$$

we say that the consumer **Comparison Shops** if her acquired distribution over posteriors is supported on a subset of ℓ^*. Figure 1 illustrates the space of valuations and the comparison shopping line. Posteriors in the red (blue) region are those at which the consumer's valuation for firm 2's (1's) good is highest. The dotted diagonal line is the comparison shopping line when the prior is the dot, which since the prior is symmetric is on the boundary of the red and blue regions. In equilibrium, the consumer only learns along the dotted, comparison shopping line.

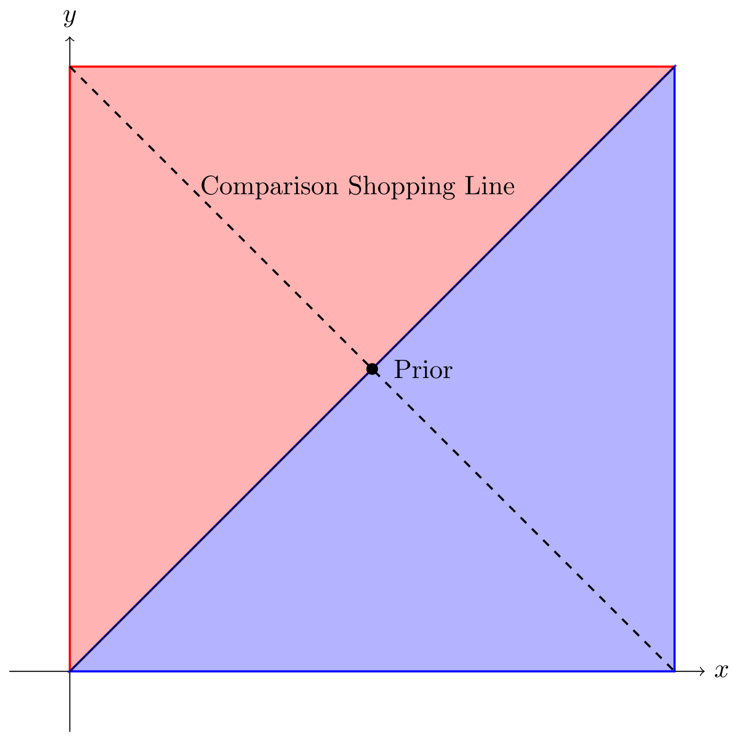

Figure 1: The Space of Valuations

Theorem 5.1. *If firms choose symmetric, atomless, distributions that admit densities with support on some closed interval* $\left[\underline{p}, \bar{p}\right]$*, the consumer comparison shops.*

The proof is in §A.7. The crucial observation behind this theorem is that the consumer's payoff as a function of her posterior is strictly concave along the vector $(1,1)$ and maximized along this vector by points on the comparison-shopping line. Bayes' plausibility then pins down the comparison shopping line $1-x$.

5.1 Solving the Information Investment Problem

In solving the consumer's problem, we conjecture its solution and use the corresponding strategies by the firms in the pricing-only game to generate the consumer's value function.

Then, we verify that the consumer's optimal learning is precisely the two- or three-point support that we conjectured in §4.

The value function for the consumer from acquiring belief (x, y) is

$$V(x,y) := \mathbb{P}(x - p_1 \geq y - p_2)\,\mathbb{E}(x - p_1 | x - p_1 \geq y - p_2) + \mathbb{P}(y - p_2 \geq x - p_1)\,\mathbb{E}(y - p_2 | y - p_2 \geq x - p_1) - \kappa c(x,y),$$

which is continuously differentiable except on $\partial[0,1]^2$ (because firms randomize continuously over prices) and is bounded above on the entire square. Accordingly, by Theorem 1 of Dworczak and Kolotilin (2024), we have weak duality and the price function solution lies weakly above the consumer's value in her information acquisition problem.[3] From there, it is easy to solve the dual problem and verify that this corresponds to a solution to the primal problem.

Intuitively, we can make use of the symmetry of the information acquisition problem and "split" the prior probabilities of $(1, 1)$ and $(0, 0)$ equally between the triangles

$$\Delta^1 := \left\{(x,y) \in [0,1]^2 : 0 \leq x \leq 1 \;\&\; 0 \leq y \leq x\right\}, \tag{6}$$

and

$$\Delta^2 := \left\{(x,y) \in [0,1]^2 : 0 \leq x \leq 1 \;\&\; 1 \geq y \geq x\right\}; \tag{7}$$

before solving two standard 3-state persuasion problems (as any simplex is homeomorphic to the standard simplex), on each of the two triangles. These problems satisfy the assumptions of Yoder (2021), whose Proposition 2 proves that the concavification approach is valid. It remains to verify that the maximum of the two concavifying planes is convex and lie every above the value function, and that the two planes either are the same or have $y = x$ as their intersection.

The solution is then as follows. If information is expensive, the price function is just a single plane; i.e., the two concavifying planes are the same plane. If information is moderately expensive, the price function is the maximum of two planes that intersect at

[3] It is important to keep in mind that we are using the Dworczak and Kolotilin (2024) terminology: the "price function" is a multiplier on the MPC/fusion constraint in the information acquisition problem and has nothing to do with the actual prices in our model set by the firms.

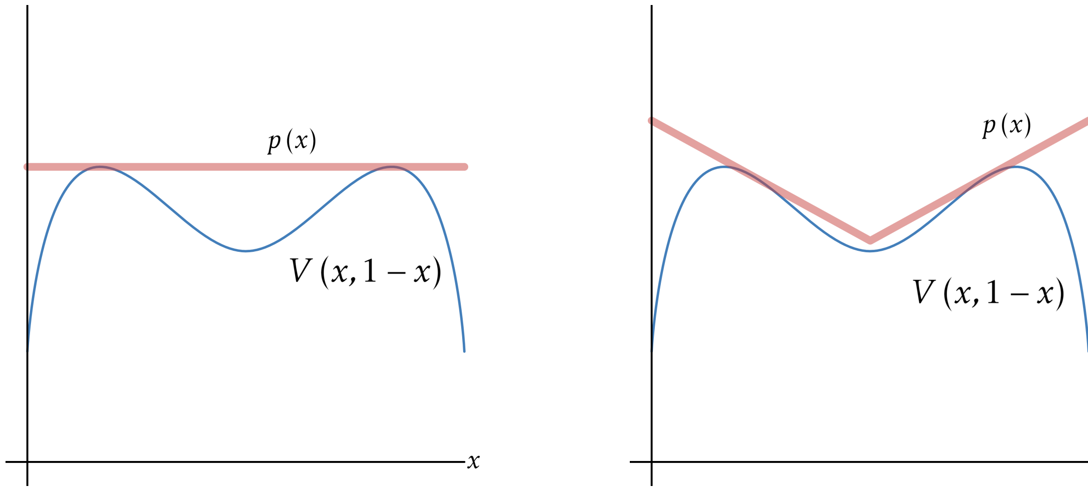

(a) Expensive Information: $\kappa \geq \bar{\kappa}$ (b) Moderately Costly Information: $\kappa \in [\underline{\kappa}, \bar{\kappa}]$

Figure 2: Two possible price functions when information is not cheap.

