

Screening with an Approximate Type Space*

Kristóf Madarász and Andrea Prat[†]

London School of Economics

June 11, 2010

Abstract

We re-visit the single-agent mechanism design problem with quasi-linear preferences, but we assume that the principal knowingly operates on the basis of only an approximate type space rather than the (potentially complex) truth. We propose a two-step scheme, the profit-participation mechanism, whereby: (i) the principal ‘takes the model seriously’ and computes the optimal menu for the approximate type space; (ii) but she discounts the price of each allocation proportionally to the profit that the allocation would yield in the approximate model. We characterize the bound to the profit loss and show that it vanishes smoothly as the distance between the approximate type space and the true type space converges to zero. Instead, we show that it is not a valid approximation to simply act as if the model was correct.

*We thank Mark Armstrong, Philippe Jehiel, Thomas Mariotti, Jim Minifie, John Moore, Hervé Moulin, Ennio Stachetti, Balázs Szentes, and seminar participants at Cambridge, the CRETA-IAS Workshop, LSE, UCL, NYU and Caltech for useful suggestions.

[†]k.p.madarasz@lse.ac.uk or a.prat@lse.ac.uk, or London School of Economics, Houghton Street, WC2A 2AE, London UK.

1 Introduction

In their path-breaking analysis of organizational decision-making, Herbert Simon and James March argue that organizations recognize the limits imposed by our cognitive abilities and might develop institutions to achieve good results in the presence of such limits:

“Most human-decision making, whether individual or organizational, is concerned with the discovery and selection of satisfactory alternatives; only in exceptional cases is it concerned with the discovery and selection of optimal alternatives.” (March and Simon, 1958, p 162).¹

When applied to a specific organizational problem, their views spur economists to ask two related questions. Given the cognitive limits that it faces, could an organization find a solution that yields a near-optimal payoff? If so, do the features of the near-optimal solution differ systematically from those of the optimal solution?

This paper attempts to answer these two questions in the well-known setting of single-agent mechanism design with quasi-linear preferences. This model – often simply referred to as the ‘screening problem’ – has found numerous applications in economics, such as taxation, regulation, or labor markets. In its most common application, often called ‘nonlinear pricing’, a multi-product monopolist faces a buyer, or a continuum of buyers, and offers a menu of product specifications (quality or quantity) at different prices.

The standard formulation of the screening problem assumes that the principal knows the space in which the agent’s type lives along with the distribution of types. In this paper, we re-visit the problem by assuming

¹In a similar fashion, Gordon (1948) writes:

"The fear of bankruptcy and the even more widespread fear of temporary financial embarrassment are probably more powerful drives than the desire for the absolute maximum in profits. [...] Given the fog of uncertainty within which [the businessman] must operate, the limited number of variables his mind can juggle at one time, and his desire to play safe, it would not be at all surprising if he adopted a set of yardsticks that promised reasonably satisfactory profits." Quoted also in Armstrong and Huck (2010) <http://else.econ.ucl.ac.uk/papers/uploaded/359.pdf>

that the principal does not use or know the true type space and distribution. For reasons that we will discuss in detail below, in contrast our principal operates on the basis of only an approximate type space. The principal is aware that her model is potentially incorrect. At the same time, she has a sense of the quality of her model: she knows an upper bound on the distance between a true type and its closest type in the approximate type space. We call this distance *the approximation index*.²

Given this information, is there a near-optimal mechanism? Can the principal guarantee herself an expected payoff that is not much lower than it would if she could optimize based on the knowledge of the true space type?

In order to discuss type approximation in a meaningful way, there must be a notion of proximity on the preferences of different types. Our type space is Euclidean and the agent's payoff function is Lipschitz-continuous in his type. We do not make further functional assumptions on the agent's payoff, the principal's cost function, or the agent's type distribution.

Finding a near-optimal solution in this strategic setting poses a challenge that is, to the best of our knowledge, absent in non-strategic environments. Although all primitives are well-behaved, the fact that the agent best-responds to the principal's menu choice creates the potential for discontinuity in the principal's expected payoff as a function of the chosen menu. The effect of this discontinuity is heightened by two elements, which we mention now in vague terms but which will become clearer shortly. First, in the exact solution of the screening problem the principal's payoff function is discontinuous exactly at the equilibrium allocation: this is because profit maximization implies that for every allocation that is offered in equilibrium there must be a binding incentive-compatibility constraint or participation constraint. Second, outside the monotonic one-dimensional case of Mussa and Rosen (1978), binding constraints are not necessarily local (McAfee and McMillan 1988, Wilson 1993, Armstrong 1996, Rochet and Choné 1998, Rochet and Stole 2003). This makes approximation difficult: if a certain type

²This assumption can be further relaxed by assuming that the principal only knows that the approximation index is satisfied with a high enough probability.

is assigned a certain allocation, a small perturbation of that type may lead to a very different allocation.

There are conceivably two ways to achieve a valid approximation in these circumstances. One could make enough assumptions to guarantee that only local constraints are binding. The resulting approximation would be useful, however, only if the principal is certain that these assumptions are satisfied, and the known environments that satisfy this condition are quite restrictive. An alternative, more general approach, which we adopt in this paper, is to look for a scheme that works even in the presence of binding non-local constraints. As we shall see, the idea is to find solutions that are robust to violations of incentive compatibility constraints, in the sense that the damage generated by such violations is bounded.

The core of the paper proposes and studies a mechanism for finding near-optimal solutions to screening problems. Given an approximate type space and its corresponding approximation index, we define the *profit-participation mechanism*, based on two steps:

- (i) We compute the optimal menu based on the set of all feasible products as if the model type space was the true type space.
- (ii) We take the menu obtained in the first step, a vector of product-price pairs, keep the product component unchanged and instead modify the price component. In particular, we offer a discount on each product proportional to the profit (revenue minus production cost) that the principal would get if she sold that product at the original price. The size of the discount, which is determined by the mechanism, depends only on the approximation index.

We prove the existence of an upper bound on the difference between the principal's payoff in the optimal solution with the true space and in the solution found by the profit-participation mechanism. Such upper bound is a smooth function of the Lipschitz constant and the approximation index. This means that, for any screening problem, the upper bound vanishes as the approximation index goes to zero.

