Screening Inattentive Agents[*]

Jeffrey Mensch[†]

July 13, 2020

Abstract

Information plays a crucial role in mechanism design problems. A potential complication is that agents may be inattentive, and so information may endogenously and flexibly depend on what options are offered. I show that it is without loss to look at *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. The reduction to contour mechanisms has significant implications for both implementation of the optimal mechanism, and the revenues that can be achieved.

[*]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. This project was generously supported by the Israel Science Foundation, Grant #798/18. Tal Shewlker provided excellent research assistance.

[†]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* in a variety of economic environments. As such, the potential buyer may 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 toward determine this. This must be accounted for when setting, as a seller, the trading mechanism which buyers face, as the mechanism will affect both the information acquisition that the buyer performs, and the eventual purchasing decision.

The fact that there are many features that determine the value to 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 in the informativeness of the signals associated with them, but 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, which can be represented by a menu of prices and trade probabilities, as in Myerson (1981). This ensures that the rules of the mechanism are known to the buyers ahead of their participation. 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

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

satisfies Bayes' rule, at a cost that is increasing 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 principal 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 approach has significant shortcomings: not only the distribution, but even the support of 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 look at reported interim values when considering incentive compatibility of mechanisms due to the additional ex-ante constraints.

Instead, I show that the set of mechanisms that can be implemented 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 agent, 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, therefore, the principal's problem reduces to one of Bayesian persuasion (Aumann and Maschler, 1995; Kamenica and Gentzkow, 2011).

When choosing this optimal distribution, the principal faces three major tradeoffs. By having the agent acquire information that leads to higher beliefs, the principal can extract more rents, as the agent 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 would lead to higher rents, its low probability of realization may not justify it. Lastly, the principal would like to extract rents from a given belief. However, the more rent is extracted, the weaker the power of the incentives to induce innformation acquisition. As the agent cannot acquire the object with probability greater than 1, this prevents sufficient incentives for the agent to acquire information towards 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, this will no longer be the case. Nevertheless, I provide sufficient conditions to guarantee that the mechanisms will be of this form.

Yet even when the results go through, the optimal reserve prices and revenues differ from previous results. For instance, as the number of bidders 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 bidders 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, one may want to change the mechanism for a given interim distribution to entice 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 defend against not just choosing a different *quantity* of information, but against different *interim values* as well, even if they occur with probability 0. This means providing enough surplus at a given interim value to discourage such deviations. By contrast, without flexible information acquisition, there is no difference regarding revenue extraction.

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 presents the model of mechanism design and information acquisition. Section 3 presents necessary and sufficient conditions for a mechanism to be implementable. Section 4 discusses optimal single-agent mechanism design by means of Bayesian persuasion. Section 5 uses the techniques of the previous section for applications, including the discussion of the posted-price mechanism. Section 6 then explores optimal mechanism design with multiple agents. Section 7 discusses connections to related literature. Section 8 concludes.

2 Model

Throughout, I restrict attention to symmetric mechanisms and strategies across all N agents, and so do not use subscripts to distinguish agents. I therefore focus on the problem as if there is a single agent, with the necessary adaptations for multiple agents introduced in Section 6.

Each agent (A) has ex-post type (value) $\theta \in \Theta$, where Θ is a finite subset of $\mathbb{R}$ of size K. The principal (P) chooses a mechanism consisting of pairs of (probability) allocations and transfers $\mathcal{M} \equiv \{(x,t)\}$, where $x \in [0,1]$. Let X be the set of allocations offered by $\mathcal{M}$, which is assumed to be compact and endowed with the Borel σ-algebra over X. The principal's ex-post payoff is $u_P(x,t) = t$, while the agent's is given by $u_A(x,t,\theta) = x\theta - t$.

Given a mechanism, the agent responds by first acquiring information about the state. There is a common prior μ_0 to the principal and the agent. The agent 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 is satisfied. The cost of information acquisition is given by the expected difference in a posterior-separable measure of uncertainty $H(\mu)$.[2] Thus, given distribution τ,

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

where H is strongly concave[3] and twice-Lipschitz-continuously differentiable on any compact set of interior beliefs μ, i.e. $\mu(\theta) > 0, \forall\theta$. I 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.[4] 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 agent 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

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

[3]This is a slightly stronger assumption than strict concavity, as it essentially requires that the second derivative is 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.

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

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

Based on the posterior μ, the agent chooses an option (x, t) from the menu, with indifferences broken in favor of the principal. By the compactness of $\mathcal{M}$, an optimal choice for the agent exists. Let $\mathbf{x}(\mu)$ be the optimal choice of allocation; the corresponding transfer is then given by $\mathbf{t}(\mathbf{x}(\mu))$.

I will focus much of the discussion in the paper on what I call *canonical* mechanisms $\mathcal{M}$, i.e. in which all $(x, t) \in \mathcal{M}$ are contained in the support of those chosen by the agent. While this object is endogenous, it is without loss of optimality to restrict attention to such mechanisms, as removing option in $\mathcal{M}$ that are never chosen by the agent does not change the choice by the agent. So, the payoff to the principal from any mechanism can replicated by the from some canonical mechanism. As such, it will be useful to help characterize mechanisms in general. One immediate observation is that in any canonical mechanism, there is a unique possible t for any given x, as otherwise, the agent would simply choose the cheaper of the two options.

As in a canonical mechanism, the interim problem (post-information acquisition) reduces to finding the optimal (x, t) for each $\mu \in \text{supp}(\tau)$, one can represent these as direct revelation mechanisms. That is, each (x, t) corresponds to some $(\mathbf{x}(\mu), \mathbf{t}(\mathbf{x}(\mu))$ for some μ. The principal's objective is then

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

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

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

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

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

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

where the first constraint is that the distribution of posteriors is Bayes-plausible; the second is that the information acquisition strategy of the agent is ex-ante optimal; the third is that the choice from the menu is interim optimal; the fourth and fifth are ex-ante and interim individual rationality constraints, respectively.[5]

3 Implementability

In this section, I characterize the set of canonical mechanisms, and illustrate how they can be used for a general theory of mechanism design. I highlight how the ability of the agent to acquire information constrains the set of possible canonical mechanisms that the principal can implement. The critical issue is that the agent can deviate not only at the interim stage, but also at the ex-ante stage. The principal must therefore defend against a larger set of possible deviations. Using this observation, I then generalize to an exogenous set of mechanisms to provide a full characterization of the set of mechanisms that are canonically implementable.

In the standard mechanism design approach, as mentioned in the description of the model, one can write the choice of the agent as a function of her interim type, here being her expected value given the signal realization she receives. By the revelation principle, it suffices to represent any mechanism as one in which the agent merely reports the type to the principal, and then the principal implements the strategy for the agent. These are often referred to as "direct revelation mechanisms."

Here, while this description technically remains correct, it is not the most useful because the distribution of interim types (given by the conditional expectation of θ given a signal) is endogenous. One cannot simply look at

[5]With multiple agents, 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 6.

mechanisms that are interim-incentive compatible and conclude that they are implementable by a canonical mechanism, as they may lead to a different information acquisition strategy. Indeed, this ex-ante incentive compatibility problem may be so severe that such μ would not appear in the support of τ at all. As a result, a direct revelation mechanism that would be incentive compatible were the distribution exogenous may lead certain interim types to occur with probability 0, and thus would not be a canonical mechanism.

Instead, as will be shown, it is generally more convenient to operate with a reversed approach: we describe the beliefs conditional on the allocation x. To see that this is without loss of generality, suppose that there were a family of signal realizations such that the allocation x is chosen in all of them. As information acquisition is costly, the agent could improve her payoff by merging all such signal realizations into one signal, which one can without loss of generality label x, for the recommended allocation. This ensures that one can express the conditional belief as a function of x, given by $\mu(\cdot|x) \in \Delta(\Theta)$ as found by Bayes' rule. This will be useful for envelope-theoretic results analogous to those in Myerson (1981).[6] I therefore introduce the following definition pertaining to implementability via canonical mechanisms.

Definition 1: A mechanism $\mathcal{M}$ is *canonically implementable* if:

1. All $x \in X$ are contained in the support of the agent's choice:

$$x \in X \iff \exists \mu \in \text{supp}(\tau) : \mathbf{x}(\mu) = x$$

2. The functions $\mathbf{t}(\cdot)$ and $\mathbf{x}(\cdot)$ are one-to-one, respectively, so that $\mu(\cdot|\mathbf{x}(\mu)) = \mu$.

Any canonically implementable mechanism thus describes triplets (x, t, μ) in which the mapping from any one element of the triple to another is one-to-one. This allows the definition of transfers as a functions of μ as well, which we write as $\tilde{\mathbf{t}}(\cdot) \equiv \mathbf{t}(\mathbf{x}(\mu))$.