$y = x$ and lie weakly above the value function on that line. Finally, if information is cheap, the price function is as in the moderate cost case, with the additional specification that it is equal to the value function at its minimum, along the line $y = x$.

5.2 High Information Costs: Limited Learning

The easiest scenario to analyze is that in which either information is expensive (κ is large). Our main result of this section is that if frictions are sufficiently large, then there is an equilibrium in which the consumer acquires a binary distribution over posteriors.

We say that a consumer **Comparison Shops With Uniform Two-point Support** if the consumer's acquired distribution over valuations is supported on

$$\left\{\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right), \left(\frac{1+\lambda}{2}, \frac{1-\lambda}{2}\right)\right\},$$

each with probability $\frac{1}{2}$.

Theorem 5.2. *If κ is sufficiently high, there is an equilibrium in which the consumer comparison shops with uniform two-point support and firms randomize over prices according to Expression 4.*

The proof is in §A.8. Notably, when κ is sufficiently large, λ is strictly increasing in κ: as frictions shrink, the consumer learns more and more in a mean-preserving spread sense. For all such κ, the price function is a single plane with zero slope. Eventually (as κ continues to shrink), unless at most one product has high value, κ hits a threshold $\bar{\kappa}$. Then, for all κ within some interval $[\underline{\kappa}, \bar{\kappa}]$ the consumer's learning is the same–as are the pricing strategies by the firms. Here, the price function is the maximum of two planes whose intersection is the line $y = x$. Both cases are depicted in Figure 2, where we have substituted in $y = 1 - x$ (thanks to Theorem 5.1).

With the equilibrium in hand, we can ask, how does consumer welfare change with the size of the information frictions? For intermediate costs, the consumer's learning is unaffected by the information cost; she learns the same for all κ, and therefore the firms' pricing is the same. However, information is becoming more expensive, so the consumer's welfare is strictly decreasing in κ on this interval.

On the flip side, when $\kappa \geq \bar{\kappa}$ the opposite relationship exists. When κ increases, the consumer learns less. By learning less, the consumer induces more intense competition, which drives down the distribution of prices. The price effect dominates, so consumer welfare is increasing in the size of the information friction.

Proposition 5.3. *For intermediate information costs, ($\kappa \in [\underline{\kappa}, \bar{\kappa}]$), the consumer's welfare is strictly decreasing in the size of the friction. For large information costs, ($\kappa \geq \bar{\kappa}$), the consumer's welfare is strictly increasing in the size of the friction.*

The proof is in §A.9.

5.3 Low Information Costs: Approaching Efficiency

The essential question is whether the RRS inefficiency persists as $\kappa \downarrow 0$. We find it does not. As information becomes cheap, the informational hold-up is significantly weakened by competition, making substantial information investment worthwhile for the consumer. The value of differentiating the products further increases relative to the cost so the consumer now learns precisely enough to distinguish the products perfectly when they truly differ in value. Moreover, the consumer finds it optimal to acquire a posterior at which her values for the two firms' products are the same. This encourages competition, as it provides a strong undercutting force. Crucially, this competition ensures that

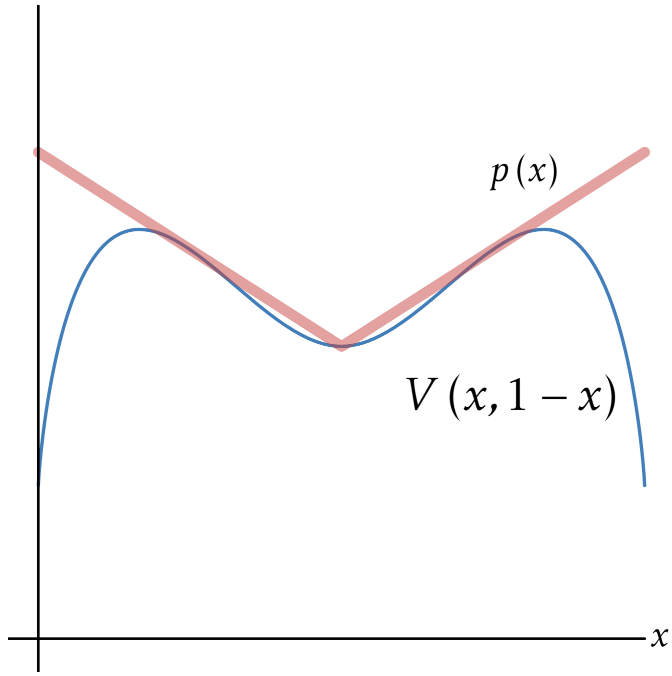

Figure 3: Cheap Information

the consumer retains enough surplus even at the extreme informative posteriors (close to $(0,1)$ and $(1,0)$) to make this high level of learning incentive-compatible, preventing the underinvestment trap seen in the monopoly case.

We say that a consumer **Comparison Shops With Occasional Indifference** if her acquired distribution over valuations has support on 3 points $\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right)$, $\left(\frac{1}{2}, \frac{1}{2}\right)$, and $\left(\frac{1+\lambda}{2}, \frac{1-\lambda}{2}\right)$.

Theorem 5.4. *If κ is sufficiently low, there is an equilibrium in which the consumer comparison shops with occasional indifference and firms randomize over prices according to Expression 5.*

Corollary 5.5. *As information costs vanish, $\kappa \downarrow 0$, there is an efficient limiting equilibrium.*

The proof for both are in §A.10. The limiting equilibrium has support on three points; $(0,1)$, $(1,0)$, and $\left(\frac{1}{2}, \frac{1}{2}\right)$.The firms price so that the consumer purchases from the advantaged firm–if there is an advantaged firm after her learning–with certainty. The consumer never purchases from a firm that is worse than the other *ex post*. The consumer's learning in this equilibrium is depicted in Figure 3.

This limiting efficiency (Corollary 5.5) directly contrasts with the RRS benchmark

(Proposition 3.2) and demonstrates that competition fundamentally alters the incentives for information investment. By ensuring the consumer benefits sufficiently from learning, competition overcomes the informational hold-up that plagues the monopoly setting.

6 Discussion and Extensions

Our analysis highlights how competition mitigates the informational hold-up, but, in order to make traction, we make several simplifying assumptions. Let us briefly explain a few of our assumptions.

Binary Values: We assume that each firm's product takes just one of two values. When frictions are large, this is inconsequential: we show in §6.1 that unless κ is too small, there is an analogous equilibrium to the one we construct in our main specification when the consumer's value for the two firms' products is distributed according to some density on $[0,1]^2$. It is when information is cheap that the problem changes: the learning with three-point support that we identify is no longer optimal for the consumer. Instead, we conjecture that the consumer now acquires a continuum of posteriors close to the prior plus possible point masses on more extreme posteriors.