Profit participation yields a valid approximation because it takes care of binding non-local incentive-compatibility constraints. By offering a profit-related discount, the principal guarantees that allocations that yield more profit in the menu computed for the approximate types become relatively more attractive for the agent. Now, a type that is close to an approximate type may still not choose the product that is meant for that approximate type. If he chooses a different product, however, this must be one that would have yielded an approximately higher profit in the original menu – the difference is bounded below by a constant that is decreasing in the discount.

While a profit-related discount is beneficial because it puts an upper bound to the profit loss due to deviation to different allocations, it also has a cost in terms of lower sale prices. The discount rate used in the profit-participation mechanism strikes a balance between the cost and the benefit. As the approximation index goes to zero, a given upper bound to the profit loss can be achieved with a lower discount and hence the optimal discount goes to zero as well.

One may wonder whether there are other, perhaps more immediate, ways of achieving a valid approximation in our class of problems? To answer this question, we restrict attention to model-based mechanisms, namely approximation schemes that begin with step (i) of the profit-participation mechanism: they first “take the model seriously.” In a second stage, they leave the set of alternatives unchanged and they modify prices according to any rule. In particular this includes the naive mechanism, whereby the principal uses the optimal menu for the approximate type space.

We prove that any model-based mechanism that violates a profit-participation condition cannot be a valid approximation scheme: the upper bound to the profit loss does not vanish as the approximation index goes to zero. This means that if there exist mechanisms that do at least as well as the profit-participation mechanism, they must either be very similar, in that they contain an element of profit participation, or radically different, because they do not begin from the exact solution for the approximate type space. The theorem implies that the naive mechanism is not a valid approximation:

the principal should not simply act as if her model was correct.

The economic insight from our result is that models can play a useful role in screening as long as the risk of mis-specification is dealt with in an appropriate manner. A principal who faces a complex screening model, but has only an imperfect model of the type space, can start by taking the model at face value and find its optimal solution. However, the resulting allocation is not robust to model mis-specification. To make sure that small errors in the model do not lead to bad allocations, the principal must act “magnanimously” by returning to the agent some of the profit that she would make if her model was true. Such magnanimity takes the form of a discount that is greater for more lucrative products.

Finally, let us ask why the principal uses an approximate type space to start with. We name three possible answers and our results have a different interpretation in each of these three cases.

Our preferred interpretation, which is also the most immediate, is that the principal – or the economist interested in modelling the problem at hand – is unsure about the agent’s preferences and has no way of resolving this uncertainty. She has, however, a model of the agent’s preferences and is willing to take a stand on at most how far her model could be from the truth: the approximation index. Our result provides comfort to the principal. Even if her model is misspecified, she can still use it to compute a menu. As long as she discounts the menu appropriately, she can place a bound on her loss. The more faith the principal has in her model, the lower is the necessary

discount and the is smaller the loss.³⁴

Under this model uncertainty interpretation, it is important to note that the principal need not know the true type space either to construct the profit-participation mechanism or to compute the upper bound on her loss. It is sufficient to know how good the approximate type space must be, as measured by the approximation index. Accordingly, there is a set of multiple prior distributions on the true type space consistent with the approximation index. The upper bound that we find in our main theorem applies to all of these.

In the second interpretation, the principal knows all the primitives of the model, but faces computation costs. Single-agent mechanism design has been proven to be an NP-complete problem even when the agent has quasilinear payoffs (Conitzer and Sandholm 2003). To reduce the heavy computation burden the principal may replace the true type space with a smaller one. By combining a method for partitioning the type space and the profit-participation mechanism, we obtain what computer scientists refer to as a polynomial-time approximation scheme (PTAS): a valid approximation of the exact solution which requires a computation time that is only polynomial in the size of the input.

The third interpretation is in terms of sampling costs. Suppose that the principal is not concerned with computation cost, but she does not know what the agent preferences only the structure of the type space. She can however sample the type space. For a fixed marketing fee, she can

³The model uncertainty interpretation is related to the notion of imprecise probability (see e.g., Walley 1991).

Related concepts are applied to macroeconomics (Hansen and Sargent 2010), but in non-strategic settings.

Bergemann and Schlag (2007) study monopoly pricing (with one good of exogenous quality) under model uncertainty, with two possible decision criteria: maximin expected utility and minimax expected regret.

Our problem is difficult to interpret as a game between the principal and an adversarial nature which can choose the agent's type. In our set-up such adversarial nature would have only a limited ability to deviate from the principal's model distribution.

⁴Although there is a potential relation between model uncertainty and quantal response equilibrium (McKelvey and Palfrey 1995), there does not appear to be a direct interpretation of our results in a QRE sense.

observe the payoff function of a particular type. By incurring this sampling cost repeatedly, she can sample as many types as she wants. The profit-participation mechanism, as stated above, supplies the principal with an approximate solution whose total sampling cost is polynomial in the input size. In this interpretation, the principal first performs a market analysis leading to the identification of a limited set of typical consumers. Then, she tailors her product range to the approximate type space and prices it ‘magnanimously’ in the sense above.

In the last two interpretations we can perform a comparative statics exercise on the cognitive limits of the principal. Suppose that the principal is constrained to solving the problem in at most N time units (in the computation time story) or sampling at most N types (in the search cost story). In both cases, the principal will select the approximate type space optimally. As a result of this we can show that as the principal’s resources decrease (N goes down): (i) The approximate type relies on a rougher categorization; (ii) The menu contains fewer alternatives; (iii) Pricing becomes more ‘magnanimous’.

The paper is structured as follows. Section 2 introduces the screening problem and defines the notion of an approximate type space. Section 3 develops profit-participation pricing and establishes an approximation bound (Lemma 1). Section 4 shows the main result of the paper, namely that the profit-participation mechanism is a valid approximation scheme (Theorem 4). In section 5 we discuss the three possible interpretations of our results. Section 6 shows that model-based mechanisms are valid approximation schemes only if they contain an element of profit participation (Theorem 5). Section 7 concludes.

1.1 Literature

To the best of our knowledge, this is the first paper to discuss near-optimal mechanisms when the principal uses an approximate type space.

Of course, there is a large body of work on approximation, in many disciplines. However, as we argued in the introduction, strategic asymmetric

information settings such as ours generate non-standard discontinuity issues.⁵ There are only a small number of papers that study approximation in mechanism design, which we attempt to summarize here.

