Playing Favorites with Inattentive Buyers

Draft

Dhongkyu Yoon*

October 6, 2025

Abstract

We study how a seller should design mechanisms when buyers can flexibly acquire costly information before purchase. In our model, two ex ante identical buyers decide not only how much information to acquire but also what kind of information to focus on. We show that the seller may optimally create endogenous exclusivity by inducing one buyer to acquire precise information while leaving the other uninformed. This reflects a core tradeoff between rent extraction and trade probability: exclusivity strengthens inspection incentives but lowers the likelihood of trade, while equal treatment maximizes coverage at the expense of weaker learning. Comparing mechanisms, we find that sequential offers are optimal when rent extraction is paramount, while symmetric simultaneous offers are optimal when trade probability is more valuable. The distinction arises only under flexible information acquisition, highlighting how timing and exclusivity interact with buyers' learning incentives to generate asymmetric treatment of symmetric buyers.

Keywords: Mechanism design, Flexible information acquisition

*Department of Economics, University of Washington dhongkyu@uw.edu

1 Introduction

Information is costly, but everywhere buyers pay it. Homebuyers hire inspectors, car buyers pay for diagnostics, investors subscribe to analyst reports, and patients rely on expensive medical trials before choosing treatment. In all of these settings, buyers face the same basic problem: before committing, they must decide how much effort to invest in discovering the truth. Modern surveys show that more than 80% of used car buyers pay for mechanical checks before purchase, and financial firms spend billions annually on research and due diligence. These examples illustrate just how ubiquitous and economically significant costly information acquisition is in markets for unfamiliar goods.

Recognizing this, sellers can shape buyers' incentives to gather such information by carefully designing the selling mechanism—through menus of prices, warranties, or allocation rules—that determine how buyers learn and commit to purchase. When making such offers, sellers face a fundamental tradeoff. On the one hand, they wish to encourage careful learning so that informed buyers will be willing to pay higher prices. On the other hand, this carries the risk that buyers may discover negative information about the product's quality and decide not to purchase at all. The presence of multiple buyers introduces further complexity: if one buyer expects that others may win the item, their incentive to incur the cost of learning diminishes. To address this, sellers may find it optimal to favor a particular buyer by granting a higher probability of allocation, or even to approach buyers sequentially rather than simultaneously, thereby restoring incentives to acquire information.

This paper studies how a seller should treat two ex ante identical buyers who, upon observing the offer, must decide how much, and what kind of information to acquire before purchasing. We show that—even when buyers are symmetric ex ante—the seller may optimally play favorites, offering stronger terms or higher allocation probabilities to one buyer to encourage information acquisition. Moreover, we examine whether the seller should offer the good simultaneously or sequentially, and how this choice interacts with buyers' incentives to acquire information.

The key tradeoff lies between rent extraction and trade probability. Since the information is endogenous, buyers will only choose to inspect if they expect to receive sufficient surplus conditional on receiving a favorable signal. If the seller prioritizes rent extraction, she will want to guarantee some form of exclusivity: by giving a buyer a stronger claim on the good, the seller raises that buyer's expected utility and thus strengthens her incentive to acquire precise information. However, more precise information also increases the likelihood of receiving a bad signal, thereby reducing the overall probability of trade. When the seller instead prioritizes trade probability over rent extraction, she will prefer to spread allocation more evenly—making the good available to multiple buyers and maximizing the likelihood of sale. Which objective the seller prioritizes—rent extraction or trade probability—depends on the relative costliness of acquiring favorable signals.

We observe how this tradeoff plays out in simultaneous and sequential mechanisms with two buyers. In the simultaneous mechanism, the resource constraint prevents the seller from guaranteeing full allocation to both buyers, so informational rents must typically be shared. If the seller prioritizes trade probability, she spreads allocation more evenly, ensuring broad coverage and a higher chance of sale. If instead the seller seeks higher rent extraction, she can loosen the allocation constraint for one buyer—offering priority or a more favorable allocation menu—thereby creating exclusivity that strengthens incentives for costly inspection. In the sequential mechanism, the seller can extend exclusivity to both buyers in turn: the first buyer receives full allocation initially, and if she declines to trade, the same exclusive opportunity is offered to the second buyer. This arrangement enhances rent extraction by reinforcing inspection incentives but lowers the overall probability of trade.

How the seller weighs these opposing effects determines which mechanism performs better. When rent extraction is more important, the sequential mechanism tends to be preferable, as it provides stronger exclusivity and sharper incentives for information acquisition. When trade probability is more valuable, the simultaneous symmetric mechanism performs better, maximizing coverage and the likelihood of sale. In short: play favorites for rents, equal treatment for trade probability.

This comparison is meaningful only in the flexible information acquisition model, where buyers not only decide how much information to acquire but also what kind of information to focus on. For example, in financial markets, an investor with inflexible information acquisition might only choose the precision of a research report—a coarse forecast versus a more detailed projection. With flexible information acquisition, however, the investor can also decide what dimension of the option to study—stress-testing downside risk (bad outcomes) versus confirming upside potential (good outcomes). In the inflexible case, where buyers choose only from a fixed set of signal precisions, the optimal asymmetric simultaneous mechanism can be implemented sequentially without changing the seller's revenue, making the two mechanisms effectively equivalent.

On the other hand, in the flexible information acquisition case, the same allocation outcome can yield different revenues depending on the timing of the mechanism. This occurs because buyers not only choose how much information to acquire, but also how precisely to learn about unfavorable outcomes. Since the signal realizes privately, the precision of the unfavorable signal determines what price and allocation probability the seller can offer at a given favorable signal. How exact this unfavorable signal can get depends on each buyer's resource constraint, which varies across mechanisms. In the simultaneous mechanism, the resource constraint is tighter: since another buyer is competing for the good, each buyer has weaker incentives to acquire precise information about bad outcomes, limiting the price the seller can charge when the signal is favorable. As a result, the price the seller can charge at a given signal depends on the intensity of competition. Flexibility in information choice is precisely what makes timing matter.

2 Literature Review

Recent work on mechanism design with buyers' endogenous information acquisition distinguishes between models of flexible and inflexible learning. In flexible settings, buyers optimally choose how much and what kind of information to acquire (Mensch (2022), Thereze (2024), and Mensch and Ravid (2022)). In contrast, inflexible acquisition models, such as Shi (2012), assume buyers can choose only the informativeness of the signal. Bergemann and Pesendorfer (2007), Kamenica and Gentzkow's (2011) assume the seller can strategically design what information to reveal to buyers in order to influence their behavior. These models highlight the power of committed disclosure policies and how the seller's informational control can serve as a substitute for or complement to price discrimination.

A second key dimension is the timing between mechanism design and information acquisition. In our model, the seller commits to a mechanism first, and buyers decide on their information strategy in response. This sequencing enhances the seller's ability to shape incentives. In contrast, Roesler and Szentes (2017) consider settings where buyers acquire information before the mechanism is designed, shifting informational advantage to the buyers. Ravid et al. (2022) study environments in which mechanism offering and information acquisition occur simultaneously, limiting both parties' ability to respond to the other, and leading to different design constraints.

Finally, the literature has investigated the optimality of asymmetric mechanisms, particularly when buyers are ex ante symmetric but differ endogenously due to information acquisition or sequential design. Gershkov et al. (2021) show that asymmetric mechanisms can improve revenue when the seller can incentivize buyer to take a favorable action. Bergemann and Pesendorfer (2007) similarly demonstrate that asymmetries can be optimal even in static settings with exogenous information. These results highlight how asymmetry can be an intentional feature of optimal design rather than a failure of fairness.

3 Preliminaries

There is one seller with one item. The value of the item is unknown to buyer(s) and the seller, who are risk neutral. The value will realize to be $v \in \{0, 1\}$ where $Pr(v = 1) = \mu$. Ex ante, this expected value μ is μ_0 and common knowledge. The seller offers to each buyer, a mechanism $\mathcal{M}$ which specifies (x, p), allocation probability and price. The seller can choose to offer however many options in a menu, among which the buyer(s) will choose after privately observing the realization of signals after acquiring information.

3.1 Flexible information acquisition

Each buyer can choose an information structure $(\mathcal{S}, \mathcal{T})$. It consists of a set of signal realizations $\mathcal{S}$ and a distribution of posterior means, $\mathcal{T} \in \Delta(M)$ such that it is Bayes consistent, i.e. $\int \mu d\mathcal{T}(\mu) = \mu_0$, where $M = [0, 1]$ is a set of possible posteriors. As in the literature, each signal realization $m \in \mathcal{S}$ has a value which represents the posterior belief it generates. Therefore a signal realizes as a value of a posterior belief $\mu \in M$, and the buyer's payoff if the trade was made will be $u_B(x, p, \mu) = x\mu - p$.

Each buyer can freely choose any information but with cost. The cost depends on how precise the signal is, i.e. the more mean preserving spread is the distribution of posteriors, the more expensive is the information acquisition. One common way to model information acquisition cost in the literature is to define as the expected difference in a posterior-separable cost function $c(\mathcal{T})$:

$$c(\mathcal{T}) = H(\mu_0) - E_{\mathcal{T}}[H(\mu)]$$

where H is strongly concave and twice Lipschitz continuously differentiable on any posteriors $\mu \in M$. The cost function is well defined for any information acquisition as long as the information structure satisfies Bayes' rule and increasing in Blackwell order. Commonly used functions for $H(\mu)$, the measure of uncertainty, would be informational entropy: $H(\mu) = \mu ln(\mu) + (1 - \mu)ln(1 - \mu)$, or quadratic function: $H(\mu) = \mu(1 - \mu) + (1 - \mu)\mu$.

In models of costly information acquisition, particularly those using the binary entropy function $H(\mu)$ to represent uncertainty, the second derivative $H''(\mu)$ captures the marginal cost of increasing posterior precision. However, the marginal cost alone does not fully describe how the buyer's incentives to acquire information evolve across different posterior levels. To better understand this dynamic, we introduce the relative curvature index $RC(\mu) := -\frac{H'''(\mu)\mu}{H''(\mu)}$. This index captures the local sensitivity or curvature of the marginal cost of uncertainty with respect to posterior precision, offering insight into how the cost of acquiring information evolves as the buyer becomes more or less certain. A lower value of the index indicates that the curvature decreases more gradually, while a higher value implies that the curvature declines more rapidly. In particular, the index helps determine whether the information environment exhibits accelerating or decelerating marginal costs as posteriors move away from ignorance (i.e., $\mu = 0.5$). The relative curvature index plays a central role throughout the analysis—governing the seller's willingness to trade off allocation risk for more informative signals and shaping the structure of optimal mechanisms.

