Screening Inattentive Buyers*

Jeffrey Mensch†

December 16, 2021

Abstract

Information plays a crucial role in mechanism design problems. A potential complication is that buyers may be inattentive, and so their information may endogenously and flexibly depend on the offered mechanism. I show that it is without loss of generality to consider *contour mechanisms*, which comprise triplets of allocation probabilities, prices, and beliefs, and are uniquely determined by a single such point. The mechanism design problem then reduces to Bayesian persuasion along the optimal contour. This reduction has significant implications for both the implementation of the optimal mechanism and the revenues that can be achieved.

*This paper was previously circulated under the title "Screening Inattentive Agents." I would like to thank Alex Gershkov, Konrad Mierendorff, Doron Ravid, Jakub Steiner, and Balazs Szentes, as well as various seminar participants at Tel Aviv University and the Technion for helpful conversations and comments. Four anonymous referees and the editor made very useful critiques and suggestions that greatly improved the paper. This project was generously supported by Israel Science Foundation grant 798/18. Tal Shewlker provided excellent research assistance. Michael Borns provided scientific editing services.

†Hebrew University; jeffrey.mensch@mail.huji.ac.il

1 Introduction

The question of optimal sales mechanisms has been widely studied. Crucially, the results of such analyses often depend on the assumption that the values that the buyers assign to the item for sale are exogenous. Yet in many cases, the factors that go into the evaluation of the item are manifold and complex. Starting with Sims (2003), a growing literature has focused on the impact of *inattention* regarding such factor. The potential buyer may therefore not immediately know how much the item for sale is worth to her: she may be inattentive regarding the true value. Instead, she must undertake some costly investigation to determine it. This inattentiveness must be accounted for by the seller when he sets the trading mechanism that buyers face, as the mechanism will affect both the information that the buyers acquire, and the eventual purchasing decisions.

The fact that there are many features that determine the value of an item for the buyer implies that uncertainty over these features enables the buyer to select over a wide range of information acquisition strategies. In particular, the buyer can decide not only *how much* information to acquire, but also *what kind* of information. Thus, the potential strategies will not automatically be completely ordered by the informativeness of the signals associated with them, but will instead be chosen *flexibly* in response to the mechanism.

I model this environment as follows. First, the seller determines the (symmetric) mechanism that the buyers face, with the aim of maximizing revenue. This ensures that the rules of the mechanism are known to the buyers prior their participation therein. One can think of this as applying to retail pricing or to auction houses, whose operating procedures are generally known to potential buyers. Upon seeing the mechanism, the buyers then choose a signal structure, which will determine the joint distribution of signal realizations and possible ex-post valuations. Unlike previous work in mechanism design,[1] I assume that information can be acquired flexibly, i.e. that any information is possible to acquire, as long as it satisfies Bayes' rule, at a cost that is increas-

[1]See Section 7 for a discussion of the related literature.

ing in the Blackwell order. After observing the signal realization, the buyers then decide which option in the menu to choose, if any.

The ability of the buyers to flexibly acquire information has several significant implications. First, it severely limits what the seller can effectively implement. In standard mechanism design, the revelation principle (Myerson, 1981) ensures that it is without loss to look at direct mechanisms, in which the buyer tells the seller her value for the item, and the seller gives a particular outcome. As a result, any monotone mechanism, in which the probability of sale is increasing in the buyer's value, would be implementable for appropriate prices. While the revelation principle is technically still valid here ex interim (as any mechanism can be expressed via direct revelation), this restriction is not sufficient here: not only the distribution, but even the support of the interim values is endogenous when the buyers can acquire information flexibly. Thus even if the mechanism were incentive compatible ex interim, it would not necessarily be so ex ante: the buyer could deviate to a different information acquisition strategy, which would eliminate some of these (interim) values from the support. This makes it insufficient to merely consider reported interim values when evaluating the incentive compatibility of mechanisms due to the additional ex-ante constraints.

As a consequence, I show that the set of outcomes that can be implemented by mechanisms is much narrower. In particular, if one knows that a particular triplet of (a) posterior beliefs, (b) probability of receiving the item, and (c) payments is in the support of the choice of the buyer, then the mechanism is effectively fully determined. That is, one can pin down exactly what payments and beliefs must occur at any other probability of receiving the item. The only remaining leeway for the seller is to determine the optimal distribution over such triplets, while respecting Bayes' rule. In effect, the seller's problem then reduces to one of Bayesian persuasion (Aumann and Maschler, 1995; Kamenica and Gentzkow, 2011).

When choosing this optimal distribution, the seller faces three major tradeoffs. By having the buyer acquire information that leads to higher beliefs, the seller can extract more rents, as the buyer will conditionally value the item more. On the other hand, by Bayes' rule, high interim expected values occur with low probabilities, and so even if such a posterior were to lead to higher rents, its low probability of realization may not justify it. Lastly, the seller would like to extract rents from a given belief. However, the more rent is extracted, the weaker the power of the incentives to induce information acquisition. As the buyer cannot acquire the object with probability greater than 1, this prevents sufficient incentives for the buyer to acquire information that would lead to posteriors with high interim values.

Since the tradeoffs with flexible information acquisition are different, the form of the optimal mechanism may be different as well. In standard mechanism design (Myerson, 1981), as with inflexible information acquisition (e.g. Shi, 2012), the optimal mechanism takes the form of a posted price with one buyer, and a second-price auction with a reserve price with multiple buyers. For many cost functions for flexible information acquisition, however, this will no longer be the case. Nevertheless, I provide sufficient conditions to guarantee that the optimal mechanisms take these respective forms.

Yet even when the results go through, the optimal reserve prices and revenues differ from previous results. For instance, as the number of buyers goes to infinity (Theorem 9), a second-price auction is indeed optimal, but with a reserve price of 0 being strictly better than any other one. Moreover, while normally an infinite number of buyers implies that revenue converges to the top of the distribution of ex-post values, the incentive to acquire information vanishes as well, and so the distribution of interim expected values converges to the point mass at the prior. This leads to a precise balance between the two effects, and so the interim expected value of the winner of the auction can be precisely calculated at a value between the prior expected value and the highest possible ex-interim value.

Lastly, the endogeneity of the interim distribution of the buyer's values affects the ability of the seller to extract revenue (Remarks 2 and 3). While with inflexible information acquisition, he may want to change the mechanism for a given interim distribution in order to induce a particular information acquisition strategy, the result here is even stronger: even for a given mechanism (that is, probability of sale for each interim value in the distribution), the revenue is lower. This is because the seller must not just defend against choosing a different *quantity* of information, but against different *interim values* as well, even if they occur with probability 0. This involves providing enough surplus at a given interim value to discourage such deviations. By contrast, without flexible information acquisition, there is no difference in revenue extraction whether or not the information of the buyer is endogenous.

Beyond the differences in the characterizations of the optimal mechanisms, the paper also contributes to the growing literature on rational inattention, which use models of flexible information acquisition to model inattentive behavior (Caplin and Dean, 2013, 2015; Matejka and McKay, 2015). To that end, the techniques developed here may be useful in other environments with rationally inattentive agents. This is especially so for other principal-agent problems besides optimal auctions.

The remainder of the paper is organized as follows. Section 2 provides a motivating example of the complications arising from information acquisition in determining what mechanisms are optimal, or can even be implemented. Section 3 presents the model of mechanism design and information acquisition, as well as some preliminary results for a sufficient class of mechanisms. Section 4 then presents necessary and sufficient conditions for a mechanism to be implementable. Section 5 discusses optimal single-buyer mechanism design by means of Bayesian persuasion. Section 6 uses the techniques of the previous section for applications, including the discussion of the posted-price mechanism. Section 7 then explores optimal mechanism design with multiple buyers. Section 8 discusses connections to related literature. Section 9 concludes.

2 A Motivating Example

Consider the case of a single buyer (she) who values the item either at 0 or at 6. Neither the buyer nor the seller (he) starts with any further information about the value, though they share a common prior that each value is equally likely. However, the buyer can acquire further information about her value at a cost proportional to the expected entropy reduction, where the formula for entropy in this case, as a function of belief μ, is

$$H(\mu) \equiv \mu(\{0\}) \ln(\mu(\{0\})) + \mu(\{6\}) \ln(\mu(\{6\}))$$

Thus, given a menu of options, the buyer will acquire information about her value, and then choose the option that maximizes her conditional expected utility based on her signal.

Normally, when solving for optimal sales mechanisms, it is without loss of generality to restrict to *direct revelation mechanisms* (Myerson, 1981). That is, the buyer reports her type to the seller, who offers a probability of sale, as well as a transfer (price), as a function thereof. This approach is complicated here, though, by the fact that the buyer's belief about her value is *endogenous*: the distribution of values, as estimated by the buyer's posterior beliefs, will depend on which mechanism is offered.

Suppose, for instance, that the seller believes that the buyer will acquire no information. Thus, he believes that she will believe that her expected value is 3 with probability 1. If one ignores information acquisition, the optimal mechanism is to offer to sell the item with probability 1 at a price of 3. However, if he were to do so, the buyer would have an incentive to change their information acquisition strategy. Consider, for instance, the alternative information acquisition in which, at one posterior, the value is believed to be 6 with probability $0.5 + \epsilon$, and at the other, is believed to be 6 with probability $0.5 - \epsilon$. Thus the buyer will purchase given the first signal realization, and refrain from doing so given the latter. The gain in expected utility from this change in strategy is $\frac{1}{2}\epsilon$, while the loss from the increase in information costs is

$$\ln(0.5) - 2[\frac{1}{2}(0.5 + \epsilon) \ln(0.5 + \epsilon) + \frac{1}{2}(0.5 - \epsilon) \ln(0.5 - \epsilon)]$$

Dividing by ϵ and taking the limit as $\epsilon \to 0$ yields a marginal gain of $\frac{1}{2}$, and a marginal loss of 0. Hence for small enough ϵ, this alternative strategy is an improvement, and so the seller could not implement such an outcome with a mechanism.

Suppose instead that the seller believes that the buyer will acquire a signal whose realizations yield expected values of 5 and 1, each with probability $\frac{1}{2}$. If this signal were exogenous, it would be optimal to offer to sell the item with probability 1 at a price of 5. Yet the buyer, when faced with such a mechanism, would instead be better off choosing to acquire no information, thereby saving on the information acquisition costs, while losing nothing from the surplus from purchasing, which was 0 anyway.

In order to implement these two posterior expected values, the buyer must not be able to improve their payoff via an alternative signal. Suppose that the seller will not sell with any probability conditional on the posterior expected value of 1, and that he sells with probability x and charges t conditional on the posterior expected value of 5. Let the posterior probabilities of having a value of 6 given expected values of 5 and 1 be μ_H and μ_L, respectively. This means that by Bayes' rule, $\frac{5}{6}$ and $\frac{1}{6}$ must be the respective values that solve the problem

$$\max_{\mu_H, \mu_L \in [0,1]} [6x\mu_H - t](\frac{0.5 - \mu_L}{\mu_H - \mu_L}) \tag{1}$$

This yields the requisite values if and only if $x = \frac{1}{6}\ln(25) \approx 0.536$ and $t = \ln(5) \approx 1.609$. Thus there is only one possible way to implement these posteriors, given the offer of $x = 0$ to μ_L. Indeed, as will be shown later in the paper, this holds more generally: given a description of what happens at a single posterior (here, μ_L), there will be (at most) a unique offer that the seller can make to induce any other posterior (if at all).

What would happen if, say, the seller were to increase x to 0.6, while continuing to try to implement the posterior $\mu_L = \frac{1}{6}$ as before? In order to incentivize this, the seller would naturally need to increase t: receiving the item with higher probability should cost more. In order to solve (1), one then needs $\mu_H \approx 0.880$ and $t \approx 1.936$.

One can already see the three major tradeoffs that were mentioned in the introduction. By offering the item with higher probability, the seller changes the incentive to acquire information, such that the buyer will now choose a posterior belief with a higher expected value of about 5.28. This allows the him to extract more rents for two reasons: naturally, if the buyer receives the item with higher probability, she will be willing to pay more. In addition, the willingness to pay for each point of additional probability of receiving the item goes up, due to the higher posterior expected value: the average increase in payoff from this increase in x is $\frac{\Delta t}{\Delta x} \approx 5.15$, while at the original μ_H, $t/x = 3$. At the same time, the probability of *receiving* a high posterior goes down: while originally the chance of receiving μ_H was 0.5, it is now approximately 0.467. Thus the probability that the seller receives a high payoff *decreases*.

To fully derive the optimal mechanism in this environment, one needs to understand exactly what the seller can implement, and then what techniques can be used to optimize over the set of such mechanisms. I do so in the next three sections.

3 Preliminaries

Throughout, I restrict attention to symmetric mechanisms and strategies across all N buyers, and so generally do not use subscripts to distinguish buyers.[2] I therefore analyze both the single and multiple buyer cases together in Sections 3 and 4, with the necessary adaptations for the full analysis of the multiple buyers case introduced in Section 7.

Each buyer (B) has ex-post type (value) $\theta \in \Theta$, where $\Theta \subset \mathbb{R}$ is of finite size K. Given a probability of receiving the item x and transfers t, the seller's ex-post payoff is $u_S(x, t) = t$, while the buyer's is $u_B(x, t, \theta) = x\theta - t$, i.e. the buyer is risk-neutral.

Before presenting the formal model, I describe the timing of the model. First, the seller offers a mechanism $\mathcal{M}$. Next, given the mechanism, the buyer will acquire information about their type. The signal is then realized, giving

[2] The sole exception is in the proof of Lemma 1, a preliminary result on the class of mechanisms that need be considered.

the buyer a posterior belief. Once all information has been acquired, the buyer chooses a strategy to play in the mechanism, which in equilibrium will lead to the buyer receiving the item with some probability x and paying some transfer t given her posterior.

Given a mechanism, the buyer responds by first acquiring information about the state. The buyer and seller share a common prior μ_0. The buyer can flexibly choose her information via some signal structure $(\pi, \mathcal{S})$, where s is a signal realization, and $\pi(s|\theta)$ is the probability that s is realized. Beliefs are then updated via Bayes' rule. This yields a distribution of posteriors $\tau \in \Delta(\Delta(\Theta))$ such that $\int \mu d\tau(\mu) = \mu_0$. Note that flexibility implies that τ is feasible as long as this equality is satisfied. The cost of information acquisition is given by the expected difference in a posterior-separable cost function $c(\tau)$, defined by the measure of uncertainty[3] $H(\mu)$. Thus, given distribution τ,

$$c(\tau) = H(\mu_0) - E_\tau[H(\mu)]$$

where H is strongly concave[4] and twice Lipschitz continuously differentiable on any compact set of interior beliefs μ, i.e. $\mu(\theta) > 0, \forall\theta$. Assume that the slope of H as μ approaches the boundary (i.e. $\mu(\theta) = 0$ for some θ) is sufficiently high so as to make it strictly optimal to choose beliefs bounded away from the boundary.[5] This allows for commonly used functions for H, such as informational entropy, i.e. $H(\mu) = -\sum_{\theta\in\Theta} \mu(\theta) \ln \mu(\theta)$.

To understand this cost function, note that by Bayes' rule, the beliefs that the buyer has upon acquiring information form a mean-preserving spread from the prior belief μ_0. Furthermore, acquiring more information in the Blackwell order will also form a mean-preserving spread. Thus c is well defined for any

[3]For further discussion of the assumption of the informational cost form, see Appendix A.

[4]This is a slightly stronger assumption than strict concavity, as it essentially requires that the second derivative be bounded away from 0 from below. More formally, there exists $m > 0$ such that, letting $\mathbf{H}(\mu)$ be the Hessian matrix of $H(\mu)$, $\mathbf{H} + mI$ is negative semi-definite.

[5]While this is endogenous, it is possible to provide primitive sufficient conditions to ensure that this is the case. I discuss this further in Appendix A.

signal structure that satisfies Bayes' rule, and is increasing in the Blackwell order.

As the set of possible mechanisms is potentially very large, it is important to reduce it to a more tractable class. The natural way to do so is via a form of the *revelation principle* (Myerson, 1981), which states that it is without loss to consider mechanisms in which players report their type. However, the set of types (in the sense of posterior beliefs) is endogenous, complicating matters. At the same time, Kamenica and Gentzkow (2011) show that it is without loss to consider *straightforward* signals, which are recommendations for actions. I combine these in a similar manner to Ravid (2020), in the following definition.

Definition 1: A mechanism $\mathcal{M}$ uses *recommendation strategies* if it consists of the smallest compact set of allocation probabilities $X \subset [0,1]$ (along with their corresponding transfers t) and, for each state $\theta \in \Theta$, the buyer has a probability measure $\pi(\cdot|\theta) \in \Delta(X)$ such that $\int_X d\pi(x|\theta) = 1$.

In words, the seller recommends that the buyer acquire information according to π, which tells them which choice of x to make. Thus the buyer does not randomize given her signal realization: there is precisely one choice of x for each realization. Conversely, by Bayes' rule, this implies that there is exactly one signal realization, and hence one posterior belief μ, for each x offered.

Lemma 1: *For every feasible mechanism $\mathcal{M}$, there is an equivalent mechanism $\mathcal{M}'$ that uses recommendation strategies in which the seller and the buyers receive the same payoffs as in $\mathcal{M}$.*

All proofs are in Appendix B.

Since it is without loss to consider recommendation strategies, one can write each of these three variables (beliefs, allocation probabilities, and transfers) as functions of each other. Notationally, I use $\mathbf{x}(\mu)$ to express the allocation as a function of μ, including what would be chosen if[6] $\mu \notin \operatorname{supp}(\tau)$, $\mathbf{t}(x)$ is the transfer as a function of the allocation, $\mu(\cdot|x)$ is the belief given x, and

[6] This is because the seller must consider what the buyer would do if she were to acquire information differently.

$\tilde{\mathbf{t}}(\mu)$ is the transfer as a function of beliefs.

I now formally describe the problem with its individual rationality and incentive compatibility constraints, given that the interim problem (post-information acquisition) for the buyer reduces to finding the optimal (x, t) for each $\mu \in \text{supp}(\tau)$. The seller's objective is thus

$$\max_{\mathcal{M},\tau} \int \mathbf{t}(\mathbf{x}(\mu))d\tau(\mu)$$

$$\text{s.t.} \quad \int \mu d\tau(\mu) = \mu_0$$

$$\tau \in \arg\max_{\sigma\in\Delta(\Delta(\Theta))} \int\int [\mathbf{x}(\mu)\theta - \mathbf{t}(\mathbf{x}(\mu))]d\mu(\theta)d\sigma(\mu) - [H(\mu_0) - \int H(\mu)d\sigma(\mu)] \qquad (IC-A)$$

$$\mathbf{x}(\mu) \in \arg\max_{(x,\mathbf{t}(x))\in\mathcal{M}} \int [x\theta - \mathbf{t}(x)]d\mu(\theta), \forall\mu \in \Delta(\Theta) \qquad (IC-I)$$

$$\int\int [\mathbf{x}(\mu)\theta - \mathbf{t}(\mathbf{x}(\mu))]d\mu(\theta)d\tau(\mu) - [H(\mu_0) - \int H(\mu)d\tau(\mu)] \geq 0 \qquad (IR-A)$$

$$\int [\mathbf{x}(\mu)\theta - \mathbf{t}(\mathbf{x}(\mu))]d\mu(\theta) \geq 0, \forall\mu \in \text{supp}(\tau) \qquad (IR-I)$$

where the first constraint is that the distribution of posteriors is Bayes-plausible, the second is that the information acquisition strategy of the buyer is ex-ante optimal, the third is that the choice from the menu is interim optimal, and the fourth and fifth are ex-ante and interim individual rationality constraints, respectively.[7]

While the above constraints are all necessary for the formal description of the seller's problem, the use of recommendation strategies allows for the simplification of the problem. In particular, the (IC-I) and (IR-A) constraints are redundant.