In economics, the closest work in terms of approximation in mechanism design is Armstrong (1999), who studies near-optimal nonlinear tariffs for a monopolist as the number of product goes to infinity, under the assumption that the agent's utility is additively separable across products. He shows that the optimal mechanism can be approximated by a simple menu of two-part tariffs, in each of which prices are proportional to marginal costs (if agent's preferences are uncorrelated across products, the mechanism is even simpler: a single cost-based two-part tariff). Clearly, there are many differences between our approach and Armstrong's. Perhaps, the most important one is that his approximation moves from a simplification of the contract space while we operate on the type space.

Nisan and Segal (2006) discuss approximation schemes for multiple-agent mechanism design. In this line of research, the designer is concerned about the communication burden of the mechanism, namely the total amount of information that must be communicated. The authors construct an approximate problem by discretizing the space of agents' valuations and they determine an upper bound to the communication burden needed to find an exact solution to the discretized problem, both in the general case and in particular cases. While communication complexity is an important issue with multiple agents, in our screening problem the communication burden of any mechanism is linear in the minimum between the size of the type space (if using a direct mechanism) and the size of the alternative space (with the indirect version).

Xu, Shen, Bergemann and Yeh (2010) study optimal screening with a one-dimensional continuous type when the principal is constrained to offer a limited number of products. They uncover a connection with quantization

⁵Note that screening problems are hard to solve exactly because of their strategic nature and the presence of asymmetric information. If we are willing to assume that either the principal maximizes total surplus (rather than profit) or that the type of the agent is observable, then the problem simplifies (See section 5 for a more formal discussion of this point)

theory and use it to bound the loss that the principal incurs from having to use coarser contracts. Our paper differs both because we look at environments where non-local constraints may be binding and because we impose restrictions on the model the principal uses rather than on the contract space.

A growing field of computer science, algorithmic mechanism design, approaches mechanism design from a computational complexity perspective (see Hartline and Karlin (2007) for a survey). Most the work in the area focuses on prior-free mechanisms, where the designer – as is the case in online mechanisms that must work for a range of environments – has no information on the agents’ prior distributions. Instead, we have in mind designers, such as most firms, that have some information about their agents and make use of it when deciding on what mechanism to use.

An exception to the prior-free principle is Chawla, Hartline, and Kleinberg (2007). They study approximation schemes for single-buyer multi-item unit-demand pricing problems. The valuation of the buyer is assumed to be independently (but not necessarily identically) distributed across items. Chawla et al. find a constant-approximation mechanism based on virtual valuations (with an approximation factor of 3). Our paper differs because we consider a general pricing problem and because our approximation operates on the type space rather than on the contract space.⁶

2 Model

From a set of available alternatives Y , the principal selects a subset of alternatives and assigns transfer prices $p \in \mathbb{R}$ to the elements of this subset. The resulting menu consists of a set of alternative-price pairs or allocations. Let’s denote a menu by $M = \{(y', p'), (y'', p'') \dots\}$. We assume that a menu always contains the outside option y_0 whose price p_0 is normalized here to be zero.

Once a menu is offered by the principal, the agent is asked to choose

⁶A few papers in management science study numerical heuristics in the context of monopoly pricing, Green and Krieger (1985) and Dobson and Kalish (1993).

exactly one item from this menu. Although we specify the model with a single agent, our setup equally applies to settings with a large number of agents.

The agent's preferences depend on his private type $t \in T$ drawn according to some probability distribution $f \in \Delta T$ with full support. In particular, the agent's payoff is his type-dependent valuation of the object y net the transfer price to the principal:

$$v(t, y, p) = u(t, y) - p \quad (1)$$

The principal's profit is the transfer price net the cost of producing the object:

$$\pi(t, y, p) = p - c(y) \quad (2)$$

The above assumption follows Rochet and Chone (1998), and much of the literature on non-linear pricing, in that the principal's payoff does not directly depend on the agent's type.

To determine the principal's expected profit from offering a menu M one has to account for the agent's choice behavior given this menu. In particular, for a fixed menu M , the agent's incentive-compatibility and participation constraints give rise to a type-dependent allocation profile $\{(y(t), p(t))\}_T$. Formally, the principal's expected profit given menu M is

$$\Pi(T, M) = \int_T [p(t) - cy(t)] f(t) dt \quad (3)$$

such that for all $t \in T$

$$u(t, y(t)) - p(t) \geq u(t, y') - p' \text{ for all } (y', p') \in M \quad (4)$$

Note, that there might be multiple profiles that satisfy the incentive compatibility constraint. Thus the above profit might not be unique. Although generically it will often be, when it is not, we define $\Pi(T, M)$ to be the maximal profit from the allocation profiles that satisfy the IC con-

straints. (The same will apply below when we define the principal's profit given an approximate type space.)

2.1 Assumptions

We make three main assumptions. First, the agent's type lives in a compact and connected set within a finite dimensional Euclidean space. Second, for any fixed alternative y , the agent's preferences are Lipschitz continuous in his type. Finally, the principal's expected payoff is bounded where the natural upper-bound equals the total surplus generated by the best possible alternative-type combination:

Condition 1 (Type Topology) $T \subset \mathbb{R}^m$ is a compact and connected set where D is the diagonal of the minimal hypercube to contain T .

Condition 2 (Lipschitz Continuity) There exists a number K such that, for any fixed $y \in Y$ and $t, t' \in T$ it follows that $\left| \frac{u(t,y) - u(t',y)}{d(t,t')} \right| \leq K$.

Condition 3 (Bounded Profit) The finite upper-bound equals

$$\Pi_{\max} = \sup_{y \in Y, t \in T} u(t, y) - c(y)$$

and without loss of generality, we set the lower bound on Π to equal 0.⁷

Given the above assumptions, we can identify an equivalence class of problems, by noting that affine transformations of the payoff functions leave our results unaffected. Hence we can normalize any two of the above three

⁷The third condition is implied by the first two plus the assumption that the cost function is bounded from below.

parameters. Specifically, for our purposes a set of problems characterized by some k

$$k = K \frac{D}{\Pi_{\max}}$$

is an equivalent class. We thus normalize $D = 1$ and $\Pi_{\max} = 1$, and refer to a problem by its Lipschitz constant K .