3.2 Example

To illustrate the mechanism in a financial setting, consider first a single buyer who is evaluating whether to purchase a call option. The buyer begins with a prior belief $\mu_0 \in (0, 1)$ that the option will finish in the money, a belief that is common knowledge between the buyer and the seller. The seller offers a simple contract $(x, p) = (1, 0.7)$, corresponding to a fixed premium of 0.7 for one unit of the option. If the buyer accepts without conducting further research, her expected payoff is $\mu_0 - 0.7$, which is negative whenever $\mu_0 < 0.7$. In this sense, a relatively high premium induces the buyer to consider acquiring information before deciding whether to trade.

Suppose the buyer can conduct costly research $\mathcal{T}$ that refines her belief into two possible posteriors, $\mu_H = 0.9$ and $\mu_L = 0.1$, each realized with equal probability. If the favorable posterior μ_H occurs, the buyer accepts and earns 0.2, whereas if the unfavorable posterior μ_L occurs, she rejects the contract. Her ex ante expected utility is then $0.1 - c(\mathcal{T})$, so inspection takes place whenever the cost is below 0.1. This example highlights the fundamental tension for the seller: setting a high premium creates incentives for information acquisition, but it must also leave sufficient expected surplus to cover the buyer's cost of research.

Extending this framework to two buyers competing for the same option clarifies how the design of the mechanism shapes incentives for research. Suppose two investors each begin with the same prior belief μ_0 that the option will end up in the money. Each investor may conduct costly analysis—such as forecasting volatility or assessing macroeconomic factors—to generate more informative posteriors before deciding whether to purchase. The seller's objective is to design the menu of contracts so as to maximize expected revenue, while recognizing how the presence of multiple potential buyers influences their willingness to undertake such costly research.

In the simultaneous case, suppose both investors face the same contract terms at the same time. Because their analyses are independent, there is no possibility of free-riding on the other's information. Nevertheless, the presence of competition alters incentives. Each investor recognizes that even if her research yields a favorable posterior, she may still fail to obtain the contract if the rival also chooses to buy. This competitive pressure reduces the expected marginal benefit of acquiring costly information, since a successful trade is no longer guaranteed. As a result, symmetric offers may lead to weaker incentives for information acquisition, even when the contract itself would have motivated research in a single-buyer setting.

By contrast, in a sequential structure, the seller approaches one investor first and offers the contract. The first investor knows that her decision alone determines whether the trade takes place. This direct link between her posterior and the outcome strengthens her incentive to conduct research: a favorable posterior allows her to purchase the option at attractive terms, while an unfavorable posterior lets her reject without risk. The second investor is only approached if the first declines, so the seller can still preserve the possibility of trade. In this way, sequential offers may better align the contract terms with buyers' incentives to generate precise information.

Overall, when the object of trade is a financial option, the timing and structure of the contract—whether simultaneous or sequential—play a central role in determining not only the seller's expected revenue but also the extent to which investors undertake costly analysis. Simultaneous competition weakens the incentive to invest in research, while sequential allocation preserves stronger informational incentives. The mechanism's design thus governs both the efficiency of option allocation and the informativeness of buyers' decisions.

4 Benchmark Model

We begin with the case of a single buyer to better observe the effect of having an additional buyer in the following section. A buyer will be offered a menu and acquire information. Then a signal will realize which assigns the buyer privately his type (which is a posterior mean), based on which he makes a choice from the menu.

Since each posterior realization will assign 'type' to the buyer and is private information, the menu choices should screen each type. When it comes to searching for optimal screening mechanism, we cannot use revelation principle (Myerson 1981) since the buyer's types are endogenous. Mensch (2022) searches among mechanisms with *recommendation strategies*: by offering a certain mechanism, the seller recommends the buyer to acquire a specific posterior distribution $\mathcal{T} \in \Delta(\mathcal{M})$, and further specifies which choice of (x, p) to make for each posterior realization. This implies there is one to one relationship between posteriors and choices of (x, p) offered.

Since the buyer can choose any posteriors ex ante (acquires information flexibly), recommendation strategies must specify which choice of (x, p) to be made for all feasible posteriors within allocation constraint, even for those realized with 0 probability in equilibrium. We can write posterior, allocation probability, and price as functions of each other. We will use $x(\mu)$ to express incentive compatible allocation probability as a function of posterior which includes off the equilibrium path posteriors too. $p(x)$ is to express the price as a function of allocation probability. Information acquisition yields a distribution of posteriors $\mathcal{T} \in \Delta(M)$ where $\mathbf{m} = supp(\mathcal{T})$. Restricting to recommendation strategies enables us to write the seller's objective as:

$$
\begin{aligned}
&\max_{x(.),p(.),\mathcal{T}} \int p(x(\mu))d\mathcal{T}(\mu) \\
\text{s.t. [IC-A]: } &\mathcal{T} \in arg \max_{\sigma\in\Delta(M)} \int (x(\mu)\mu - p(x(\mu)))\, d\sigma(\mu) \\
&- \left[H(\mu_0) - \int H(\mu)d\sigma(\mu)\right] \\
\text{[IC-P]: } &x(\mu) \in arg \max_{(w,p(w))\in\mathcal{M}} w\mu - p(w), \forall\mu \in [0,1]
\end{aligned}
$$

with individual rationality and Bayes' rule constraint. [IC-A] is incentive compatibility ex ante, and [IC-P] is incentive compatibility ex post.[1]

4.1 Characterizing optimal mechanism

We simplify the problem by considering the information structure with at most two signals only, each recommending either trade or no trade action to the buyer:

$$
|supp(\mathcal{T})| \leq 2 \tag{1}
$$

This is without loss of generality in one buyer case, with the following assumption on $H(\mu)$.

Assumption 1. *$H \in C^3, RC(\mu) < 1$ or $RC(\mu) \geq 1$ for $\mu \in [0,1]$*

Even with two buyers, since signals recommend actions, binary signals are reasonable: one signaling to buy, the other signaling not to buy. [2] We will come back to the meaning of this assumption in section 4.2.

The seller's optimal mechanism design problem is equivalent to choosing

1. which posteriors to induce: $\mathbf{m} = \{\mu_1, \mu_2\}$

2. which allocation probabilities and prices to recommend at each posterior.

Given two posteriors μ_1, μ_2, it is equivalent to Myerson's screening contract. The seller will choose the low type's choice binding at individual rationality constraint, then make the high type's choice the highest utility yielding allocation probability, extracting the highest rent possible. It means at any given two posteriors

[1] The formal description of other constraints is characterized in Mensch (2022). Even though [IC-P] is after information acquisition, the seller must guarantee incentive compatibility for off the equilibrium posteriors which might have realized if the buyer deviated ex ante.

[2] However, this is not without loss of generality in two buyers case, since more than two posteriors could do better as in Remark 4, Mensch (2022).

$\{\mu_1.\mu_2\}$, the menu will specify so that: at the low posterior, the buyer will not trade, and at the high posterior, the buyer get the item with probability 1 at the highest possible price.

Suppose not: if the buyer is offered less than 1 allocation probability at the high posterior. Fixing the low posterior, the seller can always get higher price if it increases to 1, since there is no cost to providing higher x for the seller.

From now, we will use the notation $\underline{\mu}$ for μ such that $x(\mu) = 0$. Since this posterior realization signals the buyer to not trade, we will call it *no trade signal*. Similarly we will use the notation $\bar{\mu}$ to denote μ such that $x(\mu) = 1$, and call it *trade signal*. τ denotes the probability of trade signal realizing. Then the Bayes constraint becomes: $\tau\bar{\mu} + (1-\tau)\underline{\mu} = \mu_0$.

Incentive Compatibility. When it comes to choosing $\mathbf{m}$, flexible information acquisition restricts what's incentive compatible. That is, given $\underline{\mu}$, not only the value of $\bar{\mu}$ but also the price at $\bar{\mu}$ are fixed so that the buyer does not have an incentive to deviate to any other posterior. Following from Appendix B in the proof of Lemma 4 in Mensch (2022), incentive compatible $x(\mu)$ must satisfy the following:

$$x'(\mu) = -H''(\mu) \tag{2}$$

The concavity of H captures how fast the marginal cost of information acquisition with respect to posterior is increasing. The buyer must be incentivized by the increase in allocation probability proportional to this change in cost. With the envelope condition of [IC-P]: $\frac{\partial p}{\partial x} = \mu$, we derive how much the price should marginally

$$\hat{p}'(\mu) = -H''(\mu)\mu \tag{3}$$

We've shown that $\mathbf{m} = \{\underline{\mu}, \bar{\mu}\}$. With 2 and 3, the seller's problem simplifies to choosing the value of no trade signal $\underline{\mu}$.

4.2 Rent vs Trade Probability Tradeoff

When the seller chooses $\underline{\mu}$, there is a tradeoff due to two constraints, Bayesian consistency and allocation constraint ($x \leq 1$). For lower $\underline{\mu}$, Bayes consistency gives higher probability of $\bar{\mu}$ realizing. However the value of $\bar{\mu}$ must be lowered to meet allocation constraint $x(\bar{\mu}) = 1$. We will call this *rent vs trade probability* tradeoff.

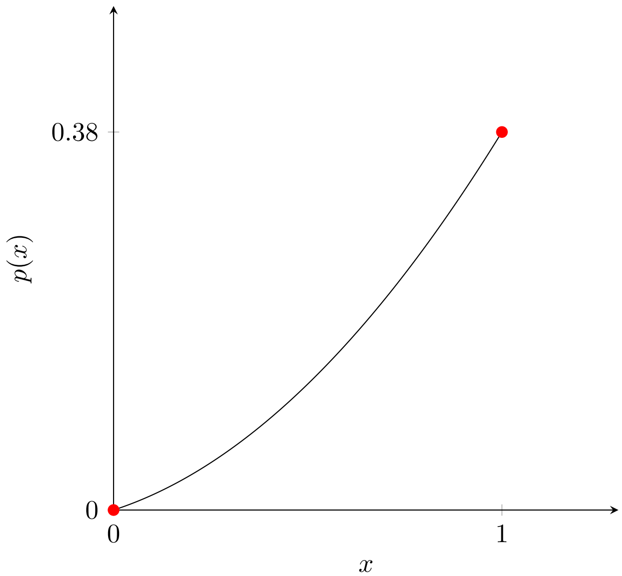

Figure 1: Supply curve when $\mu_0 = 1/2$

Here is an example that helps analyze the tradeoff.