Lemma 2: *The (IC-I) and (IR-A) constraints are implied by (IC-A) and (IR-I).*

[7]With multiple buyers, there will also be a feasibility constraint that ensures that the items can be physically allocated with probability between 0 and 1 for any vector of realization of signals. See Section 7.

We are now ready to move on to the discussion of which recommendation strategies are implementable.

4 Implementability

In this section, I characterize the set of mechanisms that use recommendation strategies, and illustrate how they can be used for a general theory of mechanism design. I highlight how the ability of the buyer to acquire information restricts the set of such mechanisms that the seller can implement. The critical issue, as seen in the motivating example, is that the buyer can deviate not only at the interim stage, but also at the ex-ante stage. The seller must therefore defend against a larger set of possible deviations. Using this observation, I then generalize to a sufficient, tractable class of mechanisms.

Assume for now that the individual rationality constraint is satisfied; I will subsequently return to it. In order for incentive compatibility to be satisfied, the buyer must be optimizing, both through information acquisition and choice of x. As it is without loss to consider mechanisms that use recommendation strategies, small perturbations of the information acquisition strategy cannot be payoff-improving. Suppose that belief $\mu(\cdot|x)$ occurs with probability $\tau(\mu(\cdot|x))$, and so state θ has total probability mass $\tau(\mu(\cdot|x))\mu(\theta|x)$ from this signal realization. Consider the perturbation that increase the probability that x is recommended by ϵ, by means of increasing $\pi(x|\theta)$ by $\frac{\epsilon}{\mu_0(\theta)}$. By Bayes' rule, this perturbation increases the conditional probability of θ given x to $\frac{\mu(\theta|x)+\epsilon}{1+\epsilon}$, while it correspondingly decreases the conditional probability of other states θ' to $\frac{\mu(\theta'|x)}{1+\epsilon}$, where $\mu(\cdot|x)$ are the respective conditional probabilities before the perturbation. By the envelope theorem, the marginal change in the consumption utility to the buyer from the choice of x as $\epsilon \to 0$ is the payoff from x at state θ, namely, $x\theta - \mathbf{t}(x)$. This must be balanced against the change in costs of information acquisition.[8] Expressing $\frac{\partial H}{\partial \mu(\theta)}$ as the partial derivative of H with respect to the probability of state θ, one can write the change in cost

[8]It must also be balanced against maintaining Bayes' rule, as this mass on state θ must be distributed across signals so as to add to $\mu_0(\theta)$. This will be addressed shortly.

from marginally changing the mass of θ at this signal as

$$h(x,\theta) \equiv H(\mu(\cdot|x)) + \frac{\partial H}{\partial \mu(\theta)}(\mu(\cdot|x))(1-\mu(\theta|x)) - \sum_{\theta' \neq \theta} \frac{\partial H}{\partial \mu(\theta')}(\mu(\cdot|x))\mu(\theta'|x) \tag{2}$$

To understand (2), the perturbation of the information acquisition strategy has three effects on the cost from recommending x, as described in the previous paragraph. First, it makes x more likely, with the associated information cost term $H(\mu(\cdot|x))$. Second, conditional on x being recommended, it is now more likely that θ is the actual state, yielding the second term in (2). Lastly, the other states $\theta' \neq \theta$ are now less likely, yielding the third term.

Of course, such a perturbation must be consistent with Bayes' rule: one cannot arbitrarily increase $\pi(x|\theta)$ without correspondingly decreasing $\pi(x'|\theta)$ for some other x'. Thus for a signal to be optimal, the buyer cannot benefit from such a marginal change in conditional probability from x to x'. It turns out that this requirement is also sufficient for optimality: due to the convexity of the cost function due to the concavity of H, one can generate a local improvement from any global improvement.[9]

Lemma 3: *π is optimal for the buyer if and only if*

$$x\theta - \mathbf{t}(x) + h(x,\theta) = x'\theta - \mathbf{t}(x') + h(x',\theta) \tag{3}$$

for all $\theta \in \Theta$ and almost all $x, x' \in X$ with respect to π.

Lemma 3 was established only for $x, x' \in X$, i.e. the support of the recommendation strategies. While it was established in Lemma 1 that it was without loss to offer mechanisms that use recommendation strategies, this may not be the most tractable way to approach the problem, as it is not clear what strategies the seller should recommend. To enable a general analysis of mechanism design with inattentive buyers, it is preferable to establish a class of mechanisms for which this issue does not apply: one can analyze it without

[9]This result parallels Lemma 3 in Caplin and Dean (2013). It should be noted that they also have an inequality for unchosen acts. This is not relevant here, because, as the mechanism uses recommendation strategies, there are no unchosen acts.

knowing which strategies are recommended.

One can do so as follows. Given a mechanism $\mathcal{M}$, consider instead a mechanism $\mathcal{M}' = \{(x, \mathbf{t}(x))\}$ such that $x \in [0,1]$ that induces τ and $\mathbf{x}$ as under $\mathcal{M}$. As such, $\mathbf{t}(x)$ is the same as in $\mathcal{M}$ when $x \in X$. If $\mathcal{M}'$ is to induce the same τ as in $\mathcal{M}$, $\mu(\cdot|x)$ and $\mathbf{t}(x)$ must be defined in such a way that the buyer does not want to place positive weight on signals that lead to $x' \notin X$. One can do so by extending the results of Lemma 3 to the rest of $x' \in [0,1]$, and assigning beliefs conditional on all such x, even if they end up being chosen with probability 0 because they were not recommended.

Lemma 4: *Suppose that* $(x^*, \mathbf{t}(x^*), \mu(\cdot|x^*))$ *is in the support of the choice of the buyer. Then there exists a unique extension* $\mathcal{M}' = \{(x, \mathbf{t}(x))\}$ *and beliefs* $\mu(\cdot|x)$ *to each* $x \in [0,1]$ *that satisfies (3). Moreover, any* τ *such that* *supp*$(\tau) \subset \{\mu : \exists x \in [0,1] : \mu = \mu(\cdot|x)\}$ *is then optimal for the buyer given* $\mathcal{M}'$*.*

Informally, the above result implies that given an initial point $(x^*, \mathbf{t}(x^*), \mu(\cdot|x^*))$, one can uniquely extend the mechanism to a triple of feasible points $(x, \mathbf{t}(x), \mu(\cdot|x))$ for all $x \in [0,1]$. By Lemma 3, we have seen that any distribution τ over points that satisfy (3) is optimal for the buyer. Thus the buyer is *indifferent* between all distributions τ over such points that satisfy Bayes' rule, as one can construct a mechanism in recommendation strategies as in Lemma 3 by only including the points in $\mathcal{M}'$ with $\mu(\cdot|x)$ in the support of[10] τ. Hence (3), in the context of Lemma 4, defines an envelope condition that guarantees incentive compatibility. In the derivation in Appendix B, I present a differentiable law of motion defining $\mathbf{t}(x)$ and $\mu(\cdot|x)$.

One can therefore use the envelope condition to derive the entire remainder of the mechanism from a single point. I therefore present one such derivation, starting from the point $x^* = 0$, which is natural to consider.

Definition 2: The *non-participation belief* is given by $\underline{\mu} \equiv \mu(\cdot|x = 0)$, and the *certain allocation belief* is given by $\bar{\mu} \equiv \mu(\cdot|x = 1)$.

[10] This will have implications for optimal mechanism design, which I explore in Section 5.

Definition 3: A *contour mechanism* $\mathcal{C}$ is given by

$$\mathcal{C} \equiv \{(x, \mathbf{t}(x), \mu(\cdot|x))\}$$

with initial conditions $(0, \mathbf{t}(0), \underline{\mu})$, and $(x, \mathbf{t}(x), \mu(\cdot|x))$ subsequently determined by (3).

Thus, any mechanism $\mathcal{M}$ can be represented by a contour mechanism $\mathcal{C}$. The upshot of this is that the buyer is then incentivized to obey a recommendation strategy corresponding to $\mathcal{M}$, as long as it lies on the same contour. This means the seller can implement any distribution of his choosing along the contour, subject to Bayes' rule, as long as the recommendation satisfies (3) and hence is optimal for the buyer.

To complete the characterization of implementable mechanisms, one must discuss the individual rationality constraints. Here, the results do not diverge much from the standard mechanism design case: the (IR-I) constraint is satisfied whenever it is satisfied at $x = 0$ (i.e. at $\underline{\mu}$). By revealed preference, since choosing $x = 0$ is always an option, the choice of the buyer of something other than $x = 0$ implies that it is preferable ex interim, and hence must give non-negative utility. The results up to now are thus summarized in the following theorem.

Theorem 1 (Implementability): *A mechanism $\mathcal{M}$ is implementable in recommendation strategies if and only if the same strategies can be implemented by a contour mechanism $\mathcal{C}$ such that $\mathbf{t}(0) \leq 0$.*

I now describe and illustrate how the contour mechanisms look in practice. For some commonly used cost functions, such as entropy-based (Sims, 2003; Matejka and McKay, 2015) or log-likelihood ratio-based (Morris and Strack, 2017; Pomatto et al., 2019) cost functions, there is always an infinite marginal cost of acquiring information that sets $\mu(\theta) = 0$ for any θ, regardless of the multiplication of costs by any coefficient $\kappa > 0$. This reflects the idea that it is prohibitively costly to fully rule out any possible type. To compensate for this, the seller would have to provide an infinite payoff conditional on reporting such beliefs, which is obviously suboptimal for the seller. The following example provides a description of the functional form of contour mechanisms for entropy-based costs, which appeared in the motivating example.

Example 1: Suppose that, for some $\kappa > 0$,

$$H(\mu) = -\kappa \sum_{\theta \in \Theta} \mu(\theta) \ln \mu(\theta)$$

Then $h(x, \theta) = -\kappa \ln \mu(\theta|x)$, and so by (3), one gets the likelihood ratio condition[11]

$$\kappa \ln \frac{\mu(\theta'|x')/\mu(\theta|x')}{\mu(\theta'|x)/\mu(\theta|x)} = (x' - x)(\theta' - \theta)$$

for any pairs x, x' and θ, θ'. As $\sum_{\theta \in \Theta} \mu(\theta|x) = 1$, $\mu(\theta|x)$ is thus pinned down by $\underline{\mu}$. Solving for $\mu(\theta|x)$ then yields

$$\mu(\theta|x) = \frac{\underline{\mu}(\theta) e^{\frac{x\theta}{\kappa}}}{\sum_{\theta' \in \Theta} \underline{\mu}(\theta') e^{\frac{x\theta'}{\kappa}}} \tag{4}$$

Meanwhile, (3) also implies that

$$\mathbf{t}(x) - \mathbf{t}(0) = x\theta - \kappa(\ln \mu(\theta|x) - \ln \underline{\mu}(\theta)) \tag{5}$$

I illustrate this for $\Theta = \{\theta_1, \theta_2, \theta_3\} = \{5, 10, 15\}$, $\kappa = 1$, and $\underline{\mu} = (0.9, 0.09, 0.01)$. In Figure 1(a), I plot the set of beliefs that are along the contour mechanism, while in Figure 1(b), I plot $\mathbf{t}(x)$ against x. As x increases, θ_2 and θ_3 become more likely, with the marginal effect on the latter value becoming stronger as x rises. The slope of the transfers with respect to an increase in x rises as well: as higher states become more likely, the marginal transfers $\mathbf{t}'(x)$ must rise as well. Note that the transfers $\mathbf{t}(x)$ here depend only on the posterior beliefs $\mu(\cdot|x)$, and not on the prior μ_0. $\square$

As the contour mechanism $\mathcal{C}$ is pinned down following some μ, it uniquely pins down the transfers for $x > \mathbf{x}(\mu)$ as well. This presents an even stronger form of the celebrated revenue equivalence result of Myerson (1981). In the

[11] This was derived in Mensch (2021).

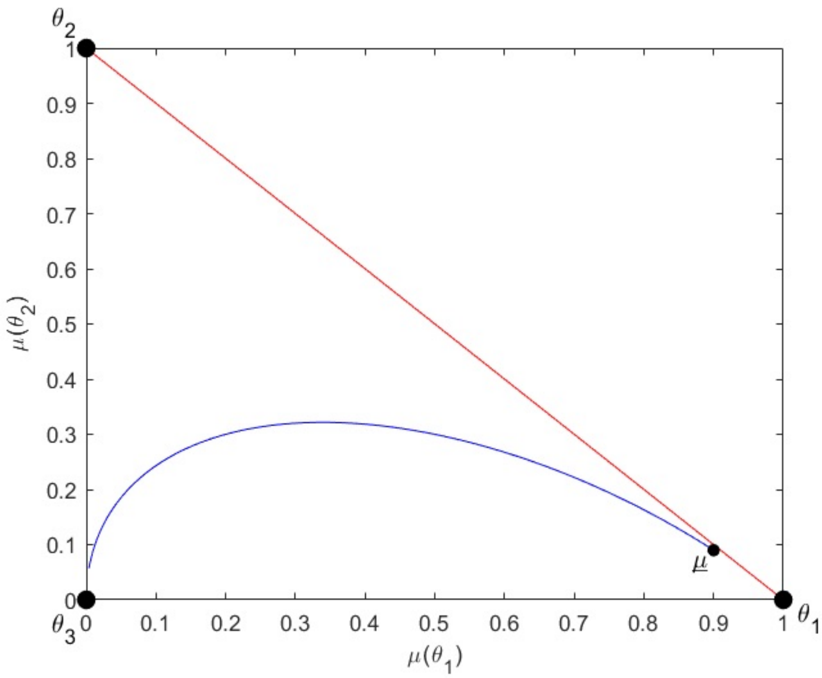

(a) Contour of beliefs

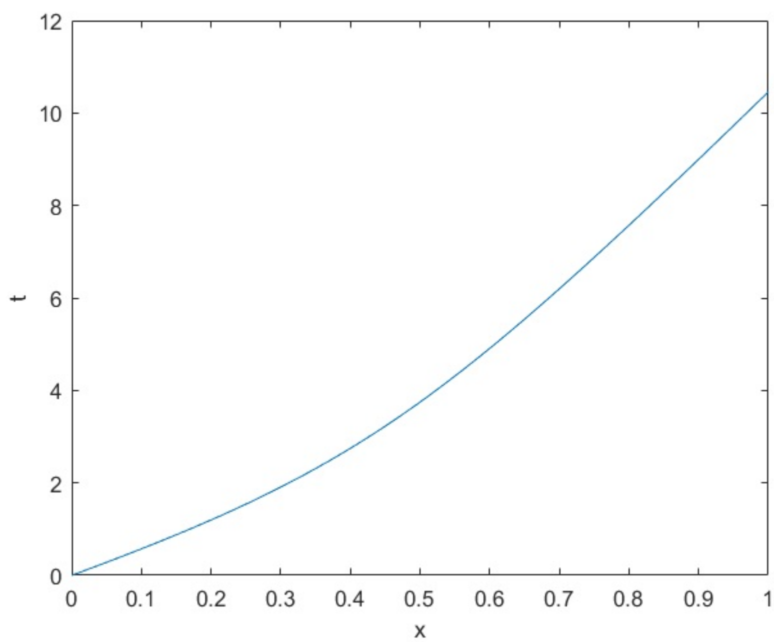

(b) Allocations and transfers

Figure 1: Entropy costs, three states

latter, for any exogenously given distribution of types, one must specify the mechanism's allocation, and based on which is specified, one can pin down the transfers (and hence the revenue) necessary to implement it. Here, however, *the mechanism itself* is also pinned down, not only by the distribution of types, but also by some initial value of μ along the contour. This means that if μ is along the contour for two different contour mechanisms, the transfers associated with $\mathbf{x}(\mu) + \delta$ must differ from $\mathbf{t}(x(\mu))$ by the same amount, and the beliefs at $\mathbf{x}(\mu) + \delta$ must agree (as long as $\mathbf{x}(\mu) + \delta \in [0, 1]$). In other words, if one starts at the same belief μ, one must proceed along the same contour of beliefs, with values of x shifted by by the same amount. I summarize this in Proposition 2.

Proposition 2: *Let* $(x, \mathbf{t}(x), \mu(\cdot|x)) \in \mathcal{C}$*, and let* $\hat{\mathcal{C}} \ni (\hat{x}, \hat{\mathbf{t}}(\hat{x}), \hat{\mu}(\cdot|\hat{x}))$ *be an implementable contour mechanism such that for some* $\delta^* > 0$*,* $x, \hat{x} \in [\delta^*, 1-\delta^*]$ *and* $\hat{\mu}(\cdot|\hat{x}) = \mu(\cdot|x)$*. Then*

$$\hat{\mathbf{t}}(\hat{x} + \delta) - \hat{\mathbf{t}}(\hat{x}) = \mathbf{t}(x + \delta) - \mathbf{t}(x)$$

$$\hat{\mu}(\cdot|\hat{x} + \delta) = \mu(\cdot|x + \delta)$$

for all $\delta \in (-\delta^*, \delta^*)$*.*

5 Optimal Mechanism Design: Bayesian Persuasion

The results of the previous section imply that in order to implement a mechanism, one must ensure that the points induced by it lie on some contour mechanism $\mathcal{C}$. However, it does not pin down the *distribution* of points along the contour. As shown in Theorem 1, any such distribution will be incentive-compatible. This leaves some leeway for the seller to choose which distribution to induce. In Figure 2, I show one such possible set of posteriors, in which the prior, denoted by μ_0, is in their convex hull. It is therefore possible to assign appropriate weights to each of the posteriors in order to make them sum to the prior, thereby satisfying Bayes' rule.

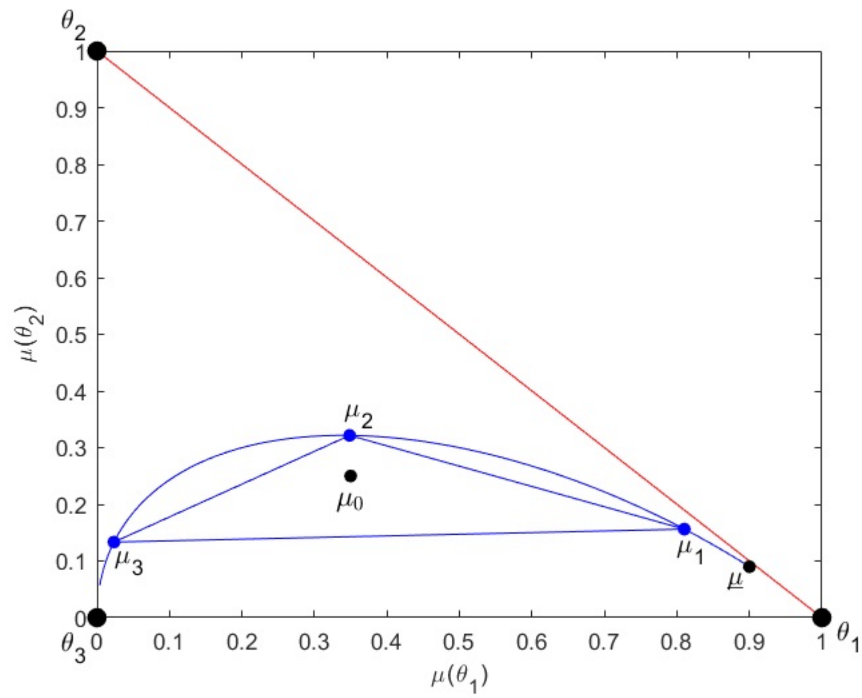

Figure 2: Possible distribution given $\mathcal{C}$

One can thus view the problem as one of optimal *Bayesian persuasion* (Kamenica and Gentzkow, 2011). Formally, for a fixed contour mechanism $\mathcal{C}$, define the *value* of a posterior to be

$$v_{\mathcal{C}}(\mu) \equiv \begin{cases} \mathbf{t}(\mathbf{x}(\mu)), & (\mathbf{x}(\mu), \mathbf{t}(\mathbf{x}(\mu)), \mu) \in \mathcal{C} \\ -\infty, & \text{otherwise} \end{cases} \tag{6}$$

The *concave closure* of $v_{\mathcal{C}}$ is then given by

$$V_{\mathcal{C}}(\mu) \equiv \sup\{z : (\mu, z) \in co(v_{\mathcal{C}})\} \tag{7}$$

where $co(v_{\mathcal{C}})$ is the convex hull of the graph of $v_{\mathcal{C}}$.