[6] By contrast, one could not do so for ex-post values θ, since θ can be in the support of multiple values of x.

In order for a mechanism to be canonically implementable, all of the incentive-compatibility and individual rationality constraints must hold. Assume for now that the individual rationality constraints are satisfied; I will subsequently return to them. In order for incentive-compatibility to be satisfied, the agent must be optimizing, both through information acquisition and choice of x. In particular, this means that 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 mass $\tau(\mu(\cdot|x))\mu(\theta|x)$ from this signal realization. Consider the following perturbation: increase the probability that x is recommended by ϵ, by increasing the $\pi(x|\theta)$ by $\frac{\epsilon}{\mu_0(\theta)}$. By Bayes' rule, this 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 $\mu(\cdot|x)$. here are the respective conditional probabilities before the perturbation. By the envelope theorem, the marginal change in the consumption utility to the agent 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.[7] Expressing $\frac{\partial H}{\partial \mu(\theta)}$ as the partial derivative of H with respective to the probability of state θ, the change in cost from marginally changing the mass of θ at this signal can be written as

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

To understand (1), 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 (1). 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. So,

[7] 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.

for a signal to be optimal, the agent cannot benefit from such a marginal change in mass from some signal realization x to another x'. It turns out that this requirement is also sufficient for optimality: one need only consider marginal perturbations to ensure optimality, as due to the convexity of the cost function due to the concavity of H, one can generate a local improvement from any global improvement.

Lemma 1: *Fix menu $\mathcal{M}$ and suppose that all $x \in X$ are in the support of those chosen by the agent. Then τ is optimal for the agent if and only if*

$$x\theta - \mathbf{t}(x) + h(x, \theta) = x'\theta - \mathbf{t}(x') + h(x', \theta) \tag{2}$$

for all θ and all x, x'.

All proofs are in Appendix B.

The previous lemma is more powerful than it may initially seem. In particular, it implies that any mechanism that satisfies (2) satisfies not only (IC-A), but also (IC-I). Suppose that it were not, and so (IC-I) were violated by τ. Then conditional on a positive measure of posteriors μ with respect to τ, the agent prefers some $x \in X$ to $\mathbf{x}(\mu)$. In this case, after combining multiple alternative signal that recommend the same action (which is a further improvement since it reduces the information acquisition cost), there are now alternative measures $(\hat{\tau}, \hat{\mu}(\cdot|x))$ that are an overall improvement. Yet by Lemma 1, no such alternative $\hat{\tau}$ exists, since such an improvement implies the presence of a local one ex-ante (as given by (2)) as well. This will mean, in turn that transfers must satisfy the standard Myersonian envelope condition ex-interim when applicable.

Lemma 1 was established only for $x, x' \in X$ when all $x \in X$ were contained in the support of those chosen by the agent. While any mechanism will have a subset X for which this is true, this is an endogenous characterization of any given (exogenous) mechanism. To enable a general analysis of mechanism design with inattentive agents, one must be able to tie this into an exogenous general framework.

One can do so as follows. Given mechanism $\mathcal{M}$, consider instead a (noncanonical) 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, $t(x)$ is given by that of $\mathcal{M}$ when $x \in X$. If $\mathcal{M}'$ is to induce the same τ as in $\mathcal{M}$, one must define $\mu(\cdot|x)$ and $\mathbf{t}(x)$ in such a way that the agent does not want to place positive weight on signals that lead to $x' \notin X$. I show that one can do so by extending the results of Lemma 1 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.

Lemma 2: *Suppose that* $(x^*, \mathbf{t}(x^*), \mu(\cdot|x^*))$ *is in the support of the choice of the agent. 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 (2). Moreover, any* τ *such that* *supp*$(\tau) \subset \{\mu : \exists x \in [0,1] : \mu = \mu(\cdot|x)\}$ *is then optimal for the agent 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 1, we have seen that any distribution τ over points that satisfy (2) is optimal for the agent. So, the agent is *indifferent* between all distributions τ over such points that satisfy Bayes' rule, as one can construct a canonical mechanism as in Lemma 1 by only including the points in $\mathcal{M}'$ with $\mu(\cdot|x)$ in the support of τ.[8] Hence (2), in the context of Lemma 2, 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 present a definition of such mechanisms as defined starting from the point $x = 0$.

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

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

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

[8]This will have implications for optimal mechanism design, which I explore in Section 4.

with initial conditions $(0, \mathbf{t}(0), \underline{\mu})$, and $(x, \mathbf{t}(x), \mu(\cdot|x))$ subsequently determined by (2).

Thus, any canonical mechanism $\mathcal{M}$ can be represented by a contour mechanism $\mathcal{C}$. The upshot of this is that the agent is then incentivized to obey a recommendation of information acquisition corresponding to $\mathcal{M}$, as long as it lies on the same contour. This means the principal can implement whatever distribution he likes along the contour, subject to Bayes' rule, as such as recommendation will satisfy (2) and hence be optimal for the agent.

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: both (IR-I) and (IR-A) are satisfied whenever the (IR-I) constraint 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 agent of something other than $x = 0$ implies that it is preferable, and hence must give nonnegative utility. I thus summarize the results of this subsection in the following theorem.

Theorem 1 (Implementability): *A canonical mechanism $\mathcal{M}$ is canonically implementable if and only if it 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 multiplication of costs by any coefficient $\kappa > 0$. This reflects the idea that it is prohibitively costly to fully rule out a possibility. In order to compensate for this, the principal would have to provide an infinite payoff conditional on reporting such beliefs, which is obviously suboptimal for the principal. I provide a description in the following example of how contour mechanisms look for entropy-based costs.

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 from (2), one gets the likelihood ratio condition[9]

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

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

$$\mu(\theta|x) = \frac{\underline{\mu}(\theta) e^{\frac{x\theta}{\kappa}}}{\sum_{\theta'\in\Theta} \underline{\mu}(\theta') e^{\frac{x\theta'}{\kappa}}} \tag{3}$$

Meanwhile, also from (2), we get that

$$\mathbf{t}(x) - \mathbf{t}(0) = x\theta - \kappa(\ln \mu(\theta|x) - \ln \underline{\mu}(\theta)) \tag{4}$$

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 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, and so 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 the prior μ_0. $\square$

As the contour mechanism $\mathcal{C}$ is pinned down following some μ, one 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 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,

[9] This was derived in Mensch (2020).