Example 1. The common cost function used in literature is the residual variance cost function $H(\mu) = 2\kappa[\mu(1-\mu)]$. Then the implementable (x, p) is $x'(\mu) = 4\kappa$, $\hat{p}'(\mu) = 4\kappa\mu$, where κ is a cost parameter.

For $\kappa = 1/2$, $H''(\mu) = -2$ and $x'(\mu) = 2$. We can invert $x(\mu)$ to get $\mu(x)$ which denotes for a given x, which posterior the buyer is induced to acquire. Then $\mu(x) = x/2 + \underline{\mu}$. We can derive $p(x) = \int \mu(x)dx = 1/4x^2 + \underline{\mu}x$. Since $p(x)$ is convex, the seller wants to spread out posteriors to μ's such that $x(\mu) = 0$ or 1. We observe the red dots to be supply points. Then subject to Bayes rule constraint, the seller will optimize to choose no trade signal. This will determine the probability of $\bar{\mu}$ realizing as it realizes with probability $2(\mu_0 - \underline{\mu})$. At the same time, $\bar{\mu} = 1/2 + \underline{\mu}$ which is increasing in $\underline{\mu}$.

Note that the supply curve $p(x)$ will be convex as in figure 1, which implies that the seller will provide at allocation probability of either 0 or 1. This is true as long as $H'(\mu)$ is continuous. We can observe the tradeoff when the seller chooses $\underline{\mu}$ in the supply curve also. The higher the $\underline{\mu}$ is, the steeper $p(x)$ is, therefore higher rent extraction at $x = 1$. However this lowers $x(\mu_0)$ implying lower expected revenue. Solving for optimal $\underline{\mu}$, $\underline{\mu} = \mu_0/2 - 1/8$. Notice we need $\mu_0 > 1/4$ to ensure the posteriors are interior.

For $\mu_0 = \frac{1}{2}$, the posteriors acquired are $\frac{1}{8}, \frac{5}{8}$. The trade is made when $\frac{5}{8}$ realizes with probability $\frac{3}{4}$ at the price of $\frac{3}{8}$. If the buyer follows recommendation and acquires information, the payoff is $\frac{3}{4} * \left(\frac{5}{8} - \frac{3}{8}\right) = \frac{3}{16}$ whereas if not and accept at the prior, the payoff is $\frac{1}{2} - \frac{3}{8} = \frac{1}{8}$.

Corner case. Due to allocation constraint, $\underline{\mu}$ can be lowered only until $\bar{\mu} = \mu_0$. When lowering $\underline{\mu}$, if the loss from the lower value of $\hat{p}(\bar{\mu})$ is always less than the gain from the higher probability of $\bar{\mu}$ realizing, the seller will lower $\underline{\mu}$ until $\bar{\mu} = \mu_0$ and $\tau = 1$, in which the buyer will not acquire any information, staying at μ_0. This is optimal when $\hat{p}(\mu)$ is concave, meaning marginally higher posterior does not give as higher rent gain, and incentivizing high posterior acquisition ends up being too costly for the seller.

Proposition 1. *In optimal mechanism, if $RC(\mu) \geq 1$ for all $\mu \in [0, 1]$, $\bar{\mu} = \mu_0$. If $RC(\mu) < 1$ for all $\mu \in [0, 1]$, $\bar{\mu} > \mu_0$.*

$RC(\mu) \geq 1$ implies the curvature of H decreases relatively fast, meaning the buyer will be tempted to choose very high $\bar{\mu}$. In this case, the loss from low trade probability will be significant for the seller. Accordingly, the seller will want to give up higher rent extraction if $RC(\mu) \geq 1$, in which case we will have a corner case where no information acquisition is recommended (Case 1 hereafter). Contrarily, $RC(\mu) < 1$ implies the curvature decreases relatively slow, so it is expensive enough for the buyer to acquire higher posteriors. In this case the loss from low trade probability will not be as significant. Accordingly, the seller still wants to extract higher rent at the sacrifice of low trade probability by inducing information acquisition (Case 2 hereafter). [3]

For example, $H(\mu) = -\mu log(\mu)$ satisfies $RC(\mu) = 1$ (Figure 2 - black curve). For any functional form satisfying $H(\mu) = -\mu log(\mu) + \epsilon\mu^{\beta}$ for small $\epsilon > 0$ and $0 < \beta < 1$, we have $RC(\mu) > 1$ (Figure 2 - red curve). The concavity decreases faster as posterior increases. In this case, the seller will offer a price where the buyer stays ignorant and accept with probability 1. Contrarily, quadratic entropy function $H(\mu) = 2\mu(1-\mu)$ satisfies $RC(\mu) = 0$ (Figure 3).

Although the seller's revenue is not directly given by $H(\mu)$, but rather by an expression involving its curvature—such as $\int -H''(m)mdm$—the second and third derivatives of the seller's revenue are proportional to those of $H(\mu)$. As a result, the curvature index $RC(\mu) = -\frac{H'''(\mu)\mu}{H''(\mu)}$ is invariant under this transformation and remains a valid descriptor of the seller's local trade-offs. Even though the seller is not literally saving, $RC(\mu)$ captures their willingness to forgo trade probability in exchange for more informative posteriors and higher rent extraction, much like a coefficient of relative prudence in standard decision theory.

4.3 Difference from inflexible information acquisition

Inflexible information acquisition differs from flexible information acquisition in the range of experiments available to the buyer. With inflexible information acquisition, the buyer must choose from a fixed set of

[3] It is possible that a given H satisfies neither. For example, it happens when $\hat{p}'(\mu)$ is discontinuous, as discussed in Example 2, Mensch (2022). We do not discuss this case here.

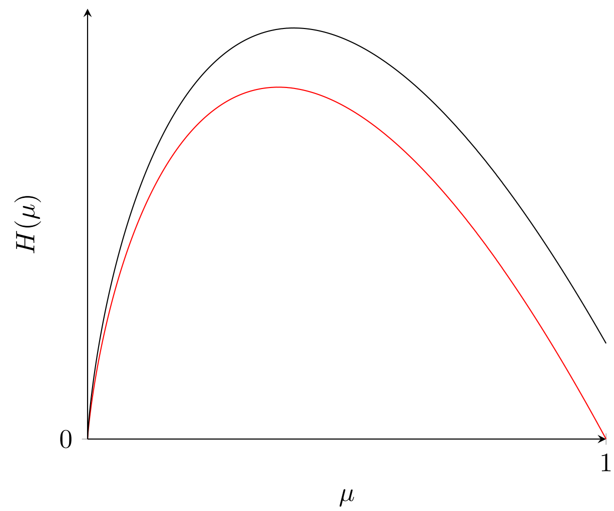

Figure 2: Case 1 ($RC(\mu) \geq 1$)

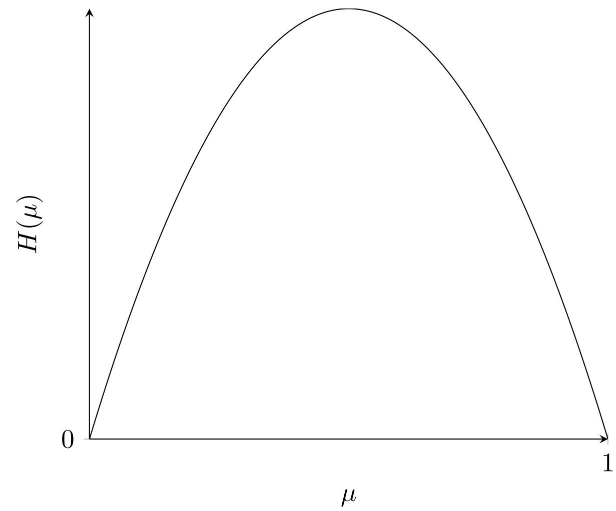

Figure 3: Case 2 ($RC(\mu) < 1$)

information structures. For example, an investor evaluating a financial option may only be able to purchase research reports that differ in their exactness, where a coarse report provides a low-precision signal about the option's payoff (e.g., predicting only whether returns will be above or below a threshold), while a more detailed report delivers a high-precision signal that refines the posterior distribution more accurately. In contrast, flexible information acquisition allows the investor to design and tailor the precision of the signal to their needs—for instance, one investor may focus on stress-testing downside risks, while another may search for conditions under which the payoff distribution is especially favorable. In this context, choosing a lower no-trade signal corresponds to acquiring a more precise (and thus more costly) report that emphasizes detecting unfavorable outcomes.

What differentiates this model from inflexible information acquisition frameworks is the endogeneity of the no-trade signal. In inflexible information acquisition models (e.g., Shi, 2012), buyers have only one dimension along which they can deviate: the overall informativeness of the signal structure. That is, once the informativeness level is chosen, the probabilities of each signal realization are fully determined. Simplifying Shi (2012) to the case of binary signals, buyers can only choose how dispersed the signals are—how distinguishable the high and low signals are from each other. In contrast, under flexible information acquisition, buyers can independently choose the precision of both the good signal ($\bar{\mu}$) and the bad signal ($\underline{\mu}$). This flexibility makes the no-trade region endogenous and restores a rent–trade probability tradeoff: buyers can strategically concentrate learning on the no-trade region, which in turn affects the likelihood of trade. Without this endogeneity, the trade probability would be independent of μ, and the fundamental tradeoff would disappear.

5 Two Buyers: Simultaneous Mechanism

Now suppose there are two potential buyers and the seller will offer menus to buyers simultaneously. Upon observing the menus, buyers will acquire information simultaneously. Then their posteriors will realize privately and independently, based on which they will decide which option to choose.

5.1 Symmetric simultaneous mechanism

The literature has studied symmetric simultaneous mechanism with N (Mensch (2022)). The mechanism is now constrained by tighter resource constraint [4]:

$$\int_{x^*}^{1} x d\mathcal{T}(\mu(x)) \leq \frac{1-\mathcal{T}(\{\mu : x(\mu) < x^*\})^N}{2}, \forall x^* \in [0,1]. \tag{4}$$

4 implies that the seller cannot promise to provide with 1 probability even under high posterior realization since the other buyer could also have the high posterior realized, in which case the buyers will get the item with less than 1 probability. This puts additional constraint on incentivizing high posteriors to buyers.