2.2 Stereotype Set and Profit

Our key point of departure is that the principal, when facing the above screening problem, does not have access to the full truth. Instead, she is constrained to operate on the basis of a model which might systematically differ from the truth. The truth in our setup is given by T and f , and the principal's model will be given by a pair S and f_S . Here S is some finite subset of the true type-space T and $f_S \in \Delta S$ is a probability distribution with full support. We refer to S as the stereotype set or equivalently as an approximate type space.

Analogously to the way the principal's expected profit was defined for T and f , we can define her expected profit if the stereotype model was true. This is a fictional object in the sense that neither S nor f_S are necessarily true. Nevertheless, as we will show in Section 4, this object plays an important intermediate role for the principal to estimate a bound on the true optimal profit. Given a fixed menu M and a resulting stereotype allocation $\{(y(t), p(t))\}_S$ the principal's expected profit under her model of S and f_S is given by:

$$\Pi(S, M) = \sum_{t \in S} f_S(t) (p(t) - c(y(t))) \quad (5)$$

such that for all $t \in S$

$$u(t, y(t)) - p(t) \geq u(t, y') - p' \text{ for all } (y', p') \in M.$$

2.3 Model Quality

The last component of our setup helps to establish a relationship between the principal’s model and the truth. In particular, we define a distance between the principal’s model and the truth that we call the true approximation index and denote it by $\varepsilon_{\text{true}}$.

To define the approximation index we have to start from T and take a partition \mathcal{P} of this set such that it satisfies two conditions: (i) each cell of this partition contains exactly one element of the stereotype set S ; (ii) for any given stereotype $\hat{t} \in S$, the probability $f_S(\hat{t})$ equals the integral of f over the set of types in T that belong to the same partition cell as \hat{t} . Holding this partition fixed, we can then determine the *maximal* Euclidean distance between a type t from T and a stereotype \hat{t} from S such that they belong to the same partition cell.

Note that multiple partitions of T will satisfy the above two conditions. Each of these will induce a potentially different maximal distance. We choose the partition with the smallest maximal distance and we call it the *best approximation partition*. Given the properties of T and f , this will always exist. The resulting minimal maximal distance is then the true approximation index $\varepsilon_{\text{true}}$. We call any ε such that $\varepsilon \geq \varepsilon_{\text{true}}$ is an approximation index.

The above structure and the approximation has a different meaning depending on which of the three interpretations of our framework – model uncertainty, computational complexity, searching the case of cost – one has in mind.

In the case of computational complexity or search costs, the principal knows the true T and f and hence can directly set $\varepsilon = \varepsilon_{\text{true}}$. She can decide on the approximation index as a function of her objectives. As we discuss it in Section 5 in detail, operating on the basis of a small S instead of T will often lead to significant reductions of complexity or search costs.

In the case of model uncertainty, the principal does not know the true T and f . Furthermore, in the sense that such uncertainty is radical, the principal does not formulate fully-specified Bayesian beliefs about these objects

either. However, she knows that her representation does not correspond to the truth and believes that relative to her model, T and f are such that $\varepsilon_{\text{true}} < \varepsilon$.

The approximation index is a form of prior knowledge, which imposes restrictions on what the true type space and distribution might be relative to the principal’s model. Importantly, to construct the Profit Participation Mechanism, the principal need not know T and f , only her model as given by S and f_S in addition to an upper bound on K and a valid approximation index ε . Importantly, while the approximation index ε does restrict the relationship between the truth and the principal’s model, it does not pin down the truth given the principal’s model. Instead, given an S and f_S pair, it allows for a great degree of flexibility about what the true type-space and type-generating probability distribution might be.

Finally, in this model uncertainty interpretation, our analysis can be easily extended to situations where the principal is not 100% certain that $\varepsilon_{\text{true}} \leq \varepsilon$. If the principal thinks that there is a small probability δ that $\varepsilon_{\text{true}} > \varepsilon$, we can simply modify the upper bound to the loss by adding a worst-case scenario (a profit of zero) that occurs with probability δ .

3 Profit-Participation Pricing

In this section we introduce the key component of our solution method. After defining such profit-participation pricing we prove a key intermediate result. This result is stated somewhat generally, because we will need to apply it twice, in a somewhat different fashion, in the proof of our main theorem.

Fix the principal’s model S with associated probability distribution f_S , the agent’s payoff function u (which need only be defined for stereotypes), the cost function c , a Lipschitz-constant K , and an approximation index ε . We define profit-participation pricing as follows:

Definition 1 *For any menu $M = ((y^1, p^1), \dots, (y^k, p^k))$ profit-participation*

pricing transforms this menu to be $\tilde{M} = ((y^1, \tilde{p}^1), \dots, (y^k, \tilde{p}^k))$, where

$$\tilde{p}^i = (1 - \tau)p^i + \tau c(y^i), \text{ for } i = 1, \dots, k$$

and

$$\tau = \sqrt{2K\varepsilon}.$$

The next lemma puts a bound on the profit loss when the principal replaces a particular menu with its profit-participation discounted version.

For any partition $\tilde{\mathcal{P}}$ of the type space T , let $\mathcal{S}(\tilde{\mathcal{P}})$ be the class of stereotype sets such that each partition cell contains exactly one stereotype.

Lemma 1 *Fix T , f , S , f_S , as well as the associated best approximation partition \mathcal{P} . Take any feasible menu M . Let \tilde{M} be the menu derived through profit-participation pricing. Take any stereotype set $S' \in \mathcal{S}(\mathcal{P}')$ where \mathcal{P}' is a partition of T that is at least as fine as \mathcal{P} . Then:*

$$\Pi(S', \tilde{M}) - \Pi(S, M) \geq -2\sqrt{2K\varepsilon}$$

Proof. Take any menu M and compute the discounted menu \tilde{M} . Consider two types \hat{t} and t that they belong to the same cell of \mathcal{P} . Suppose that $\hat{t} \in S$ and $t \in S'$. We have to distinguish between two cases.