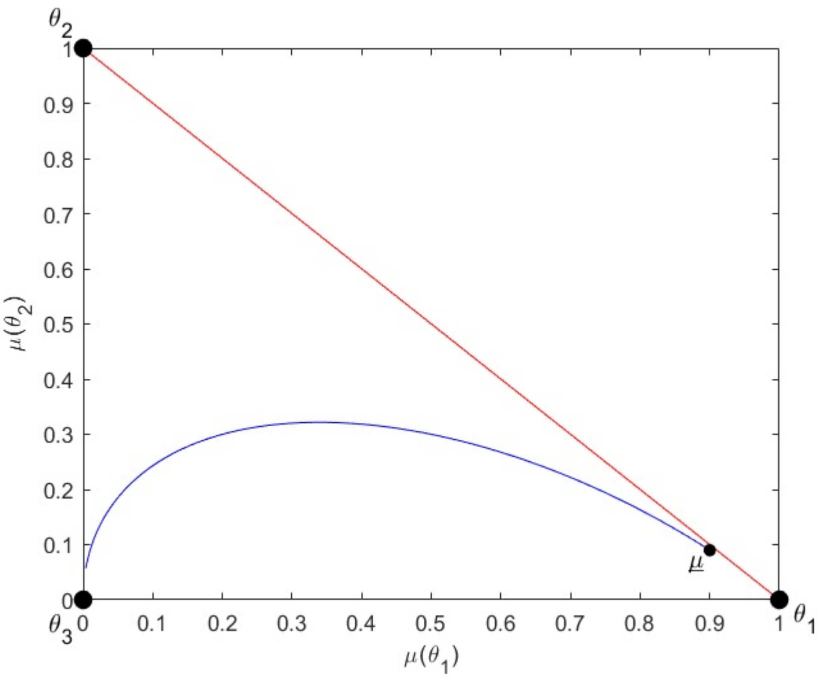

(a) Contour of beliefs

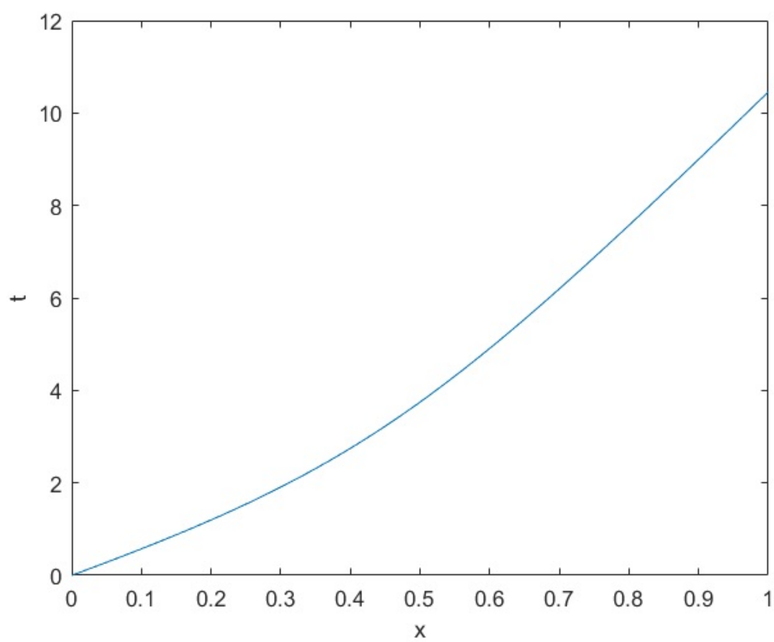

(b) Allocations and transfers

Figure 1: Entropy costs, three states

the mechanism itself is also pinned down, not only by the distribution of types, but 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 $\mathbf{x}(\mu)+\delta \in [0,1]$). I summarize this in Proposition 2.

Proposition 2: *For any given* $(x, \mathbf{t}(x), \mu(\cdot|x)) \in \mathcal{C}$*, any implementable contour mechanism* $\hat{\mathcal{C}} \ni (\hat{x}, \hat{\mathbf{t}}(\hat{x}), \hat{\mu}(\cdot|\hat{x}))$ *such that* $1-\delta_2 \geq x, \hat{x} \geq \delta_1$ *and* $\hat{\mu}(\cdot|\hat{x}) = \mu(\cdot|x)$ *must satisfy*

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

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

for all $\delta \in (-\delta_1, \delta_2)$*.*

4 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 principal to choose which distribution to induce. In Figure 2, I show one such possible set of posteriors, in which the prior, given by μ_0, is in their convex hull. It is therefore possible to find appropriate weights on each posterior to make them sum to the prior, satisfying Bayes' rule.

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

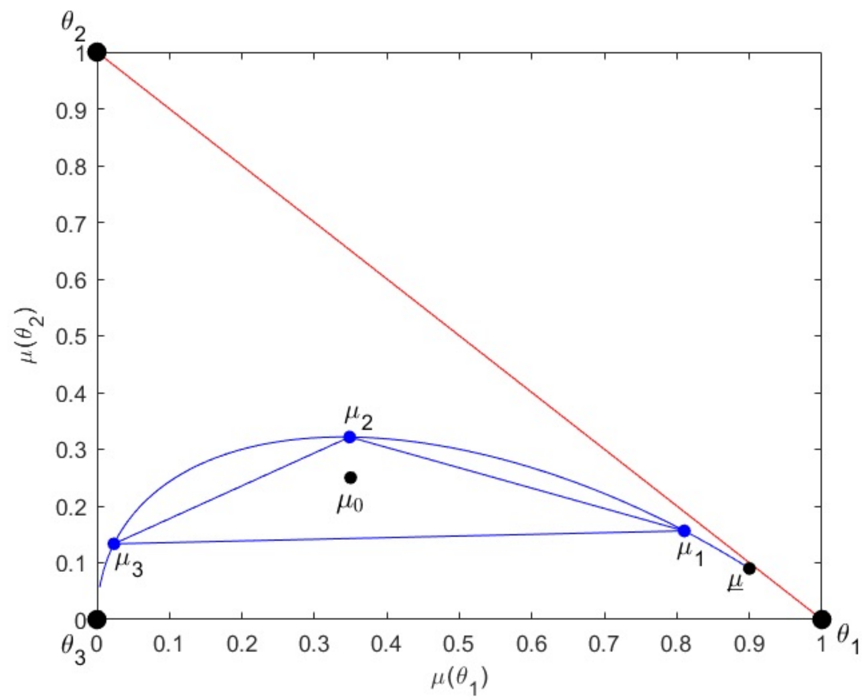

Figure 2: Possible distribution given $\mathcal{C}$

define the *value* of a posterior to be

$$v_{\mathcal{C}}(\mu) \equiv \begin{cases} \mathbf{t}(\mathbf{x}(\mu)), & (\mathbf{x}(\mu), \mathbf{t}(\mathbf{x}(\mu)), \mu) \in \mathcal{C} \\ -\infty, & \text{otherwise} \end{cases} \tag{5}$$

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

$$V_{\mathcal{C}}(\mu) \equiv \sup\{z : (\mu, z) \in co(v_{\mathcal{C}})\} \tag{6}$$

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*

$$\max_{\mathcal{C}} V_{\mathcal{C}}(\mu_0) \tag{7}$$

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 import results from that literature about optimal persuasion mechanisms. For instance, one need only restrict attention to mechanisms with at most K signal realizations.[10] This allows for a much smaller search over possible mechanisms in order to find the optimum, and enables one to use first-order conditions as a tool to do so.

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

Remark 1: As the utility of the agent 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 distributions of posteriors for which the prior is a convex combination thereof. 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, as based on H, would need to be measurable in the posterior mean, in contrast to the present assumption of strong concavity of H in the posteriors. □

While I will press the analysis of the form of the optimal mechanisms further for the case of binary states in Section 5, it is possible to provide a partial characterization for the general case, which I do in the following result.

Proposition 5: *In the optimal single-agent 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 canonical mechanism that sells the item to the agent with probability 1. The intuition is similar to that in 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.

[10] This follows from the standard reasoning based on Carathéodory's Theorem.

One may wonder if one can say more 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, entailing a certain chance of the item not being sold for prices other than the highest possible.

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 we were to take these signals as exogenously given, the agent would be indifferent between this allocation and that associated with the lower realization (in the former case, that given by non-participation), as otherwise the seller could increase the price. Yet here, ex-interim, the agent strictly prefers their assigned allocation.

This can be inferred from Figure 1. As described there, the contour mechanism maps out a continuum of possible values of x that increase as one moves along the contour of beliefs, as in Figure 1(a). In Figure 1(b), one sees that as one does so, the transfers associated with x are increasing. So, implicitly, one is offering 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 agent must be strictly increasing for intervening values of μ.

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

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. As will be shown, this will make it much easier to explicitly solve for them in those cases.

5 Optimal Mechanisms: Single Agent, $|K| = 2$

5.1 General characterization

Up until now, the results have held generally for an arbitrary number of states for a given agent. 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 subsequent persuasion problem, along that contour.

Let $\Theta = \{\theta_1, \theta_2\}$ with $\theta_2 > \theta_1$. This means that beliefs are one-dimensional. Ordering them by the probability they place on θ_2, and with some abuse of notation treating μ as standing for $\mu(\theta_2)$, then as derived in Appendix B in the proof of Lemma 2, the marginal change in $\mathbf{x}$ as $\mu(\theta_2)$ increases is

$$\mathbf{x}'(\mu) = -\frac{\frac{\partial^2 H}{\partial \mu(\theta_2)^2}(\mu) - 2\frac{\partial^2 H}{\partial \mu(\theta_2)\partial \mu(\theta_1)}(\mu) + \frac{\partial^2 H}{\partial \mu(\theta_1)^2}(\mu)}{\theta_2 - \theta_1} \tag{8}$$

This is positive due to the 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 concavity of H can be interpreted as the rate at which the marginal cost of information increases. In order to incentivize the agent to acquire information despite this cost, the allocation provided at posterior μ must change sufficiently in order to do so. At the same time, a larger difference between states, $\theta_2 - \theta_1$, incentivizes the agent to acquire more information, and so a smaller change in allocation is needed to reach μ.