Proposition 3, Mensch (2022) shows that for convex $\hat{p}(\mu)$ and concave $x(\mu)$, the seller will allocate the item in the following way [5]:

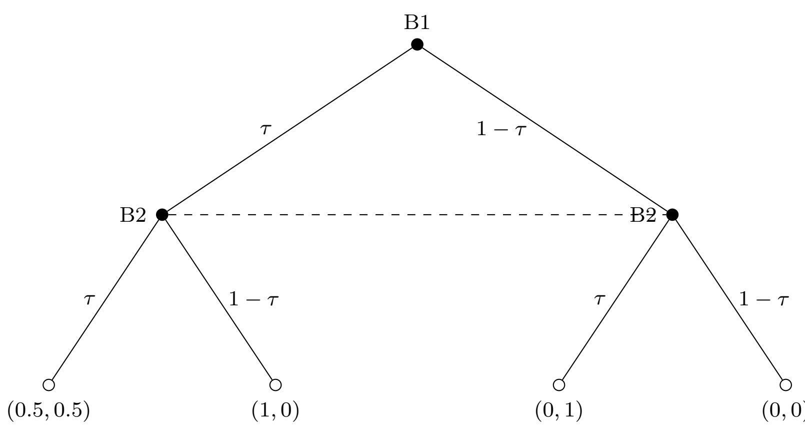

Allocation probability when simultaneously meeting buyers with symmetric mechanism

5.2 Dropping symmetric constraint

While symmetric and simultaneous mechanisms are widely used as tractable benchmarks in the literature, they may fail to capture important strategic dynamics that arise in real-world settings. In many practical environments—especially those involving only two buyers—the assumption of symmetry is often violated. Sellers frequently have access to information that enables them to differentiate among buyers, whether based

[4]Border (1991, Proposition 3.2)
[5]If the seller did not induce information acquisition, the buyers will have equal probability of getting the item at their prior.

on observable characteristics, prior interactions, or relationship history. Such asymmetric treatment can be strategically advantageous. For instance, offering different pricing menus or using sequential bargaining can improve surplus extraction or create competitive pressure (Gershkov et al. 2021; Bergemann and Pesendorfer 2007). These examples suggest that rigid symmetry, while analytically convenient, may overlook valuable seller strategies rooted in differentiation and timing. This motivates a closer examination of mechanisms that permit asymmetric treatment of buyers.

In the one-buyer case, since allocating a higher probability to a single buyer carries no opportunity cost, the seller can freely offer high allocation to incentivize information acquisition. However, with two buyers, allocation becomes a constrained resource. Now, granting a higher allocation probability to one buyer comes at the opportunity cost of allocating less to the other. As a result, the seller must decide how to distribute a single unit of the good between two buyers in a way that induces just enough information to extract revenue, without incurring excessive allocation costs.

We show that the seller may strategically favor one buyer by guaranteeing higher allocation probability when a high posterior realizes. This is profitable when higher posteriors are not too expensive to incentivize. In particular, a buyer who receives a high signal can be motivated with a smaller marginal increase in allocation probability, while the rent extracted from such a buyer increases more than proportionally. Thus, the seller can reallocate a small amount of allocation from the lower-posteriors (or unfavored) buyer to the high-posteriors (favored) buyer, and the gain in rent exceeds the loss in allocation value. At the same time, high posteriors cannot be too cheap, or the marginal rent gain would be insufficient. This trade-off—between the rent from one buyer and the forgone allocation value from the other—is central to the seller's incentive to induce asymmetry.

We formalize this intuition in the following extensive form analysis:

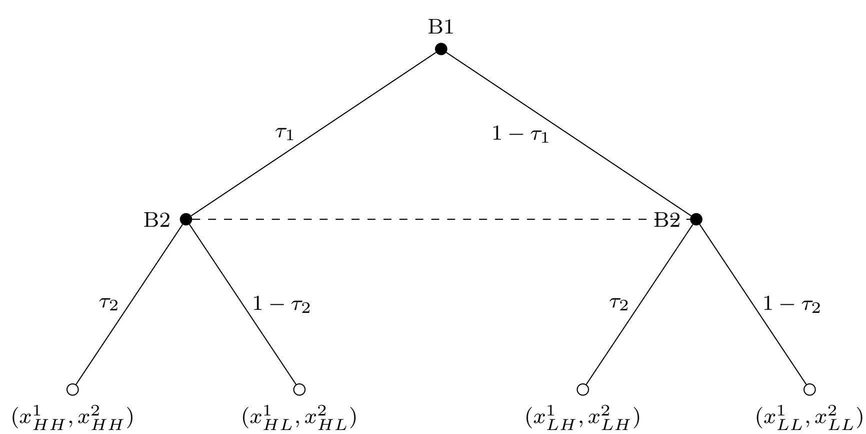

Allocation probability when simultaneously meeting buyers

The seller will maximize

$$\tau_1 p_1(x_H^1) + (1-\tau_1)p_1(x_L^1) + \tau_2 p_2(x_H^2) + (1-\tau_2)p_2(x_L^2) \tag{5}$$

Subject to resource constraints:

$$x_{HH}^1 + x_{HH}^2 \leq 1$$
$$x_{HL}^1 + x_{HL}^2 \leq 1$$
$$x_{LH}^1 + x_{LH}^2 \leq 1$$
$$x_{LL}^1 + x_{LL}^2 \leq 1$$

where $x_H^1 = \tau_2 x_{HH}^1 + (1-\tau_2)x_{HL}^1$, $x_L^1 = \tau_2 x_{LH}^1 + (1-\tau_2)x_{LL}^1$, $x_H^2 = \tau_1 x_{HH}^2 + (1-\tau_1)x_{LH}^2$ and $x_L^2 = \tau_1 x_{HL}^2 + (1-\tau_1)x_{LL}^2$, with Bayesian consistency constraints.

We now turn to characterizing the optimal simultaneous mechanism in each case. Recall that in Case 1, the seller prefers to deter information acquisition, as the expected loss from the resulting reduction in trade probability outweighs the potential benefit from better price discrimination. By contrast, in Case 2, the seller wishes to encourage information acquisition as much as possible, since the additional rent that can be extracted from informed buyers more than compensates for the lower probability of trade.

Case 1. We show that even without symmetric constraint, the seller prefers to allocate symmetrically.

Proposition 2. *In optimal simultaneous mechanism with two buyers, the seller will implement no information acquisition to both buyers and offer the item with equal allocation probability, if the following holds:*

$$RC(\mu) \geq 1 \tag{6}$$

Intuitively, the only reason to break symmetry would be to give one buyer stronger incentives to acquire costly information. However, when the seller prefers to deter information acquisition altogether, this motive disappears. In that case, the seller gains nothing from creating asymmetry and instead maximizes expected revenue by treating buyers symmetrically.

Case 2. We show that the seller will end up allocating as following:

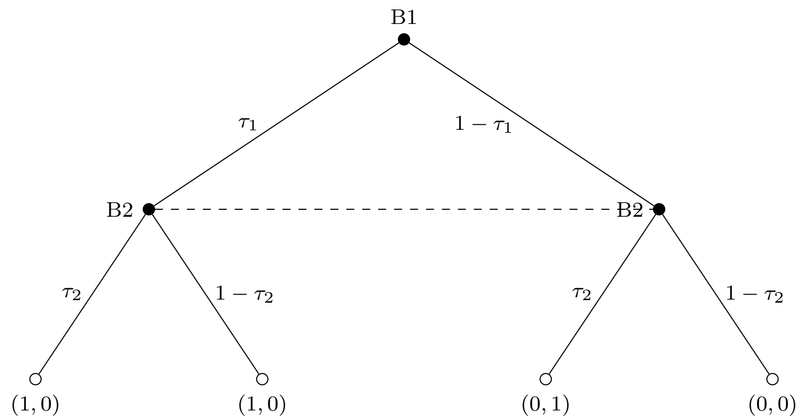

Allocation probability when simultaneously meeting buyers

Two buyers will be offered the item at the same time, but B1 will be favored to have the item for sure if his posterior realizes to be high. B2 has a chance only if B1's posterior was realized to be low. [6] The favor that B1 gets incentivizes B1 to get higher posterior which enables the seller to charge higher price.

The following theorem summarizes this result.

Theorem 1. *In optimal simultaneous mechanism with two buyers, the seller will implement posteriors* $(\underline{\mu}_1, \bar{\mu}_1), (\underline{\mu}_2, \hat{\mu}_2)$ *to each buyer so that* $x_H^1 = 1$, $x_H^2 = 1 - \tau_1$, *if the following holds:*

$$0 \leq RC(\mu) < 1 \tag{7}$$

where $\hat{\mu}_2$ *is such that* $x_2(\hat{\mu}_2) = 1 - \tau_1$ *with* $\tau_1 \bar{\mu}_1 + (1 - \tau_1)\underline{\mu}_1 = \mu_0$.

Notice 7 is derived by putting together $RP(\mu) < 1$ and $-H'''(\mu) \leq 0$. Intuitively, $RP(\mu) < 1$ ensures that the seller is not excessively averse to risk in the face of uncertain trade outcomes. When offering a higher price to B1, the seller must accept a lower probability of trade—a form of risk-taking that is only optimal when the seller's effective precautionary motive is sufficiently weak. Although the seller is not literally engaging in precautionary saving, the trade-off is analogous: the seller is deciding how much allocation risk to tolerate in order to extract more rent from informative posteriors. The seller's revenue is constructed from $H''(\mu)$ $(\tau \int -H''(m)mdm)$, and the relevant curvature index $RC(\mu)$ remains invariant. As such, $RC(\mu)$ plays a role similar to the coefficient of relative prudence, capturing the seller's local willingness to trade off certainty in allocation for informational gains.

The concavity of $x(\mu)$ $(-H'''(\mu) \leq 0)$ implies that the marginal cost of increasing posteriors is lower at higher values of μ. As a result, shifting a fixed amount of informational effort toward B1 increases her

[6] This is still different from sequential offers where the offers are made sequentially, which will be discussed in the following section

posteriors more than it decreases B2's. This makes such a reallocation beneficial, as the seller can generate more value from the same amount of allocation. Therefore, the seller optimally directs the item toward buyers where allocation can be used most efficiently—maximizing informational returns.