1. When \tilde{M} is offered, t chooses the allocation $y(\hat{t})$ meant for \hat{t} . Here, the only loss for the principal is due to the price discount determined by τ :

$$\tilde{p}(\hat{t}) - c(y(\hat{t})) = (1 - \tau)(p(\hat{t}) - c(y(\hat{t})))$$

2. When \tilde{M} is offered, t chooses the allocation y' different from $y(\hat{t})$. By the Lipschitz condition and the ε distance limit (the two types belong to the same cell of \mathcal{P} , and hence cannot be more than ε away from each other), we know that

$$\begin{aligned} |u(\hat{t}, y(\hat{t})) - u(t, y(\hat{t}))| &\leq K\varepsilon \\ |u(\hat{t}, y') - u(t, y')| &\leq K\varepsilon \end{aligned}$$

When combining these inequalities they imply that utility differentials for t and \hat{t} cannot be too different:

$$u(t, y(\hat{t})) - u(t, y') \geq u(\hat{t}, y(\hat{t})) - u(\hat{t}, y') - 2K\varepsilon$$

As the next step of the proof, we consider a revealed preference argument. From the incentive compatibility constraints we know: (i) that t prefers y' to $y(\hat{t})$ and hence

$$\tilde{p}(y(\hat{t})) - \tilde{p}(y') \geq u(t, y(\hat{t})) - u(t, y')$$

and (ii) that \hat{t} preferred $\hat{y}(\hat{t})$ to y' and hence:

$$u(\hat{t}, y(\hat{t})) - u(\hat{t}, y') \geq p(y(\hat{t})) - p(y')$$

If we combine the last two inequalities with the Lipschitz bound, we obtain

$$\tilde{p}(y(\hat{t})) - \tilde{p}(y') \geq p(y(\hat{t})) - p(y') - 2K\varepsilon$$

Substituting the definition of the discounted price \tilde{p} , we get:

$$\tau(p(y') - c(y') - (p(y(\hat{t})) - c(y(\hat{t})))) \geq -2K\varepsilon$$

This inequality guarantees that a non-served type, $t \notin S$, will never choose an allocation y' that is much worse than $y(\hat{t})$ for the principal.

In total there are two sources of profit loss: (i) the price discount, (ii) the deviation of types from the chosen option of their stereotypes. The former is bounded by:

$$\pi(t, y', \tilde{p}(y')) - \pi(t, y', p(y')) \leq -\tau\pi(t, y', p(y')) \leq -\tau$$

and the latter is bounded by:

$$\pi(t, y', p(y')) - \pi(t, y(\hat{t}), p(y(\hat{t}))) \leq -\frac{2K\varepsilon}{\tau}$$

To minimize total profit loss we can choose the discount rate τ such that for any type $t \in S'$,

$$\begin{aligned} \pi(t, y', \tilde{p}(y')) - \pi(t, y(\hat{t}), p(y(\hat{t}))) &\geq -\tau - \frac{2K\varepsilon}{\tau} \\ &= -2\sqrt{2K\varepsilon} \end{aligned}$$

This concludes the proof. ■

The lemma is the crucial result of the paper. It shows that profit participation puts an upper bound to the loss that the principal suffers if the type space is not what she thought it was. The existence of this bound is based on a trade-off introduced by profit-participation pricing. First offering a price discount leads to a loss to the principal proportional to τ . At the same time, profit-based discounts also loosen the agent's incentive compatibility constraints in a particular way. In the discounted menu allocations that generate higher profit to the principle become now relatively more attractive to the agent. Hence if a type from the refined S' chooses differently than its assigned stereotype from S it not only follows that this type realizes a utility similar to its assigned stereotype, by virtue of the Lipschitz condition, but also that this deviation cannot hurt the principal too much. Furthermore, the greater is the profit-based discount, the smaller is the potential loss that the principal might need to suffer due to a deviation. Setting $\tau = \sqrt{2K\varepsilon}$ optimizes on this trade-off between the loss from lower prices and the loss from deviations and establishes the above upper bound.

4 Profit Participation Mechanism

In the previous section we did not mention optimality. Neither the set of alternatives offered to the agent nor the prices of these alternatives satisfied were chosen with expected profit in mind. We now introduce a solution method that combines finding the optimal menu for a stereotype set and then adapting this menu to the true type set through profit-participation pricing:

Definition 2 *The profit-participation mechanism (PPM) consists of the following steps:*

- (i) *Compute an optimal menu \hat{M} for the screening problem defined by $S, f_S, \varepsilon, Y, u, \pi$;*
- (ii) *Apply profit-participation pricing on \hat{M} to obtain a discounted menu \tilde{M} .*

PPM takes the pricing problem described in Section 2 as its input. It produces an output that consists of a menu \tilde{M} . Our focus is the difference between the optimal profit that the principal could achieve by solving the screening problem for the true type space and the profit that she can expect when she offers \tilde{M} to the true type space. This comparison gives rise to a notion of the approximation loss.

Definition 3 *The PPM loss is the difference between the expected profit in the optimal solution of the true type space and the expected profit if the menu found through PPM is offered (to the true type space).*

We can now state the main result of this paper:

Theorem 4 *The PPM loss is bounded above by $4\sqrt{2K\varepsilon}$.*

Proof. Step 1. Define the optimal mechanism

$$M^* = \arg \max_M \Pi(T, M).$$

to be an allocation vector that maximizes the principal's expected profit subject to the IC constraints and contains the outside option. Let's denote this optimal profit by $\Pi(T, M^*)$.

The optimal mechanism M^* and hence the maximal profit $\Pi(T, M^*)$ are unknown objects and they remain unknown in our approach. In fact, all the sets and menus in the proof are not known to the principal, except the ones found through PPM.