By (IC-I), the Myersonian envelope condition $\mathbf{t}'(x) = E_{\mu(\cdot|x)}[\theta]$, and so with slight abuse of notation,

$$\tilde{\mathbf{t}}'(\mu) = [\theta_2\mu(\theta_2) + \theta_1(1-\mu(\theta_2))]\mathbf{x}'(\mu) \tag{9}$$

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

$$v_{\mathcal{C}}(\mu) = \tilde{\mathbf{t}}(\mu) = \int_{\underline{\mu}}^{\mu} \tilde{\mathbf{t}}'(\mu)d\mu \tag{10}$$

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, this means that 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 (8) and (9) 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 principal 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 three economic tradeoffs that the principal makes when choosing the optimal mechanism. 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 principal 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 will also decrease $\tau(\mu_2)$ if one holds μ_1 fixed. So, 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 principal 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 principal is constrained by the allocation probability x needing to be at most 1. Clearly, this means that in the optimal mechanism, $\mathbf{x}(\mu_2) = 1$. Beyond this, though, there is now an additional incentive to decrease μ_2, as this relaxes the aforementioned allocation constraint. This is distinct from that in the previous paragraph, which is connected to Bayes' rule on posteriors, rather than allocations. Even if (μ_1, μ_2) are optimal for a given $\underline{\mu}$, it may be an improvement for the principal to lower $\underline{\mu}$ in order to extract more rents; this may lead to a different choice of posteriors being optimal overall.

5.2 Posted price mechanisms

A classic result from the literature on pure allocation mechanisms is that posted-price mechanisms are optimal: the agent can either pay a certain t and receive the item with probability 1, or not pay anything and receive nothing with certainty. Here, this result will not automatically hold. Viewing the problem as one of Bayesian persuasion makes the reasoning clear. There are two possibilities for posted-price mechanisms: either the agent 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[11]

$$\tilde{\mathbf{t}}'(\mu)=\begin{cases}1, & \mu(\theta_2)<0.4\\ \epsilon, & \mu(\theta_2)\in[0.4,0.5]\\ \frac{0.07}{\epsilon}, & \mu(\theta_2)\in(0.5,0.5+\epsilon]\\ 0.4, & \mu(\theta_2)>0.5+\epsilon\end{cases}$$

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 one does at lower values of $\mu(\theta_2)$. More formally, if

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

[11]Recall from the discussion following (9) and (10) that one can generate any monotone $\tilde{\mathbf{t}}$ with the appropriate concave H. While technically not differentiable at the transition points, one can uniformly approximate this by a differentiable function. This would lead to the solution presented here being approximately correct for appropriately differentiable H.

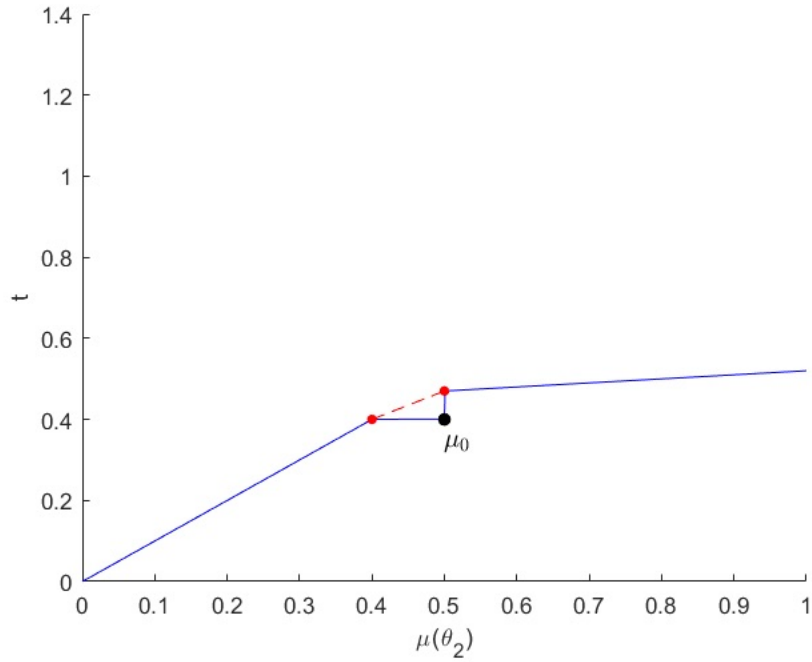

Figure 3: Posted price not optimal

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 yet nonzero probability. □

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 principal's optimal mechanism design problem. In order for the principal to find it optimal for the agent 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 for the principal to sell to the agent at belief $\mu > \mu_0$ due to higher rent extraction must significantly outweigh the lower probability with which they occur due to Bayes' rule. So, ideally, the principal 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 principal 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 principal to make the higher of the two posteriors more likely, without violating the allocation probability constraint. Below $\mu(\theta_2) = 0.4$, the principal must increase the rate of change of allocation probability. This does not allow placing enough additional weight (by Bayes rule) on $0.5 + \epsilon$ to outweigh the loss of revenue from decreasing the allocation probability at μ, 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 principal wants the agent to acquire as much information as possible, given the incentives of the contour mechanism. This means one ends up at the extreme points of the contour mechanism, i.e. corresponding to $x = 0$ and $x = 1$, as shown in Figure 5. Of course, this is by definition a posted price mechanism.

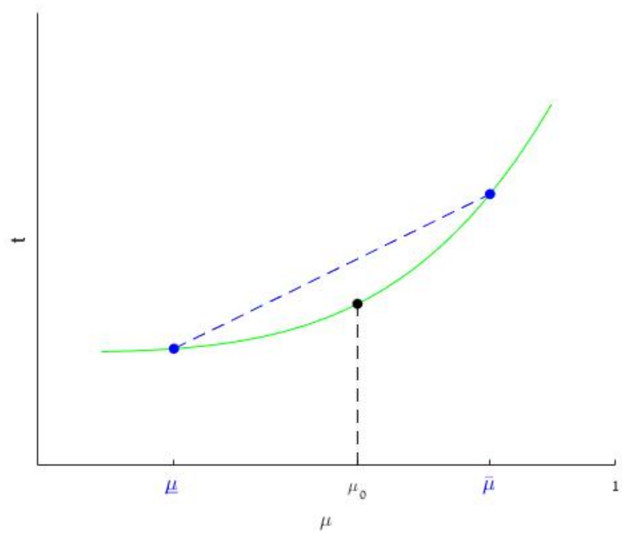

Figure 4: Posted price when $\tilde{\mathbf{t}}$ convex

Alternatively, if $\tilde{\mathbf{t}}$ is concave, as in Figure 6, then the principal wants the agent to acquire no information. Thus the agent's posteriors will remain at the prior μ_0. Given this, 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.

Theorem 6: *When Θ is binary, if $\tilde{\mathbf{t}}(\mu)$ is either concave or convex in μ for all choices of $\underline{\mu}$,*[12] *then a posted price mechanism is optimal.*

[12] As derived in (8) and (9), this is a function of primitives H and Θ.

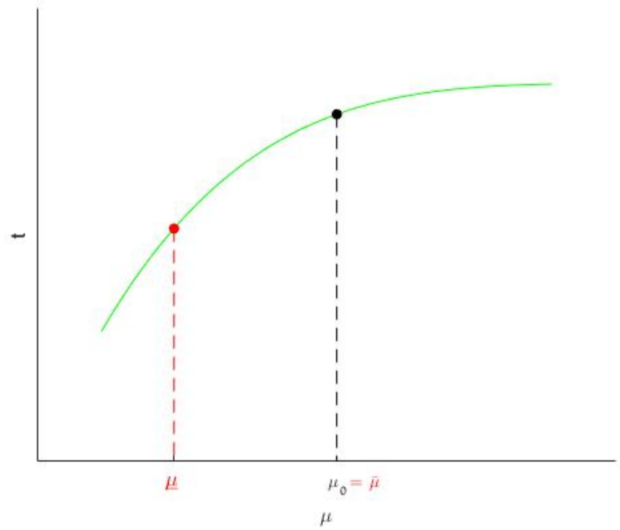

Figure 5: Posted price when $\tilde{\mathbf{t}}$ concave

Example 3: Suppose that H is proportional to residual variance, i.e.[13]

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

Then

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

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.[14] $\square$

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

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

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

acquisition worthwhile, the principal 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 agent to acquire high beliefs, but cheap to acquire low ones. The principal can be sure that the agent will not acquire beliefs that are so high that they occur with too low probability to make their rent extraction gains worthwhile, while the cheapness of low beliefs means that the principal wants the agent to acquire information in that direction so as to make the high-rent outcome more likely. Hence the principal 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 principal does not find it worthwhile to attempt to extract higher rents by inducing higher beliefs, since offering a higher allocation leads the agent to a posterior so much higher that it is very unlikely to occur, outweighing the rent gains. Meanwhile, low beliefs are too costly relative to the gains in probabilities of the agent receiving the high signal. The principal therefore optimally induces no information acquisition.