5.3 Difference from costly investment model

Both our model and Gershkov et al. (2021) show that it can be optimal for the seller to favor one buyer over another, even when buyers are ex ante symmetric. In both settings, favoritism functions as a tool to influence buyers' endogenous actions. By tilting allocation probabilities toward a particular buyer, the seller strengthens that buyer's incentive to take an action that ultimately benefits the seller—whether by acquiring information in our framework or by making a costly investment in theirs. In both cases, favoritism may reduce efficiency, since the unfavored buyer could turn out to have the higher valuation ex post, yet the seller finds it worthwhile to bias allocation in order to induce the desired behavior.

The two models differ, however, in the nature of endogeneity and the underlying tradeoff. In Gershkov et al., buyers make costly investments that deterministically lower the seller's cost, so allocation serves as a reward for investment. Because types shift deterministically and signals are absent, Bayesian consistency is not required, and allocation decisions do not depend on posteriors. By contrast, our model features flexible information acquisition: buyers choose the precision of their signals, and the seller leverages favoritism to stimulate more precise information in order to extract informational rents. Here, the tradeoff lies between higher rent extraction and lower trade probability, since allocation must be incentive compatible and therefore tied to buyers' posteriors. Favoritism is optimal only when the marginal rent gain outweigh the loss from reduced trade.

Remark *Without flexible information acquisition setting, the characterized optimal asymmetric simultaneous mechanism can be implemented by a sequential one, as shown by Gershkov et al. (2021). That is, the seller can replicate the allocation rule and achieve the same revenue by offering the item sequentially. However, this equivalence does not hold in our setting. Since buyers endogenously acquire information in response to the offered menu, the seller's revenue depends not only on the allocation probability but also on when and how that allocation is delivered. In particular, collecting payment conditional on allocation with probability 1 and trade occurring with probability* $1 - \tau$ *is not equivalent to collecting payment with allocation probability* $1 - \tau$ *for sure. The difference arises because buyers' information acquisition incentives depend on their expected surplus, which is shaped by the full menu. As a result, whether the seller prefers the sequential or simultaneous mechanism depends on the incentive to spread out posteriors and how costly it is to induce higher posteriors under the allocation constraint (we will return to this trade-off in detail* *in Section 7). This distinction highlights why the sequential mechanism must be analyzed separately in the flexible information acquisition setting. Unlike in Gershkov et al. (2021), where sequential allocation simply replicates the simultaneous outcome, here the timing of allocation interacts directly with buyers' incentives to acquire information. To capture these effects, the next section introduces and characterizes the sequential mechanism under flexible information acquisition.*

6 Two Buyers: Sequential Mechanism

We observed that favoring one buyer can improve the seller's revenue. The seller can do even better by taking advantage of the timing structure—resetting the resource constraint when moving from one buyer to the next. Specifically, the seller can first offer the item to buyer 1 (B1); if trade does not occur, the seller then offers the item to buyer 2 (B2). Because B2 is approached only after B1 has declined to trade, the resource constraint effectively resets—allowing the seller to offer the item to B2 with allocation probability 1. We assume B1 and B2 are ex ante identical in terms of both their priors and information cost functions, and that their signal realizations are independent. The timing of the interaction proceeds as follows:

1. B1 is offered a menu
2. B1 decides their information acquisition strategy. The signal realizes and B1 chooses whether to buy or not
3. If the item was unsold, B2 is offered a menu
4. B2 decides their information acquisition strategy. The signal realizes and B2 chooses whether to buy or not

The extensive form representation will be:

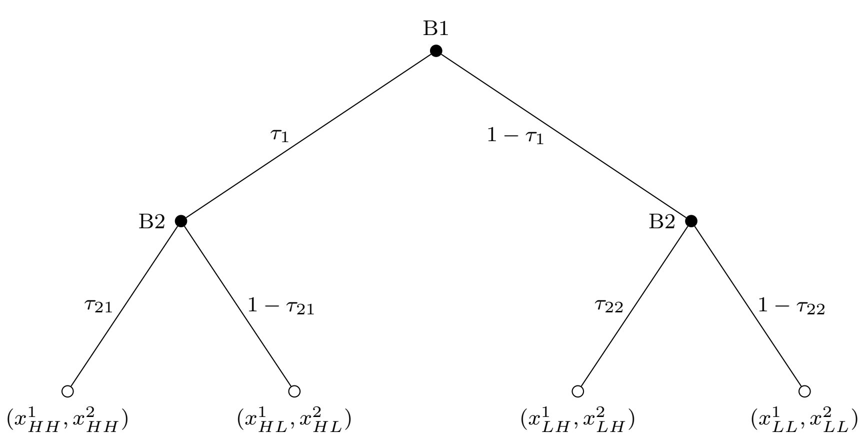

Allocation probability when sequentially meeting buyers

B2 will be recommended to acquire information as in one buyer case. Then we can write the seller's problem as

$$\begin{aligned} &\max_{\mathcal{T}} E_{\mathcal{T}} \hat{V}_1(\mu) \\ &\text{s.t. } \int_{\mathcal{T}} \mu d\mathcal{T}(\mu) = \mu_0 \\ &\text{where } \hat{V}_1(\mu) = \hat{p}_1(\mu) + (1 - x_1(\mu))V_2, \end{aligned}$$

and V_2 is the expected transfer from B2. [7]

6.1 Effect of having an additional buyer

The effect of having an additional buyer manifests differently in the simultaneous and sequential mechanisms. In the simultaneous mechanism, the impact is primarily in terms of how much allocation probability the seller can guarantee to each buyer, since the resource constraint applies jointly across both buyers. In the sequential mechanism, however, the seller can take greater risk with B1—offering a higher price and thus inducing a higher posterior—because any loss from a lower trade probability with B1 can be recovered by subsequently offering the item to B2. This sequential structure grants the seller additional flexibility in balancing rent extraction with trade probability, a feature that is most clearly illustrated in the Case 2 result.

Case 2. When $RC(\mu) < 1$, we show that both buyers acquire the maximum amount of information that is incentive-compatible: posterior beliefs are dispersed as much as possible. The seller's incentive depends on the rent–versus–trade-probability trade-off. As shown in Proposition 1, the seller prefers to spread out allocation probabilities to generate steeper posteriors—up to the point where the allocation constraint binds. This holds even with two buyers, which is summarized in the next proposition.

Proposition 3. *Suppose the seller induces as much information as possible with one buyer (B2) and there is another buyer (B1) with the same cost function to whom the seller can offer the item before offering to B2. Then the seller offers each buyer the item on a take-it-or-leave-it basis, i.e. with allocation probability either 1 or 0. Moreover, B1's no trade signal value is always higher than B2's and the price offered to B1 is higher than that to B2.*

Proposition 3 also implies that the seller induces a more optimistically skewed signal for B1—information

[7] Since x is allocation probability, conditional on the item was not sold to B1, the problem is the same as one buyer case. If we interpret x as quantity instead of allocation probability, $\hat{V}_1$ will be different. Residual quantity after trading with B1 will be a constraint to B2's available quantity. We discuss this further in Section 5.

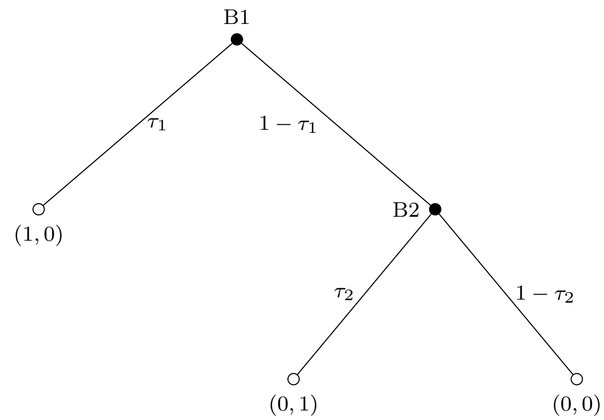

Case 2: Allocation probability when sequentially meeting buyers

that is more informative about the good-quality realization—in order to justify charging a higher price, even though trade may occur with lower probability. The seller is willing to accept this risk because B2 serves as a fallback: if B1 does not trade, the seller can still offer the item to B2.

While both buyers acquire the maximum incentive-compatible amount of information, they differ in the direction of the information they prioritize. This asymmetry arises from differences in their no-trade signals, which determine how posterior dispersion is distributed around the prior. If the lower posterior lies farther from the prior than the upper posterior, the signal is more informative about the low-value state, reflecting the buyer's attempt to avoid a bad outcome. Conversely, if the upper posterior lies farther from the prior, the signal is more informative about the high-value state, reflecting the buyer's attempt to confirm a good outcome. Thus, even though both buyers achieve the same overall level of posterior dispersion, their learning strategies diverge: one acquires information that is relatively more revealing about downside risks, while the other acquires information that is relatively more revealing about upside potential.

A direct corollary of this analysis is that, although both buyers acquire information in an incentive-compatible way, the amount of information induced may differ across them. The difference does not come from one buyer's information being inherently more revealing about the high-value state or the low-value state, but instead from the marginal curvature of the information cost function. When the cost function is more convex at a given prior, the seller optimally scales back the precision of signals for that buyer, while for the buyer facing a less convex region of the cost function, the seller induces more precise information. Thus, heterogeneity in induced information reflects the shape of the cost function rather than any asymmetry in the informativeness of particular states.