Symmetric Firms: Like the previous assumption, this is for tractability. The equilibrium of the pricing-only game becomes quite difficult to construct when firms are asymmetric.

Parametric Assumption on the Prior: We make a rather cryptic stipulation that the positive correlation between the consumer's values for the two firms' goods cannot be too high. This is again due to the challenges in constructing an equilibrium in the pricing game between the firms: it is much easier to construct an equilibrium in the pricing-only game when the probability of the "tie" belief (when the consumer's distribution has symmetric three-point support) is sufficiently high. Our parametric assumption, thus, guarantees this is true in the consumer's optimal learning when frictions are sufficiently low.

Private Learning Before Trading: We assume that the consumer learns before she observes the firms' prices and that, moreover, this learning is private.[4] The timing in our main environment (learning before trading) is realistic in many environments: in particular, learning about a service provider's reputation seems especially fitting. Our timing assumption is also that made by RRS, which allows us to identify the effects of competition cleanly. In §6.2, we allow for public learning and show that a hold-up problem emerges, which leads to zero information acquisition by the consumer.

Cost Function: The consumer's cost of acquiring information is a linear functional of the distribution over posteriors. We assume this posterior-separable form for tractability. Moreover, we specify that the convex function that is integrated has unbounded slope at the boundaries of the unit square. This is done to ensure an "interior" solution in the consumer's information acquisition problem. Importantly, this specification only makes our convergence result more difficult to attain: if the slope were bounded our results would go through with the modification that the results would no longer be limit results but would hold for sufficiently small positive κ.

Full Market Coverage: Our main specification assumes that the consumer's outside option is negligible, so she always purchases from one of the firms. Clearly, when κ is large, the equilibrium we construct remains an equilibrium even with a (potentially) relevant outside option of, say, 0. Naturally, this is not the case as κ vanishes.

6.1 Extension to a Prior with a Density

An exact analog of Theorem 5.2 holds when the consumer's valuations for the two products are symmetrically distributed with nonzero density h on the unit square and the consumer's utility is affine in her valuation for the purchased product and additively separable in her valuation, the price, and the cost of acquiring information (which is posterior-mean measurable).

That is, suppose each firm is selling a product whose value to the consumer is a random variable Z_i with full support on $[0,1]$. Random vector (Z_1, Z_2) is distributed

[4] Matějka and McKay (2012) study a related scenario in which firms set prices before the consumer learns.

on $[0,1]^2$ according to continuous density $h(z_1,z_2)$, which is symmetric around the diagonal $y = x$, i.e., $h(z_1,z_2) = h(z_2,z_1)$ for all $z_1,z_2 \in [0,1]$. The prior expected value is $\frac{1}{2} = \int_0^1 \int_0^1 af(a,b)\,db\,da$.

The consumer may acquire any fusion $G \in \mathcal{F}_H$ of the prior at cost $C(G) = \kappa \int c\,dG$ where we maintain the assumptions from the model section above. Then,

Proposition 6.1. *If κ is sufficiently high, there is an equilibrium in which the consumer comparison shops with uniform two-point support and firms randomize over prices according to Expression 4.*

The proof is in §A.11. The comparative statics from Proposition 5.3 also carry over:

Proposition 6.2. *For intermediate information costs, ($\kappa \in [\underline{\kappa}, \bar{\kappa}]$), the consumer's welfare is strictly decreasing in the size of the friction. For large information costs, ($\kappa \geq \bar{\kappa}$), the consumer's welfare is strictly increasing in the size of the friction.*

The intuition is also the same: in the intermediate-friction region, the consumer's optimal learning; and, therefore, the firms' behavior, stays the same as κ dwindles. The consumer accrues all of the benefits of cheaper information. When κ is large, the consumer's learning is affected and so firms raise their prices (on average) to take advantage of their greater market power. This (negative, for the consumer) force is dominant, and so cheaper information makes the consumer worse off.

When information is cheap yet still costly, our equilibrium construction no longer applies. The issue is that in any symmetric equilibrium, firms must be randomizing over prices–this is because that if the consumer anticipates a deterministic market price, her optimal learning has symmetric binary support; but then from Lemma 4.1, if her learning has symmetric binary support, firms must randomize over prices. But when firms are randomizing over prices and information is cheap, it becomes tricky to characterize the equilibrium. Furthermore, when information is cheap, the consumer's learning will not have binary or ternary support (the pricing strategies constructed above leave the value function convex near indifference, making the consumer want to obtain a distribution over posteriors with uncountable support). Determining the equilibrium when κ is positive but vanishing thus remains an open problem for future work.

6.2 Observable Learning Benchmark

A crucial assumption in our model is that learning is private. A natural comparison is the case in which the firms observe the consumer's acquired posterior $(x,y) \in [0,1]^2$ before posting prices. The first step in characterizing the equilibrium is to characterize equilibria in the pricing-only game between the two firms for an arbitrary vector (x,y).

Lemma 6.3. *For all (x,y), a pure-strategy equilibrium of the pricing-only game exists. If $x < y$, there exists an equilibrium in which the vector of prices is $(0, y-x)$ and the consumer always purchases from firm 2. For any (x,y) with $x = y$, the unique equilibrium is the Bertrand outcome: both firms price at marginal cost, $p_1 = p_2 = 0$.*

It is straightforward to check that the equilibria constructed in Lemma 6.3 are particularly bad for the consumer unless $x = y$: her expected payoff at any (x,y) with $y \neq x$ is strictly negative. Moreover, it is clear that the consumer's net payoff at any (x,y), in any equilibrium, must be as follows:

Lemma 6.4. *In any equilibrium of the pricing-only game with $y > x$, the expected net payoff to the consumer is weakly less than x.*

The proof is in §A.12. Working backward, we now conclude that the consumer will not learn.

Proposition 6.5. *If $\kappa > 0$, the unique equilibrium with observable learning is for the consumer to acquire no information: she chooses the degenerate distribution on the prior $\left(\frac{1}{2}, \frac{1}{2}\right)$.*

The proof is in §A.13. Although this equilibrium is quite inefficient–no matter how cheap information is, the consumer does not learn–it is good for the consumer. By committing to not learn, she forces the firms to compete intensely and the consumer acquires all of the surplus in the market.

7 Conclusion

A central finding in the literature on markets with pre-contracting investments is the potential for inefficiency driven by hold-up problems. Ravid, Roesler, and Szentes (2022)

(RRS) demonstrate this starkly for a monopoly setting with informational investments, where the seller's ability to appropriate informational rents discourages the consumer's investment in learning, leading to inefficient outcomes even when information becomes virtually free. This paper asks whether competition fundamentally alters this result.