The simple form of the optimal posted price mechanism makes it straightforward to compute the optimal beliefs. As $\tilde{\mathbf{t}}'(x)$ is positive and increasing in x, 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

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

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

6 Optimal Mechanisms: Multiple Agents

A natural follow-up to the results for single-agent environments is to extend the analysis to the N-agent case. I assume that values are private and independently, symmetrically distributed, and restrict attention to symmetric static mechanisms.[15] Given the initial motivation for the question, the latter restriction is appropriate, as one can view the symmetry restriction as coming from the same reason that leads to the order of play overall: the principal attracts participation in the first place by offering a menu (e.g. a catalog, or a description of the auction). To give an example, the auction format is known before the agents appear on the day of the auction, so there is no opportunity to treat them asymmetrically beforehand. As a result, he cannot directly control which agents participate or what information they acquire, as they do so after the mechanism has been chosen. Alternatively, as noted in Deb and Pai (2017), there may be legal restrictions that prevent asymmetric treatment of the agents.

The techniques developed for the single-agent case can be used in the the multi-agent case as well. From the perspective of each individual agent, the problem is exactly the same. That is, the tradeoff that each faces with respect to the cost of information acquisition and allocation probabilities is identical. Thus the same envelope condition as in Lemmas 1 and 2 hold here as well.

The key difference, though, is that the presence of N agents imposes a resource constraint that was not present before. If the principal allocate the object to one of the agents, he cannot physically allocate it to another one. When combined with the interim incentive compatibility constraint, the results of Border (1991, Theorem 2) imply that a given mechanism is implementable

[15] In a separate, but related environment, Gershkov et al. (2019) examine optimal mechanisms when the agents' 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.

ex-interim if and only if it satisfies

$$\int_{x^*}^{1} x d\tau(\mu(\cdot|x)) \leq \frac{1-\tau(\{\mu : \mathbf{x}(\mu) < x^*\})^N}{N} \tag{11}$$

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 agent, this implies that a mechanism will be canonically implementable if and only if the induced τ along the contour satisfies the extra constraint in (11).

6.1 General N-agent case, $K = 2$

The presence of the constraint in (11) prevents the straightforward use of concavification that one could use in the single-agent 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 object was sold with probability 1. This violates the Border constraint, since this entails a positive probability event of allocating the same object to multiple agents with probability 1, an obvious impossibility. This issue requires developing techniques that address this constraint.

The key insight is that the principal 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 so much as to violate the Border constraint. Thus, as long as there is a feasible mean-preserving spread (i.e. one that does not violate the Border constraint) from μ such that the average value of $\tilde{\mathbf{t}}$ from the spread is greater than at the original μ, the principal would want to do so. This makes the solution to the problem a form of constrained concavification.

How does this approach square with the results of the single-agent case regarding 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 relax (11) at μ^*. The principal could then potentially extract more rent by lowering the non-participation belief. This presents an extra incentive to induce less information acquisition in the multi-agent environment, rather than always inducing extremal information acquisition as found in the single-agent case.

Conversely, if $\mathbf{x}(\cdot)$ is concave (which it will be whenever $\tilde{\mathbf{t}}$ is concave, by (9)), then pooling μ over an interval to some μ^* potentially leads to a violation of (11) at μ^*. Thus the principal may want to spread beliefs further out (by inducing information acquisition) even when $\tilde{\mathbf{t}}$ is concave, which he would never want to do in the single-agent 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 incentives to spread out beliefs as much as possible are present both from concavification and from the Border constraint. Hence given $\underline{\mu}$, the principal will place as much weight as possible on the highest possible beliefs (while satisfying (11) with equality), and place the rest of the weight on $\underline{\mu}$. This can be implemented by, for instance, a second-price auction with a reserve price. I state this formally in Lemma 3 and Proposition 7.

Lemma 3: *Suppose that Θ is binary. If $\mathbf{x}(\cdot)$ is concave, then for any τ and $\underline{\mu}$ satisfying (11), there exists $\hat{\tau}_{\underline{\mu}}$ which is a mean-preserving spread of τ, such that for some μ^*, (11) 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 Θ is binary and 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 which satisfies the criteria of Proposition 7 is that of residual variance; another example is given below.

Example 4: Suppose that $\tilde{\mathbf{t}}$ is linear in μ on the set of implementable beliefs, and so for some constant $\alpha > 0$, $\mathbf{x}'(\mu) = \frac{\alpha}{\theta_2\mu(\theta_2)+\theta_1(1-\mu(\theta_1))}$.[16] 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 which sets the lowest possible value of $\underline{\mu}$, which by Lemma 3, will correspond to placing $\tau(\underline{\mu}) = 0$. In other words, the reserve price will be 0, as the object is allocated with probability 1 to some player, i.e. with probability $\frac{1}{N}$ to each. $\square$

Since the optimal mechanism is of this particular form, it is much simpler to find the corresponding non-participation belief. By Lemma 3, one can trace out 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 $\underline{\mu}$.

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 classical intuition. In classic results, the second-price auction is optimal because one has a fixed distribution, and the principal 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 principal 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 agent(s) receiving a low interim expected value. This makes the conditions for optimality of the second-price auction with reserve price more stringent.

As with exogenous signals, the optimal reserve price depends on the tradeoff between greater participation, and extraction from those who do partic-

[16] By (8), 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.

ipate. With flexible information acquisition, this comes down to how much the principal gains from persuasion versus exploitation. 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-agent problem. However, this additional persuasion prevents exploitation that could be achieved through lowering $\underline{\mu}$.

Remark 3: Similarly to Remark 2, 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 agent who has this cutoff value will receive positive rents. The reasoning is analogous: one 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 to implement the same allocation in the auction ex-interim, the principal's revenue will be lower. This again contrasts with inflexible information acquisition, e.g. Shi (2012). □

6.2 Quadratic costs, K arbitrary

Many of the techniques from the binary-state environment carry over to the case where H is quadratic. This generalizes the result seen in Example 3, which examined when H was given by residual variance for a single agent.

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 extremal mean-preserving spreads 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 states environment. One can trace out 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 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.

6.3 $N \to \infty$, K arbitrary

Lastly, I consider the case of large auctions, as the number of bidders approaches infinity. The interplay of the incentives for information acquisition and the Border constraint make the question of what happens in the infinite-bidder limit more complicated. On the one hand, as is standard in the literature, the winning bidder will be in the upper tail of the distribution of bidders' values. On the other hand, with flexible information acquisition, this distribution will be endogenous. While for each N, the bidders 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 bidder's chance of winning vanishes in the limit, and so the incentive to acquire such information vanishes as well for each bidder. Nevertheless, it is not immediately clear that inducing such information is asymptotically ineffective, since it remains possible that with flexible information acquisition, the agent 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 bidders may yield such high signals with high enough probability to increase the principal'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 μ_N be the beliefs for*

N agents satisfying (11). 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

$$\lim_{N\to\infty} \max_{\tau_N, \mathcal{C}_N} N \int_0^1 \mathbf{t}_N(x) d\tau_N(\mu(\cdot|x)) = \int_0^1 \frac{\mathbf{t}_{\mu_0}(x)}{x} dx \tag{12}$$

with $\mathbf{t}_{\mu_0}(x)$ derived from (2) 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 which 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 agent who ends up receiving the item is one whose belief 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 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 just manipulating 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 extremal mean preserving spreads of distributions over x, which will be those generated from second-price auctions (by reasoning akin to that in Lemma 3 and Proposition 7).

Remark 4: 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 bidders with 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 principal would be leaving money on the table by essentially censoring lower parts of the distribution of x. $\square$

One intriguing possibility, that I do not explore here, is that there may be potential improvements for the principal from sequential mechanisms. Since the incentive to acquire information is vanishing precisely because the chances of one receiving the item, even with information acquisition, are so small, eliminating this issue restores the incentives. For instance, the principal could use a serial dictatorship mechanism, where he selects a particular agent, and depending on the signal they get, offers to sell with a given probability. If the item remains unsold, the principal can then move on to the next agent without having to worry about feasibility constraints from allocations among multiple agents. More generally, the principal could condition the mechanism offered to the second agent on the signal realization that the first agent received. Thus, if the first agent received a signal realization corresponding to a low posterior μ_1, the principal may want the second agent to focus on high posteriors in order to potentially outbid the first, and thereby sell at a higher price. Conversely, if the first agent received a high posterior μ_2, there may be no purpose to having the second acquire any information, since one is likely to sell the item to the first agent anyway. The ability to thereby correlate the choice of information acquisition across agents may prove relevant to revenue maximization; the principal cannot do so via a simultaneous, symmetric mechanism in which agents move independently.

7 Related Literature

This paper is the first to consider a model of optimal sales mechanisms with agents who can subsequently flexibly acquire costly information. However, it is by no means the first to consider information acquisition in general in such environments. Below, I discuss previously known results, and their connections and differences to the results here.