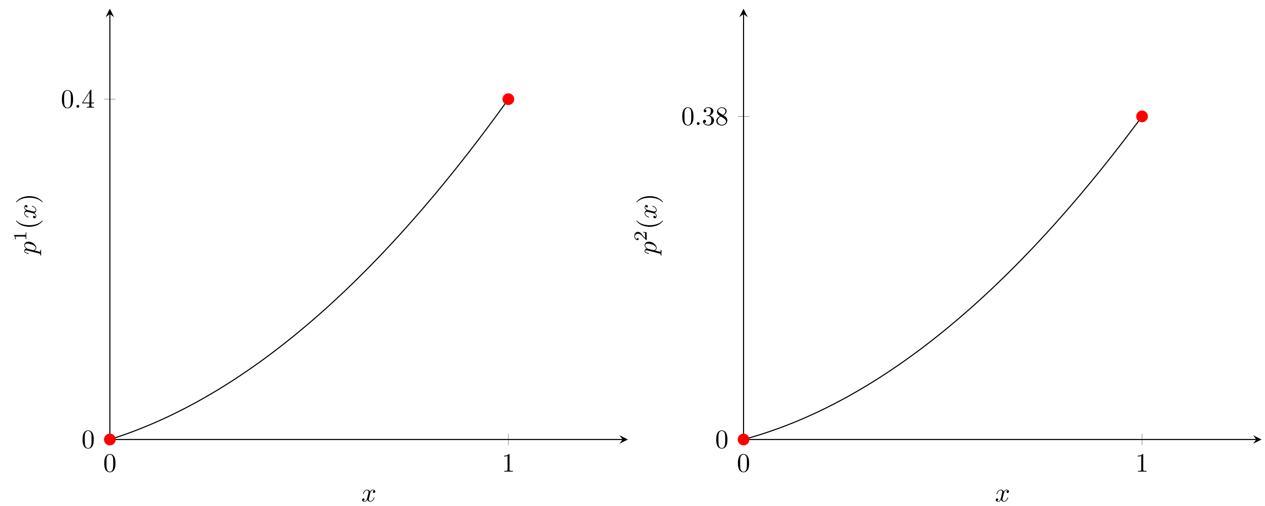

Figure 4: Supply curves for sequential mechanism

Corollary 1. *Suppose $RC(\mu) < 1$. The seller induces more information to B1 than B2 if*

$$-H'''(\mu) < 0.$$

The seller induces less information to B1 if

$$-H'''(\mu) > 0.$$

That is, when $-H'''(\mu) < 0$, the curvature of $H(\mu)$ is decreasing, making it relatively cheaper to expand posterior dispersion for B1, so the seller induces more information. Conversely, when $-H'''(\mu) > 0$, the curvature is increasing, raising the cost of dispersion for B1, and the seller correspondingly scales back her induced information.

Recall from Example 1 (single-buyer case) that the seller prefers a higher no-trade signal, as it raises the price that can be charged, but doing so comes at the cost of a lower probability that the signal will realize. With an additional buyer, the seller is less concerned about this trade-off—B2 provides a fallback—so she is more willing to take this risk with B1 to induce steeper posteriors. This logic is confirmed in the numerical example from Example 1: solving for each buyer's posteriors and prices, we find that the supply curve is steeper for B1, driven by the higher non-participation cutoff. For a common prior $\mu_0 = 1/2$, the no-trade signal values are 0.145 for B1 and 0.125 for B2. Consequently, the probability of reaching the high posterior differs: 0.71 for B1 versus 0.75 for B2.

Case 1. This risk-taking behavior might persist even when the seller prefers safe trade probability (i.e. $RC(\mu) \geq 1$): we observe cases where the presence of an additional buyer leads the seller to induce information acquisition from one buyer, even though doing so would not have been optimal in the single-buyer case. The availability of a fallback option allows the seller to tolerate lower trade probability and pursue more aggressive rent extraction.

Theorem 2. *Suppose the seller does not induce information with one buyer (B2) and there is another buyer (B1) with the same cost function and prior to whom the seller can offer the item before offering to B2. The seller induces information acquisition to B1 if the following conditions hold for all* $\mu \in [\underline{\mu}, 1]$

$$1 \leq RC(\mu) \leq 1 + \frac{V_2 H'''(\mu)}{-H''(\mu)} \tag{8}$$

where $V_2 = \int_{\underline{\mu}}^{\mu_0} -H''(\mu)\mu d\mu$, $\underline{\mu} = G(1 + H'(\mu_0))$ *with* $G(x) = (H')^{-1}(x)$.

Notice 8 is derived by putting together

$$RC(\mu) \geq 1 \tag{9}$$

$$H'''(\mu)V_2 - H''(\mu) - \mu H'''(\mu) \geq 0 \tag{10}$$

Specifically, 9 ensures a corner solution in the single-buyer case, while 10 implies that the seller's value function is convex in B1's posterior distribution. Together, these conditions suggest that the seller prefers to induce a spread in B1's posteriors away from the prior, despite the associated risk. Intuitively, this requires the relative curvature index $RC(\mu)$ to be sufficiently low, so that the seller is willing to tolerate risk in exchange for informativeness. This incentive becomes stronger as the expected revenue from B2 (V_2) increases.

In Section 3.1, we identified a corner solution in the single-buyer case, where the seller opts not to induce information because the cost from the Bayesian plausibility constraint outweighs the gain from shifting posteriors and extracting surplus. However, the presence of an additional buyer can shift this trade-off: even if information provision is suboptimal in the single-buyer setting, the seller may now find it beneficial to induce informative signals from B1 in order to increase the overall expected revenue.

In equilibrium,

1. The seller induces a binary signal $\{\underline{\mu}, \bar{\mu}\}$ for B1, with an interior distribution over the posteriors.

2. If $\bar{\mu}$ is realized, the item is sold to B1. If $\underline{\mu}$ is realized, the seller moves on to B2.

3. B2 is offered a price $\hat{p}(\mu_0)$ which is accepted with probability 1.

In this mechanism, B1 is offered a higher price than B2 but trades with less than full probability. For 10 to hold, the seller's fallback value V_2 from offering to B2 must be sufficiently high to justify the risk taken in recommending higher posteriors to B1. This result stands in contrast to the two-bidder mechanism in Gershkov et al. (2019), where the seller incentivizes a costly investment by granting B1 higher allocation probability. In our model, however, the allocation probability at the low posterior $\underline{\mu}$ is fixed by the consistency constraint, so the seller cannot directly reward the acquisition of higher posteriors. Instead, the seller is more tolerant of the risk that the high posterior may not realize, because if it does not, the seller can still trade with B2. The presence of B2 as a backup fundamentally shifts the seller's willingness to induce information acquisition in B1, even in settings where it would otherwise be suboptimal.

6.2 Difference from free information model

In Bergemann and Pesendorfer (2007) (BP), the seller has full flexibility to choose any signal structure for the buyers without cost. In contrast, in our model, the seller effectively chooses the information structure indirectly by offering a menu before buyers acquire information, which in turn shapes their endogenous learning decisions. Despite this difference in implementation, both models feature the seller inducing binary information for the first buyer to increase their willingness to pay, but no information for the second. However, a key distinction lies in how rents are handled: in BP, the seller designs the mechanism so that no information rents are left to either buyer, whereas in our setting, the seller must leave positive information rent to both buyers. This difference arises because in our model, buyers choose whether to acquire costly information in response to the menu, rather than freely receiving the information.

6.3 Generally solving using the convexity

Mathematically, the convexity of the seller's value function $\hat{V}_1(\mu)$ determines whether the seller finds it optimal to induce information acquisition for B1. This follows from the logic of Bayesian persuasion. The seller spreads out B1's posterior distribution only when doing so increases expected revenue, and the convexity of $\hat{V}_1$ guarantees that binary posteriors strictly improve the seller's payoff relative to remaining at the prior.

Example 1. In the case of two buyers with the residual variance cost function, the seller's value function takes the form:

$$\hat{V}_1(\mu) = 2\kappa(\mu^2 - \underline{\mu}^2) + (1 - 4\kappa(\mu - \underline{\mu}))V_2$$

where $V_2 = \frac{1}{2}\left(\mu_0 + \frac{1}{4}\right)^2$. Since $\hat{V}_1(\mu)$ is convex in μ, the seller optimally induces as much information as possible for B1. The only term that could alter the convexity of $\hat{V}_1$ is the one involving $(1 - x(\mu))$, but because $x(\mu)$ is linear in this example, the overall function remains convex.

Example 2. Suppose now that $H(\mu) = -\mu[log(\mu)]$, which gives $x'(\mu) = 1/\mu$, $\hat{p}'(\mu) = 1$. Then we have

$$\hat{V}_1(\mu) = \mu - \underline{\mu} + \left(1 - log\left(\frac{\mu}{\underline{\mu}}\right)\right) V_2$$

where $V_2 = \mu_0(1 - 1/e)$. Since $(1 - x(\mu))$ is convex and $\hat{p}(\mu)$ is linear, $\hat{V}_1(\mu)$ remains convex. In the single-buyer case, the seller would have been indifferent between inducing or not inducing information acquisition because the linearity of $\hat{p}$ implies no strict incentive to spread posteriors. However, with the presence of a second buyer, the seller strictly prefers to induce information acquisition in B1. Intuitively, concave $x(\mu)$ implies that higher posteriors are less effective at incentivizing buyers. But with a second buyer available for fallback trade (via $1 - x(\mu)$), the seller is more willing to induce riskier, more informative signals for B1.

This comparison highlights that even though B1 and B2 are ex ante symmetric, their recommended information structures may differ depending on the shape of $\hat{p}(\mu)$ and $x(\mu)$. We summarize the possibilities below:

	$\hat{p}(\mu)$	$x(\mu)$	B1's vs B2's recommendation
Case 1	concave	convex	same
Case 2	convex	concave	same
Case 3	concave	concave	could differ
Case 4	convex	convex	could differ

The most interesting cases are Case 3 and Case 4, where the seller may recommend information acquisition for B1 but not for B2, despite their symmetry. For example, in Case 4 with both $\hat{p}(\mu)$ and $x(\mu)$ convex, the addition of the second buyer can lead to asymmetric information provision, even though no such asymmetry would arise in the single-buyer case. This asymmetry arises because the convexity of $\hat{p}$, which depends on both μ and $x'(\mu)$, can dominate the impact of the added fallback option, maintaining or amplifying the incentive to spread out B1's posteriors.

Example 3. Consider $H(\mu) = \log(\mu)$, implying $x'(\mu) = \frac{1}{\mu^2}$ and $\hat{p}'(\mu) = \frac{1}{\mu}$. The seller's value function is then:

$$\hat{V}_1(\mu) = log\left(\frac{\mu}{\underline{\mu}}\right) + \left(1 - \left(-\frac{1}{\mu} + \frac{1}{\underline{\mu}}\right)\right) V_2$$

where $V_2 = \log\left(\mu_0/\underline{\mu}(\mu_0)\right) = \log(\mu_0 + 1)$. Again, $\hat{V}_1(\mu)$ is convex, so the seller prefers to spread out B1's posteriors and induce information acquisition. Meanwhile, B2 remains uninformed and is offered a deterministic trade at the prior. This example shows how the presence of a second buyer can lead the seller to take targeted informational risks with B1, even while leaving B2 uninformed.