We show that competition indeed serves as a powerful corrective mechanism. Our main contribution is demonstrating that duopoly competition restores *ex post* efficiency as information costs vanishIn the equilibrium we characterize, consumers are incentivized to acquire precise enough information to always choose the superior product in the limit, ensuring efficient matching. Technical difficulties arise because the equilibrium requires solving a multidimensional information design problem on top of an equilibrium pricing game that involves a distribution of prices. The value of posterior beliefs is endogenous and depends on the firms' pricing, which is random. We prove that the consumer only wants to learn the relative values, which we call comparison shopping.

The key mechanism is competition's mitigation of the informational hold-up. Unlike a monopolist who can tailor their price to extract nearly all expected surplus based on anticipated learning, competing firms pricing simultaneously without observing the consumer's private knowledge cannot fully capture the value created by the information investment. This inability to perfectly expropriate informational rents ensures that the consumer retains sufficient returns from learning about relative product fit, preserving the incentive to invest in acquiring accurate information even as costs approach zero. This finding highlights a specific benefit of competition arising from the interplay between price rivalry and the nature of consumer learning considered here.

While competition solves the limiting inefficiency, our analysis also reveals that a residual hold-up persists when information costs are positive. This leads to a non-monotonic relationship between information costs and consumer welfare, where making information cheaper can sometimes harm consumers by enabling firms to extract more surplus, even as matching improves.

Overall, our results underscore that market structure is a critical determinant of learning incentives and market efficiency. When consumers must invest in information before trading, the degree of competition significantly influences whether that investment is undertaken efficiently and whether the market ultimately allocates goods effectively. Our work suggests that, concerning informational hold-up, competition can indeed be a potent force for restoring efficiency.

References

Mark Armstrong and Jidong Zhou. Consumer information and the limits to competition. *American Economic Review*, 112(2):534–77, February 2022.

Gary Biglaiser, Jiadong Gu, and Fei Li. Information acquisition and product differentiation perception. *American Economic Journal: Microeconomics*, 2024.

Fernando Branco, Monic Sun, and J Miguel Villas-Boas. Optimal search for product information. *Management Science*, 58(11):2037–2056, 2012.

Fernando Branco, Monic Sun, and J Miguel Villas-Boas. Too much information? information provision and search costs. *Marketing Science*, 35(4):605–618, 2016.

Carlo M Cusumano, Francesco Fabbri, and Ferdinand Pieroth. Competing to commit: Markets with rational inattention. *American Economic Review*, 114(1):285–306, 2024.

Piotr Dworczak and Anton Kolotilin. The persuasion duality. *Theoretical Economics*, 19 (4):1701–1755, 2024.

Leonardo Felli and Kevin Roberts. Does Competition Solve the Hold-up Problem? *Economica*, 83(329):172–200, 2016.

Vasudha Jain and Mark Whitmeyer. Search and competition with flexible investigations. *Mimeo*, 2021.

Andreas Kleiner, Benny Moldovanu, Philipp Strack, and Mark Whitmeyer. The extreme points of fusions. *Mimeo*, 2025.

Ruitian Lang. Try before you buy: A theory of dynamic information acquisition. *Journal of Economic Theory*, 183:1057–1093, 2019.

Daniel Martin. Strategic pricing with rational inattention to quality. *Games and Economic Behavior*, 104:131–145, 2017.

Filip Matějka and Alisdair McKay. Simple market equilibria with rationally inattentive consumers. *American Economic Review*, 102(3):24–29, 2012.

Giuseppe Moscarini and Marco Ottaviani. Price competition for an informed buyer. *Journal of Economic Theory*, 101:457–493, 2001.

Marilyn Pease. Shopping for information: Implications of consumer learning for optimal pricing and product design. *The Journal of Industrial Economics*, 71(3):883–923, 2023.

Nicola Persico. Information acquisition in auctions. *Econometrica*, 68(1):135–148, 2000.

Doron Ravid. Ultimatum bargaining with rational inattention. *American Economic Review*, 110(9):2948–63, September 2020.

Doron Ravid, Anne-Katrin Roesler, and Balázs Szentes. Learning before trading: on the inefficiency of ignoring free information. *Journal of Political Economy*, 130(2):346–387, 2022.

Xianwen Shi. Optimal auctions with information acquisition. *Games and Economic Behavior*, 74(2):666–686, 2012.

Martin Vaeth. Rational voter learning, issue alignment, and polarization. *Mimeo*, 2023.

Asher Wolinsky. True Monopolistic Competition as a Result of Imperfect Information. *The Quarterly Journal of Economics*, 101(3):493–511, 08 1986.

Nathan Yoder. Extended real-valued information design. *Mimeo*, 2021.

A Omitted Proofs and Derivations

A.1 Proposition 3.1 Proof

Proof. First, we argue that our construction constitutes an equilibrium. We need to show that a firm has no profitable deviation. The firm obtains a constant profit for any price $p \in [\underline{x}, \bar{x}]$:

$$\Pi(p) = p[1 - F(p)] = \underline{x}.$$

Its profit is strictly less than $\underline{x}$ for any price strictly below $\underline{x}$ and is 0 for any price strictly above $\bar{x}$.

It is also easy to check that the consumer has no profitable deviation: for any value $x \in [\underline{x}, \bar{x}]$ the consumer's payoff is the affine function

$$-\kappa c'(\underline{x})x + \kappa(\underline{x}c'(\underline{x}) - c(\underline{x})).$$

Viz., we have

$$\int_{\underline{x}}^{x} (x - p)\, dG(p) - \kappa c(x) = -\kappa c'(\underline{x})x + \kappa(\underline{x}c'(\underline{x}) - c(\underline{x})).$$

As the consumers payoff is continuous on the interior of $(0, 1)$, strictly concave for all $x \in (0, \underline{x})$ and $x \in (\bar{x}, 1)$, and linear on $[\underline{x}, \bar{x}]$, we conclude that this distribution is also a best response for the consumer. Finally, Expression 3 is pinned down by $G(1) = 1$.

Uniqueness follows from the proof of Theorem 5.2 in Jain and Whitmeyer (2021). ■

A.2 Proposition 3.2 Proof

Proof. As $\kappa \downarrow 0$, from the identity $1 = \kappa(c'(\bar{x}) - c'(\underline{x}))$, we see that we must have $\bar{x} \uparrow 1$ and/or $\underline{x} \downarrow 1$. We claim that we cannot have $\underline{x} \downarrow 0$. In fact, $\underline{x}$ must always be strictly above a strictly positive number as follows.