Step 2. For the rest of the proof, fix \mathcal{P} to be the best approximation partition (which by definition achieves a true approximation index not greater than ε). Among all possible stereotype sets $\mathcal{S}(\mathcal{P})$, pick $S_{\max} \in \mathcal{S}(\mathcal{P})$ to maximize the principal's expected profit given that M^* is offered. Formally,

$$S_{\max} \in \arg \max_{S \in \mathcal{S}} \Pi(S, M^*)$$

The principal's profit when the agent's type is restricted to S_{\max} must be better than the optimal profit:

$$\Pi(S_{\max}, M^*) \geq \Pi(T, M^*)$$

Step 3. We now apply Lemma 1 for the first time. We begin with menu M^* offered to $S_{\max} \in \mathcal{S}(\mathcal{P})$. We discount the menu according to profit-participation pricing, thus obtaining a new menu M' . The inequality in the lemma holds for any partition \mathcal{P}' which is at least as fine as \mathcal{P} and for any $S \in \mathcal{S}(\mathcal{P}')$; so in particular it holds for $S \in \mathcal{S}(\mathcal{P})$. So we conclude that for $S \in \mathcal{S}(\mathcal{P})$:

$$\Pi(S, M') - \Pi(S_{\max}, M^*) \geq -2\sqrt{2K\varepsilon}.$$

Step 4. Take any stereotype set $S \in \mathcal{S}(\mathcal{P})$ and pick a menu \hat{M} that is optimal for that stereotype:

$$\hat{M} \in \arg \max_M \Pi(S, M)$$

By definition given stereotype set S , this menu \hat{M} is better for the principal than using menu M' which we defined in Step 3. Hence

$$\Pi(S, \hat{M}) \geq \Pi(S, M')$$

Step 5. Let us apply Lemma 1 for the second time. We now take the partition \mathcal{P}' to be the finest possible partition, namely T . We discount \hat{M} through profit-participation pricing to become \tilde{M} . The lemma guarantees

that:

$$\Pi(T, \tilde{M}) - \Pi(S, \hat{M}) \geq -2\sqrt{2K\varepsilon}$$

Summing up the above five steps:

$$\Pi(T, M^*) = [\text{max profit}] \quad (\text{Step 1})$$

$$\Pi(S_{\max}, M^*) \geq \Pi(T, M^*) \quad (\text{Step 2})$$

$$\Pi(S, M') \geq \Pi(S_{\max}, M^*) - 2\sqrt{2K\varepsilon} \quad (\text{Step 3})$$

$$\Pi(S, \hat{M}) \geq \Pi(S, M') \quad (\text{Step 4})$$

$$\Pi(T, \tilde{M}) \geq \Pi(S, \hat{M}) - 2\sqrt{2K\varepsilon} \quad (\text{Step 5})$$

and hence the profit-loss due to using \tilde{M} instead of the optimal M^* is bounded by:

$$\Pi(T, \tilde{M}) \geq \Pi(T, M^*) - 4\sqrt{2K\varepsilon}$$

■

The theorem constructs a bound to the PPM loss by applying Lemma 1 twice. The first application bounds the difference between the maximal profit with the true type space and the stereotype profit with any stereotype set satisfying the approximation index ε given the optimal menu for the true type space. The second application bounds the difference between the maximal stereotype profit and the profit that the principal obtains in the true type space if she uses the discounted version of the best menu for the stereotype set. Taken together, the two steps bound the difference between the maximal profit and the profit obtained with the discounted version of the optimal menu based on the principal's model.⁸

5 Interpretation and Comparative Statics

So far, we have not discussed in detail why the principal might operate only on the basis of an approximate model. Why is she not just using the true

⁸When using PPM, the principal actually only uses the second application of Profit-Participation Pricing. The first application is only used to prove the bound in the proof.

type space? In this section, we propose three answers to this question.

5.1 Model Uncertainty

The first interpretation – which is also the one we prefer – is immediate. The principal faces non-probabilistic uncertainty about the model (Walley 1991). She has a model of the agent’s type space, the stereotype set S with the stereotype distribution function. She knows that the model may be wrong but she does not know in which direction, however she knows the upper bound to the approximation index ε . Theorem 4 allows the principal to achieve a minimum expected profit and guarantees that, as ε vanishes, this lower bound converges to the expected profit she would have if she knew the true type space. Thus the PPM supplies a solution that is valid for a whole range of true $\{T, f\}$ pairs.

The rest of this section discusses in more detail the other two interpretations: computational complexity and search cost. We conclude this section by providing a simple comparative static property of PPM in light of constraints on thinking/searching costs for the principal.

5.2 Computation Cost

To understand our contribution, it is useful to provide some background information on the computational complexity of screening.

Unless we have a ‘clever algorithm’ (to be discussed below), finding an exact solution to our screening problem requires a computation time that is exponential in the size of the input. To see this, assume for now that T is discrete (even though computational complexity notions can be extended to uncountable). Note that under the Revelation Principle we can solve the problem in two stages: (i) For each possible allocation of alternatives to types, we see if it is implementable and, if it is, we compute the profit-maximizing price vector; (ii) Given the maximized profit values in (i), we choose the allocation with the highest profit. While given a finite T , each step (i) is a linear program, the number of allocations that we must consider in (i) is as high as $(\#Y)^{\#T}$. The number of steps we must perform can

then grow exponentially in the size of the input.⁹ This means that finding the exact solution could take too long even for relatively modest instances. The time or resources necessary to find the optimal solution to the screening problem might well become unbounded given this brute-force algorithm.

Before trying to improve on the brute force algorithm, it is useful to note that the complexity of screening problems depends on two joint assumptions: asymmetric information and conflict of interest. If either of these assumptions is missing, we can find an exact solution in polynomial time. If there were no asymmetric information and the principal could condition contracts on the agent's type, she would simply offer agent t the surplus-maximizing allocation

$$y^*(t) \in \arg \max_y u(t, y) - c(y)$$

at price

$$p^*(t) = u(t, y^*(t)) - u(t, y_0)$$

This would involve just $\#Y \cdot \#T$ steps.

If there were no conflict of interest – namely, if we wanted to maximize the surplus $u(t, y) - c(y)$ – it would be even simpler. The principal would offer all alternatives, each of them at the cost of production ($p(y) = c(y)$). The agent would select $y^*(t) \in \arg \max_y u(t, y) - c(y)$. Solving this problem would involve just $\#Y$ steps.

Obviously, the time required to find the exact solution would be lower if there existed a clever algorithm. Conitzer and Sandholm (Theorem 4, 2003) have shown that the problem of finding an exact solution to single-agent mechanism design with quasilinear utility is NP-complete.¹⁰

In the case of computational complexity, the principal knows all prim-

⁹If there are less products than types, it may be quicker to compute an indirect mechanism rather than invoke the Revelation Principle and compute the direct mechanism. To achieve the exponential bound, assume for instance that $\#Y = a\#T$, where $a > 1$, and increase both the number of types and products.