Theorem 3: *A mechanism is optimal if and only if it solves*

$$\max_{\mathcal{C}} V_{\mathcal{C}}(\mu_0) \tag{8}$$

such that $\mathbf{t}(0) = 0$.

The reduction of the problem to that of Bayesian persuasion aids in the finding of the solution, as one can then import results pertaining to optimal persuasion mechanisms. For instance, one need only restrict attention to mechanisms with at most K signal realizations.[12] As the resulting mechanism is finite-dimensional, this limits the search for the optimal mechanism to a much smaller set of possibilities, and enables one to use first-order conditions to search for the optimum.

Corollary 4: *When there is one buyer, there exists an optimal mechanism such that* $|supp(\tau)| \leq K$.

Given that one need only consider distributions of size $\leq K$, one can solve for the optimal mechanism by deriving the values of $\tilde{\mathbf{t}}(\mu)$ and $\mu(\cdot|x)$ for a given $\underline{\mu}$ and then using numerical techniques to find the best K points along the contour.

Remark 1: As the utility of the buyer from the mechanism is linear in θ, the expected utility of a given $(x, \mathbf{t}(x))$ is a function of the posterior mean of θ. Thus one may be tempted to use the results of Dworczak and Martini (2019), who provide additional tools for solving persuasion problems in this class using a price-theoretic approach. However, this is not immediately applicable here, as their results depend on the ability to implement any distribution of posteriors which contain the prior in their convex hull. This will generally not be possible due to the requirement that beliefs must lie along the contour mechanism. Moreover, the cost function of information acquisition, based on H, would need to be measurable in the posterior mean, in contrast to the present (weaker) assumption of strong concavity of H in the posteriors. $\square$

While I will press the analysis of the form of the optimal mechanisms

[12] This follows from the standard reasoning based on Carathéodory's theorem.

further for the case of binary states in Section 6, it is possible to provide a partial characterization for the general case as follows.

Proposition 5: *In the optimal single-buyer mechanism, for some* $\mu \in supp(\tau)$*,* $\mathbf{x}(\mu) = 1$*, while at the other posteriors (if any),* $\mathbf{x}(\mu) \in [0, 1)$*.*

In words, there is always an option in the optimal mechanism that sells the item to the buyer with probability 1. The intuition is similar to that for the standard case of exogenous interim values: one leaves money on the table by having all allocation probabilities be less than 1, since one can raise them without changing the information acquisition incentives. One can then raise transfers as well, increasing revenue.

One may well wonder if one can say something about the allocation probabilities at the other posteriors. For instance, with exogenous distributions, a posted-price mechanism is always optimal (e.g. Riley and Zeckhauser, 1984). The answer, as we will see in the next section, is no: there will be cases where non-posted-price mechanisms are optimal. In the parlance of Riley and Zeckhauser, it will be optimal to "haggle" over the price by making offers that are neither accepted nor rejected with certainty.

Remark 2: Consider the transfers that are associated with the optimal mechanism. By Corollary 4, the optimal mechanism will contain at most K signal realizations, and by Proposition 5 there will be one signal realization at which the item is allocated with probability 1. If one were to take these signals as exogenously given, the buyer would be indifferent between this allocation and that associated with the lower realization (in the case of posted price mechanisms, that given by non-participation), as otherwise the seller could increase the price. Yet here, ex interim, the buyer strictly prefers their assigned allocation.

This can be inferred from Figure 1. In Figure 1(a), the contour mechanism maps out a continuum of possible values of x that increase as one moves along the contour of beliefs. In Figure 1(b), as a result of the changes in beliefs as x increases, the transfers associated with x are increasing as well. Implicitly, the seller must offer a strictly increasing allocation for the intervening possible posteriors μ, even if they occur with probability 0, in order to maintain ex-ante incentive compatibility. By the envelope theorem, as in Myerson (1981), this means that the ex-interim utility to the buyer must be strictly increasing for intervening values of μ.

By contrast, in the case of exogenous signals, the seller does not need to worry about such posteriors, as they occur with probability 0. The seller therefore keeps the allocation constant in the gaps between the interim expected values in the support, and so the ex-interim utility remains constant. Thus the seller must leave more rent to the buyer in the case of endogenous distribution of interim types compared to the case of exogenous distribution, even when allocating to them with the same probabilities. □

It is well known that Bayesian persuasion problems can be difficult to solve in general, especially when there are at least three possible states. The additional level of complication of finding the right non-participation belief makes a general explicit solution even more difficult. In the remainder of the paper, I explore more particular environments, in which it is possible to say more about the optimal mechanisms.

6 Optimal Mechanisms: Single Buyer, Binary Types

6.1 General characterization

Up until now, the results have held generally for an arbitrary number of states for a given buyer. In this section, I focus on the case of two states, which has more structure, in order to derive further results. In particular, this ensures that there is essentially one contour mechanism possible, and so reduces to the choice of a non-participation belief, and the subsequent persuasion problem, along that contour.

Let $\Theta = \{\theta_1, \theta_2\}$ with $\theta_2 > \theta_1$; then beliefs μ are one-dimensional, allowing one to order them by the probability they place on θ_2. Then as derived in

Appendix B in the proof of Lemma 4, once can write the marginal change in $\mathbf{x}$ as $\mu(\theta_2)$ increases as

$$\mathbf{x}'(\mu) = -\frac{\frac{\partial^2 H}{\partial\mu(\theta_2)^2}(\mu) - 2\frac{\partial^2 H}{\partial\mu(\theta_2)\partial\mu(\theta_1)}(\mu) + \frac{\partial^2 H}{\partial\mu(\theta_1)^2}(\mu)}{\theta_2 - \theta_1} \tag{9}$$

This is positive due to the strong concavity of H, which makes the numerator negative. Intuitively, higher allocations must be given to higher interim types, which correspond to higher $\mu(\theta_2)$. The degree of concavity of H can be interpreted as the rate at which the marginal cost of information increases. To incentivize the buyer to acquire information despite this cost, the allocation provided at posterior μ must change sufficiently. At the same time, a larger difference between states, $\theta_2 - \theta_1$, incentivizes the buyer to acquire more information, and so a smaller change in allocation is needed to reach μ.

By (IC-I), the Myersonian envelope condition $\mathbf{t}'(x) = E_{\mu(\cdot|x)}[\theta]$ holds, and so with some abuse of notation treating μ as standing for $\mu(\theta_2)$ and $\tilde{\mathbf{t}}$ as a function of that single variable,

$$\tilde{\mathbf{t}}'(\mu) = [\theta_2\mu + \theta_1(1-\mu)]\mathbf{x}'(\mu) \tag{10}$$

Thus, for fixed $\underline{\mu}$, the persuasion problem is defined by

$$v_{\mathcal{C}}(\mu) = \tilde{\mathbf{t}}(\mu) = \int_{\underline{\mu}}^{\mu} \tilde{\mathbf{t}}'(\mu)d\mu \tag{11}$$

where $\mathbf{x}(\underline{\mu}) = 0 = \mathbf{t}(\mathbf{x}(\underline{\mu}))$.

An interesting fact that emerges is that any shape of $\tilde{\mathbf{t}}(\cdot)$ (respectively, $\mathbf{x}(\cdot)$) can be generated by the appropriate choice of H, as long as it is monotone and sufficiently smooth, by manipulating the degree of concavity of H at μ. If one scales H by κ while holding fixed the values of θ, one scales $\mathbf{x}'$ by κ. If one scales the values of θ by κ as well, one returns to the original values of $\mathbf{x}'$, but scales $\tilde{\mathbf{t}}'$ by κ. This allows one to arbitrarily restrict the feasible interval $[\underline{\mu}, \bar{\mu}]$ as well. In tandem, one can have arbitrary combinations of $\tilde{\mathbf{t}}'$ and intervals $[\underline{\mu}, \bar{\mu}]$. This is useful in exploring properties of the optimal mechanism, as will be done when examining the optimality of the posted-price mechanism.

Equations (9) and (10) reveal that in the binary state environment, there is essentially a unique implementable mechanism: for a given belief, both the allocations and transfers must evolve in a fixed way. The additional structure imposed allows one to represent beliefs, allocations, and transfers by a single one-dimensional variable. The only degrees of freedom for the seller along this contour mechanism, then, are the non-participation belief, and the posteriors which receive positive weight (of which, by Corollary 4, there need not be more than two). These narrow constraints allow one to straightforwardly solve for the optimal mechanism.

These equations highlight the three economic tradeoffs that the seller makes when choosing the optimal mechanism alluded to in the motivating example. First, by the envelope theorem, the marginal benefit of increasing the allocation at belief μ is $\theta_2\mu(\theta_2) + \theta_1(1 - \mu(\theta_2))$. As this is increasing in $\mu(\theta_2)$, the seller would like to provide higher allocations at higher beliefs, rather than sell with high probability x near μ_0, all things being equal.

Second, the Bayesian constraint on beliefs provides a counteracting incentive. Suppose that we label the optimal (binary) posteriors μ_1 and μ_2, with $\mu_2(\theta_2) > \mu_1(\theta_2)$. While, as argued in the previous paragraph, increasing μ_2 provides higher rent extraction opportunities, it also decreases $\tau(\mu_2)$ if one holds μ_1 fixed. Thus, there is a tradeoff between the amount of rent one can extract for a given posterior, and the probability with which that posterior occurs.

Lastly, for any $\{\mu_1, \mu_2\}$, the seller would like to extract as much rent as possible, by selling with the highest possible probabilities. This is done by setting $\underline{\mu}$ as low as possible. However, the seller is constrained by the allocation probability x needing to be at most 1. It is straightforward that in the optimal mechanism, $\mathbf{x}(\mu_2) = 1$. Beyond this, though, there is now an additional incentive to decrease μ_2 in order to relax the aforementioned allocation constraint. This incentive is distinct from that in the previous paragraph, which arises from the constraint from Bayes' rule on posteriors; here, it arises from the constraint on allocation probabilities. Hence even if (μ_1, μ_2) are optimal for a given $\underline{\mu}$, it may be an improvement for the seller to lower $\underline{\mu}$ in order to extract more rents; this may lead to a different choice of posteriors being optimal overall.

6.2 Posted-price mechanisms

A classic result from the literature on pure allocation mechanisms is that posted-price mechanisms are optimal: the buyer can either pay a certain t and receive the item with probability 1, or not pay anything and receive nothing. Here, this result does not automatically hold. Viewing the problem as one of Bayesian persuasion makes the reasoning behind this clear. There are two possibilities for posted-price mechanisms: either the buyer always purchases the item (i.e. $\tau(\mu(\cdot|x=1)) = 1$), or only does so some of the time ($\tau(\mu(\cdot|x=1)) < 1$). In the former case, there will be one signal realization (i.e. no information acquired); in the latter, there will be two. As seen above, in the optimal mechanism, the item is sold with probability 1 at the higher of the two induced beliefs (if only one is induced, the item is sold with probability 1 as well). Yet the exact optimal persuasion mechanism will depend on $\tilde{\mathbf{t}}$. This will lead to some situations where it will not be optimal to have a posted price. I illustrate this in the following example.

Example 2: Consider the case, as illustrated in Figure 3, where[13]

$$\tilde{\mathbf{t}}'(\mu) = \begin{cases} 1, & \mu < 0.4 \\ \epsilon, & \mu \in [0.4, 0.5] \\ \frac{0.07}{\epsilon}, & \mu \in (0.5, 0.5+\epsilon] \\ 0.4, & \mu > 0.5+\epsilon \end{cases}$$

[13]Recall from the discussion following equations (9-11) that one can generate any monotone $\tilde{\mathbf{t}}$ with the appropriate concave H. While the present $\tilde{\mathbf{t}}$ is technically not differentiable at the transition points, one can uniformly approximate this by a differentiable function. This leads to the solution presented here approximating the correct one for appropriately differentiable H.

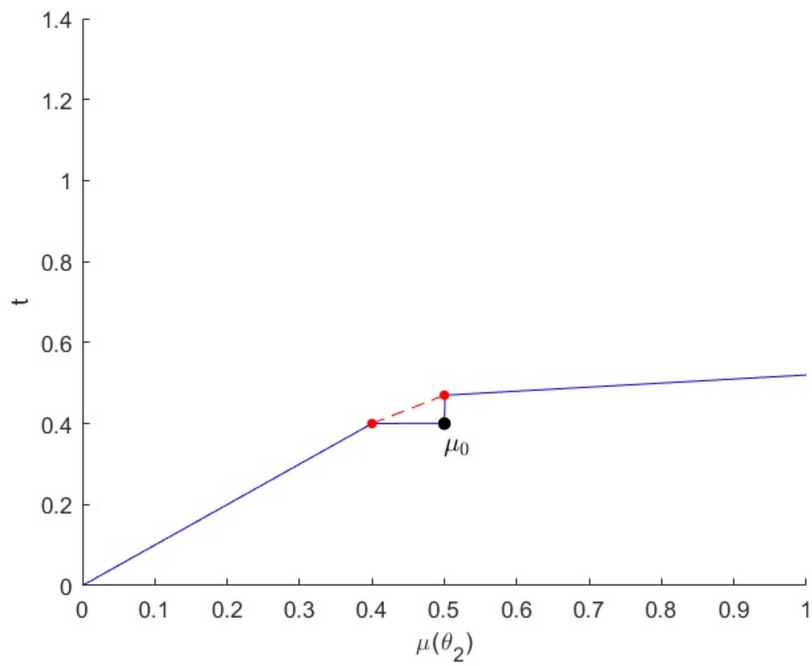

Figure 3: Posted price not optimal

Consider $\mu_0 = 0.5$ in the limit as $\epsilon \to 0$. To see that the posted-price mechanism is not optimal for some values of κ (a constant to scale H) and θ_1, θ_2, we can choose these such that $\int_{0.4}^{0.5+\epsilon} x'(\mu)d\mu < 1$. By the concavification results from Bayesian persuasion, the optimal persuasion mechanism will have support on the two points $\{0.4, 0.5+\epsilon\}$ if they are both feasible for some $\underline{\mu}_1 \leq 0.4$. We must also check that it is not better to instead set $x(\mu_0) = 1$ for some other $\underline{\mu}_2 < 0.4$. The reason that the former is preferable is that, by the envelope theorem, one needs to increase x by less to achieve the same increase of t for $\mu(\theta_2) \in [0.5, 0.5+\epsilon]$ than for lower values of $\mu(\theta_2)$. Formally, if

$$\int_{\underline{\mu}_2}^{\underline{\mu}_1} x'(\mu)d\mu = \int_{0.5}^{0.5+\epsilon} x'(\mu)d\mu$$

then

$$\int_{\underline{\mu}_2}^{\underline{\mu}_1} \tilde{\mathbf{t}}'(\mu)d\mu < \int_{\underline{\mu}_2}^{\underline{\mu}_1} x'(\mu)[0.4\theta_2 + 0.6\theta_1]d\mu$$

$$< \int_{0.5}^{0.5+\epsilon} x'(\mu)[0.5\theta_2 + 0.5\theta_1]d\mu$$

$$< \int_{0.5}^{0.5+\epsilon} \tilde{\mathbf{t}}'(\mu)d\mu$$

Moreover, for $\underline{\mu}$ such that $\mu = 0.5 + \epsilon$ is a feasible posterior,

$$\lim_{\epsilon \to 0} V_{\mathcal{C}}(0.5) = v_{\mathcal{C}}(0.5) + 0.07$$

$$\lim_{\epsilon \to 0} \int_{0.5}^{0.5+\epsilon} x'(\mu)d\mu = \frac{0.07}{0.5\theta_2 + 0.5\theta_1}$$

Thus one can ensure that $0 < \underline{\mu}_2 < \underline{\mu}_1 < 0.4$ for the right choices of $\{\theta_1, \theta_2\}$, i.e. $0.5\theta_2 + 0.5\theta_1 > 0.07$ but not too large. The limit difference in payoffs between the two mechanisms is then

$$\int_{0.5}^{0.5+\epsilon} \tilde{\mathbf{t}}'(\mu)d\mu - \int_{\underline{\mu}_2}^{\underline{\mu}_1} \tilde{\mathbf{t}}'(\mu)d\mu > 0$$

Hence the optimal mechanism will involve one option of buying with probability one, and another of buying with a lower but nonzero probability. $\square$

To understand the economic intuition for the suboptimality of the posted-price mechanism in Example 2, one must consider the tradeoffs mentioned above for the seller's optimal mechanism design problem. In order for the seller to find it optimal for the buyer to acquire any information at all (rather than simply focus on the third incentive, i.e. to extract as much rent as possible from μ_0), the incentive to sell to the buyer at belief $\mu > \mu_0$ due to higher rent extraction must significantly outweigh the lower probability with which μ occurs due to Bayes' rule. Thus, ideally, the seller would like to induce high beliefs that occur with high probability. The rapid increase in marginal cost, due to the high concavity between $\mu(\theta_2) = 0.5$ and $0.5 + \epsilon$, enables the seller to do so. Beyond $\mu(\theta_2) = 0.5 + \epsilon$, the marginal cost rises more slowly, leading to the second effect outweighing the first.

On the other hand, the marginal cost rises slowly as one moves μ down from μ_0, until one reaches $\mu(\theta_2) = 0.4$, after which it rises more rapidly. This enables the seller to make the higher of the two posteriors more likely, without violating the allocation probability constraint. Below $\mu(\theta_2) = 0.4$, the seller must increase the rate of change of the allocation probability. This does not allow him to assign enough additional weight (by Bayes rule) to posterior $0.5+\epsilon$ to outweigh the loss of revenue from decreasing the allocation probability at $\underline{\mu}$, and so the lower of the two beliefs is optimally set at $\mu(\theta_2) = 0.4$. Since the difference $\mathbf{x}(0.5+\epsilon) - \mathbf{x}(0.4) < 1$, this implies a non-posted-price mechanism.

Just as the Bayesian persuasion perspective illustrates cases where the posted-price mechanism is suboptimal, it also provides sufficient conditions for the posted-price mechanism to be optimal. Using the intuition from this perspective, if $\tilde{\mathbf{t}}$ is convex, then the seller wants the buyer to acquire as much information as possible, given the incentives of the contour mechanism. This involves inducing posterior beliefs that lie at the extreme points of the contour mechanism, i.e. corresponding to $x = 0$ and $x = 1$, as shown in Figure 5. Of course, such a policy is by definition a posted-price mechanism.