Claim A.1. *For all $\kappa > 0$, $\underline{x} \geq a$, where a solves*

$$1 - \log a - \frac{\mu}{a} = 0.$$

Proof. Differentiating

$$1 + \log\frac{\bar{x}}{\underline{x}} - \frac{\mu}{\underline{x}}$$

with respect to $\bar{x}$ and $\underline{x}$, in turn, we get $\frac{1}{\bar{x}} > 0$ and $\frac{\mu}{\underline{x}^2} - \frac{1}{\underline{x}} > 0$ ($\underline{x} < \mu$). Accordingly, by the implicit function theorem, $\underline{x}'(\bar{x}) < 0$. Thus, $\underline{x}$ is minimized when $\bar{x} = 1$, which produces the specified equation. The unique solution to that equation $a \in (0, \mu)$ is strictly increasing in μ and is approximately .19 when $\mu = \frac{1}{2}$ and does not equal 0 for all $\mu > 0$. ■

Thus, we have shown that $\underline{x} \downarrow a$ and $\bar{x} \uparrow 1$ as $\kappa \downarrow 0$. Finally, pick an arbitrary $p \in [\underline{x}(0), \bar{x}(0))$ and observe that

$$\lim_{\kappa \downarrow 0} G(p) = \lim_{\kappa \downarrow 0} \kappa (c'(p) - c'(\underline{x})) = 0,$$

and so we must have $G \to \delta_1$ (in distribution) as $\kappa \downarrow 0$. ■

A.3 Lemma 4.1 Proof

Proof. For a contradiction, assume a symmetric pure-strategy equilibrium exists. Then a firm's demand is (locally) perfectly inelastic–if it raises its price slightly, the consumer will purchase from it with the same probability. This deviation yields strictly higher profits, which contradicts our original pure-strategy being an equilibrium. ■

A.4 Proposition 4.2 Proof

Proof. We are looking for a symmetric equilibrium in which each firm chooses an atomless distribution over prices $\Gamma(p)$ with support on $\left[\underline{p}, \underline{p} + 2\lambda\right]$. We guess further that Γ can be written as $\Gamma(p) = \Gamma_L(p)$ for $p \in \left[\underline{p}, \underline{p} + \lambda\right]$ and $\Gamma(p) = \Gamma_H(p)$ for $p \in \left[\underline{p} + \lambda, \underline{p} + 2\lambda\right]$. The profit for a firm is

$$\Pi(p) = \begin{cases} \frac{p}{2}\left[2 - \Gamma_H(p+\lambda)\right], & \underline{p} \leq p \leq \underline{p} + \lambda \\ \frac{p}{2}\left(1 - \Gamma_L(p-\lambda)\right), & \underline{p} + \lambda \leq p \leq \underline{p} + 2\lambda \end{cases}.$$

For any on-path p a firm's payoff must equal some constant k. Imposing this and the conditions for Γ to be a CDF, we get the functional form specified in Expression 4.

Finally, we need to verify that firms do not want to choose a price outside of the conjectured region. If a firm chooses a price $p \in \left[\underline{p} - \lambda, \underline{p}\right]$, its payoff is

$$\frac{p}{2}\left[2 - \Gamma_L(p+\lambda)\right] = \frac{p}{2}\left[\left(\frac{\sqrt{2}\lambda + \lambda}{p + 2\lambda}\right) + 1\right].$$

The derivative of this with respect to p is

$$\frac{2\lambda\left(\frac{\sqrt{2}\lambda+\lambda}{p+2\lambda}\right)}{p+2\lambda}+1>0\;,$$

whence we conclude a firm does not want to deviate to a price in this region (we have implicitly assumed that $\underline{p} \geq \lambda$, but this is fine since a firm obviously does not want to deviate to a negative price). Evidently, if a firm chooses any price $p \leq \underline{p}-\lambda$ its payoff is just p, which is obviously strictly increasing in p and hence equals $\underline{p}+\lambda$, which we just established is not an improvement for the firm. The last case is that in which a firm chooses a price $p \in \left[\underline{p}+2\lambda, \underline{p}+3\lambda\right]$. In that case, a firm's profit is

$$\frac{p}{2}\left(1-\Gamma_H\left(p-\lambda\right)\right)=\frac{p}{2}\left(\frac{\sqrt{2}\lambda+3\lambda-p}{p-2\lambda}\right)\;,$$

which is strictly decreasing in p.

The uniqueness argument is analogous to that argued for the atomless equilibrium of Proposition 3 in Moscarini and Ottaviani (2001). ■

A.5 Lemma 4.3 Proof

Proof. Because the consumer is indifferent between the two firms with strictly positive probability, a standard under-cutting argument eliminates any symmetric equilibria in which a firm sets some price with strictly positive probability. ■

A.6 Proposition 4.4 Proof

Proof. If one firm, firm 2, say, chooses Φ, firm 1's profit as a function of p is $(1-q)\frac{q}{1-2q}\lambda$, a constant, for all $p \in \left[\frac{q}{1-2q}\lambda, \frac{q}{1-2q}\lambda+\lambda\right]$. For all $p \in \left[\frac{q}{1-2q}\lambda+\lambda, \frac{q}{1-2q}\lambda+2\lambda\right]$, firm 1's profit is

$$pq\left(1-\Phi\left(p-\lambda\right)\right),$$

which is strictly decreasing in p. For all $p \geq \frac{q}{1-2q}\lambda+2\lambda$ firm 1's profit is 0. Finally, for all $p \in \left[0, \frac{q}{1-2q}\lambda+\lambda\right]$, firm 1's profit is $p\left(1-q\Phi\left(p+\lambda\right)\right)$. For all

$$q \leq \frac{\frac{\sqrt[3]{9\sqrt{93}-47}}{\sqrt[3]{2}}-\frac{11\sqrt[3]{2}}{\sqrt[3]{9\sqrt{93}-47}}+5}{9} \approx .406,$$

this function is strictly increasing on this interval. ■

A.7 Theorem 5.1 Proof

Proof. By symmetry we restrict attention without loss of generality to the case $y \geq x$. We assume that the firms each choose the distributions over prices F with support on $\left[\underline{p}, \bar{p}\right]$. Defining $\lambda := \bar{p} - \underline{p}$, the consumer's payoff from posterior (x, y), is

$$V(x,y) = \begin{cases} y - \mathbb{E}[p] - \kappa c(x,y), & y \geq x + \lambda \\ y - \mathbb{E}[p] + U(z) - \kappa c(x,y), & x + \lambda \geq y \geq x \end{cases}, \tag{A1}$$

where $z := y - x$ and

$$\mathbb{E}[p] = \int_{\underline{p}}^{\underline{p}+\lambda} p\,dF(p) \ ,$$

and

$$U(z) := \int_{\underline{p}+z}^{\underline{p}+\lambda} (p - z) F(p - z)\,dF(p) - \int_{\underline{p}}^{\underline{p}+\lambda-z} p\,(1 - F(p + z))\,dF(p) \ .$$

Directly,

$$V_{xx}(x,y) = \begin{cases} -\kappa c_{xx}(x,y), \\ \int_{\underline{p}+z}^{\underline{p}+\lambda} f(p-z)\,dF(p) - \kappa c_{xx}(x,y) \end{cases} \qquad V_{yy}(x,y) = \begin{cases} -\kappa c_{yy}(x,y), \\ \int_{\underline{p}+z}^{\underline{p}+\lambda} f(p-z)\,dF(p) - \kappa c_{yy}(x,y) \end{cases},$$

and

$$V_{xy} = -\int_{\underline{p}+z}^{\underline{p}+\lambda} f(p-z)\,dF(p) - \kappa c_{xy}(x,y).$$

The directional second derivative in the direction of $(1,1)$ is $-\kappa c_{xx}(x,y) - \kappa c_{yy}(x,y) - 2\kappa c_{xy}(x,y) < 0$, by the strict convexity of c.