Several papers consider auction frameworks with information acquisition. Persico (2000) allows for bidders to acquire signals about values with common components in both first- and second-price auction. Compte and Jehiel (2007) and Hernando-Veciana (2009) compare information acquisition incentives in sealed bid vs. dynamic bid or open-bid environments, respectively. Bobkova (2019) considers the incentives to learn about private vs. common components in auctions. These differ from the present model in two regards: first, the auction format is taken as given, and so abstracts from the additional complications inherent in mechanism design. Second, the information acquisition technologies are not flexible, at the very least requiring that the signals acquired be ordered in terms of some measure of informativeness (such as Lehmann, 1988).

There is also a literature on comparing revenues for given auction mechanisms based on their dependence on the information structure of the agents. Milgrom and Weber (1982) discuss the revenue effects of information in auctions with common-value components. Ganuza and Penalva (2012) introduce a dispersion-based information order for comparing information structures in private-value auctions, and discusses their effects on the seller's revenue. Bergemann et al. (2017) characterizes the lower bound of bids in the first-price auction that could arise across the bidders' information structures. Sorokin and Winter (2019) derive optimal information design for fixed auction rules.

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 agents have access to information acquisition technologies that are perfectly ordered. In particular, Shi (2012), the most closely related due to its examination of optimal auctions, imposes several additional assumptions, such as the information being rotationally ordered, regularity of the conditional distributions, and, for some results, a supermodularity condition on the hazard rates. Aside from the restrictive nature of the technologies available to the agents, the inflexibility of the technologies does 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 principal not only controls the mechanism, but also the information available to the agents. Here, of course, the information is in the hands of the agent.

Limitations on the ability of the principal to implement monotone mechanisms, as well as the insight of manipulating the distributions of the types of the agent has also been used in the literature on mechanism design with limited commitment. In particular, Bester and Strausz (2001), Skreta (2006), and Doval and Skreta (2019) all consider the set of feasible conditional distributions for a given choice from a menu. Indeed, Doval and Skreta (2019) formulate how this can be viewed as a problem of Bayesian persuasion, with the principal essentially persuading himself at a later stage of the game (since he lacks commitment, he can be viewed as two separate players). The approach here differs in that the principal manipulates the beliefs of the agent, whereas with limited commitment, it is generally assumed that the agent is fully informed about her own type.

More generally, there has been a burgeoning literature on flexible information acquisition in decision problems. Starting with Sims (2003), there has been a move to reexamine many standard results while allowing for information acquisition with costs proportional to entropy reduction. This was expanded in Matejka and McKay (2015) to allow for flexible information acquisition, which provides a basis for the multinomial logit. Caplin and Dean (2013, 2015) examine this from a revealed preference framework, and extend 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 Bayesian persuasion. This paper relies on similar intuition in Sections 3 and 4 in deriving the set of implementable mechanisms, and then showing that not only the agent's problem, but the principal's problem as well reduces to one of Bayesian persuasion. The most important methodological distinction to be drawn here is that the action space and payoffs endogenously depend on the mechanism (in contrast to their results which rely on an exogenously given set of actions and payoffs), which are then tied together by envelope-theoretic considerations.[17]

Flexible information acquisition has been used in applied frameworks as well. Yang (2015), Morris and Yang (2016), and Denti (2019) consider global games models where agents can flexibly acquire information. Ravid (2020) examines a bargaining model with inattentive buyers. Yang (2020) examines optimal financial security design when buyers are inattentive. Lipnowski et al. (2020) explore optimal information provision to an inattentive receiver. This paper complements these others by applying the flexible information acquisition framework to the natural setting of optimal auctions.

Very closely related is the work of Roesler and Szentes (2017) and Ravid, Roesler and Szentes (2019), who also consider an environment involving optimal price setting by a seller, and flexible information acquisition by a buyer. Crucially, their environments differs in the timing of the model. The model here first sets the mechanism, and then the agents acquire information (as in all models mentioned in previous paragraphs). Instead, the former paper first has the buyer publicly acquire information, and then the monopolist chooses a posted-price mechanism (which will be optimal by standard results from mechanism design). Moreover, there is no cost of information acquisition in their model. In the latter, they have the players move simultaneously, and allow

[17]There are a few additional technical differences. Most importantly, the principal is the one concavifying over beliefs here, rather than the inattentive decision maker.

for information to be private and costly, focusing on the limit case as costs vanish. While these models fit well where there is a particular single buyer whom the seller can identify, and for whom there is merely uncertainty about their value, it is less so for the framework mentioned in the introduction here, where the principal sets a pricing policy for a large, anonymous population: in the latter, the principal's choice of mechanism must be observable to the buyers beforehand. The order also changes the techniques, as the seller in their mechanism can take the distribution of values as given (making it a standard mechanism design problem from the seller's perspective), whereas the seller here must consider the effects of the mechanism on the information acquisition strategy. Lastly, they only consider posted-price mechanisms, which are not without loss of optimality in my environment.

Georgiadis and Szentes (2020) also use Bayesian persuasion techniques in a principal-agent model with information acquisition. There, however, it is the principal who acquires information about the effort of the agent, using signals arising from Brownian motion whose drift depends on the (binary) effort level. They find that the principal's payoff can be reduced to a scoring rule on the likelihood of the signals received, and that the optimal stopping rule is then a pair of thresholds of certainty found by concavifying the principal's payoffs over the scoring rule.

8 Conclusion

This paper provided new tools for analyzing mechanism design with information acquisition, by considering the possibility of the agent to acquire information flexibly. This allowed using 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 second-price auction with reserve prices are 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 principal 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 principal.

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 those of 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 that may allow the principal to better incentivize the information acquisition of the agent. 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 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.

Bester, H., and Strausz, R. (2001): "Contracting with imperfect commitment and the revelation principle: the single agent case." *Econometrica*, 69, 1077-1098.

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.

Doval, L., and Skreta, V. (2019): "Mechanism design with limited commitment." 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.

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.

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.

—— (2020): "Rational inattention and the monotone likelihood ratio." Working paper.

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.

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.

Skreta, V. (2006): "Sequentially optimal mechanisms." *Review of Eco-*

nomic Studies, 73, 1085-1111.

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

$$c(\alpha\tau + (1-\alpha)\tau') \leq \alpha c(\tau) + (1-\alpha)c(\tau') \tag{13}$$

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, equality should follow when considering a decision maker valuing information. Their preferences can be represented by the expected utility from actions conditional on the state, and so they would value these lotteries equally, as they provide equivalent information in expectation.[18] 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,

[18] 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 looking it up. 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 is posterior-separable, this environment would satisfy the assumptions of the model.

Another justification comes from Pomatto et al. (2019). 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 (20) holding with equality: 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 which generates the same distribution of posteriors as this randomization. They justify this through the 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.[19] 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 the inequality (20) would also go 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 found from revealed preference representations of decision makers with costly information acquisition. In 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 one cannot improve her utility by redistributing attention across decision problems. This latter property rules out unintuitive behavior such as acquiring more information when the stakes are lower.

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

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 do so, 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 within decision problems as well by reallocating single posteriors. Intuitively, since one can then think about optimization 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 for "posterior separability," as I consider a fixed prior, and do not consider the counterfactual of what would happen if the prior were to change.[20] While there are some additional technical assumptions that are needed, the main property that must be satisfied is their axiom of "Separability." Roughly, it states the following. Suppose some posteriors $\{\mu\}$ are the ones that occur 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 otherwise choosing the same pairs (μ, a) as before. This means that the optimality of posterior μ for 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

[20]Some of the pitfalls of doing so are discussed in Mensch (2018). In particular, if one thinks of information as coming from a 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 A and state space Ω. Her ex-post utility $u(a, \omega)$, 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*

$$\max_{a \in A, \omega \in \Omega} u(a, \omega) - \min_{\hat{a} \in A, \hat{\omega} \in \Omega} u(\hat{a}, \hat{\omega}) \leq C$$

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, $\mu(\theta) \geq \epsilon$.[21]

Proof: Suppose 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 keeping actions 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 measure over x defined by the pushforward measure from $\mathbf{x}(\mu)$ and τ for the sets M_1

[21] Caplin and Dean (2013) prove a similar yet slightly weaker result in 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.

and M_2, respectively. Consider the probability integral transform 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 original respective 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

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

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

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

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

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 coming from the distance between any two elements in M_1 and M_2 as discussed above; and the second inequality coming from the strong convexity of H, i.e. $\mathbf{H} + mI$ is negative semi-definite for some $m > 0$. Meanwhile, the marginal loss in payoff from the changes in conditional choice are bounded by $C\tau(M_1)$. Hence for κ sufficiently large, the former will be larger than the latter, and hence an improvement for the decision maker. $\square$