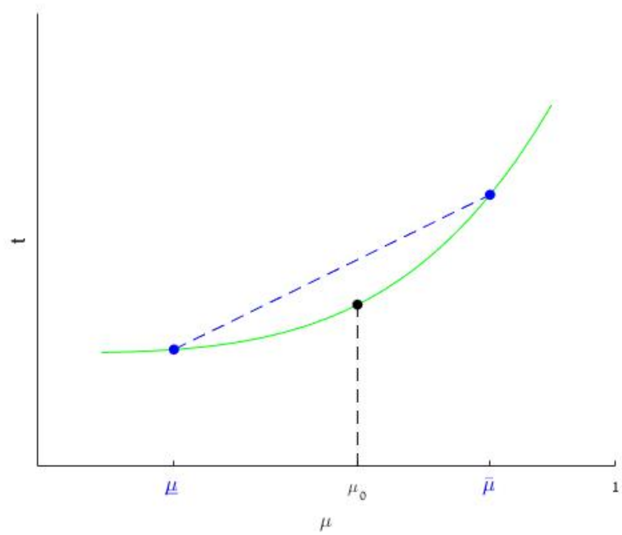

Figure 4: Posted price when $\tilde{\mathbf{t}}$ is convex

Alternatively, if $\tilde{\mathbf{t}}$ is concave, as in Figure 6, then the seller wants the buyer to acquire no information. Thus the buyer's posteriors will remain at the prior μ_0. Given such a policy, it is optimal to set x as high as possible for μ_0, i.e. $\mathbf{x}(\mu_0) = 1$. Thus it leads to a posted price being optimal as well. I state this formally in the following theorem.

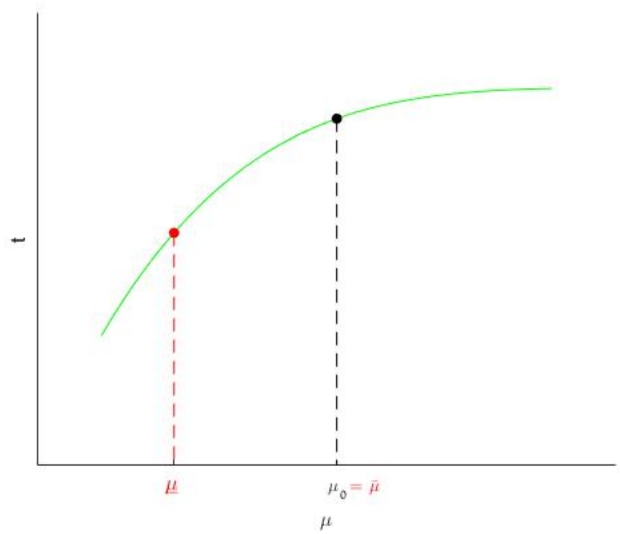

Figure 5: Posted price when $\tilde{\mathbf{t}}$ is concave

Theorem 6: *When Θ is binary, if $\tilde{\mathbf{t}}(\mu)$ is either concave or convex in μ for all choices of $\underline{\mu}$, then a posted-price mechanism is optimal.*

Example 3: Suppose that H is proportional to residual variance, i.e.[14]

$$H(\mu) = \kappa[\mu(\theta_2)(1 - \mu(\theta_2)) + \mu(\theta_1)(1 - \mu(\theta_1))]$$

Then

$$\mathbf{x}'(\mu) = \frac{4\kappa}{\theta_2 - \theta_1}$$

[14]This formula has been used as a cost function for information design by Gentzkow and Kamenica (2014) and Ely et al. (2015).

and so

$$\tilde{\mathbf{t}}'(\mu) = \frac{4\kappa}{\theta_2 - \theta_1}(\theta_2\mu(\theta_2) + \theta_1(1 - \mu(\theta_2)))$$
$$\implies \tilde{\mathbf{t}}''(\mu) = 4\kappa > 0$$

Hence the posted-price mechanism is optimal.[15] □

The economic intuition is as follows. Recall that any statement about the shape of $\tilde{\mathbf{t}}$ is implicitly a statement about H, i.e. the costs of information acquisition. From the discussion of Example 2, in order to make information acquisition worthwhile, the seller must see sufficient gains in rents from extracting from higher types, and these must occur with high enough probability. Convexity of $\tilde{\mathbf{t}}$ essentially means that it is expensive for the buyer to acquire high beliefs, but cheap to acquire low ones. The seller can be sure that the buyer will not acquire beliefs that are so high that they occur with such low probability to make their rent extraction gains not worthwhile. Conversely, the cheapness of low beliefs means that the seller can easily incentivize low posteriors in order to make the high-rent outcome more likely. Hence the seller optimally induces as much information acquisition as possible, given the allocation probability constraint.

For concavity, the intuition runs the other way: it is cheap to acquire high beliefs, and expensive to acquire low ones. Hence the seller does not find it worthwhile to attempt to extract higher rents by inducing higher beliefs, since offering a higher allocation leads the buyer to a posterior so much higher that it is very unlikely to occur; this would outweigh the rent gains from higher posteriors. Meanwhile, low beliefs are too costly relative to the gains in probabilities of the buyer receiving the high signal. The seller therefore optimally induces no information acquisition.

The simple form of the optimal posted-price mechanism makes it straightforward to compute the optimal beliefs. By standard envelope-theoretic reasoning ex interim, the marginal increase in transfer from an increase in the probability of receiving the item is equal to the expected value of the buyer.

[15] In fact, I show in Section 6 that when costs are defined by residual variance for an arbitrary number of states and buyers, a second-price auction with a reserve price is optimal.

Thus, increasing p must correspond to increasing the interim expected values at which the item is sold with probability 1 (namely, $\bar{\mu}$), and hence increasing the non-participation belief ($\underline{\mu}$) as well to ensure that $\int_{\underline{\mu}}^{\bar{\mu}} \mathbf{x}'(\mu) d\mu = 1$, as $\mathbf{x}'(\mu) > 0$. As $\tilde{\mathbf{t}}'(x)$ is positive and increasing in x, it follows that for any posted price p, there will be a unique pair of beliefs $(\underline{\mu}, \bar{\mu})$ such that

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

and

$$\int_{\underline{\mu}}^{\bar{\mu}} \tilde{\mathbf{t}}'(\mu) d\mu = p$$

Therefore, p will implement posteriors with support on $\{\underline{\mu}, \bar{\mu}\}$.

7 Optimal Mechanisms: Multiple Buyers

I now extend the results for single-buyer environments to the N-buyer case. I assume that values are private and independently, symmetrically distributed. The mechanism design literature has long focused on symmetric mechanisms, going back to Myerson (1981) and Maskin and Riley (1984). Additionally, as has been noted by Bergemann and Pesendorfer (2007) and Deb and Pai (2017), there may be legal or fairness restrictions that prevent asymmetric treatment of the buyers. I therefore primarily restrict attention to symmetric mechanisms as a first step to approaching this problem. However, this restriction may not be without loss of optimality. Indeed, in the case of H quadratic (Section 7.2), I show that for two buyers, one can improve over the optimal symmetric mechanism by a simple asymmetric one.

The main techniques developed for the single-buyer case can be used in the the multi-buyer case as well. From the perspective of each individual buyer, the problem is exactly the same. That is, the tradeoff between information acquisition and allocation probabilities is identical. Thus the same envelope condition as in Lemmas 3 and 4 hold here as well.

The key difference, though, is that the presence of N buyer imposes a resource constraint that was not present before. If the seller allocates the item to one of the buyers, he cannot physically allocate it to another one. When combined with the interim incentive compatibility constraint, the results of Border (1991, Proposition 3.2) imply that a given mechanism is implementable ex interim if and only if it satisfies

$$\int_{x^*}^{1} x d\tau(\mu(\cdot|x)) \leq \frac{1 - \tau(\{\mu : \mathbf{x}(\mu) < x^*\})^N}{N} \tag{12}$$

for all x^*. Since, by Theorem 1, any distribution τ is implementable as long as it lies along the same contour in the case of a single buyer, this implies that a mechanism will be implementable if and only if the induced τ along the contour satisfies the extra constraint in (12).

7.1 General N-buyer case, $K = 2$

The presence of the constraint in (12) prevents the straightforward use of concavification that one could use in the single-buyer case with $K = 2$, as the resultant distribution will no longer necessarily be implementable. For instance, in Example 2, the optimal mechanism induced a high posterior with positive probability, at which the item was sold with probability 1. This violates the Border constraint, since it entails a positive probability event of allocating the same item to multiple buyers with probability 1, an obvious impossibility. This issue requires developing techniques that incorporate this constraint.

The key insight is that the seller would still like to concavify; he is merely constrained from doing so by feasibility considerations. So, he would essentially like to concavify "as much as possible." That is, he would like to spread out the posteriors in such a way that exploits the possible gains in rents from information acquisition, without promising any buyer to receive the item with high enough probability to violate the Border constraint. Thus, he will implement any feasible mean-preserving spread from μ such that the average value of $\tilde{\mathbf{t}}$ from the spread is greater than at the original μ. This makes the solution to the problem a form of constrained concavification.

How does this approach square with the results of the single-buyer case regarding the concavity/convexity of $\tilde{\mathbf{t}}$? As before, ignoring the Border constraint, one wants to pool the probabilities in the former case, and spread them out in the latter. However, one also has an incentive to manipulate the Border constraint as well. If $\mathbf{x}(\cdot)$ is a convex function, then pooling the probabilities for $\mu > \mu^*$ decreases x on average, and so relaxes (12) at μ^*. The seller could then potentially extract more rent by lowering the non-participation belief. This presents an extra incentive to the seller to induce less information acquisition in the multi-buyer environment, rather than always inducing as much information acquisition as possible (i.e. the most extreme mean-preserving spread) as found in the single-buyer case.

Conversely, if $\mathbf{x}(\cdot)$ is concave (which it will be whenever $\tilde{\mathbf{t}}$ is concave, by (10)), then pooling μ over an interval to some μ^* potentially leads to a violation of (12) at μ^*. Thus the seller may want to spread out the beliefs even further (by inducing information acquisition) even when $\tilde{\mathbf{t}}$ is concave, which he would never want to do in the single-buyer case.

Notice that if $\tilde{\mathbf{t}}$ is convex and $\mathbf{x}(\cdot)$ is concave (as it is, for instance, when the cost is given by residual variance), the seller has incentives to spread out beliefs as much as possible both from concavification and from the Border constraint. Hence, given $\underline{\mu}$, the seller will assign as much weight as possible to the highest possible beliefs (while satisfying (12) with equality), and assign the rest of the weight to $\underline{\mu}$. This can be implemented by, for instance, a second-price auction with a reserve price. I state this formally in Lemma 5 and Proposition 7.

Lemma 5: *Suppose that Θ is binary. If $\mathbf{x}(\cdot)$ is concave, then, for any τ and $\underline{\mu}$ satisfying (12), there exists a mean-preserving spread of τ, denoted by $\hat{\tau}_{\underline{\mu}}$, such that for some μ^*, (12) holds with equality for all $\mu \geq \mu^*$, and $\hat{\tau}_{\underline{\mu}}(\underline{\mu}) = 1 - \hat{\tau}_{\underline{\mu}}(\mu \geq \mu^*)$. Moreover, $\hat{\tau}_{\underline{\mu}}(\underline{\mu})$ is increasing in $\underline{\mu}$.*

Proposition 7: *Suppose that Θ is binary and that boundary beliefs are not implementable given μ_0. If $\mathbf{x}(\cdot)$ is concave and $\tilde{\mathbf{t}}(\cdot)$ is convex, the optimal* *mechanism can be implemented by a second-price auction with a reserve price.*

A simple example of a cost function that satisfies the criteria of Proposition 7 is that of residual variance; another example is the following.

Example 4: Suppose that $\tilde{\mathbf{t}}$ is linear in μ on the set of implementable beliefs, and so for some constant[16] $\alpha > 0$, $\mathbf{x}'(\mu) = \frac{\alpha}{\theta_2\mu(\theta_2)+\theta_1(1-\mu(\theta_1))}$. Notice that persuasion has no effect here: given $\underline{\mu}$, $E_\tau[\tilde{t}(\mu)]$ is immutable. Thus, the only way to affect payoffs is to change $\underline{\mu}$, which at the optimum will be as low as possible. This is done by the second-price auction that sets the lowest possible value of $\underline{\mu}$, which, by Lemma 5, corresponds to setting $\tau(\underline{\mu}) = 0$. In other words, the reserve price will be 0, as the item is sold with probability 1 to some player, i.e. with probability $\frac{1}{N}$ to each. □

Since the optimal mechanism is of this particular form, it is much simpler to find the corresponding non-participation belief. By Lemma 5, one can infer the corresponding reserve price for values of $\underline{\mu}$ as we decrease it from μ_0, such that Bayes' rule is satisfied, as there will be a unique such reserve price for each such $\underline{\mu}$ as defined by (9) and (10).

While the optimal mechanism here is, as in classic results like Myerson (1981), a second-price auction with a reserve price, the intuition here is essentially a dual to the classic intuition. In classic results, the second-price auction is optimal because one has a fixed distribution, and the seller benefits most from selling with highest probability to the types with highest value. Here, one considers a fixed set of possible probabilities of sale, and considers the optimal distribution over interim expected values. Since the seller wants to assign the highest possible values to the highest probabilities of sale, he spreads out beliefs in order to achieve this. However, such a spread of beliefs is only beneficial overall if the requisite convexity/concavity conditions are satisfied, since it also increases the probability of the buyer(s) receiving a low interim

[16] By (9), this would mean that

$$\frac{\partial^2 H}{\partial\mu(\theta_2)^2}(\mu) - 2\frac{\partial^2 H}{\partial\mu(\theta_2)\partial\mu(\theta_1)}(\mu) + \frac{\partial^2 H}{\partial\mu(\theta_1)^2} = -\frac{\alpha(\theta_2-\theta_1)}{\theta_2\mu(\theta_2)+\theta_1(1-\mu(\theta_1))}$$

from which it is possible to integrate the function to derive H.

expected value. This makes the conditions for optimality of the second-price auction with a reserve price more stringent.

As with exogenous signals, the optimal reserve price depends on the trade-off between greater participation and greater extraction of rents from those buyers who do participate. With flexible information acquisition, this comes down to how much the seller gains from persuasion versus exploitation of the induced distribution of posteriors. Higher reserve prices induce less weight on intermediate posteriors (i.e. those below μ^*), and so can be thought of as a form of "more" persuasion for the purpose of extracting rents from high beliefs. This is the analogue of the first two incentives that were present in the single-buyer problem. However, this additional persuasion prevents exploitation that could be achieved through lowering $\underline{\mu}$.

Remark 3: Similarly to Remark 2, here the reserve price for a given cutoff interim value $E_{\mu^*}[\theta]$ will be lower than if the distribution were exogenously given, as long as it is above $E_{\underline{\mu}}[\theta]$. This means that the buyer who has this cutoff value will receive positive rents. The reasoning is analogous: the seller must provide positive rents to this interim value in order to deter deviation to values between $E_{\underline{\mu}}[\theta]$ and $E_{\mu^*}[\theta]$. So, even if the seller were to implement the same allocation in the auction ex interim, his revenue would be lower. This again contrasts with inflexible information acquisition, e.g. Shi (2012), where the reserve price is precisely equal to the value of the cutoff type when there are also types present just below it. Indeed, even if one were to consider a binary information decision of whether to acquire information according to τ or not to acquire any information, then as long as the buyers would get sufficient rent on average, it would not be necessary to modify the reserve price at all if it is above the ex-ante mean expected value (and so they would rather acquire the information to be able to purchase the item with positive probability at all ex ante). The ability of the buyer to deviate flexibly on the margin in the choice of information aquisition thus drives the reserve price down. $\square$

7.2 Quadratic costs, K arbitrary

Many of the techniques from the binary-state environment carry over to the case where H is quadratic, i.e.

$$H(\mu) = \sum_{\theta,\theta' \in \Theta} a_{\theta,\theta'} \mu(\theta)\mu(\theta')$$

for some set of coefficients $\{a_{\theta,\theta'}\}_{\theta,\theta' \in \Theta}$. This generalizes the case examined in Example 3, where H was given by residual variance for a single buyer.[17] These cost functions are natural to consider for several reasons. First, it is the simplest example of a function of the form that I consider (i.e. strongly concave and sufficiently differentiable). Since the second derivative of H plays a major role in the analysis, examining the case where the second derivative is constant is a natural place to start. In addition, the residual variance function, as noted by Ely et al. (2015), can describe the amount of private information in insider trading models (Kyle, 1985; Ostrovsky, 2012).

The following result extends the posted-price result seen in Example 3.

Proposition 8: *Let H be quadratic. Then the optimal mechanism is implementable by a second-price auction with a reserve price.*

The intuition for Proposition 8 is that it turns out that the contour mechanisms $\mathcal{C}$ for such H are relatively simple: the beliefs $\mu(\cdot|x)$ must all be linear in x. As a result, it is straightforward to consider mean-preserving spreads of beliefs, allowing for manipulation analogous to that in Lemma 3. As transfers $\mathbf{t}(x)$ will be convex in x, the most extreme mean-preserving spreads possible will be optimal. Of course, these are implemented by second-price auctions with reserve prices, as was similarly found for the case of binary states.

This makes finding the non-participation belief/reserve price pairs similar to the binary-state environment. One can infer possible values of the non-participation belief along the line of possible beliefs that intersects μ_0. For each such possible non-participation belief, there will be a unique Bayes-plausible

[17] Recall that any terms that are affine in $\mu(\theta)$ wash out in expectation. Hence it is without loss to consider H that has only quadratic terms, and none of lower degree.

reserve price. One then calculates the optimal non-partipation belief along this line by checking which of these maximizes revenue given its associated reserve price. This greatly simplifies the maximization problem to a single value on a given line.

Remark 4: The simple structure of the contour mechanisms for H quadratic allows for a straightforward construction of an asymmetric mechanism that would do better than a (symmetric) second price auction with a reserve price. Consider the case of $N = 2$. Let μ^* be the cutoff belief, such that for all μ along the contour, if $E_\mu[\theta] < E_{\mu^*}[\theta]$, then either $\mu \notin \text{supp}(\tau)$ or $\mu = \underline{\mu}$. Thus μ^* is the posterior belief at which the buyer pays the reserve price conditional on winning.