7 Comparing Simultaneous and Sequential Mechanism

The tradeoff between rent extraction and trade probability persists even when an additional buyer is introduced. If the seller prioritizes rent extraction, the sequential mechanism is more attractive. By meeting buyers one at a time, the seller can induce higher posteriors and achieve stronger price discrimination. This is because the resource constraint is reset after a failed trade with B1, and the seller can still offer the item to B2 with certainty. This structure enables the seller to extract payments from both buyers, though at the cost of a lower overall probability of trade. If instead the seller values maximizing the probability of trade, the simultaneous mechanism becomes optimal. By offering the item to both buyers at once, the seller cannot guarantee allocation to any single buyer, but can ensure a higher overall likelihood of trade relative to the sequential mechanism. The following theorem formalizes this tradeoff between extracting higher rents sequentially and securing higher trade probabilities simultaneously.

Theorem 3. *Suppose there are two ex-ante symmetric buyers. The seller wants to offer the item simultaneously rather than sequentially if the following holds for all* $\mu \in [0, 1]$

$$RC(\mu) \geq 1 \tag{11}$$

Contrarily, the seller wants to offer the item sequentially rather than simultaneously if

1. *the following holds for all* $\mu \in [0, 1]$

$$0 \leq RC(\mu) < 1 \tag{12}$$

2. *the following holds for all* $\mu \in [\mu_0, 1]$

$$-H''(\mu) < \frac{1}{2\mu^2} \tag{13}$$

Condition 12 ensures that the seller prefers to induce dispersion in posteriors under the simultaneous mechanism. However, sequential provision becomes more attractive when the seller can reset B2's resource constraint—offering the item to B2 only if B1 does not trade—allowing for more targeted rent extraction.

Conversely, when the seller has a strong precautionary motive—effectively valuing a higher trade probability—she may prefer to discourage information acquisition altogether. Doing so requires offering buyers enough surplus to deter them from acquiring information. This becomes more difficult when the resource constraint less tight, as buyers are more willing to acquire information when high allocation probabilities are available.

Finally, condition 13 ensures that higher posteriors are not excessively costly for buyers. If they were, buyers would require disproportionately higher allocation probabilities to justify acquiring such information. In that case, the seller would prefer to allocate the item to the other buyer at a lower posterior, making simultaneous provision more favorable.

8 Discussions

8.1 Allocation probability vs quantity

In one buyer case, whether to interpret x as allocation probability or quantity did not matter. In sequential mechanism, the revenue is different for the seller. If it is interpreted as quantity, it is equivalent to having two buyers at each posterior realization of B1 where the seller has the resource constraint of $(1-x_1)$ quantity. Therefore, we can conclude depending on whether the seller wants to spread out the posteriors as much or not (convexity of $\hat{p}$), it is better to choose whether to allow the option of splitting the item or not.

For convex $\hat{p}$, the seller wants to spread out the signals as much, therefore in equilibrium the result will be the same whether the item can be split or not. However if buyers have discontinuous cost function so that in equilibrium the seller wants to induce interior allocation probabilities (as in Ex2 from Mensch 2022), the seller will prefer one way to the other.

8.2 Sequential mechanism with correlated values

We have previously assumed that there is no dependence between two buyers' posteriors. Now we assume that they have common values: B1's posterior realization is observed to B2 which is updated as the new expected value. Timing can be summarized as following:

1. Seller offers $(x_1, p_1(x))$ to B1 first
2. B1 acquires information
3. B1's posterior μ_1 realizes and decides whether to buy or not

4. If the item was not sold, the seller offers B2 mechanism $(x, p(x))$ with μ_1 as prior

5. B2 acquires information

6. μ_2 realizes, B2 decides whether to buy or not

Now B2 will be recommended to acquire as in one buyer case, but now as a function of realized μ_1. Then the seller's problem of B1's recommendation will be:

$$
\begin{aligned}
&\max_{\mathcal{T}} E_{\mathcal{T}} \hat{V}_1(\mu) \\
&\text{s.t.} \int \mu d\mathcal{T}(\mu) = \mu_0 \\
&\text{where } \hat{V}_1(\mu) = \hat{p}_1(\mu) + (1 - x_1(\mu))V_2(\mu^M(\mu))
\end{aligned}
$$

Unlike in private value case, the seller does not want to take risk anymore when 9 and 10 hold, since the value from B2 now depends on B1's signal, therefore B2 cannot serve as a backup. In other words, the incentive to fully extract rent from B1 always overwhelms the incentive to leave the allocation probability to the second buyer. The integrated value function is no longer convex since $(1-x(\mu_1))$'s convexity is multiplied by $V_2(\mu_1)$ which is concave.

8.3 Heterogeneous buyers

We have assumed two buyers are symmetric ex-ante: they have the same prior expected value. The seller could randomly select any buyer to offer the item sequentially. If one of the buyers have higher expected value, in sequential mechanism, which buyer would the seller want to offer the item first? The seller wants the stronger buyer first since the seller can sell it at a higher price. In fact the seller will be better with heterogeneous buyers since the optimal posteriors can be induced with less cost.

References

Bergemann and Pesendorfer. 2007. "Information Structures in Optimal Auctions." *Journal of Economic Theory* 137: 580-609.

Bulow and Klemperer. 1996. "Auctions Versus Negotiations." *American Economic Review* 86: 180-194.

Gershkov et al. 2021. "A Theory of Auctions with Endogenous Valuations." *Journal of Political Economy* 129 (4): 981-1318.

Kamenica and Gentzkow. 2011. "Bayesian Persuasion." *American Ecomonic Review* 101 (6): 2590 - 2615.

Kim and Koh. 2022. "Auctions with flexible information acquisition." *Games and Economic Behavior* 133: 256-281.

Mensch. 2022. "Screening inattentive buyers." *American Economic Review* 112 (6): 1949-1984.

Mensch and Ravid. 2022. "Monopoly, Product Quality, and Flexible Learning." Working paper

Myerson. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6(1), 58-73.

Ravid. 2020. "Ultimatum Bargaining with Rational Inattention." *American Economic Review* 110(9), 2948-63.

Ravid, Roesler, and Szentes. 2022. "Learning Before Trading: On the Inefficiency of Ignoring Free Information." *Journal of Political Economy* 130(2), 346-387.

Roesler and Szentes. 2017. "Buyer-optimal Learning and Monopoly Pricing." *American Economic Review* 107(7), 2072-80.

Thereze. Working Paper. "Screening Costly Information"

Zhang. 2018. "Auctions vs. Negotiations: Optimal Selling Mechanism with Endogenous Bidder Values." Working paper, Dept. Econ., Univ. Bonn.

Appendix A. Proofs

Proof of Proposition 1.

The first part follows from strict convexity of $p(x)$:

$$
\begin{aligned}
p'(x) &= \mu(x) \text{ follows from envelope theorem from [IC-P]} \\
p''(x) &= \mu'(x) = (x^{-1})'(\mu) = \frac{1}{x'(\mu)} = \frac{1}{-H''(\mu)}
\end{aligned}
$$

Since $-H''(\mu) > 0$, $p''(x) > 0$. Then the seller wants to spread out allocation as much, choosing $x = 0$ and $x = 1$.

The second part follows from Bayesian Persuasion: if $\hat{p}(\mu)$ is convex, the seller wants to spread out posteriors as much, whereas if $\hat{p}(\mu)$ is concave, the seller wants to stay at a prior.

$$
\begin{aligned}
(\hat{p})'(\mu) &= x'(\mu)\mu = -H''(\mu)\mu \\
(\hat{p})''(\mu) &= -H'''(\mu)\mu - H''(\mu)
\end{aligned}
$$

$\hat{p}$ is concave if and only if $RP(\mu) \geq 1$.

Proof of Proposition 2.

We define $p^{\dagger}(k) = \hat{p}(\underline{\mu}(k))$ where $\underline{\mu}(k)$ is such that $x(\mu_0) = k$. Then from

$$
\begin{aligned}
p(\mu_0) &= \int_{\underline{\mu}}^{\mu_0} -H''(m) m dm \\
&= H(\mu_0) - \mu_0 H'(\mu_0) - H(\underline{\mu}) + \underline{\mu} H'(\underline{\mu}) \\
x(\mu_0) &= \int_{\underline{\mu}}^{\mu_0} -H''(m) dm \\
&= -H'(\mu_0) + H'(\underline{\mu}) = k,
\end{aligned}
$$

we can write

$$
\begin{aligned}
p^{\dagger}(k) &= H(\mu_0) - \mu_0 H'(\mu_0) - H(\underline{\mu}(k)) + \underline{\mu}(k)(k + H'(\mu_0)) \\
(p^{\dagger})'(k) &= (-H'(\underline{\mu}(k)) + k + H'(\mu_0)) \frac{d\underline{\mu}}{dk} + \underline{\mu} \\
(p^{\dagger})''(k) &= (-H''(\underline{\mu}) \frac{d\underline{\mu}}{dk} + 1) \frac{d\underline{\mu}}{dk} + (-H'(\underline{\mu}) + k + H'(\mu_0)) (\frac{-1}{H'''(\underline{\mu})} \frac{d\underline{\mu}}{dk} + \frac{d\underline{\mu}}{dk}
\end{aligned}
$$

Since $\frac{d\underline{\mu}}{dk} = \frac{1}{H''(\underline{\mu})}$ and $-H'(\underline{\mu}) + k + H'(\mu_0) = 0$, $(p^{\dagger})''(k) = \frac{d\underline{\mu}}{dk} < 0$. With $p^{\dagger}(0) = 0$, $p^{\dagger}$ is concave. It follows that

$$
0.5\{p(x) + p(1-x)\} \leq p(0.5), \forall x \in [0,1]
$$

For example, assuming $-H(\mu)'' = \mu^a$ ($a < -1$ for $\hat{p}(\mu)$ to be concave),

$$
\begin{aligned}
\hat{p}(\mu_0) &= \frac{1}{a+2}(\mu_0^{a+2} - \underline{\mu}^{a+2}) \\
\underline{\mu}(k) &= (\mu_0^{a+1} - (a+1)k)^{\frac{1}{a+1}} \\
p^\dagger(k) &= \frac{1}{a+2}(\mu_0^{a+2} - (\mu_0^{a+1} - (a+1)k)^{\frac{a+2}{a+1}} \\
(p^\dagger)'(k) &= (\mu_0^{a+1} - (a+1)k)^{\frac{1}{a+1}} \\
(p^\dagger)''(k) &= -(\mu_0^{a+1} - (a+1)k)^{-\frac{a}{a+1}} < 0
\end{aligned}
$$

Proof of Theorem 1.