Appendix B: Proofs
For Online Publication

Proof of Lemma 1: Given that all actions are in the support of the agent's choice, that the equality must necessarily hold is trivial. To show sufficiency, suppose that $(\tau, \mu(\cdot|x))$ is suboptimal, and that some $(\hat{\tau}, \hat{\mu}(\cdot|x))$ is better for the buyer. For any $\alpha \in (0,1)$, define the measures $(\tau_\alpha, \mu_\alpha(\cdot|x))$ given by

$$\mu_\alpha(\cdot|x) = \frac{(1-\alpha)\mu(\cdot|x)d\tau(\mu(\cdot|x)) + \alpha\hat{\mu}(\cdot|x)d\tau(\hat{\mu}(\cdot|x))}{(1-\alpha)d\tau(\mu(\cdot|x)) + \alpha d\tau(\hat{\mu}(\cdot|x))}$$

$$d\tau_\alpha(\mu_\alpha(\cdot|x)) = (1-\alpha)d\tau(\mu(\cdot|x)) + \alpha d\tau(\hat{\mu}(\cdot|x))$$

This is also an improvement over $(\tau, \mu(\cdot|x))$, since

$$\int_X \sum_{\theta\in\Theta}[x\theta - \mathbf{t}(x) + H(\mu(\cdot|x))]\mu(\theta|x)d\tau(\mu(\theta|x))$$

$$< (1-\alpha)\int_X \sum_{\theta\in\Theta}[x\theta - \mathbf{t}(x) + H(\mu(\cdot|x))]\mu(\theta|x)d\tau(\mu(\theta|x))$$

$$+\alpha\int_X \sum_{\theta\in\Theta}[x\theta - \mathbf{t}(x) + H(\hat{\mu}(\cdot|x))]\hat{\mu}(\theta|x)d\hat{\tau}(\hat{\mu}(\theta|x))$$

$$\leq \int_X \sum_{\theta\in\Theta}[x\theta - \mathbf{t}(x) + H(\mu_\alpha(\cdot|x))]\mu_\alpha(\theta|x)d\tau_\alpha(\mu_\alpha(\theta|x))$$

where the second inequality is from merging signals between τ and $\hat{\tau}$ that recommend the same x, and the fact that $\tau \neq \hat{\tau}$. Dividing by α and taking the limit as $\alpha \to 0$, this becomes the Fréchet derivative taken at (τ, μ) in the direction of $(\hat{\tau}, \hat{\mu})$:

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

yielding that for some x, x',

$$\sum_{\theta\in\Theta}[x\theta - \mathbf{t}(x) + h(x,\theta)] < \sum_{\theta\in\Theta}[x'\theta - \mathbf{t}(x') + h(x',\theta)]$$

and so, for some θ,

$$x\theta - \mathbf{t}(x) + h(x,\theta) < x'\theta - \mathbf{t}(x') + h(x',\theta)$$

$\square$

Proof of Lemma 2: 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 solution. I then verify that the necessary and sufficient conditions of Lemma 1 are satisfied.

I start by deriving a differentiable law of motion that satisfies (2), which will be used to show sufficiency. Thus I show that there exists a differentiable locus of points on which the agent's choice has its support; one can then convert it to a canonical mechanism by dropping the values of x which are not in the support, and thus invoke Lemma 1. First, to define $\mathbf{t}'(x)$, any solution that is optimal for the agent must satisfy (IC-I). It is well known from Myerson (1981) that in order to do so,

$$\lim_{\epsilon\to 0}\frac{\mathbf{t}(x+\epsilon) - \mathbf{t}(x)}{\epsilon} = E_{\mu(\cdot|x)}[\theta] \tag{14}$$

So, one can define

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

This thereby defines the law of motion of beliefs from $\mu(\cdot|x)$. From (1), for $\mu(\cdot|x)$ to be differentiable,

$$\frac{\partial h}{\partial x}(x,\theta) = \sum_{\theta''\in\Theta}\frac{\partial^2 H}{\partial\mu(\theta'')\partial\mu(\theta)}(\mu(\cdot|x))\frac{\partial\mu}{\partial x}(\theta''|x)(1-\mu(\theta|x))$$

$$-\sum_{\theta''\in\Theta}\sum_{\theta'\neq\theta}\frac{\partial^2 H}{\partial\mu(\theta'')\partial\mu(\theta)}(\mu(\cdot|x))\frac{\partial\mu}{\partial x}(\theta''|x)\mu(\theta'|x) \tag{15}$$

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

$$\sum_{\theta'\in\Theta}\frac{\partial^2 H}{\partial\mu(\theta')\partial\mu(\theta)}(\mu(\cdot|x))\frac{\partial\mu}{\partial x}(\theta'|x) = -(\theta + C_{\mu(\cdot|x)}), \forall\theta \tag{16}$$

is a solution to (15). Since H is strongly concave, the Hessian $\mathbf{H}(\mu(\cdot|x))$ is negative definite, and so once can write

$$\begin{pmatrix} \frac{\partial\mu}{\partial x}(\theta_1|x) \\ \vdots \\ \frac{\partial\mu}{\partial x}(\theta_K|x) \end{pmatrix} = -\mathbf{H}^{-1}(\mu(\cdot|x))\begin{pmatrix} \theta_1 + C_{\mu(\cdot|x)} \\ \vdots \\ \theta_K + C_{\mu(\cdot|x)} \end{pmatrix} \tag{17}$$

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 $\mathbf{H}^{-1}_{(i,j)}$,

$$C_{\mu(\cdot|x)} = -\frac{\sum_{i=1}^K\sum_{j=1}^K\theta_j\mathbf{H}^{-1}_{(i,j)}(\mu(\cdot|x))}{\sum_{i=1}^K\sum_{j=1}^K\mathbf{H}^{-1}_{(i,j)}(\mu(\cdot|x))} \tag{18}$$

It now remains to be shown that the system of differential equations defined by (14) and (17) 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 (18), $C_{\mu(\cdot|x)}$ 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 (2) 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 agent given prior $\mu_0 = \int d\tau(\mu(\cdot|x))$ by Lemma

1, and so (14) and (17) are sufficient, with

$$\frac{\partial}{\partial x}\{E_{\mu(\cdot|x)}[\theta]\} = -\sum_{\theta,\theta'\in\theta}[\frac{\partial^2 H}{\partial\mu(\theta)\partial\mu(\theta')}(\mu(\cdot|x))]\frac{\partial\mu}{\partial x}(\theta'|x)\frac{\partial\mu}{\partial x}(\theta|x) > 0 \tag{19}$$

as is easily be derived from (14) and (17), which is positive due to the negative-definiteness of the Hessian matrix.[22]

To see that one can set $[x-a, x+b] = [0,1]$, suppose that the supremal value of a were less than x. $\mu(\cdot|x-a)$ must still be in the interior of the simplex by Lemma A since $x+b-\mathbf{t}(x+b)-(x-a)+\mathbf{t}(x-a) \le b-a+\max\{\theta\in\Theta\}$. So, 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 agent's problem must be identical to that given above. To do so, fix x^*, and suppose that there exists $\hat{\tau}$ which 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 (14) and (17). Consider the distribution $\tilde{\tau}$ over $\{\mu(\cdot|x)\}$ whose pushforward measure over $x \in [0,1]$ is uniform. Then by Lemma 1, $\alpha\hat{\tau} + (1-\alpha)\tilde{\tau}$ is optimal for the agent for any $\alpha \in (0,1)$ given prior $\tilde{\mu}_0 = \alpha\mu_0 + \int_{\{\mu(\cdot|x)\}} d\tilde{\tau}(\mu(\cdot|x))$. Immediately, in order to satisfy (IC-I), the transfers conditional on x must be the same under both the mechanism that generates $\hat{\tau}$ and $\tilde{\tau}$, respectively. Thus by (1) and (2),

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

$$= H(\mu(\cdot|x)) + \frac{\partial H}{\partial\mu(\theta)}(\mu(\cdot|x))(1-\mu(\theta|x)) - \sum_{\theta'\neq\theta}\frac{\partial H}{\partial\mu(\theta')}(\mu(\cdot|x))\mu(\theta'|x) \tag{20}$$

Multiplying the above by $\hat{\mu}(\theta|x)$ and $\mu(\theta|x)$, then summing over $\theta \in \Theta$ and

[22] As remarked following Lemma 1, any set of triplets $(x, \mathbf{t}(x), \mu(\cdot|x))$ that satisfies (2) and on which τ has its support is incentive compatible, and so the monotonicity of $E_{\mu(\cdot|x)}[\theta]$ is implied anyway.