As can be easily derived from (12) and the linearity of $\mu(\cdot|x)$ in x, the posterior distribution of interim values must be uniform over the interval $[E_{\mu(\cdot|x^*)}[\theta], E_{\bar{\mu}}[\theta]]$. Then the alternative distribution of posterior values

$$\hat{\tau}(\mu) = \begin{cases} \frac{1}{4}(1 - \tau(\underline{\mu})), & \mu = \mu^* \\ \frac{1}{2}(1 - \tau(\underline{\mu})), & \mu = \frac{1}{2}(\mu^* + \bar{\mu}) \\ \frac{1}{4}(1 - \tau(\underline{\mu})), & \mu = \bar{\mu} \\ \tau(\underline{\mu}), & \mu = \underline{\mu} \end{cases}$$

is a mean-preserving spread of τ. Now suppose that the seller induces buyer 1 to acquire information according to

$$\tau_1(\mu) = \begin{cases} (1 - \tau(\underline{\mu})), & \mu = \frac{1}{2}(\mu^* + \bar{\mu}) \\ \tau(\underline{\mu}), & \mu = \underline{\mu} \end{cases}$$

and buyer 2 to acquire information according to

$$\tau_2(\mu) = \begin{cases} \frac{1}{2}(1 - \tau(\underline{\mu})), & \mu = \mu^* \\ \frac{1}{2}(1 - \tau(\underline{\mu})), & \mu = \bar{\mu} \\ \tau(\underline{\mu}), & \mu = \underline{\mu} \end{cases}$$

Thus $\frac{1}{2}\tau_1+\frac{1}{2}\tau_2 = \hat{\tau}$. Moreover, since $\tilde{\mathbf{t}}(\mu)$ is convex, this alternative information structure yields higher profits to the seller. This means that the seller could improve his payoff by using a hierarchical mechanism,[18] in which:

- the seller sells to buyer 1 if she receives a posterior of $\bar{\mu}$;
- otherwise, he sells to buyer 2 if she receives a posterior of $\frac{1}{2}(\mu^* + \bar{\mu})$;
- if neither hold, then he sells to buyer 1 if she received a posterior of μ^*; and
- the seller does not sell the item in any other case. $\square$

The intuition for why the optimal mechanism is asymmetric differs from that of Bergemann and Pesendorfer (2007, Lemma 3). In the latter, the seller would prefer to break any symmetry by awarding the item to a particular buyer i with the same probability for adjacent posterior expected values, and then merging them; this would increase the virtual value associated with the merged posterior, and hence the revenue that can be extracted for it, without changing the probability of receiving the item. By contrast, here, x and t are determined uniquely for a given posterior by $\underline{\mu}$, regardless of the rest of the distribution τ; thus the virtual values cannot be manipulated in this way. Instead, it is the nonlinearity of the transfer function $\tilde{\mathbf{t}}$ which may cause mean-preserving spreads/contractions of the distribution to improve the seller's payoff. As the set of possible mean-preserving spreads/contractions is expanded by allowing for asymmetry, it may be optimal to use an asymmetric mechanism.

7.3 $N \to \infty$, K arbitrary

Lastly, I consider the case of large auctions, as the number of buyers approaches infinity. The interplay of the incentives for information acquisition and the Border constraint make the question of what happens in the infinite-buyer limit more complicated. On the one hand, as is standard in the literature, the winning buyer will be in the upper tail of the distribution of buyers'

[18] Such mechanisms play a prominent role in the proofs in Border (1991).

values. On the other hand, with flexible information acquisition, this distribution will be endogenous. While for each N, the buyers can still be incentivized to acquire information that yields beliefs μ with positive probability such that $E_\mu[\theta]$ is well above $E_{\mu_0}[\theta]$, each individual buyer's chance of winning vanishes in the limit, and so the incentive to acquire such information vanishes as well for each buyer. Nevertheless, it is not immediately clear that inducing such information is asymptotically ineffective, since it remains possible that with flexible information acquisition, the buyer acquires signal realizations with significantly higher values than those at the prior, albeit with vanishing probabilities. Therefore, depending on the rate of vanishing, aggregating across the buyers may yield such high signals with high enough probability to increase the seller's revenue. The question, then, is how exactly the two effects interact, what mechanism is optimal, and what information is chosen as a result.

It turns out that the incentives balance in such a way that one can calculate the optimal mechanism, which is implemented by a second-price auction with no reserve price. One can then generate a formula for the expected revenue.

Theorem 9: *Let τ_N be any distribution of posteriors and $\underline{\mu}_N$ be the beliefs for N buyers satisfying (12). Then in the limit as $N \to \infty$,*

(i) For all $\epsilon > 0$,

$$\lim_{N\to\infty} \tau_N(\{\mu : E_\mu[\theta] - E_{\mu_0}[\theta] < \epsilon\}) = 1$$

(ii) The non-participation belief approaches the prior, i.e.

$$\lim_{N\to\infty} \underline{\mu}_N = \mu_0$$

(iii) Maximal revenue is generated by a second-price auction with a reserve price of 0, yielding revenue

$$\lim_{N\to\infty} \max_{\tau_N, \mathcal{C}_N} N \int_0^1 \mathbf{t}_N(x) d\tau_N(\mu(\cdot|x)) = \int_0^1 \frac{\mathbf{t}_{\mu_0}(x)}{x} dx \tag{13}$$

with $\mathbf{t}_{\mu_0}(x)$ derived from (3) for $\underline{\mu} = \mu_0$.

Parts (i) and (ii) of the theorem follow from the fact that information remains costly, but the chance of winning vanishes. The cost of information chosen must vanish as well, which can only be done by beliefs that converge to the prior. Since if this involved a value of x greater than 0 in the limit, it would not be physically feasible, the associated value of x must be 0 as well, i.e. beliefs approach the non-participation belief.

Part (iii) is established by examining the probability that the buyer who ends up receiving the item has a posterior belief that is associated with getting the item with probability at least x. Recall from the discussion at the beginning of this section in the binary-state environment with convex $\tilde{\mathbf{t}}$ that there is a tradeoff between acquiring more information and relaxing the Border constraint. Here, though, the analogous tradeoff disappears in the limit. From part (ii), the limit value of $\underline{\mu}$ is μ_0. So, any manipulation of the limit distribution of beliefs essentially is equivalent to a manipulation of the limit distribution of x, without altering $\mathbf{t}(x)$ and $\mu(\cdot|x)$. Thus, relaxing the Border constraint offers no advantage, as one cannot thereby change the values of $\underline{\mu}$ and $\mathbf{t}$. As $\mathbf{t}(\cdot)$ is a convex function of x, it is optimal to generate mean-preserving spreads of distributions over x that are as extreme as possible, which will be those generated from second-price auctions (by reasoning akin to that in Lemma 5 and Proposition 7).

Remark 5: One might think that the reserve price is of no importance here, because the chances of getting a signal anywhere other than at the top of the distribution vanishes. So, leaving out buyers with an interim value under any reserve price would not affect revenue. While this is true when the distribution of values is exogenous, it is no longer so when the distribution is endogenous as is the case here. As derived in the proof, the density function of the interim values as $N \to \infty$ will be approximately proportional to $\frac{1}{Nx}$. So, the lower signals will be much more likely, making them still relevant for the seller. Indeed, if one were to set reserve price $r > 0$, and defined x^* to solve $\mathbf{t}_{\mu_0}(x^*) = r$, then the limit revenue would only be $\int_{x^*}^{1} \frac{\mathbf{t}_{\mu_0}(x)}{x} dx$. If the reserve price were anything other than 0, the seller would be leaving money on the table by essentially censoring lower parts of the distribution of x. □

One intriguing possibility, which I do not explore here, is that there may be potential improvements for the seller from sequential mechanisms, in which information acquisition is possible during the execution of the mechanism itself, rather than just beforehand once the mechanism has been set. Since the incentive to acquire information is vanishing precisely because the chances of a given buyer receiving the item, even with information acquisition, are so small, eliminating this reason restores the incentives. For instance, the seller could offer the item to the buyers sequentially, where he selects a particular buyer, and depending on the signal she gets, offers to sell with a given probability. If the item remains unsold, the seller can then move on to the next buyer and recommend her an information acquisition strategy, without having to worry about feasibility constraints of allocations among multiple buyers. More generally, the seller could condition the mechanism offered to the second buyer on the signal realization that the first buyer received. Thus, if the first buyer received a signal realization corresponding to a low posterior μ_1, the seller may want the second buyer to focus on high posteriors in order to potentially outbid the first, thereby allowing him to sell at a higher price. Conversely, if the first buyer received a high posterior μ_2, there may be no purpose to having the second acquire any information, since he is likely to sell the item to the first buyer anyway. The ability thereby to correlate the choice of information acquisition across buyers may prove relevant to revenue maximization; the seller cannot do so via a simultaneous, symmetric mechanism in which buyers move independently.[19] In such a case, the relevant solution concept would be Bayes correlated equilibrium (Bergemann and Morris, 2016).

[19] In a related environment, Gershkov et al. (2019) examine optimal mechanisms when the buyers' valuations are endogenous due to potential investments. There are situations in that framework in which sequential mechanisms may be optimal because of this endogeneity. As the values are endogenous here (albeit for the different reason of information acquisition), similar reasoning may apply.

8 Related Literature

Several papers consider auction frameworks with information acquisition, including Persico (2000), Compte and Jehiel (2007), Hernando-Veciana (2009), and Bobkova (2019). These differ from the present model in that the auction format is given (they are not mechanism design problems), and information acquisition is not flexible.

There is also a literature on comparing revenues for given auction mechanisms based on their dependence on the information structure of the buyers, such as Milgrom and Weber (1982), Ganuza and Penalva (2012), Bergemann et al. (2017), and Sorokin and Winter (2019). However, there is no endogenous information acquisition in these models.

Papers that allow for mechanism design as well as information acquisition include Bergemann and Valimaki (2002), Szalay (2009), and Shi (2012). These again all share the feature that the buyers have access to information acquisition technologies that are perfectly ordered. As a result, their models do not allow for the additional insights available from Bayesian persuasion.

Another closely related paper is that of Bergemann and Pesendorfer (2007), who examine optimal Bayesian persuasion in the context of optimal auctions. In their model, the seller controls not only the mechanism, but also the information available to the buyers. Here, of course, the information is in the hands of the buyer.

More generally, starting with Sims (2003) and Matejk and McKay (2015), there has been a burgeoning literature on entropy reduction/flexible information acquisition in decision problems. Caplin and Dean (2013, 2015) examine this using a revealed preference framework, and extend it to other cost functions. They provide necessary and sufficient conditions for optimality of choices by an inattentive decision maker, and also note that her problem is analogous to one of Bayesian persuasion. Many of the techniques for analyzing the buyer's choice, analyzed in Sections 4 and 5, modify their methods to the present environment.

Flexible information acquisition has been used in applied frameworks as well, including Yang (2015), Morris and Yang (2016), Denti (2019), Georgiadis and Szentes (2020), Ravid (2020), Yang (2020), and Lipnowski et al. (2020). Very closely related is the work of Roesler and Szentes (2017) and Ravid, Roesler and Szentes (2021), who also consider an environment involving optimal price setting by a seller, and flexible information acquisition by a buyer. However, these models differ in the timing of the model, and focus on the zero-cost limits of information acquisition. This simplifies the mechanism design problem, as the seller need not worry about the effects of the mechanism on the subsequent incentive to acquire information; hence the optimal mechanism in those models is a posted price as usual.

9 Conclusion

This paper provides new tools for analyzing mechanism design with information acquisition, by considering the possibility of the buyer acquiring information flexibly. This allows the use of techniques from Bayesian persuasion, as the design problem effectively becomes one of implementing a Bayes-plausible distribution of posteriors. This insight allows several additional observations, such as whether the standard mechanisms like a posted price or a second-price auction with a reserve price is optimal.

Several possible extensions present themselves. First, one could consider the case of optimal monopoly quality provision, as in Mussa and Rosen (1978). In such an environment, the seller would offer a menu of qualities and prices, analogous to the allocation probabilities and prices here. Implementability would be again given by contour mechanisms, but the optimal mechanism would need to take into account the cost of production. This would potentially lead to making it less desirable to induce extreme beliefs, as a convex cost of production would lead to large losses for the seller.

Additionally, the use of contour mechanisms is not limited to allocation mechanisms. A similar technique should also be applicable to any sort of principal-agent problem, such as that of contracting under moral hazard. It would be of interest to see if further insights can be generated in these environments as well.

Lastly, I briefly mentioned that the optimal mechanism may be dynamic, since it may allow the seller to better incentivize the information acquisition of the buyer. In particular, dynamic mechanisms provide an additional tool, since they allow responses of later movers to the signals received by the earlier agents. As the benefits of such dynamic mechanisms are not limited to this environment (see Gershkov et al., 2019), this should be explored more generally.

References

Aumann, R. J., and Maschler, M. (1995): *Repeated Games with Incomplete Information.* With the collaboration of R. Stearns. MIT Press, Cambridge, MA.

Bergemann, D., Brooks, B., and Morris, S. (2017): "First-price auctions with general information structures: Implications for bidding and revenue." *Econometrica*, 85, 107-143.

Bergemann, D., and Morris, S. (2016): "Bayes correlated equilibrium." *Theoretical Economics*, 11, 487-522.

Bergemann, D., and Pesendorfer, M. (2007): "Information structures in optimal auctions." *Journal of Economic Theory*, 137, 580-609.

Bergemann, D., and Valimaki, J. (2002): "Information acquisition and efficient mechanism design." *Econometrica*, 70, 1007-1033.

Blackwell, D. (1951): "Comparison of experiments." *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 93-102.

——— (1953): "Equivalent comparisons of experiments." *Annals of Mathematical Statistics*, 24, 265-272.

Bobkova, N. (2019): "Information choice in auctions." Working paper.

Border, K. (1991): "Implementation of reduced-form auctions: A geometric approach." *Econometrica*, 59, 1175-1187.

Caplin, A., and Dean, M. (2013): "Behavioral implications of rational inattention with Shannon entropy." NBER Working Paper 19318.

—— (2015): "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review*, 105, 2183-2203.

Caplin, A., Dean, M., and Leahy, J. (2019): "Rationally inattentive behavior: Characterizing and generalizing Shannon entropy." NBER Working Paper No. 23652.

Coddington, E. A., and Levinson, N. (1972): *Theory of Ordinary Differential Equations.* New York: McGraw-Hill Inc.

Compte, O., and Jehiel, P. (2007). "Auctions and information acquisition: sealed bid or dynamic formats?" *Rand Journal of Economics*, 38, 355-372.

De Oliveira, H., Denti, T., Mihm, M., and Ozbek, K. (2017): "Rationally inattentive preferences and hidden information costs." *Theoretical Economics*, 12, 621-654.

Deb, R., and Pai, M. (2017): "Discrimination via symmetric auctions." *American Economic Journal: Microeconomics*, 9, 275-314.

Denti, T. (2019): "Unrestricted information acquisition." Working paper.

—— (2020): "Posterior-separable cost of information." Working paper.

Dworczak, P., and Martini, G. (2019): "The simple economics of optimal persuasion." *Journal of Political Economy*, forthcoming.

Ely, J., Frankel, A., and Kamenica, E. (2015): "Suspense and surprise." *Journal of Political Economy*, 123, 215-260.

Frankel, A., and Kamenica, E. (2019): "Quantifying information and uncertainty." *American Economic Review*, 109, 3650-3680.

Ganuza, J. J., and Penalva, J. S. (2010): "Signal orderings based on dispersion and the supply of private information in auctions." *Econometrica*, 78, 1007-1030.

Georgiadis, G., and Szentes, B. (2020): "Optimal monitoring design." *Econometrica*, forthcoming.

Gentzkow, M., and Kamenica, E. (2014): "Costly persuasion." *American Economic Review (Papers and Proceedings)*, 104, 457-462.

Gershkov, A., Moldovanu, B., Strack, P., and Zhang, M. (2019): "A theory of auctions with endogenous valuations." Working paper.

Hernando-Veciana, A. (2009): "Information acquisition in auctions: Sealed bids vs. open bids." *Games and Economic Behavior*, 65, 375-405.

Kamenica, E., and Gentzkow, M. (2011): "Bayesian persuasion." *American Economic Review*, 101, 2590-2615.

Kyle, A.S. (1985): "Continuous auctions and insider trading." *Econometrica*, 53, 1315-1335.

Lehmann, E. (1988): "Comparing location experiments." *Annals of Statistics*, 16, 521-533.

Lipnowski, E., Mathevet, L., and Wei, D. (2020): "Attention management." *American Economic Review: Insights*, 2, 17-32.

Maskin, E., and Riley, J. (1984): "Optimal auctions with risk-averse buyers." *Econometrica*, 52, 1473-1518.

Matejka, P., and Mckay, A. (2015): "Rational inattention to discrete choices: A new foundation for the multinomial logit." *American Economic Review*, 105, 272-298.

Mensch, J. (2018): "Cardinal representations of information." Working paper.

—— (2021): "Rational inattention and the monotone likelihood ratio." *Journal of Economic Theory*, 196.

Milgrom, P., and Weber, R. (1982): "A theory of auctions and competitive bidding." *Econometrica*, 50, 1089-1122.

Morris, S., and Strack, P. (2017): "The Wald problem and the equivalence of sequential sampling and static information costs." Working paper.

Morris, S., and Yang, M. (2016): "Coordination and the relative cost of distinguishing nearby states." Working paper.

Mussa, M., and Rosen, S. (1978): "Monopoly and product quality." *Journal of Economic Theory*, 18, 301-317.

Myerson, R. (1981): "Optimal auction design." *Mathematics of Operations Research*, 6, 58-73.

Ostrovsky, M. (2012): "Information aggregation in dynamic markets with strategic traders." *Econometrica*, 80, 2595-2647.

Persico, N. (2000): "Information acquisition in auctions." *Econometrica*, 68, 135-148.

Pomatto, L., Strack, P., and Tamuz, O. (2019): "The cost of information." Working paper.

Ravid, D. (2020): "Bargaining with rational inattention." Working paper.

Ravid, D., Roesler, A. K., and Szentes, B. (2019): "Learning before trading: On the inefficiency of ignoring free information." Working paper.

Riley, J., and Zeckhauser, R. (1984): "Optimal selling strategies: When to haggle, when to hold firm." *Quarterly Journal of Economics*, 98, 267-289.

Roesler, A. K., and Szentes, B. (2017): "Buyer-optimal learning and monopoly pricing." *American Economic Review*, 107, 2072-2080.

Shannon, C. (1948): "A mathematical theory of communication." *Bell System Technical Journal*, 27.

Shi, X. (2012): "Optimal auctions with information acquisition." *Games and Economic Behavior*, 74, 666-686.

Sims, C. (1998): "Stickiness." *Carnegie-Rochester Conference Series on Public Policy*, 49, 317-356.

—— (2003): "Implications of rational inattention." *Journal of Monetary Economics*, 50, 665-690.

Sorokin, C., and Winter, E. (2019). "Pure information design in classic auctions." Working paper.

Szalay, D. (2009): "Contracts with endogenous information." *Games and Economic Behavior*, 65, 586-625.

Torgersen, E. (1991): *Comparison of Statistical Experiments.* Cambridge University Press: Cambridge, UK.

Yang, M. (2015): "Coordination with flexible information acquisition." *Journal of Economic Theory*, 158, 721-738.

—— (2020): "Optimality of debt under flexible information acquisition." *Review of Economic Studies*, 87, 487-536.

Appendix A: Discussion of Information Acquisition Costs

As mentioned in Section 7, there is a large literature on flexible information acquisition, in which it is generally assumed that the cost of information acquisition is posterior-separable. There are several reasons for doing so. First, it captures the idea that information acquisition is flexible, and that the cost of information is increasing in the Blackwell order. Second, it is relatively tractable, allowing for use of concavification techniques to solve for the decision maker's optimal choice for a given decision problem (e.g. Caplin and Dean, 2013).

It remains to provide a microeconomic foundation for this form of cost function. As discussed in De Oliveira et al. (2017), a cost function for information, in addition to being increasing in the Blackwell order, should satisfy the inequality

$$c(\alpha\tau + (1-\alpha)\tau') \leq \alpha c(\tau) + (1-\alpha)c(\tau') \tag{14}$$

for any two distributions τ, τ', since the agent can always randomize between two signals to achieve their convex combination. Representability via the functional form in Section 2 is equivalent to assuming that this inequality holds with equality (Torgersen, 1991).