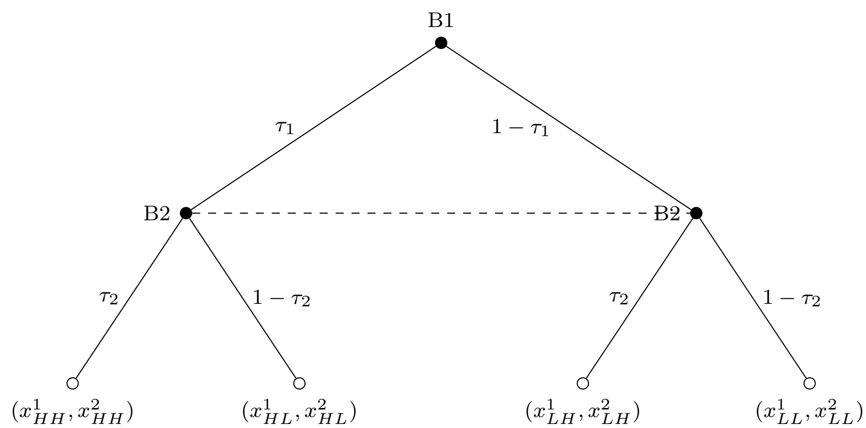

Allocation probability when simultaneously meeting buyers

We show that any (x^1_{HH}, x^2_{HH}) not $(1, 0)$ cannot be optimal. Suppose $x^1_{HH} \geq x^2_{HH}$. Consider increasing x^1_{HH} by ϵ and decreasing x^2_{HH} by ϵ. Since $x(\mu)$ is concave, $d\mu$ is greater for B1. The gain from giving B1 more allocation from B2 in HH state is greater than the loss of revenue from B2. The seller will continue to induce higher posterior by giving higher allocation to B1 until $x^1_{HH} = 1$.

Proof of Proposition 3.

First, to show the value of no trade signal is higher for B1 - comparing the seller's problem to one buyer case, with the second buyer,

$$\max \tau \hat{p}(\mu_1) + (1-\tau)V_2$$

the second term is added. Since the second term, plugging in the Bayes rule constraint, is increasing in $\underline{\mu}$, $\underline{\mu}$ should be higher with this additional term.

If we solve for $\underline{\mu}$, $\underline{\mu} = (\mu_0 - 1/4)/(2 - V_2)$. From the example in 3.1, it is the no trade signal for B2 multiplied by $2/(2-V_2)$ which is always greater than 1.

Next is to show price is higher for B1. We show $\frac{\partial \hat{p}(\bar{\mu})}{\partial \underline{\mu}} > 0$. Note that $\hat{p}(\bar{\mu})$ is integral which includes $\underline{\mu}$. Letting $P(m) = \int \hat{p}(m)dm = \int x'(m)mdm$,

$$
\begin{aligned}
\frac{\partial \hat{p}(\bar{\mu})}{\partial \underline{\mu}} &= P'(\bar{\mu})\frac{d\bar{\mu}}{d\underline{\mu}} - P'(\underline{\mu}) \\
&= -H''(\underline{\mu})(\bar{\mu} - \underline{\mu}) > 0
\end{aligned}
$$

Since from resource constraint

$$1 = \int_{\underline{\mu}}^{\bar{\mu}} x'(\mu) dm$$

we have

$$\frac{d\bar{\mu}}{d\underline{\mu}} = \frac{H''(\underline{\mu})}{H''(\bar{\mu})}$$

Proof of Theorem 2.

For concave $\hat{p}(.) = -\mu H''(\mu)$, and τ as the probability of high signal realizing, $\hat{V}_1(\mu) = \hat{p}_1(\mu) + (1 - x_1(\mu))V_2(\mu^M)$ is convex if and only if

$$H'''(\mu)(V_2 - \mu) + H''(\mu) \geq 0$$

where $V_2 = \hat{p}(\mu_0) - \hat{p}(\underline{\mu})$, $\underline{\mu} = G(1 + H'(\mu_0))$ with $G(x) = (H')^{-1}(x)$.

Therefore, if

1. $V_2 > \mu$ and $H'''(\mu) \geq -H''(\mu)/(V_2 - \mu)$
2. $V_2 \leq \mu$ and $-H'''(\mu) \geq H''(\mu)/(V_2 - \mu)$

The seller induces information acquisition to B1 and no information acquisition to B2.

Proof of Corollary 1.

$-H'''(\mu) < 0$ implies $x(\mu)$ is concave. Then with higher $\underline{\mu}$, $\mu \in \{0 \leq x(\mu) \leq 1\}$ lies in less concave region, implying the interval length will also be greater than when $\underline{\mu}$ is lower.

On the contrary when $-H'''(\mu) > 0$ so that $x(\mu)$ is convex, higher $\underline{\mu}$ implies the interval length of the set $\mu \in \{0 \leq x(\mu) \leq 1\}$ will be less than when $\underline{\mu}$ is lower.

Proof of Theorem 3.

1. For the first part, we show a simultaneous mechanism can yield higher revenue than the optimal sequential mechanism. Consider the optimal sequential mechanism when $\hat{p}(\mu)$ is concave. Then when the seller approached B2, the seller will not induce information acquisition as in one buyer case. Suppose the trade with B1 was made with τ_1 probability (whether the seller chose to induce information or not to B1 depends on the condition in Theorem 2). If we compare at the revenues from B2 in each mechanism:
   - Sequential: $(1 - \tau_1)\hat{p}^*(\mu_0)$ where $x^*(\mu_0) = 1$.
   - Simultaneous: $\hat{p}^2(\mu_0)$ where $x^2(\mu_0) = 1 - \tau_1$.

   From Proposition 2, $p^\dagger$ is concave. Then $kp^\dagger(1) < p^\dagger(k)$ for $k \in (0, 1)$.
2. From Theorem 1, the optimal simultaneous mechanism will implement posteriors $(\underline{\mu}_1, \bar{\mu}_1)$ and $(\underline{\mu}_2, \hat{\mu}_2)$. We show a sequential mechanism can yield higher revenue than the optimal simultaneous mechanism. If we look at the revenues from each mechanism:
   - Simultaneous: $\hat{\tau}_1 \hat{p}^1(\bar{\mu}_1) + \hat{\tau}_2 \hat{p}^2(\hat{\mu}_2)$ where $x^2(\hat{\mu}_2) = 1 - \hat{\tau}_1$
   - Sequential: $\tau_1 \hat{p}^1(\bar{\mu}_1) + (1 - \tau_1)V_2$

First we write the revenue as $V(k) := T(\underline{\mu}(k), k)P(\underline{\mu}(k), k)$ where given k as the resource constraint on B2 so that $\underline{\mu}(k)$ and $\mu(k)$ are defined from $x_2(\mu(k)) = \int_{\underline{\mu}(k)}^{\mu(k)} -H''(m)dm = k$,

$$
\begin{aligned}
\underline{\mu}(k) &= argmax T(\underline{\mu}, k)P(\underline{\mu}, k) \\
T(\underline{\mu}(k), k) &= \frac{\mu_0 - \underline{\mu}(k)}{\mu(k) - \underline{\mu}(k)} \\
P(\underline{\mu}(k), k) &= \int_{\underline{\mu}(k)}^{\mu(k)} -H''(m)mdm
\end{aligned}
$$

The revenue from B2 in simultaneous mechanism is $V(1-\hat{\tau}_1) = T(\underline{\mu}(1-\hat{\tau}_1), 1-\hat{\tau}_1)P(\underline{\mu}(1-\hat{\tau}_1), 1-\hat{\tau}_1)$ and we show that $(1-\hat{\tau}_1)V(1) = (1-\hat{\tau}_1)T(\underline{\mu}(1), 1)P(\underline{\mu}(1), 1)$ is greater which is implementable in sequential mechanism, by showing $V''(k) > 0$ for $0 < k < 1$.

By envelope theorem, we can ignore indirect effect $\frac{\partial \underline{\mu}}{\partial k}$ when observing $V'(k)$. Simplifying $T(\underline{\mu}(k), k)$ and $P(\underline{\mu}(k), k)$ just as a function of k, and the notation $\underline{\mu}(k), \mu(k)$ to $\underline{\mu}, \mu$,

$$
\begin{aligned}
T'(k) &= -\frac{\mu_0 - \underline{\mu}}{(\mu - \underline{\mu})^2}\frac{d\mu}{dk} = \frac{T(k)}{H''(\mu)(\mu - \underline{\mu})} < 0 \\
T''(k) &= \frac{2T'(k)}{H''(\mu)(\mu - \underline{\mu})} + \frac{T'(k)H'''(\mu)}{(H''(\mu))^2)} \\
P'(t) &= \mu(t) \\
P''(t) &= \mu'(t) = -\frac{1}{H''(\mu)}
\end{aligned}
$$

Then we have

$$
\begin{aligned}
V''(k) &= T''(k)P(k) + 2T'(k)P'(k) + T(k)P''(k) \\
&= \frac{T'(k)}{H''(\mu)}\left(\frac{2}{\mu - \underline{\mu}} + \frac{H'''(\mu)}{H''(\mu)}\right) + 2T'(k)\mu + \frac{T(k)}{-H''(\mu)} \\
&> T'(k)\left(\frac{2}{H''(\mu)(\mu - \underline{\mu})} + \frac{H'''(\mu)}{(H''(\mu))^2} + 2\mu\right)
\end{aligned}
$$

Since $T'(k) < 0$, $V''(k) > 0$ if

$$
\frac{H'''(\mu)}{(H''(\mu))^2} + 2\mu < \frac{2}{-H''(\mu)(\mu - \underline{\mu})}
$$

Since $RP(\mu) \leq 1$ and $\mu > \underline{\mu}$,

$$
\begin{aligned}
\frac{H'''(\mu)\mu}{-H''(\mu)} + 2\mu^2(-H''(\mu)) \leq 1 + 2\mu^2(-H''(\mu)) < 2 < \frac{2\mu}{\mu - \underline{\mu}} \\
-H''(\mu) < \frac{1}{2\mu^2}
\end{aligned}
$$

Since $V(0) = 0$, along with $V''(k) > 0$, $V(1-\hat{\tau}_1) > (t - \hat{\tau}_1)V(1)$ for $0 < \hat{\tau}_1 < 1$.