Let us now evaluate the function $c(x, a + x)$, where a is a parameter taking values in $[-1, 1]$. We have already (just) shown that it is strictly concave. Directly,

$$\frac{\partial}{dx} c(x, a + x) = \varphi'(x) - \varphi'(1 - x) + \varphi'(a + x) - \varphi'(1 - a - x).$$

By direct substitution, we see that this equals 0 when $y = 1 - x$, i.e., $x = \frac{1-a}{2}$. Accordingly, for all a in the specified interval, $c(x, a + x)$ is maximized at $y = 1 - x$; *viz.*, on the the comparison shopping line. Thus, any price function of the value function restricted to the comparison shopping line, $V(x,1-x)$, must correspond to a price function that lies everywhere weakly above $V(x,y)$, and so learning along the comparison-shopping line is optimal. ■

A.8 Theorem 5.2 Proof

Proof. Recall that firms are pricing according to CDF

$$\Gamma(p) = \begin{cases} \Gamma_L := \frac{p-\sqrt{2}\lambda}{\lambda+p}, & \sqrt{2}\lambda \leq p \leq \left(1+\sqrt{2}\right)\lambda \\ \Gamma_H := \frac{(3+\sqrt{2})\lambda-2p}{\lambda-p}, & \left(1+\sqrt{2}\right)\lambda \leq p \leq \left(2+\sqrt{2}\right)\lambda. \end{cases}$$

Given pricing, the consumer is merely facing a standard information acquisition problem. We need to show that there is an optimal solution to this in which she obtains the binary distribution over posteriors anticipated by the firms. To that end, we begin by explicitly collecting some important objects. The first is the consumer's value function: her expected payoff as a function of her vector of beliefs (x,y), given her optimal behavior at the beliefs–which is, of course, to purchase the highest net-expected-value good.

For convenience, define $\underline{p} := \sqrt{2}\lambda$, $\tilde{p} := \underline{p}+\lambda$, and $\bar{p} := \underline{p}+2\lambda$. The consumer's value function is (restricting attention to $y \geq x$ by symmetry)

$$V(x,y) = \begin{cases} y - \mathbb{E}[p] - \kappa c(x,y), & y \geq x+2\lambda \\ y - \mathbb{E}[p] + T_1(z) - \kappa c(x,y), & x+2\lambda \geq y \geq x+\lambda \\ y - \mathbb{E}[p] + T_2(z) - \kappa c(x,y), & x+\lambda \geq y \geq x \end{cases}, \tag{A2}$$

where $z := y-x$ and

$$\mathbb{E}[p] = \int_{\underline{p}}^{\tilde{p}} p d\Gamma_L(p) + \int_{\tilde{p}}^{\bar{p}} p d\Gamma_L(p) = \left(\left(\sqrt{2}+1\right)\log\left(\sqrt{2}+1\right)+\sqrt{2}-1\right)\lambda ,$$

$$T_1(z) := \int_{\underline{p}}^{\bar{p}-z} (1-\Gamma_H(p+z))\Gamma_L(p)\,dp ,$$

and

$$T_2(z) := \int_{\tilde{p}}^{\bar{p}-z} (1-\Gamma_H(p+z))\Gamma_H(p)\,dp + \int_{\tilde{p}-z}^{\tilde{p}} (1-\Gamma_H(p+z))\Gamma_L(p)\,dp + \int_{\underline{p}}^{\tilde{p}-z} (1-\Gamma_L(p+z))\Gamma_L(p)\,dp .$$

From Theorem 5.1, we may restrict attention to learning along the line $y = 1 - x$. Thus, the second object we need to compute is the directional derivative of the value function along vector $(1, -1)$, evaluated at all points of the form $(x, 1-x)$. It is

$$D(x) = \begin{cases} -\kappa c_x(x, 1-x) + \kappa c_y(x, 1-x) - 1, & x \leq \frac{1}{2} - \lambda \\ 2P_1(1-2x) - \kappa c_x(x, 1-x) + \kappa c_y(x, 1-x) - 1, & \frac{1}{2} - \lambda \leq x \leq \frac{1-\lambda}{2} \\ 2P_2(1-2x) - \kappa c_x(x, 1-x) + \kappa c_y(x, 1-x) - 1, & \frac{1-\lambda}{2} \leq x \leq \frac{1}{2} \end{cases}, \qquad (A3)$$

where

$$P_1(z) := \int_{\underline{p}}^{\bar{p}-z} \gamma_H(p+z)\Gamma_L(p)\,dp\ ,$$

and

$$P_2(z) := \int_{\tilde{p}}^{\bar{p}-z} \gamma_H(p+z)\Gamma_H(p)\,dp + \int_{\tilde{p}-z}^{\tilde{p}} \gamma_H(p+z)\Gamma_L(p)\,dp + \int_{\underline{p}}^{\tilde{p}-z} \gamma_L(p+z)\Gamma_L(p)\,dp\ .$$

Direct substitution yields $D\left(\frac{1}{2}\right) = 0$. Moreover, by the symmetry and convexity of c, and since $c\left(\frac{1}{2}, \frac{1}{2}\right) = 0$, for $x \leq \frac{1}{2}$, $c_y(x, 1-x) - c_x(x, 1-x) \geq 0$ with equality at $x = \frac{1}{2}$.

Now, thinking of the price function corresponding to the consumer's optimal learning, there are two possible ways in which a symmetric distribution with binary support can be optimal. First, the price function on the square can be a single plane–when restricted to the line $y = 1 - x$, a line. Moreover, by the symmetry of the problem, this line must have zero slope. Second, the price function can be v-shaped, the minimum of two lines with slopes of opposite sign that intersect at $x = 1/2$.