On the one hand, (14) should hold with equality when considering a decision maker maximizing expected utility. The agent's preferences can be represented by the expected utility from actions conditional on the state, and so she would value these lotteries equally, as they provide equivalent information in expectation.[20] As a cost function, though, it is slightly less immediate why they should be equal, as physically, these are different experiments. Thus further justification is needed.

There are several possible answers. The first, going back to Sims (1998,

[20]It is fairly straightforward, by the supporting hyperplane theorem, to construct a decision problem for which the expected utility from the optimal choice conditional on belief μ is given by $-H(\mu)$.

2003) is that the cost is not one of acquiring, but of processing information. Suppose that all information is available in some database; the only difficulty is in accessing it. As has been known since Shannon (1948), the optimal way to encode information that is available through a limited channel (in terms of flow per unit of time) is given by informational entropy. As this functional form is posterior-separable, this environment would satisfy the assumptions of the model.

Pomatto et al. (2019) provide another justification. In addition to assuming that the cost of running two conditionally independent experiments is equal to the sum of each of their costs, they also provide an axiom closely related to (14) holding with equality: namely, that the cost of randomizing between experiment τ with probability α, and no information with probability $1-\alpha$, should be equal to that of a single experiment that generates the same distribution of posteriors as this randomization. They justify this through a scenario where the decision maker has access to a large number of independent draws of the same experiment, of which the decision maker can sample as many as she likes.[21] In this case, the decision maker could continue to draw until the experiment yielded some information, which in expectation would be $\frac{1}{\alpha}$ draws. Thus inequality (14) also goes in the other direction. Their representation yields a particular functional form of H based on log-likelihood ratios.

Further arguments in favor of such a representation can be made using revealed preference representations of decision makers with costly information acquisition. Per Caplin and Dean (2015), one can represent the decision maker as maximizing her utility subject to such a constraint if and only if it satisfies two properties. The first, "No Improving Action Switches" (NIAS), states that the decisions conditional on the resultant posteriors must be optimal (by revealed preference). The second, "No Improving Attention Cycles" (NIAC), states that the decision maker cannot improve her utility by redistributing attention across decision problems. This latter property rules out counterintuitive behavior such as acquiring more information when the stakes

[21] Indeed, a major justification for their axiomatization of their cost function is based on sampling problems. See also Morris and Strack (2017).

are lower.

There have been a couple of approaches to extend these revealed-preference results to a representation that is posterior-separable. This would allow for testing whether the cost is actually of this form. To this end, Denti (2019) strengthens NIAC to "No Improving Posterior Cycles," which states that the decision maker cannot improve her utility by redistributing attention not only across decision problems, but also within decision problems as well by redistributing between individual posteriors. Intuitively, since optimization occurs posterior by posterior, the problem admits a posterior-separable cost representation.

Another approach is that of Caplin et al. (2019), who provide axioms for different forms of cost representations. For the purposes of this paper, the relevant one is that of "posterior separability," as I consider a fixed prior, and do not consider the counterfactual of what would happen if the prior were to change.[22] While there are some additional technical assumptions that are needed, the main property that must be satisfied is their axiom, "Separability." Roughly, it states the following. Suppose some posteriors $\{\mu\}$ are optimal when the actions are chosen from some subset $A = \{a\}$ for some decision problem. If those posteriors are still feasible (by Bayes' rule) in combination with some other set of posteriors $\{\mu'\}$, there must be some other subset of actions $A' = \{a'\}$ such that the optimal attention strategy is to choose pairs (μ', a'), while pairing the same (μ, a) as before. This means that the optimality of posterior μ for a given action a does not depend on what other posteriors are chosen with positive probability, and so is a form of separability.

As a final remark, I claimed in Section 2 that all posteriors will be in the interior of the probability simplex if the slopes of the derivatives of H are sufficiently high. I formally state and prove a slightly stronger version of this now, which states that this bound away from the boundary is uniform for all decision problems where the maximum payoff difference is less than a given

[22]Some of the pitfalls associated with this are discussed in Mensch (2018). In particular, if one thinks of information costs as coming from the physical cost of a single experiment, the representation H *cannot* remain the same across priors.

bound.

Suppose that a decision maker faces compact choice set $X \subset [0,1]$ and state space Θ. Her ex-post utility is $u(x,\theta)$, and her cost of information acquisition is given by the posterior-separable representation from H as given in Section 2.

Lemma A: *For any $C > 0$, there exist $\epsilon, \kappa > 0$ such that if*

$$\max_{x \in X}[\max_{\theta \in \Theta} u(x,\theta) - \min_{\hat{\theta} \in \Theta} u(x,\hat{\theta})] \leq C$$

and

$$c(\tau) = \kappa(H(\mu_0) - E_\tau[H(\mu)])$$

then for all posteriors μ in the support of those chosen by the decision maker,[23] *$\mu(\theta) \geq \epsilon$.*

Proof: Suppose that there is no such $\epsilon > 0$. Then for every $\epsilon > 0$ and some θ, one can consider the sets $M_1 \equiv \{\mu : \mu(\theta) < \epsilon\}$ and $M_2 \equiv \{\mu : \mu(\theta) > \mu_0(\theta)\}$. Both sets must have positive measure under τ by Bayes' rule. Moreover, the Euclidean distance between any posteriors in M_1 and any in M_2 must be at least some $d > \mu_0(\theta) - \epsilon$. By the fundamental theorem of calculus and the intermediate value theorem,

$$\frac{\partial H}{\partial \mu(\theta')}(\mu_1) - \frac{\partial H}{\partial \mu(\theta')}(\mu_2) = \int_0^1 \sum_{\theta'' \in \Theta} \frac{\partial^2 H}{\partial \mu(\theta')\partial \mu(\theta'')}(\alpha\mu_1 + (1-\alpha)\mu_2)\frac{\mu_1(\theta'') - \mu_2(\theta'')}{|\mu_1 - \mu_2|}d\alpha$$

$$= \sum_{\theta'' \in \Theta} \frac{\partial^2 H}{\partial \mu(\theta')\partial \mu(\theta'')}(\hat{\mu})\frac{\mu_1(\theta'') - \mu_2(\theta'')}{|\mu_1 - \mu_2|}$$

for some $\hat{\mu} \equiv \alpha\mu_1 + (1-\alpha)\mu_2$, for some $\alpha \in [0,1]$.

Now consider the Bayes-plausible change where, while actions are kept fixed, beliefs in the former set are moved closer to those in the latter, and vice versa, as defined as follows. Let σ_1 and σ_2 be the probability measures over x

[23] Caplin and Dean (2013) prove a similar yet slightly weaker result in demonstrating their Theorem 1, showing that posteriors must lie on the interior if the partial derivatives of H are unbounded at the boundaries of the probability simplex.

defined by the pushforward measures from $\mathbf{x}(\mu)$ and τ for the sets M_1 and M_2, respectively. Consider the probability integral transforms from σ_1 and σ_2 to the uniform distribution over $s \in [0,1]$, and let ν_1 and ν_2 be the mappings back to the respective original beliefs $\mu_1(\cdot|x)$ and $\mu_2(\cdot|x)$. Then for each $s \in [0,1]$, consider changes in the beliefs in M_1 in the direction of $\nu_2(\cdot|s) - \nu_1(\cdot|s)$, and for M_2 in the direction of $\nu_1(\cdot|s) - \nu_2(\cdot|s)$. The marginal decrease in cost for this is given by the Fréchet derivatives

$$\tau(M_1)\int_0^1 \sum_{\theta'\in\Theta} \kappa \frac{\partial H}{\partial\mu(\theta')}(\nu_1(\cdot|s))[\nu_2(\theta'|s) - \nu_1(\theta'|s)]ds$$

$$+\frac{\tau(M_1)}{\tau(M_2)}\tau(M_2)\int_0^1 \sum_{\theta'\in\Theta} \kappa \frac{\partial H}{\partial\mu(\theta')}(\nu_2(\cdot|s))[\nu_1(\theta'|s) - \nu_2(\theta'|s)]ds$$

$$= \tau(M_1)\int_0^1 \sum_{\theta',\theta''\in\Theta} \kappa \frac{\partial^2 H}{\partial\mu(\theta')\partial\mu(\theta'')}(\hat{\nu}(\cdot|s))\frac{\nu_1(\theta''|s) - \nu_2(\theta''|s)}{|\nu_1(\cdot|s) - \nu_2(\cdot|s)|}[\nu_2(\theta'|s) - \nu_1(\theta'|s)]ds$$

$$\geq \tau(M_1)\kappa m d > 0$$

where the coefficient in front of the second integral is due to the changes in beliefs being scaled by $\frac{\tau(M_1)}{\tau(M_2)}$ relative to those in the first, the first inequality originates from the distance between any two elements in M_1 and M_2 as discussed above, and the second inequality originates from the strong convexity of H, i.e. $\mathbf{H} + mI$ is negative semi-definite for some $m > 0$. Meanwhile, the marginal loss in the payoff due to the changes in conditional choice are bounded by $2C\tau(M_1)$. Hence, for κ sufficiently large, the decrease in cost will be larger than the decrease in payoff, and hence is an improvement for the decision maker. □

Appendix B: Proofs

Proof of Lemma 1: The first step, analogous to Myerson (1981), establishes that given the information acquisition of the buyers, it is sufficient for them to report their posteriors. Let Y be the action space in $\mathcal{M}$, and τ be the distribution of posteriors that the buyer acquires in equilibrium. For each $\mu \in \text{supp}(\tau)$, the buyer will choose some strategy $\xi : \mu \to \Delta(Y)$. Let $\hat{\mathbf{x}}(\xi(\mu_1), ..., \xi(\mu_N))$ be the vector of probabilities that buyers receive the item by playing according to strategy ξ; similarly, define $\hat{\mathbf{t}}(\xi(\mu_1), ..., \xi(\mu_N))$ to be the vector of expected transfers. One can then define the direct revelation mechanism $\mathcal{M}'$ where each buyer reports her posterior μ_i, and the probabilities of receiving the item and transfers are given by

$$\mathbf{x}(\mu_1, ..., \mu_N) = \hat{\mathbf{x}}(\xi(\mu_1), ..., \xi(\mu_N))$$

$$\mathbf{x}(\mu_1, ..., \mu_N) = \hat{\mathbf{x}}(\xi(\mu_1), ..., \xi(\mu_N))$$

Hence each buyer receives the same expected utility as in $\mathcal{M}$ for each possible report of posterior; since ξ was an equilibrium strategy in $\mathcal{M}$, it is optimal in $\mathcal{M}'$ to report one's true posterior.

Similarly, any distribution of posteriors τ' will yield a weakly lower payoff than τ, as the same set of payoffs is feasible in $\mathcal{M}'$ as from acquiring τ' in mechanism $\mathcal{M}$ and then choosing $\xi(\mu)$ for each $\mu \in \text{supp}(\tau')$. Hence it will be optimal to acquire τ in $\mathcal{M}'$.

The above shows that it is without loss to consider mechanisms in which the seller recommends that the buyer acquire τ, and report their posterior μ; there will then be a unique x for each reported μ. It is also clear that for each x, there must be a unique t, since otherwise the buyer could misreport her type μ in order to get a lower t. To complete the proof, one must show that for each x, there is a unique $\mu \in \text{supp}(\tau)$ that receives the item with probability x. Suppose otherwise; let $1_x(s)$ be the indicator function on the signal space that takes the value 1 if, upon receiving signal s, the buyer receives the item with probability x, and 0 otherwise. This is a measurable function with respect to

π, and so the buyer's ex-ante payoff is given by

$$\sum_{\theta\in\Theta}\int_{\mathcal{S}}\int_0^1 (x\theta - \mathbf{t}(x))1_x(s)\mu_0(\theta)dxd\pi(s|\theta) - H(\mu_0) + \int_{\Delta(\Theta)} H(\mu)d\tau(\mu)$$

where $\mathbf{t}(x)$ is the transfer associated with x. If the set of signal realizations for which the same x is chosen is of measure greater than 0 with respect to π, then there exists $\hat{\pi}$ in which all signal realizations s for which x is chosen are merged into one signal $\hat{s}$, upon whose reception the buyer again chooses x. If $\mu(\cdot|s)$ is not the same almost everywhere for all such s, then the cost of information acquisition is strictly lower, and hence an improvement for the buyer. Hence it is without loss that there is a unique μ for which x is chosen almost everywhere. $\square$

Proof of Lemma 2: To see that (IR-A) is implied by the other constraints, let $\underline{x}^* \equiv \min\{x \in X\}$. By standard single-crossing arguments from (IC-I), $E_{\mu(\cdot|x)}[\theta]$ is increasing in x. Thus, for all $x \in X$,

$$\underline{x}^* E_{\mu(\cdot|x)}[\theta] - \mathbf{t}(\underline{x}^*) \geq \underline{x}^* E_{\mu(\cdot|\underline{x}^*)}[\theta] - \mathbf{t}(\underline{x}^*) \geq 0$$

Furthermore, the buyer can acquire no information, which is costless. Therefore, by (IC-I),

$$\begin{aligned}\int\int [\mathbf{x}(\mu)\theta - \mathbf{t}(\mathbf{x}(\mu))]d\mu(\theta)d\tau(\mu) - [H(\mu_0) - \int H(\mu)d\tau(\mu)] \geq \\ \int\int [\underline{x}^*\theta - \mathbf{t}(\underline{x}^*)]d\mu(\theta)d\tau(\mu) - [H(\mu_0) - H(\mu_0)] \\ = \int [\underline{x}^*\theta - \mathbf{t}(\underline{x}^*)]d\mu_0(\theta) \\ \geq 0\end{aligned}$$

To see that (IC-I) is implied by (IC-A), suppose that for some subset of allocations $Y = \{x\}$ that are recommended with positive probability according to π, there is some action $\hat{\mathbf{x}}(x)$ that the buyer strictly prefers, i.e.

$$\sum_{\Theta}\int_{Y}[\hat{\mathbf{x}}(x)\theta-\mathbf{t}(\hat{\mathbf{x}}(x))]\mu_0(\theta)d\pi(x|\theta)>\sum_{\Theta}\int_{Y}[x\theta-\mathbf{t}(x)]\mu_0(\theta)d\pi(x|\theta)$$

This same ex-interim payoff could be achieved by using the recommendation strategy $\hat{\pi}(x|\theta)$ where, instead of recommending x, $\hat{\mathbf{x}}(x)$ is recommended, i.e.

$$d\hat{\pi}(x|\theta)=\begin{cases}0, & x\in Y\\ d\pi(x|\theta)+\int_{y\in Y:\hat{\mathbf{x}}(y)=x}d\pi(y|\theta), & x\notin Y\end{cases}$$

Moreover, since H is concave, the information cost is reduced because the buyer no longer distinguishes between the cases where x was recommended and $\{y\in Y:\hat{\mathbf{x}}(y)=x\}$ was recommended, and instead generates a single posterior that is the weighted average (according to τ) of $\mu(\cdot|x)$ and $\{\mu(\cdot|y):\hat{\mathbf{x}}(y)=x\}$. Thus the buyer could improve her expected payoff at least as much by an ex-ante deviation. $\square$

Proof of Lemma 3: By Lemma A, $\mu(\theta|x)>\epsilon,\forall\theta,x$. Hence $H(\mu(\cdot|x))$ and $\frac{\partial H}{\partial\mu(\theta)}(\mu(\cdot|x))$ are bounded. By Bayes' rule and Fubini's theorem, the buyer's objective can be written as the linear operator of $\pi(\cdot|\theta)$,

$$F(\pi)\equiv\sum_{\theta\in\Theta}\int_X[x\theta-\mathbf{t}(x)+H(\frac{d\pi(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta}d\pi(x|\theta')\mu_0(\theta')})]d\pi(x|\theta)\mu_0(\theta)\tag{15}$$

where $\frac{d\pi(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta}d\pi(x|\theta')\mu_0(\theta')}$ is the Radon-Nikodym derivative of the measure $d\pi(x|\theta)\mu_0(\theta)$ with respect to $\sum_{\theta'\in\Theta}d\pi(x|\theta')\mu_0(\theta')$. By assumption, since π is a valid signal (i.e. it generates posteriors via Bayes' rule), the measures $\{\pi(\cdot|\theta)\}_{\theta\in\Theta}$ are mutually absolutely continuous and so this Radon-Nikodym derivative is well defined.

Consider the set of finite signed measures $\{\{\hat{\pi}(\cdot|\theta)\}_{\theta\in\Theta}\}$ that are absolutely continuous with respect to π, and endow it with the norm

$$\|\hat{\pi}\|=[\sum_{\theta\in\Theta}\int(\frac{d\hat{\pi}(x|\theta)}{d\pi(x|\theta)})^2d\pi(x|\theta)\mu_0(\theta))]^{(\frac{1}{2})}$$

Thus $\{\{\hat{\pi}(\cdot|\theta)\}_{\theta\in\Theta}\}$ constitutes a normed vector space. Of particular interest are those $\hat{\pi}$ such that $\hat{\pi}(\cdot|\theta)$ is a conditional probability measure. For such $\hat{\pi}$, consider the vector $\epsilon(\hat{\pi}-\pi)$. As the linear operator

$$A(x,\theta) = x\theta - \mathbf{t}(x) + h(x,\theta)$$

is bounded, in the limit,

$$\lim_{\epsilon\to 0}\frac{1}{\epsilon\|\hat{\pi}-\pi\|}[F(\pi+\epsilon(\hat{\pi}-\pi)) - F(\pi) - \epsilon\sum_{\theta\in\Theta}\int_X A(x,\theta)d(\hat{\pi}-\pi)(x|\theta)\mu_0(\theta)] = 0$$

and so F is Fréchet differentiable. Hence in order to be optimal, one must have that for all conditional probability measures $\hat{\pi}$,

$$\sum_{\theta\in\Theta}\int_X A(x,\theta)d(\hat{\pi}-\pi)(x|\theta)\mu_0(\theta) = 0$$

and so $A(x,\theta) = A(x',\theta)$ almost everywhere with respect to π. Thus (3) is necessary.