¹⁰Their proof refers to the case where the seller's cost function may depend on the buyer's type ($c(y, t)$ in our notation), while we assume that it depends only on product specification ($c(y)$), in line with Rochet and Choné (1998). A polynomial-time exact algorithm for the case with $c(y)$ is not known. However, if it existed, then our approximation scheme would still offer a time saving that is polynomial in the size of the input.

itives of the model, but pays a cost $c > 0$ for each unit of computational time. As we argued before, absent a smart algorithm, solving the computational time needed for the solution with the true type space is $\#Y^{\#T}$. The computation time for PPM is of the order of

$$\#Y^{\#S}$$

and thus it is polynomial in the number of possible alternatives, $\#Y$, independent of the number of types $\#T$, and exponential in the size of the stereotype space $\#S$. The above theorem implies that through profit-participation pricing, such a reduction in computational time can be achieved at an approximation cost of $4\sqrt{2K}\varepsilon$. This means that our mechanism is particularly successful in reducing the complexity of the type space. Once the principal is satisfied with, say, a 1% profit loss, her computation cost is independent of the complexity of the type space.

We can formalize these properties of PPM. To do so, we adopt here the definition whereby an algorithm is a polynomial-time approximation scheme (PTAS) if it returns a solution that is within a factor ε of being optimal (AS) and for every ε , the running time of the algorithm is a polynomial function of the input size (PT).

Proposition 1 *PPM yields a polynomial-time approximation scheme that is constant in T and polynomial in Y .*

Proof. Consider $\mathcal{S}(\mathcal{P})$ and pick a stereotype set S such that the cardinality of the stereotype set is minimal while the partition still satisfies the ε maximal distance property. Let $Q(\varepsilon)$ stand for the smallest cardinality of such a stereotype set S . To find an upper-bound on $Q(\varepsilon)$, let us partition the type space into identical m -dimensional hypercubes with diagonal length ε . Given such a partition, the maximal number of stereotypes we need is:

$$\bar{Q}(\varepsilon) = \left(\frac{1}{2\varepsilon}\right)^m \tag{6}$$

Note that this upper bound is tight if types are uniformly distributed on the

type space and the number of true types goes to infinity.

We can now prove that the profit participation scheme is an approximation scheme (AS). This is true because

$$\lim_{\varepsilon \rightarrow 0} 4\sqrt{2K\varepsilon} = 0$$

To prove that PPM is polynomial in time (PT), fix an $\varepsilon > 0$ and note that the cardinality of the minimal stereotype set S here is

$$\#S = \bar{Q}(\varepsilon) = \left(\frac{1}{2\varepsilon}\right)^m$$

Thus, the total computation time of PPM is proportional to the number of steps needed to compute the optimal mechanism for the stereotype set S . The Revelation Principle guarantees that this number is bounded above by

$$\#Y^{\#S}$$

Hence, for any given ε , the dimension of the stereotype space $\#S$ is fixed, and the computation time of PPM is polynomial in the input size $\#Y * \#T$.¹¹

■

The proposition is proven by showing that, for any ε , it is possible to construct a stereotype set such that every type is at most ε away from a stereotype. This bounds the exponent of the term $\#Y^{\#S}$. The computation time then becomes polynomial in $\#Y$ and constant in $\#T$. The stereotype set is constructed by partitioning the whole type space in hypercube and selecting the mid point of each cube as a stereotype.

¹¹A more stringent notion of approximation quality, fully polynomial-time approximation scheme (FPTAS), requires the computation time to be polynomial not only in the input size but also in the quality of the approximation, namely in $\frac{1}{\varepsilon}$. It is easy to see that this requirement fails here. A designer who wants to move from a 1% approximation to a 0.5% approximation, a 0.25% approximation, and so on, will face an exponentially increasing computation time. However, it is known that many problems – all the “strongly” NP-complete ones (Garey and Johnson, 1974) – do not have an FPTAS.

5.3 Sampling cost

Our analysis has an alternative interpretation in terms of sampling cost. Suppose that the principal knows the set of possible types, T , and the set of possible alternatives, Y , but does not know the payoff function of the agent: $u : T \times Y \rightarrow \Re$ (but she knows that u satisfies the Lipschitz condition for K). The principal can choose to sample as many types as she wants, but each sampling operation entails a fixed cost γ . Sampling is simultaneous, not sequential. The principal chooses a sampling set S ex ante. By equating the sampling set S with the stereotype set, we can apply PPM, as defined above. Theorem 1 guarantees that the resulting pricing scheme is an ε -approximation of the optimal pricing scheme.

5.4 Comparative Statics on the Principal's Cognitive Resources

Assume now that there is some binding constraint either on computational time or on total sampling cost, and let's denote this constraint by N . Then using PPM to solve the non-linear pricing problem gives rise to the following comparative static results with respect to N .

Corollary 1 *If the principal uses PPM, then, as N decreases:*

- (i) $\#\hat{Y}$ decreases (there are fewer items on the menu)
- (ii) $\#S$ decreases (the type model is based on a rougher categorization)
- (iii) τ increases (the principal prices alternatives in more magnanimous way).

In words, as the thinking or searching cost rises in our framework, a principal who uses PPM to solve our problem will use mechanisms that are simpler as measured in the number of distinct items offered to the agent. At the same time, the principal will derive this mechanism based on a rougher categorization of true types into stereotypes, but will offer the items at a greater discount when compared to the prices optimal for the stereotypes.

6 Other Approximation Schemes

While the profit-participation mechanism is a valid approximation scheme, can we guarantee that there are no other mechanisms that ‘do better’? The performance of any approximation scheme depends on the class of problems to which it is applied. According to the No Free Lunch Theorem of Optimization (Wolpert and McCreedy 1997), an elevated performance over one class of problems tends to be offset by performance over another class. In our case, the more prior information the principal has, the more tailored the mechanism can be. For more restrictive classes of problems (e.g. one-dimensional problems with the standard regularity conditions), it is easy to think of mechanisms that perform better than the profit-participation mechanism. But a more pertinent question is whether there are other valid mechanisms for our same general class of problems.

As indicated by the No Free Lunch Theorem of Optimization (Wolpert and McCreedy 1997), the performance of an approximation scheme should be evaluated for a given set of problems. In this paper, we focus on a large class of screening problems, limited only by the Lipschitz constant K and the approximation index ε . In this section, we ask whether there are other mechanisms, besides PPM, that work for the class of problems under consideration.