We tackle the zero-slope case first, which corresponds to large κ. We have

$$D\left(\frac{1-\lambda}{2}\right) = D\left(\frac{1+\lambda}{2}\right) = 0,$$

i.e.,

$$\tau(\kappa, \lambda) := 2P_1(\lambda) - \kappa c_x\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) + \kappa c_y\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) - 1 = 0. \qquad (A4)$$

Note that

$$P_1(\lambda) = \int_{\underline{p}}^{\tilde{p}} \gamma_H(p+\lambda)\Gamma_L(p)\,dp = -\frac{\left(2^{\frac{5}{2}}+6\right)\log\left(\sqrt{2}+2\right) - \left(2^{\frac{7}{2}}+12\right)\log\left(\sqrt{2}+1\right) + \left(2^{\frac{3}{2}}+3\right)\log(2) + 2}{2},$$

which is evidently independent of the parameters and is approximately $\frac{1}{10}$. Directly, $\tau'(\kappa) > 0$ and

$$\tau'(\lambda) = \frac{\kappa}{2}\left[c_{xx}\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) - 2c_{xy}\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) + c_{yy}\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right)\right] > 0.$$

By the implicit function theorem $\lambda'(\kappa) < 0$. Moreover, $\lim_{\lambda\to 1}\tau = \infty$ and $\tau(\kappa, 0) < 0$. Accordingly, there is a unique solution $\lambda^* = \lambda(\kappa)$ to this equation, which is strictly decreasing in κ. Moreover, $\lim_{\kappa\uparrow\infty}\lambda^*(\kappa) = 0$ and $\lim_{\kappa\downarrow 0}\lambda^*(\kappa) = 1$.

We also need to check the following, which ensures that the zero-slope line tangent to the value function at $x = \frac{1-\lambda^*}{2}$ is, indeed, a valid price function. In short, the value function obtains a unique maximum on $x \leq \frac{1}{2}$ at the specified point $x = \frac{1-\lambda^*}{2}$. We will verify that it is first increasing in x, has zero slope at $\frac{1-\lambda^*}{2}$, then decreasing in x before having zero slope again at $x = \frac{1}{2}$. Recalling that D is the value function's derivative, equivalently, we will verify that D is strictly positive for $x < \frac{1-\lambda^*}{2}$, equal to zero at $\frac{1-\lambda^*}{2}$ and $\frac{1}{2}$, and negative for x between $\frac{1-\lambda^*}{2}$ and $\frac{1}{2}$. Formally,

Claim A.2. *$D(x) > 0$ for all $x < \frac{1-\lambda^*}{2}$ and $D(x) \leq 0$ for all $x \in \left[\frac{1-\lambda^*}{2}, \frac{1}{2}\right]$. Moreover, $V\left(\frac{1-\lambda^*}{2}, \frac{1+\lambda^*}{2}\right) \geq V\left(\frac{1}{2}, \frac{1}{2}\right)$.*

Proof. Directly, we differentiate the function D (from Expression A3) with respect to x. This yields

$$D'(x) = \underbrace{-\kappa c_{xx}(x, 1-x) - \kappa c_{yy}(x, 1-x)}_{=:\kappa\rho(x)<0} + \tau(x),$$

where

$$\tau(x) = \begin{cases} 0, & x \leq \frac{1}{2} - \lambda \\ 4R(1-x), & \frac{1}{2} - \lambda \leq x \leq \frac{1-\lambda}{2} \\ 4M(1-x), & \frac{1-\lambda}{2} \leq x \leq \frac{1}{2} \end{cases},$$

where

$$R(z) := \int_{\underline{p}}^{\bar{p}-z} \gamma_H(p+z)\gamma_L(p)\,dp\ ,$$

and

$$M(z) := \int_{\tilde{p}}^{\bar{p}-z} \gamma_H(p+z)\gamma_H(p)\,dp + \int_{\tilde{p}-z}^{\tilde{p}} \gamma_H(p+z)\gamma_L(p)\,dp + \int_{\underline{p}}^{\tilde{p}-z} \gamma_L(p+z)\gamma_L(p)\,dp\ .$$

It is straightforward to check that $\tau(x) \geq 0$ (strictly if $x > \frac{1}{2} - \lambda$) and that it is strictly increasing in x (for all $x > \frac{1}{2} - \lambda$). Moreover,

$$\rho'(x) = c_{yyy}(x, 1-x) - c_{xxx}(x, 1-x) = 2\left(\varphi'''(1-x) - \varphi'''(x)\right) \geq 0.$$

Accordingly, D' has at most one sign change, from negative to positive. As

$$D\left(\frac{1-\lambda^*}{2}\right) = D\left(\frac{1}{2}\right) = 0,$$

this establishes the claim. ■

For this distribution to be feasible (a fusion/MPC of the prior) we need $\frac{\lambda^*}{2} \leq \omega$. Define $\bar{\kappa} \geq 0$ to be the value of κ such that $\lambda^* = 2\omega$. Observe that $\bar{\kappa} = 0$ if and only if $\omega = \frac{1}{2}$, which is the perfect negative correlation case. Directly, $D\left(\frac{1}{2} - \omega\right)$ is strictly increasing in κ.

We now finish the proof of the theorem by tackling the v-shaped price function case, in which $\lambda = 2\omega$.

Claim A.3. *There exists an interval of κs, $[\underline{\kappa}, \bar{\kappa}]$, where $\underline{\kappa} \in [0, \bar{\kappa}]$ for which the equilibrium $\lambda = 2\omega$.*

Proof. Directly, $\frac{\partial}{\partial \kappa} D(x) > 0$. By construction, when $\kappa = \bar{\kappa}$, the equilibrium $\lambda^* = 2\omega$. Moreover, as we noted in Claim A.2, $V\left(\frac{1-\lambda^*}{2}, \frac{1+\lambda^*}{2}\right) \geq V\left(\frac{1}{2}, \frac{1}{2}\right)$. If this is an equality then $\underline{\kappa} = \bar{\kappa}$. If this inequality is strict then there is an interval of κs ($[\underline{\kappa}, \bar{\kappa}]$) for which the line tangent to $V(x-\omega, x+\omega)$ lies above $V(x, 1-x)$ at $\frac{1}{2}$.[5] ■

This completes the proof of the theorem. ■

[5]Here is a heuristic explanation of what is going on. For high κ, we are in the first case, with the flat price function. The support points generated by this solution are moving further and further apart (symmetrically, along the line $y = 1-x$) as κ diminishes. At some point, the constraint that the distribution is an MPC of the prior binds. $\bar{\kappa}$ is the exact information cost at which this occurs. At this point, the price function is still flat. If we continue to reduce $\bar{\kappa}$ from this point, we are now rotating the price function down, preserving the tangency at $\lambda = 2\omega$ (indeed we cannot have a more spread apart pair of points). For it to be a valid price function, it must lie everywhere above the value function, but we know that this must be true as long as κ is not decreased too much below $\bar{\kappa}$, whenever $V\left(\frac{1-\lambda^*}{2}, \frac{1+\lambda^*}{2}\right) > V\left(\frac{1}{2}, \frac{1}{2}\right)$.