taking the difference between the former and the latter, one gets

$$\sum_{\theta\in\Theta}(\frac{\partial H}{\partial\mu(\theta)}(\mu(\cdot|x)) - \frac{\partial H}{\partial\mu(\theta)}(\hat{\mu}(\cdot|x)))(\mu(\theta|x) - \hat{\mu}(\theta|x)) = 0 \tag{21}$$

By the intermediate value theorem, there exists some $\alpha \in [0, 1]$ such that for $\tilde{\mu} \equiv \alpha\mu(\cdot|x) + (1-\alpha)\hat{\mu}(\cdot|x)$,

$$\frac{\partial H}{\partial\mu(\theta)}(\mu(\cdot|x)) - \frac{\partial H}{\partial\mu(\theta)}(\hat{\mu}(\cdot|x)) = \sum_{\theta'\in\Theta} \frac{\partial^2 H}{\partial\mu(\theta)\partial\mu(\theta')}(\tilde{\mu})(\mu(\theta'|x) - \hat{\mu}(\theta'|x)) \tag{22}$$

Putting together (21) and (22), this means that

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

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

Proof of Theorem 1: By Lemmas 1 and 2, 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:

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

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

Hence all four constraints are satisfied. □

Proof of Proposition 2: Immediate from (14) and (17) 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 $\mathbf{t}(0) < 0$, 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 2 in equations (14) and (17), 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 (2). Therefore τ is an incentive compatible distribution from Lemma 1. This implies that the set of feasible payoffs to the principal is compact, and so a maximum exists.

Given the existence of an optimal mechanism, 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 $\mathcal{C}$, it follows that the support of $co(v_{\mathcal{C}})$ must be contained in $\mathcal{C}$ with probability 1. Hence the mechanism given by (7) 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 such that $x^* \equiv \sup\{x : \exists \mu \in \text{supp}(\tau) : \mathbf{x}(\mu) = x\} < 1$. 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 (2). So, τ is optimal, with the choice of x under $\hat{\mathcal{C}}$, $\hat{\mathbf{x}}(\mu)$ equal to $\mathbf{x}(\mu) + 1 - x^*$, and $\mathbf{t}(x) = \hat{\mathbf{t}}(x)$ by Proposition 2. By Lemma 2, one can then complete $\hat{\mathcal{C}}$ to apply to values of $x < 1 - x^*$. Since by (14), $\hat{\mathbf{t}}'(x) > 0$, one can then increase $\hat{\mathbf{t}}$ by $\int_0^{1-x^*} \hat{\mathbf{t}}'(x)dx$ while maintaining (2). □

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 $\mathcal{C}$. $\square$

Proof of Lemma 3: Fix τ, and suppose 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 (11) 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

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

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

$$\implies d\tau(\mu) = \frac{1}{N-1} [\mathbf{x}(\mu)]^{\frac{1}{N-1}-1} \mathbf{x}'(\mu) d\mu \tag{23}$$

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 (23) 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 (11), there exists a mean-preserving spread that satisfies (11); by Zorn's lemma, there will then be a maximal element, which must be of the form of the lemma. First, suppose there is an atom at some $\mu_* \in (\underline{\mu}, \mu^*)$. Then there exists $\delta > 0$ such that for sufficiently small ϵ, (11) does not hold with equality at $\hat{\mu}, \forall \hat{\mu} \in (\mu_*, \mu_* + \epsilon)$ or else (11) 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 $\mu \in [\mu_* - \epsilon^2, \mu_* + \epsilon]$, places all mass at $\{\mu_* - \epsilon^2, \mu_* + \epsilon\}$, 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 (11) at $\mu_* + \epsilon$. For all $\mu \leq \mu_* - \epsilon^2$, the right-hand side of (11) 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 (11) 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 $\mu_* \in \text{supp}(\tau)$ such that $\mu_* \in (0, \mu^*)$ and (11) does not hold with equality. By hypothesis, such a point exists. Then for sufficiently small ϵ, (11) 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 (11) here as well.

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. $\square$

Proof of Proposition 7: By Jensen's inequality, any mean-preserving spread of any τ is a weak improvement for the principal. By Lemma 3, any τ has a feasible mean-preserving spread unless it satisfies (11) 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). $\square$

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 (11) as

$$\int_{x^*}^{1} x d\sigma(x) \leq \frac{1 - \sigma(x < x^*)^N}{N}, \forall x^* \in [0, 1] \tag{24}$$

Lemma B: *For any σ satisfying (24), there exists a mean-preserving spread $\hat{\sigma}$ over $x \in [0, 1]$ which*

(i) satisfies (24) with equality between some x^* and 1;

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

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

Proof: Suppose that (24) is satisfied for all $x \geq x^*$. As in the proof of Lemma 3, it is easy to show that to find a mean-preserving spread, one only need consider the distribution between 0 and x^*, since (24) for $x > x^*$ does not depend on the exact distribution of lower values, but only 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, (24) does not hold with equality at $\hat{x}$, $\forall \hat{x} \in (x_*, x_* + \epsilon)$, or else (24) 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 $x \in [x_* - \epsilon^2, x_* + \epsilon]$, places all mass at $\{x_* - \epsilon^2, x_* + \epsilon\}$, while preserving $E_{\hat{\sigma}^\epsilon}[x] = E_\sigma[x]$. By Bayes' rule,

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

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

and so $\hat{\sigma}^\epsilon$ does not violate (24) at $x_* + \epsilon$. For all $x \leq x_* - \epsilon^2$, the right-hand side of (24) is the same as under σ, while $\int_x^1 s d\hat{\sigma}^\epsilon(s) = \int_x^1 s d\sigma(s)$. So, (24) 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 (24) does not hold with equality. By hypothesis, such a point exists. Then for sufficiently small ϵ, (24) 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 (24) here as well.

By Zorn's Lemma, there then exists a maximal mean-preserving spread, which must satisfy (i)-(iii). $\square$

Proof of Proposition 8: Since H is quadratic, $\mathbf{H}$ is independent of μ. By (17) and (18), this means that $\frac{\partial \mu}{\partial x}(\theta|x)$ is constant, not depending on x or $\underline{\mu}$. So, for any contour mechanism $\mathcal{C}$, all values of $\mu(\cdot|x)$ are linear in x. By (19) so is $E_{\mu(\cdot|x)}[\theta]$, and as a result by (14) $\mathbf{t}$ 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 principal'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 (11) 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. $\square$

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 agent's probability of winning converges to 0, their expected utility converges to 0 as well. So (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 (11), $E_{\tau_N}[\mathbf{x}_N(\mu)] \to 0$. By Proposition 2, $\mathbf{x}'(\mu)$ is determined for any μ regardless of $\underline{\mu}$. By (3) and (4), $\frac{\partial \mu}{\partial x}(\theta|x = 0)$ 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 $|\tau_N - \underline{\mu}_N| \to 0$ in the weak* topology. 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 (19), $\mathbf{t}(x)$ 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 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

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

where $\mu(\cdot|0) = \mu$. Clearly $\phi_N(\mu) = \mu$ 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 $\mu \in \bar{B}_\delta(\mu_0)$, for N large enough, $|\mu - \int_0^1 \mu(\cdot|x) d\hat{\tau}_N(\mu(\cdot|x))| < \delta$ by (11) and (17) for all $\mu \in \bar{B}_\delta(\mu_0)$, as $\tau \to \mu$ by (ii). Hence by the triangle inequality,

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

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

and so $\phi_N(\mu) \in \bar{B}_\delta(\mu_0)$. Since $\phi_N(\mu)$ is continuous, by the Brouwer fixed-point theorem there exists $\mu \in \bar{B}_\delta(\mu_0)$ such that $\phi_N(\mu) = \mu$, which implies that $\int_0^1 \mu(\cdot|x) d\hat{\tau}_N(\mu(\cdot|x)) = \mu_0$ as required. Thus there exists 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

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

noting that the last equality holds (rather than as an inequality) because either $\hat{\sigma}$ is absolutely continuous (if $x^* = 0$) or $\hat{\sigma}([0,x^*)) = \hat{\sigma}(x = 0)$. So, $\hat{\sigma}$ 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,

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

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

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

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

by the dominated convergence theorem, assuming that $\lim_{N\to\infty} \int_0^1 N\mathbf{t}_{\mu_0}(x)d\hat{\sigma}_N(x)$ is finite. Differentiating (24) when it holds with equality at x,

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

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

Indeed,

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

Since from (14),

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

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 before, 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 revenue given by (12). □