For the sufficiency of (3), suppose that π is suboptimal, and that instead some $\hat{\pi}$ is better for the buyer. First, the conditional distribution $\hat{\mu}(\cdot|x)$ must be weak* continuous with respect to x almost everywhere: suppose not, and that there exists some point x^* around which there exists $\epsilon > 0$ such that, for every $\delta > 0$, the open ball $B_\delta(x^*)$ contains two subsets of positive measure $X_1^\epsilon, X_2^\epsilon$ such that $|\mu(\cdot|x_1) - \mu(\cdot|x_2)| > \epsilon$, for all $x_i \in X_i^\epsilon$, respectively. Then for sufficiently small δ, the alternative signal that recommends x^* instead of any other $x \in B_\delta(x^*)$ will be an improvement, as the information cost will be strictly lower by the strong concavity of H, while by the compactness of $\mathcal{M}$, the loss from recommending x^* instead vanishes as $\delta \to 0$. That is, indicating this alternative recommendation by $\tilde{\pi}_\delta$, for small enough δ,

$$\sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\tilde{\pi}_\delta(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\tilde{\pi}_\delta(x|\theta')\mu_0(\theta')})]d\tilde{\pi}_\delta(x|\theta)\mu_0(\theta)$$

$$-\sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\hat{\pi}(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\hat{\pi}(x|\theta')\mu_0(\theta')})]d\hat{\pi}(x|\theta)\mu_0(\theta)$$

$$= \sum_{\theta\in\Theta} \hat{\pi}(B_\delta(x^*)|\theta)[x\theta - \mathbf{t}(x) + H(\frac{\int_{B_\delta(x^*)} d\tilde{\pi}_\delta(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta}\int_{B_\delta(x^*)} d\tilde{\pi}_\delta(x|\theta')\mu_0(\theta')})]$$

$$-\sum_{\theta\in\Theta}\int_{B_\delta(x^*)} [x\theta - \mathbf{t}(x) + H(\frac{d\hat{\pi}(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\hat{\pi}(x|\theta')\mu_0(\theta')})]d\hat{\pi}(x|\theta)\mu_0(\theta)$$

$$> 0$$

Next, consider the case where $\hat{\pi}$ is absolutely continuous with respect to π. For any $\alpha \in (0,1)$, consider the conditional probability measures $(1-\alpha)\pi+\alpha\hat{\pi}$. This will also be an improvement for the buyer over π, since

$$\sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\pi(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\pi(x|\theta')\mu_0(\theta')})]d\pi(x|\theta)\mu_0(\theta) \tag{16}$$

$$< (1-\alpha)\sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\pi(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\pi(x|\theta')\mu_0(\theta')})]d\pi(x|\theta)\mu_0(\theta)$$

$$+\alpha\sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\hat{\pi}(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\hat{\pi}(x|\theta')\mu_0(\theta')})]d\hat{\pi}(x|\theta)\mu_0(\theta)$$

$$\leq \sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{((1-\alpha)d\pi + \alpha d\hat{\pi})(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta}((1-\alpha)d\pi + \alpha d\hat{\pi})(x|\theta')\mu_0(\theta')})]((1-\alpha)d\pi + \alpha d\hat{\pi})(x|\theta)\mu_0(\theta) \tag{17}$$

where the second inequality is from merging recommendations of the same x, and the fact that $\pi \neq \hat{\pi}$ and H is concave. Subtracting (16) from (17), dividing by α, and taking the limit as $\alpha \to 0$, this becomes the Fréchet derivative as above in the direction of $\hat{\pi} - \pi$:

$$0 < \sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + h(x,\theta)](d\hat{\pi} - d\pi)(x|\theta)\mu_0(\theta)$$

yielding that for some positive measure of x with respect to π and some positive measure of $\hat{x}$ with respect to both $\pi, \hat{\pi}$,

$$\sum_{\theta\in\Theta}[x\theta - \mathbf{t}(x) + h(x,\theta)] < \sum_{\theta\in\Theta}[\hat{x}\theta - \mathbf{t}(\hat{x}) + h(\hat{x},\theta)]$$

and so, for some θ,

$$x\theta - \mathbf{t}(x) + h(x,\theta) < \hat{x}\theta - \mathbf{t}(\hat{x}) + h(\hat{x},\theta)$$

contradiction.

Now suppose that $\hat{\pi}$ is singular with respect to π. Since π is a recommendation strategy, for any $x \in X$, the open ball of radius ϵ has measure $\pi(B_\epsilon(x)|\theta) > 0$. Then construct the alternative measure $\hat{\pi}_\epsilon$ defined by partitioning $[0,1]$ into intervals I of length between $\epsilon/2$ and ϵ whose endpoints are not mass points of $\hat{\pi}$, and set, for all $x \in I$,

$$d\hat{\pi}_\epsilon(x|\theta) = \frac{\int_{I\cap X} d\hat{\pi}(\hat{x}|\theta)}{\int_{I\cap X} d\pi(\hat{x}|\theta)} d\pi(x|\theta)$$

Clearly, $\hat{\pi}_\epsilon$ is absolutely continuous with respect to π. By the compactness of $\mathcal{M}$ and the Portmanteau theorem,

$$\lim_{\epsilon\to 0}\sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\hat{\pi}_\epsilon(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\hat{\pi}_\epsilon(x|\theta')\mu_0(\theta')})] d\hat{\pi}_\epsilon(x|\theta)\mu_0(\theta)$$

$$\geq \lim_{\epsilon\to 0}\sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\hat{\pi}(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\hat{\pi}(x|\theta')\mu_0(\theta')})] d\hat{\pi}_\epsilon(x|\theta)\mu_0(\theta)$$

$$= \sum_{\theta\in\Theta}\int_X [x\theta - \mathbf{t}(x) + H(\frac{d\hat{\pi}(x|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} d\hat{\pi}(x|\theta')\mu_0(\theta')})] d\hat{\pi}(x|\theta)\mu_0(\theta)$$

But for low enough ϵ, that would mean that $\hat{\pi}_\epsilon$ is also better than π, which we saw was impossible for any measure that is absolutely continuous with respect to π. $\square$

Proof of Lemma 4: I define a system of partial differential equations defining the motion of $(x, \mathbf{t}(x), \mu(\cdot|x))$, and show that they have a unique solution.

I then verify that the necessary and sufficient conditions of Lemma 3 are satisfied.

I start by deriving a differentiable law of motion that satisfies (3), which will be used to show sufficiency. Thus I show that there exists a differentiable locus of points on which the buyer's choice has its support; one can then convert it to a mechanism in recommendation strategies by dropping the values of x that are not in the support, and invoking Lemma 3 on the remaining values of x to verify that it is optimal for the buyer. First, to define $\mathbf{t}'(x)$, any solution that is optimal for the buyer must satisfy (IC-I). It is well known from Myerson (1981) that in order to do so,

$$\lim_{\epsilon \to 0} \frac{\mathbf{t}(x+\epsilon) - \mathbf{t}(x)}{\epsilon} = E_{\mu(\cdot|x)}[\theta] \tag{18}$$

So, one can define

$$\frac{\partial h}{\partial x}(x,\theta) \equiv \lim_{\epsilon \to 0} \frac{h(x+\epsilon,\theta) - h(x,\theta)}{\epsilon} = E_{\mu(\cdot|x)}[\theta] - \theta \tag{19}$$

This implicitly defines the law of motion of beliefs from $\mu(\cdot|x)$. By (2), for $\mu(\cdot|x)$ to be differentiable,

$$\begin{aligned}\frac{\partial h}{\partial x}(x,\theta) = & \sum_{\theta'' \in \Theta} \frac{\partial^2 H}{\partial \mu(\theta'')\partial \mu(\theta)}(\mu(\cdot|x))\frac{\partial \mu}{\partial x}(\theta''|x)(1-\mu(\theta|x)) \\ & - \sum_{\theta'' \in \Theta}\sum_{\theta' \neq \theta} \frac{\partial^2 H}{\partial \mu(\theta'')\partial \mu(\theta)}(\mu(\cdot|x))\frac{\partial \mu}{\partial x}(\theta''|x)\mu(\theta'|x)\end{aligned} \tag{20}$$

Thus, for any constant $C_{\mu(\cdot|x)}$,

$$\sum_{\theta' \in \Theta} \frac{\partial^2 H}{\partial \mu(\theta')\partial \mu(\theta)}(\mu(\cdot|x))\frac{\partial \mu}{\partial x}(\theta'|x) = -(\theta + C_{\mu(\cdot|x)}), \forall \theta \tag{21}$$

is a solution to (20), as by plugging these values into (18), (19) is satisfied.

Since H is strongly concave, the Hessian $\mathbf{H}(\mu(\cdot|x))$ is negative definite, and so

$$\begin{pmatrix} \frac{\partial \mu}{\partial x}(\theta_1|x) \\ \vdots \\ \frac{\partial \mu}{\partial x}(\theta_K|x) \end{pmatrix} = -\mathbf{H}^{-1}(\mu(\cdot|x)) \begin{pmatrix} \theta_1 + C_{\mu(\cdot|x)} \\ \vdots \\ \theta_K + C_{\mu(\cdot|x)} \end{pmatrix} \tag{22}$$

Lastly, in order to be a probability distribution, $\sum_{\theta\in\Theta} \frac{\partial \mu}{\partial x}(\theta|x) = 0$, which means that, indicating the $(i,j)^{th}$ entry of $\mathbf{H}^{-1}$ by $\mathbf{H}^{-1}_{(i,j)}$,

$$C_{\mu(\cdot|x)} = -\frac{\sum_{i=1}^{K}\sum_{j=1}^{K} \theta_j \mathbf{H}^{-1}_{(i,j)}(\mu(\cdot|x))}{\sum_{i=1}^{K}\sum_{j=1}^{K} \mathbf{H}^{-1}_{(i,j)}(\mu(\cdot|x))} \tag{23}$$

It now remains to be shown that the system of differential equations defined by (18) and (22) has a solution, in order to demonstrate that the assumption of differentiability yields a valid solution. Since H is twice Lipschitz continuously differentiable and strongly concave, $\mathbf{H}(\mu)$ is Lipschitz continuous in μ and bounded away from 0, and so $\mathbf{H}^{-1}$ is Lipschitz continuous as well. Lastly, by (23), $C_{\mu(\cdot|x)}$ is defined by the ratio of Lipschitz continuous functions, and so C_μ is itself Lipschitz continuous in μ. By the Picard-Lindelöf theorem (Coddington and Levinson, Theorem 5.1), there exists an interval $[x-a, x+b]$ on which the system $(x, \mathbf{t}(x), \mu(\cdot|x))$ has a unique solution.

By the fundamental theorem of calculus, it then follows that (3) is satisfied for all pairs $x, x' \in [x-a, x+b]$. Hence any distribution τ over $\{\mu(\cdot|x) : x \in [x-a, x+b]\}$ is optimal for the buyer given prior $\mu_0 = \int d\tau(\mu(\cdot|x))$ by Lemma 3, and so (18) and (22) are sufficient for (IC-A) to be satisfied, with

$$\frac{\partial}{\partial x}\{E_{\mu(\cdot|x)}[\theta]\} = -\sum_{\theta,\theta'\in\theta}[\frac{\partial^2 H}{\partial\mu(\theta)\partial\mu(\theta')}(\mu(\cdot|x))]\frac{\partial\mu}{\partial x}(\theta'|x)\frac{\partial\mu}{\partial x}(\theta|x) > 0 \tag{24}$$

as is easily derived from (18) and (22), which is positive due to the negative-definiteness of the Hessian matrix.[24]

To see that one can set $[x-a, x+b] = [0,1]$, suppose that the maximal such value of a were less than x. Beliefs $\mu(\cdot|x-a)$ must still be in the interior of the simplex by Lemma A since $x+b-\mathbf{t}(x+b)-(x-a)+\mathbf{t}(x-a) \leq b-a+\max\{\theta \in \Theta\}$. Thus, the conditions of the Picard-Lindelöf theorem are still satisfied, and so this cannot be the supremum. The same reasoning applies to b.

For necessity, one must show that any incentive-compatible solution to the buyer's problem must be identical to that given above. To do so, fix x^*, and suppose that there exists $\hat{\tau}$ that places positive measure, for some subset of allocations $\{x\}$, on beliefs $(\hat{\mathbf{t}}(x), \hat{\mu}(\cdot|x)) \neq (\mathbf{t}(x), \mu(\cdot|x))$, where the beliefs on the right-hand side are those derived from (18) and (22). Consider the distribution $\tilde{\tau}$ over $\{\mu(\cdot|x)\}$ whose pushforward measure over $x \in [0,1]$ is uniform. Then, by Lemma 3, $\alpha\hat{\tau} + (1-\alpha)\tilde{\tau}$ is optimal for the buyer for any $\alpha \in (0,1)$ given prior $\tilde{\mu}_0 = \alpha\mu_0 + \int_{\{\mu(\cdot|x)\}} d\tilde{\tau}(\mu(\cdot|x))$. It is immediate that in order to satisfy (IC-I), the transfers conditional on x must be the same under the mechanisms that generate $\hat{\tau}$ and $\tilde{\tau}$, respectively. Thus, by (2) and (3),

$$H(\hat{\mu}(\cdot|x)) + \frac{\partial H}{\partial \mu(\theta)}(\hat{\mu}(\cdot|x))(1-\hat{\mu}(\theta|x)) - \sum_{\theta' \neq \theta} \frac{\partial H}{\partial \mu(\theta')}(\hat{\mu}(\cdot|x))\hat{\mu}(\theta'|x)$$

$$= H(\mu(\cdot|x)) + \frac{\partial H}{\partial \mu(\theta)}(\mu(\cdot|x))(1-\mu(\theta|x)) - \sum_{\theta' \neq \theta} \frac{\partial H}{\partial \mu(\theta')}(\mu(\cdot|x))\mu(\theta'|x) \quad (25)$$

Multiplying the above by $\hat{\mu}(\theta|x)$ and $\mu(\theta|x)$, then summing over $\theta \in \Theta$ and taking the difference between the former and the latter, one gets

$$\sum_{\theta \in \Theta} (\frac{\partial H}{\partial \mu(\theta)}(\mu(\cdot|x)) - \frac{\partial H}{\partial \mu(\theta)}(\hat{\mu}(\cdot|x)))(\mu(\theta|x) - \hat{\mu}(\theta|x)) = 0 \quad (26)$$

By the intermediate value theorem, there exists some $\alpha \in [0,1]$ such that for

[24] As remarked in the discussion following Lemma 3, any set of triplets $(x, \mathbf{t}(x), \mu(\cdot|x))$ that satisfies (3) and on which τ has its support is incentive compatible, and so the monotonicity of $E_{\mu(\cdot|x)}[\theta]$ is implied anyway.

$\tilde{\mu} \equiv \alpha\mu(\cdot|x) + (1-\alpha)\hat{\mu}(\cdot|x)$,

$$\frac{\partial H}{\partial \mu(\theta)}(\mu(\cdot|x)) - \frac{\partial H}{\partial \mu(\theta)}(\hat{\mu}(\cdot|x)) = \sum_{\theta' \in \Theta} \frac{\partial^2 H}{\partial \mu(\theta)\partial \mu(\theta')}(\tilde{\mu})(\mu(\theta'|x) - \hat{\mu}(\theta'|x)) \quad (27)$$

Combining (26) and (27), one gets

$$\sum_{\theta \in \Theta}\sum_{\theta' \in \Theta} \frac{\partial^2 H}{\partial \mu(\theta)\partial \mu(\theta')}(\tilde{\mu})(\mu(\theta'|x) - \hat{\mu}(\theta'|x))(\mu(\theta|x) - \hat{\mu}(\theta|x)) = 0$$

But by the negative-definiteness of $\mathbf{H}$, the left-hand side must be negative, contradiction. □

Proof of Theorem 1: By Lemmas 3 and 4, the contour mechanism satisfies (IC-A) and (IC-I). Since $t(0) < 0$ and (IC-I) is satisfied, (IR-I) is satisfied by standard arguments (e.g. Myerson, 1981). Lastly, (IR-I) implies (IR-A) by revealed preference:

$$\tau \in \arg\max_{\sigma \in \Delta(\Delta(\Theta))} \int\int [x(\mu)\theta - t(x(\mu))] d\mu(\theta) d\sigma(\mu) - [H(\mu_0) - \int H(\mu) d\sigma(\mu)]$$

$$\Longrightarrow \int\int [x(\mu)\theta - t(x(\mu))] d\mu(\theta) d\tau(\mu) - [H(\mu_0) - \int H(\mu) d\tau(\mu)] \geq -t(0)$$

$$\geq 0$$

Hence all four constraints are satisfied. □

Proof of Proposition 2: Immediate from (18) and (22) defining an autonomous system of differential equations. □

Proof of Theorem 3: I first establish that an optimal mechanism exists. It is clear that any contour mechanism's revenue can be increased if $\mathbf{t}(0) < 0$, and so it is without loss of optimality to restrict attention to ones with $\mathbf{t}(0) = 0$. Within this set, let $\{\mathcal{C}_m\}_{m=1}^{\infty}$ be a sequence of such contour mechanisms, and let τ_m be the corresponding distributions over posteriors. By Lemma A, there exists $\epsilon > 0$ such that for all m, $\mu(\theta|x) \geq \epsilon$. As shown in the proof of Lemma 4 in equations (18) and (22), the functions $\mathbf{t}'(x)$ and $\frac{\partial \mu}{\partial x}(\cdot|x)$ are Lipschitz continuous on any compact set in the interior of the simplex, no matter what $\mu(\cdot|x)$ is, and so $\{\mathbf{t}_m\}$ and $\{\mu_m(\cdot|x)\}$ are equi-Lipschitz continuous. Therefore, by the Arzelà-Ascoli theorem, there exists a subsequence of $\{(\mathcal{C}_m, \tau_m)\}_{m=1}^{\infty}$ such that $\mathcal{C}_m \to \mathcal{C}$ uniformly and $\tau_m \to \tau$ in the weak* topology, with support within the same compact set. By Coddington and Levinson, Theorem 7.1, the solutions of differential equations for a sequence of starting points converge uniformly to a solution of the differential equations for the limit point as well, so the limit values of $(\mathbf{t}(x), \mu(\cdot|x))$ in $\mathcal{C}$ satisfy (3). Therefore τ is an incentive-compatible distribution by Lemma 3. This implies that the set of feasible payoffs to the seller is compact, and so a maximum exists.

Given the existence of an optimal mechanism, it follows that by Theorem 1, any implementable mechanism can be expressed by some $\mathcal{C}$. As $v_{\mathcal{C}}(\mu) = -\infty$ for all μ not contained in $\mathcal{C}$, the support of $co(v_{\mathcal{C}})$ must be contained in $\mathcal{C}$ with probability 1. Hence the mechanism satisfying (8) is optimal if and only if it is optimal overall. That $\mathbf{t}(0) = 0$ follows from being able to increase $\mathbf{t}(x)$ by some $\epsilon > 0$ without violating either (IC-A) or (IR-I) for $\underline{\mu}$ otherwise. □

Proof of Corollary 4: This follows immediately from Kamenica and Gentzkow (2011, Proposition 4 in their Online Appendix). □

Proof of Proposition 5: Suppose that, given $\mathcal{C}$, some τ is optimal such that $x^* \equiv \sup\{x : \exists \mu \in \operatorname{supp}(\tau) : \mathbf{x}(\mu) = x\} < 1$. Then the mechanism $\hat{\mathcal{C}}$ in which $1 - x^*$ is added to all values of $x \leq x^*$, and all triplets corresponding to $x > x^*$ are excluded, also satisfies (3). Thus τ remains optimal, where the choice of x under $\hat{\mathcal{C}}$, $\hat{\mathbf{x}}(\mu)$ equals $\mathbf{x}(\mu) + 1 - x^*$, and $\mathbf{t}(x) = \hat{\mathbf{t}}(x)$, by Proposition 2. By Lemma 4, one can then complete $\hat{\mathcal{C}}$ to apply to values of $x < 1 - x^*$. Since, by (18), $\hat{\mathbf{t}}'(x) > 0$, one can then increase $\hat{\mathbf{t}}$ by $\int_0^{1-x^*} \hat{\mathbf{t}}'(x)dx$ while maintaining (3). □

Proof of Theorem 6: For each choice of $\mathcal{C}$, there will either be as much information revelation as possible in the case of convex $\tilde{\mathbf{t}}$, or none in the case of concave $\tilde{\mathbf{t}}$, by Kamenica and Gentzkow (2011, Proposition 2). Thus it must also be true for the optimal $\mathcal{C}$. □

Proof of Lemma 5: Fix τ, and suppose that it is not of the form described in the statement of the lemma. The first step is to show that there is a mean-preserving spread of this form. With binary states, one can rewrite (12) as

$$\int_{\hat{\mu}}^{1} \mathbf{x}(\mu) d\tau(\mu) = \frac{1 - [1 - \tau(\mu < \hat{\mu})]^N}{N}$$

Differentiating this when it holds with equality, one gets

$$-\mathbf{x}(\hat{\mu}) d\tau(\hat{\mu}) = -[\tau(\mu < \hat{\mu})]^{N-1} d\tau(\hat{\mu})$$

$$\implies \tau(\mu < \hat{\mu}) = [\mathbf{x}(\hat{\mu})]^{\frac{1}{N-1}}$$

$$\implies d\tau(\mu) = \frac{1}{N-1}[\mathbf{x}(\mu)]^{\frac{1}{N-1}-1} \mathbf{x}'(\mu) d\mu \tag{28}$$

with boundary condition $\tau(\mu \leq \bar{\mu}) = 1$, where $\bar{\mu} \equiv \bar{\mu}$. Let

$$\mu^* \equiv \inf\{\hat{\mu} : \tau(\mu < \tilde{\mu}) = [\mathbf{x}(\tilde{\mu})]^{\frac{1}{N-1}}, \forall \tilde{\mu} > \hat{\mu}\}$$

Note that (28) does not depend on the exact distribution below μ. Thus, to find a mean-preserving spread, one need only consider the distribution between $\underline{\mu}$ and μ^*.