A.9 Proposition 5.3

Proof. Thanks to the discussion in the text, we need only consider the case $\kappa \geq \bar{\kappa}$. In this case, the consumer's payoff at equilibrium is

$$\underbrace{\frac{1}{2} + \frac{\lambda}{2} - \mathbb{E}[p] + \int_{\underline{p}}^{\tilde{p}} (1 - \Gamma_H(p+\lambda))\Gamma_L(p)\,dp}_{S(\lambda)} - \kappa c\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right).$$

The derivative of this with respect to λ is $S'(\lambda) - \frac{1}{2}\kappa\left(c_y\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) - c_x\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right)\right)$. However, by the Optimality Equation A4, $c_y\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) - c_x\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) = 2P_1(\lambda) - 1$. Summing everything up and simplifying, we obtain $-1 - \sqrt{2} < 0$. ∎

A.10 Theorem 5.4 and Corollary 5.5 Proofs

Proof. For convenience, define $\underline{p} := \frac{q}{1-2q}\lambda$. The consumer's payoff as a function of the realized posterior belief x is (restricting attention to $y \geq x$ by symmetry)

$$V(x,y) = \begin{cases} y - \mathbb{E}[p] - \kappa c(x,y), & y \geq x + \lambda \\ y - \mathbb{E}[p] + U(z) - \kappa c(x,y), & x + \lambda \geq y \geq x \end{cases},$$

where $z := y - x$ and

$$\mathbb{E}[p] = \int_{\underline{p}}^{\underline{p}+\lambda} p\,d\Phi(p)\ ,$$

and

$$U(z) := \int_{\underline{p}+z}^{\underline{p}+\lambda} (p-z)\Phi(p-z)\,d\Phi(p) - \int_{\underline{p}}^{\underline{p}+\lambda-z} p(1-\Phi(p+z))\,d\Phi(p)\ .$$

For this to be an equilibrium, we need for there to be a line (the price-function) $\alpha x + \beta$ lying everywhere above $V(x, 1-x)$ on $0 \leq x \leq \frac{1}{2}$, and intersecting $V(x, 1-x)$ at $\frac{1-\lambda}{2}$ and $\frac{1}{2}$. Removing $\mathbb{E}[p]$ since it is a constant, we compute

$$V\left(\frac{1}{2}, \frac{1}{2}\right) = \frac{1}{2} - \frac{\lambda(1-q)q\left(\log\left(\frac{q}{1-q}\right) - 4q + 2\right)}{(1-2q)^3}.$$

Moreover, $\alpha = \kappa c_y\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) - \kappa c_x\left(\frac{1-\lambda}{2}, \frac{1+\lambda}{2}\right) - 1$.

We need

$$\left(\kappa c_y\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)-\kappa c_x\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)-1\right)\left(\frac{1-\lambda}{2}\right)+\beta=\frac{1+\lambda}{2}-\kappa c\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right),$$

or

$$\beta=1-\kappa c\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)-\left(\kappa c_y\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)-\kappa c_x\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)\right)\left(\frac{1-\lambda}{2}\right).$$

We also need

$$\left(\kappa c_y\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)-\kappa c_x\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)-1\right)\frac{1}{2}+\beta=\frac{1}{2}-\frac{\lambda(1-q)q\left(\log\left(\frac{q}{1-q}\right)-4q+2\right)}{(1-2q)^3},$$

or

$$\beta=1-\frac{\lambda(1-q)q\left(\log\left(\frac{q}{1-q}\right)-4q+2\right)}{(1-2q)^3}-\frac{1}{2}\left(\kappa c_y\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)-\kappa c_x\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right)\right).$$

Equating the βs, we get

$$\underbrace{\omega\frac{(1-q)\left(\log\left(\frac{q}{1-q}\right)-4q+2\right)}{(1-2q)^3}}_{=:t(\lambda)}+\kappa\underbrace{\left[\lambda d'(\lambda)-d(\lambda)\right]}_{v(\lambda)}=0, \tag{A5}$$

where

$$d(\lambda):=c\left(\frac{1-\lambda}{2},\frac{1+\lambda}{2}\right),$$

and where we used the fact that $t\equiv t(\lambda)$ (as $q=\frac{\omega}{\lambda}$).

Furthermore, note that we must have $2\omega\le\lambda\le 1$. Directly, $t'(\lambda)<0$ and $t(\lambda)<0$ for all $\lambda\in[2\omega,1]$, and $\lim_{\lambda\downarrow 2\omega}t(\lambda)=0$. Moreover, by the strict convexity of c, $v(\lambda)>0$. Likewise, $v'(\lambda)=\lambda d''(\lambda)>0$, and $\lim_{\lambda\uparrow 1}v'(\lambda)=\infty$. Continuing along these lines, it is easy to compute that $t'(\lambda)$ is bounded for all $\lambda\in(2\omega,1]$. Finally, it is straightforward to check that $\lim_{\lambda\uparrow 1}v(\lambda)=\infty$.

From the observations in the previous paragraph, we conclude the following:

(i) If κ is sufficiently small, then a unique solution $\lambda^*(\kappa)$ to Equation $A5$ exists.

(ii) In this unique solution, λ^* is strictly decreasing in κ.

(iii) As $\kappa\downarrow 0$, $\lambda^*\uparrow 1$.

This last item is the stated corollary (5.5). ■

A.11 Proposition 6.1 Proof

Proof. This result is an immediate implication of the fact that in the proof of Theorem 5.2, we show that as κ increases, λ decreases and in the limit goes to 0. For any prior as specified in this section, there is a threshold $\lambda > 0$ such that the comparison shopping with uniform two-point support distribution is a fusion of the prior. That we can use the same price-function approach follows from Dworczak and Kolotilin (2024). Alternatively, as Kleiner, Moldovanu, Strack, and Whitmeyer (2025) establish, these distributions are exposed (in their parlance, "strongly exposed") points in the set of finitely-supported fusions of the prior. When frictions are high ($\kappa \geq \bar{\kappa}$), the associated polyhedral subdivision is the trivial one consisting of a single element ($[0,1]^2$). When frictions are moderate ($\kappa \in [\underline{\kappa}, \bar{\kappa}]$), the associated subdivision is convex partitional, with two elements, the triangles Δ^1 and Δ^2 (Expressions 6 and 7). ■

A.12 Lemma 6.4 Proof

Proof. Consider an arbitrary equilibrium and let $v \geq 0$ be the infimum of the support of firm 1's distribution over prices. Naturally, then, $v+\lambda$ must be the infimum of the support of firm 2's distribution over prices. Thus, the consumer's net payoff is weakly less than $\max\{y-v-\lambda, x-v\} = x-v \leq x$. ■

A.13 Proposition 6.5 Proof

Proof. By Lemma 6.4, the consumer's payoff at any (x,y) is bounded above by $\min\{x,y\}$, which is weakly concave. For $\kappa > 0$ this function is strictly concave. ■