I show that for any other τ satisfying (12), there exists a mean-preserving spread that satisfies (12); by Zorn's lemma, there will then be a maximal element, that must be of the form of the lemma. First, suppose that there is an atom at some $\mu_* \in (\underline{\mu}, \mu^*)$. Then there exists $\delta > 0$ such that for sufficiently small ϵ, (12) does not hold with equality at $\hat{\mu}, \forall \hat{\mu} \in (\mu_*, \mu_* + \epsilon)$ or else (12) would be violated at μ_*. Moreover,

$$\lim_{\epsilon \to 0} \tau(\mu \in (\mu_* - \epsilon, \mu_* + \epsilon)) = \tau(\mu_*)$$

Consider the following mean-preserving spread: replace τ by $\hat{\tau}^\epsilon$ which, for all $\mu \in [\mu_* - \epsilon^2, \mu_* + \epsilon]$, asigns all mass to $\{\mu_* - \epsilon^2, \mu_* + \epsilon\}$, while preserving $E_{\hat{\tau}^\epsilon}[\mu] = \mu_0$. By Bayes' rule,

$$\lim_{\epsilon \to 0} \tau(\mu \in [\underline{\mu}, \mu_* - \epsilon^2]) + \frac{1}{1+\epsilon}\tau(\mu_*) \leq \lim_{\epsilon \to 0} \hat{\tau}^\epsilon(\mu < \mu_* + \epsilon) \leq \lim_{\epsilon \to 0} \tau(\mu \in [\underline{\mu}, \mu_* + \epsilon) \backslash \{\mu_*\}) + \frac{1}{1+\epsilon}\tau(\mu_*)$$

$$\implies \lim_{\epsilon \to 0} \hat{\tau}^{\epsilon}(\mu < \mu_* + \epsilon) = \lim_{\epsilon \to 0} \tau(\mu < \mu_* + \epsilon)$$

and so $\hat{\tau}^{\epsilon}$ does not violate (12) at $\mu_* + \epsilon$. For all $\mu \leq \mu_* - \epsilon^2$, the right-hand side of (12) is the same as under τ, while by Jensen's inequality,

$$\int_{\mu}^{1} x(s) d\hat{\tau}^{\epsilon}(s) \leq \int_{\mu}^{1} x(s) d\tau(s)$$

Hence (12) is satisfied everywhere by $\hat{\tau}^{\epsilon}$ for ϵ sufficiently small.

Alternatively, suppose that there are no such atoms. Then τ is continuous for $\mu \in (\underline{\mu}, \mu^*)$. Consider $\mu_* \in \operatorname{supp}(\tau)$ such that $\mu_* \in (0, \mu^*)$ and (12) does not hold with equality. By assumption, such a point exists. Then for sufficiently small ϵ, (12) does not hold with equality for all $\mu \in (\mu_* - \epsilon^2, \mu_* + \epsilon)$. Thus the construction of the previous paragraph can be used to create a mean-preserving spread that does not violate (12) here either.

Finally, note that for a fixed $\underline{\mu}$, $E[\mu]$ is decreasing in μ^*. There is therefore a unique μ^* for which $E_{\tau}[\mu] = \mu_0$. If one increases $\underline{\mu}$, then if $\tau(\underline{\mu})$ does not increase as well, the new resultant distribution $\hat{\tau}_{\underline{\mu}}$ will strictly first-order stochastically dominate τ. As this implies $E_{\hat{\tau}_{\underline{\mu}}}[\mu] > \mu_0$, this is impossible. □

Proof of Proposition 7: By Jensen's inequality, any mean-preserving spread of any τ is a weak improvement for the seller. By Lemma 3, any τ has a feasible mean-preserving spread unless it satisfies (12) with equality above some μ^*, and no other posterior aside from $\underline{\mu}$ is in the support. Hence some such τ will be optimal. That this can be implemented by a second-price auction with a reserve price r can be seen by setting $r = \int_{\underline{\mu}}^{\mu^*} \tilde{t}'(\mu) d\mu$ and using the revenue equivalence theorem (Myerson, 1981). □

Before presenting the proofs of Proposition 8 and Theorem 9, I introduce some additional notation and a useful lemma, analogous to Lemma 3. Consider the pushforward measure σ as generated by $\mathbf{x}(\mu)$ where μ is distributed according to τ. One can then write (12) as

$$\int_{x^*}^{1} x d\sigma(x) \leq \frac{1 - \sigma(x < x^*)^N}{N}, \forall x^* \in [0, 1] \tag{29}$$

Lemma B: *For any σ satisfying (29), there exists a mean-preserving spread $\hat{\sigma}$ over $x \in [0,1]$ that*

(i) satisfies (29) with equality between some x^* and 1;

(ii) sets $\sigma((0,x^*)) = 0$; and

(iii) has an atom at $x = 0$.

Proof: Suppose that (29) is satisfied for all $x \geq x^*$. As in the proof of Lemma 3, it is easy to show that in order to find a mean-preserving spread, one need only consider the distribution between 0 and x^*, since (29) for $x > x^*$ does not depend on the exact distribution of lower values, but only on their cumulative distribution up to x.

If there is an atom at some $x_* \in (0, x^*)$, then there exists $\delta > 0$ such that for sufficiently small ϵ, (29) does not hold with equality at $\hat{x}$, $\forall \hat{x} \in (x_*, x_* + \epsilon)$, or else (29) would be violated at x_* itself. Moreover,

$$\lim_{\epsilon \to 0} \sigma(x_* - \epsilon, x_* + \epsilon) = \sigma(x_*)$$

Consider the following mean-preserving spread: replace σ with $\hat{\sigma}^\epsilon$, which, for all $x \in [x_* - \epsilon^2, x_* + \epsilon]$, assigns all mass to $\{x_* - \epsilon^2, x_* + \epsilon\}$, while preserving $E_{\hat{\sigma}^\epsilon}[x] = E_\sigma[x]$. By Bayes' rule,

$$\lim_{\epsilon \to 0} \sigma([0, x - \epsilon^2)) + \frac{1}{1+\epsilon}\sigma(x_*) \leq \lim_{\epsilon \to 0} \hat{\sigma}^\epsilon([0, x_* + \epsilon)) \leq \lim_{\epsilon \to 0} \sigma([0, x_* + \epsilon) \backslash \{x_*\}) + \frac{1}{1+\epsilon}\sigma(x_*)$$

$$\implies \lim_{\epsilon \to 0} \hat{\sigma}^\epsilon([0, x_* + \epsilon)) = \lim_{\epsilon \to 0} \sigma([0, x_* + \epsilon))$$

and so $\hat{\sigma}^\epsilon$ does not violate (29) at $x_* + \epsilon$. For all $x \leq x_* - \epsilon^2$, the right-hand side of (29) is the same as under σ, while $\int_x^1 s d\hat{\sigma}^\epsilon(s) = \int_x^1 s d\sigma(s)$. Thus, (29) is satisfied everywhere for $\hat{\sigma}^\epsilon$ for ϵ sufficiently small.

Now suppose instead that there are no such atoms. Then σ is continuous for $x \in (0, x^*)$. Consider $x_* \in \text{supp}(\sigma)$ such that $x_* \in (0, x^*)$ and (29) does not hold with equality. By assumption, such a point exists. Then, for sufficiently small ϵ, (29) does not hold with equality for all $x \in (x_* - \epsilon^2, x_* + \epsilon)$. Thus the construction of the previous paragraph can be used to create a mean-preserving spread that does not violate (29) here either.

By Zorn's lemma, there then exists a maximal mean-preserving spread, which must satisfy (i)-(iii). □

Proof of Proposition 8: Since H is quadratic, $\mathbf{H}$ is independent of μ. By (22) and (23), this means that $\frac{\partial \mu}{\partial x}(\theta|x)$ is constant, i.e. not dependent on x or $\underline{\mu}$. Thus, for any contour mechanism $\mathcal{C}$, all values of $\mu(\cdot|x)$ are linear in x. By (24), so is $E_{\mu(\cdot|x)}[\theta]$, and as a result by (18) $\mathbf{t}$ is quadratic in x, with initial conditions $\mathbf{t}(0) = 0$ and $\mathbf{t}'(0) = E_{\underline{\mu}}[\theta]$. Letting σ be the pushforward measure over X defined by τ and $\mathbf{x}(\mu)$, any mean-preserving spread $\hat{\sigma}$ over X also defines a mean-preserving spread $\hat{\tau}$ over μ given $\mathcal{C}$, and vice versa. Any such mean-preserving spread increases the seller's expected payoff due to $\mathbf{t}(x)$ being quadratic in x (and hence convex). By Lemma B, a maximal mean-preserving spread places an atom at $x = 0$ while satisfying (12) with equality for all $x > x^*$ for some x^*, while placing measure 0 on $x \in (0, x^*)$. By the revenue equivalence theorem of Myerson (1981), this can be implemented by a second-price auction with a reserve price. □

Proof of Theorem 9: (i) The information acquisition cost is given by

$$c(\tau_N) = \int [H(\mu_0) - H(\mu)] d\tau_N(\mu)$$

Since the buyer's probability of winning converges to 0, her expected utility converges to 0 as well. Thus (with some abuse of notation), $\tau_N \to \mu_0$ in the weak* topology, and therefore $E_\mu[\theta] \to E_{\mu_0}[\theta]$.

(ii) By (12), $E_{\tau_N}[\mathbf{x}_N(\mu)] \to 0$. By Proposition 2, $\mathbf{x}'(\mu)$ is determined for any μ regardless of $\underline{\mu}$. By (4) and (5), $\frac{\partial \mu}{\partial x}(\theta|x = 0)$ is continuous in $\underline{\mu}$ since H is twice continuously differentiable, and so $\mathbf{x}'(\mu)$ is uniformly continuous on any closed ball B around μ_0 such that B is in the interior of the simplex. As shown above, for sufficiently large N, $\tau_N(\mu \in B) \to 1$, and so $|\tau_N - \delta_{\underline{\mu}_N}| \to 0$ in the weak* topology, where $\delta_{\underline{\mu}_N}$ is the Dirac measure that places probability 1 on $\underline{\mu}_N$. By the triangle inequality from (i), this means that $\underline{\mu}_N \to \mu_0$.

(iii) Fix function $\mathbf{t}(x)$. Since $E_\mu(\cdot|x)[\theta]$ is strictly increasing in x by (24),

$\mathbf{t}(x)$ will be a strictly convex function. Hence by Jensen's inequality, for any σ, there exists $\hat{\sigma}$ that satisfies the properties in Lemma B such that $\int_0^1 \mathbf{t}(x)d\hat{\sigma}(x) > \int_0^1 \mathbf{t}(x)d\sigma(x)$. As in the proof of Proposition 7, any σ that satisfies these properties can be implemented by a second-price auction with reserve price $r = \mathbf{t}(x^*)$ by the revenue equivalence theorem of Myerson (1981).

Next, for any fixed $\mathbf{t}$, the distribution σ satisfying the properties in Lemma B that maximizes $\int_0^1 \mathbf{t}(x)d\sigma(x)$ is that which sets $x^* = 0$, as for any other value, the distribution over $x \in [x^*, 1]$ would remain unchanged by setting x^* instead. Since $\mathbf{t}$ is a strictly increasing function and the new distribution first-order stochastically dominates the old one, this increases $\int_0^1 \mathbf{t}(x)d\sigma(x)$. Thus, for fixed $\mathbf{t}(\cdot)$, a second-price auction with a reserve price of 0 is optimal.

I now show that in the limit as $N \to \infty$, there is a unique limit value $\mathbf{t}(x)$ of any implementable sequence of $\{\mathbf{t}_N(x)\}_{N=1}^{\infty}$, and so one will be able to invoke the above result to conclude that this form of auction is optimal. First, consider the sequence of distributions $\{\tau_N\}$ and their pushforward measures $\{\sigma_N\}$. For sufficiently high N, there exists Bayes-plausible $\hat{\tau}_N$ such that its pushforward measure $\hat{\sigma}_N$ satisfies (a)-(c) and is a mean-preserving spread of σ_N, with some corresponding value of x^*. To see this, by Coddington and Levinson, Theorem 7.6, for any $\epsilon > 0$ there exists $\delta > 0$ such that if $\mu \in \bar{B}_\delta(\mu_0)$ (the closed ball of radius δ around μ_0 in the simplex), then the solutions for $(\mathbf{t}(x), \mu(\cdot|x))$ under $\underline{\mu} = \mu$ differ from those under $\underline{\mu} = \mu_0$ by at most ϵ in the Euclidean topology. Consider the function

$$\phi_N(\underline{\mu}) = \underline{\mu} + \frac{1}{2}[\mu_0 - \int_0^1 \mu(\cdot|x)d\hat{\tau}_N(\mu(\cdot|x))]$$

Clearly, $\phi_N(\underline{\mu}) = \underline{\mu}$ if and only if $\int_0^1 \mu(\cdot|x)d\hat{\tau}_N(\mu(\cdot|x)) = \mu_0$. As $\mu(\cdot|x)$ is uniformly continuous in $\underline{\mu} \in \bar{B}_\delta(\mu_0)$, it follows that for N large enough, $|\underline{\mu} - \int_0^1 \mu(\cdot|x)d\hat{\tau}_N(\mu(\cdot|x))| < \delta$ by (12) and (22) for all $\underline{\mu} \in \bar{B}_\delta(\mu_0)$, as τ converges to the Dirac measure on $\underline{\mu}$ by (ii). Hence, by the triangle inequality,

$$|\mu_0 - \phi_N(\underline{\mu})| \leq \frac{1}{2}|\mu_0 - \underline{\mu}| + \frac{1}{2}|\int_0^1 \mu(\cdot|x)d\hat{\tau}_N(\mu(\cdot|x)) - \underline{\mu}|$$

$$\leq \frac{1}{2}\delta + \frac{1}{2}\delta = \delta$$

and so $\phi_N(\underline{\mu}) \in \bar{B}_\delta(\mu_0)$. Since $\phi_N(\underline{\mu})$ is continuous, by the Brouwer fixed point theorem there exists $\underline{\mu} \in \bar{B}_\delta(\mu_0)$ such that $\phi_N(\underline{\mu}) = \underline{\mu}$, which implies that $\int_0^1 \mu(\cdot|x)d\hat{\tau}_N(\mu(\cdot|x)) = \mu_0$ as required. Thus, given τ_N and σ_N, there exist such $\hat{\tau}_N$ and $\hat{\sigma}_N$, respectively, for high enough N.

Let $\mathbf{t}_N$ and $\hat{\mathbf{t}}_N$ be the corresponding transfer functions. Consider any subsequence such that $\sigma_N \to \sigma$ and $\hat{\sigma}_N \to \hat{\sigma}$ in the weak* topology. For any y, by the Portmanteau theorem,

$$\int_0^y \sigma([0,x))dx \leq \liminf \int_0^y \sigma_N([0,x))dx \leq \liminf \int_0^y \hat{\sigma}_N([0,x))dx = \int_0^y \hat{\sigma}([0,x))dx$$

where the last holds with equality because either $\hat{\sigma}$ is absolutely continuous (if $x^* = 0$) or $\hat{\sigma}([0,x^*)) = \hat{\sigma}(x=0)$. Thus, $\hat{\sigma}$ is a mean-preserving spread of σ. Moreover, by the Lipschitz continuity of $\mathbf{H}$, both $\mathbf{t}_N \to \mathbf{t}_{\mu_0}$ and $\hat{\mathbf{t}}_N \to \mathbf{t}_{\mu_0}$ uniformly on $[0,1]$, where $\mathbf{t}$ is defined for the contour starting at $\underline{\mu} = \mu_0$ (Coddington and Levinson, Theorem 7.1). Since $\mathbf{t}$ is also continuous, by the Portmanteau theorem and the dominated convergence theorem,

$$\lim_{N\to\infty} \int_0^1 N\mathbf{t}_{\mu_0}(x)d\sigma_N(x) = \lim_{N\to\infty} \int_0^1 N\mathbf{t}_N(x)d\sigma_N(x)$$

$$\leq \lim_{N\to\infty} \int_0^1 N\mathbf{t}_N(x)d\hat{\sigma}_N(x)$$

$$= \lim_{N\to\infty} \int_0^1 N\mathbf{t}_{\mu_0}(x)d\hat{\sigma}_N(x)$$

$$= \lim_{N\to\infty} \int_0^1 N\hat{\mathbf{t}}_N(x)d\hat{\sigma}_N(x)$$

assuming that $\lim_{N\to\infty} \int_0^1 N\mathbf{t}_{\mu_0}(x)d\hat{\sigma}_N(x)$ is finite. Differentiating (29) when it holds with equality at x yields

$$x = [\hat{\sigma}_N((0,x))]^{N-1}$$

$$\implies \frac{d\hat{\sigma}_N}{dx}(x) = \frac{(x)^{\frac{2-N}{N-1}}}{N-1} \leq \frac{2}{Nx}$$

Indeed,

$$\lim_{N\to\infty} N\frac{d\hat{\sigma}_N}{dx}(x) = \frac{1}{x}$$

Since, by (18),

$$x \cdot \min\{\theta \in \Theta\} \leq \mathbf{t}(x) \leq x \cdot \max\{\theta \in \Theta\}$$

by the dominated convergence theorem we have (even for $x^* = 0$, by defining for each N at the limit as $x^* \to 0$)

$$\int_{x^*}^{1} N\mathbf{t}_{\mu_0}(x)d\hat{\sigma}_N(x) \leq \int_{x^*}^{1} 2\max\{\theta \in \Theta\}dx$$

$$\implies \lim_{N\to\infty} \int_{x^*}^{1} N\mathbf{t}_{\mu_0}(x)d\hat{\sigma}_N(x) = \int_{x^*}^{1} \frac{\mathbf{t}_{\mu_0}(x)}{x}dx$$

As observed earlier, for fixed $\mathbf{t}(\cdot)$, setting $x^* = 0$ is optimal. Therefore, any mechanism in the limit is dominated by a second-price auction with reserve price 0, which yields the revenue as given in (13). $\square$