We begin by defining the class of mechanisms that take the model seriously and then modify prices:

Definition 4 *A mechanism is model-based if it can be represented as a two-step process where first one performs step (i) of the PPM and then, leaving the alternative vector unchanged, modifies the price vector according to some function*

$$\tilde{p}(y) = \Psi(p(y), c(y), K, \varepsilon).$$

The function Ψ obviously does not operate on the price of the outside option y_0 , which is a primitive of the problem. We focus our attention on mechanisms that return minimal exact solutions, namely solutions where all alternatives offered are bought with positive probability.

The function Ψ can encompass a number of mechanisms. In the naive one, the principal takes the model seriously tout court, without modifying prices.

Example 1 *In the naive mechanism,*

$$\Psi(p(y), c(y), K, \varepsilon) = p(y)$$

In the flat discount mechanism, the principal acts magnanimously by discounting prices, but her generosity is not related to stereotype profits:

Example 2 *In the flat discount mechanism*

$$\Psi(p(y), c(y), K, \varepsilon) = p(y) - \delta$$

for some $\delta > 0$, which may depend on K and ε .

Finally, we can also represent the PPM in this notation:

Example 3 *In PPM*

$$\Psi(p(y), c(y), K, \varepsilon) = (1 - \tau)p(y) + \tau c(y)$$

for some $\tau > 0$, which may depend on K and ε .

The following definition is aimed at distinguishing between mechanisms depending on whether they contain an element of profit participation or not.

Definition 5 *A model based mechanism violates profit participation if for an $\varepsilon_{true} > 0$, there exists $\bar{p} > 0$ and $\bar{c} > 0$ such that for all $p' < p'' \leq \bar{p}$ and $c \leq \bar{c}$*

$$p'' - \Psi(p'', c, K, \varepsilon) \leq p' - \Psi(p', c, K, \varepsilon)$$

The principal lets the agent participate in her profits if she gives her a price discount that is strictly increasing in the profit (that she would made had she sold that alternative at the full price). Profit participation is

violated if there is a price/cost region, which includes the origin, such that this condition is violated, namely an increase in the principal's profit does not translate into a strict increase in the absolute value of the price discount.

It is easy to see that both the naive price mechanism and the flat-discount mechanism violate profit participation (indeed, they violate it for all values of $p'' > p'$ and c). Instead, with PPM, we have

$$p'' - \Psi(p'', c) = \tau p'' + \tau c > \tau p' + \tau c = p' - \Psi(p', c) \text{ for all } p'' > p'$$

Hence PPM never violates profit participation.

We now show that unlike the PPM, mechanisms that violate profit participation are not valid approximation schemes for the class of problems considered here.

Theorem 5 *The upper bound to the profit loss generated by a model-based mechanism that violates profit participation does not vanish as $\varepsilon \rightarrow 0$.*

Proof. It is useful to make the dependence of the mechanism on K and ε explicit. Hence in the proof we continue to write $\Psi(p, c; K, \varepsilon)$.

Suppose that the mechanism is model-based but violates profit participation for some $\bar{p} > 0$ and $\bar{c} > 0$. Select

$$\begin{aligned} p' &= \frac{1}{2}\bar{p} \\ p'' &\in \left(\frac{1}{2}\bar{p}, \min \left(\frac{1}{2}\bar{p} + \frac{1}{6}K, \bar{p} \right) \right) \end{aligned}$$

Suppose that $c = 0$. For the p' and p'' chosen above it must be that:

$$p'' - \Psi(p'', 0; K, \varepsilon) \leq p' - \Psi(p', 0; K, \varepsilon) \tag{7}$$

Define $h = p'$ and $q = p'' - p'$. Consider the following problem:

$$\begin{aligned}
T &= [0, 2] \\
f(t) &= \frac{1}{2} \text{ for all } t \in [0, 2] \\
Y &= [1, 2] \cup \{\bar{y}\} \cup y_0 \\
u(t, y) &= \begin{cases} h + q(t - 1 - 2|y - t|) & \text{if } y \in [1, 2] \\ h & \text{if } y = \bar{y} \\ 0 & \text{if } y = y_0 \end{cases} \\
c(y) &= 0 \text{ for all } y
\end{aligned}$$

In this screening problem, types below $t = 1$ prefer a “generic” alternative \bar{y} . Types above $t = 1$ prefer a “personalized” alternative $y = t$.

It is easy to see that in the optimal solution of this screening problem types below 1 buy \bar{y} at price h and each type $t > 1$ is offered a personalized alternative $\hat{y}(t) = t$ at price $h + q(t - 1)$. The principal’s expected profit is $h + \frac{1}{4}q$.

Note that the K -Lipschitz condition is satisfied by this problem. To see this, note that $u(t, y)$ is continuous in t for all y and that $\lim_{\tilde{t} \rightarrow t} \left| \frac{\partial}{\partial t} u(t, y) \right|$ reaches a maximum when $y > t > 1$, in which case, it is $3q$. This means that u satisfies a Lipschitz condition for $3q$. Given the normalization that $D = 1$, T must be halved to $[0, 1]$, implying a Lipschitz condition with $K = 6q$. This is always satisfied because, given the definition of q ,

$$6q = 6(p'' - p') \leq 6 \left(\min \left(\frac{1}{2}\bar{p} + \frac{1}{6}K, \bar{p} \right) \right) - 6\frac{1}{2}\bar{p} \leq K.$$

To show that the mechanism Ψ does not yield a valid approximation, we consider the following sequence of stereotype sets with associated stereotype

probability distributions:

$$\begin{aligned}
& \left\{ \begin{array}{l} S_0 = \{0, 1, 2\} \\ f_{S_0}(0) = f_{S_0}(2) = \frac{1}{4}; f_{S_0}(1) = \frac{1}{2} \end{array} \right. \\
& \left\{ \begin{array}{l} S_1 = \{0, \frac{1}{2}, 1, \frac{3}{2}, 2\} \\ f_{S_1}(0) = f_{S_1}(2) = \frac{1}{8}; f_{S_1}(\frac{1}{2}) = f_{S_1}(1) = f_{S_1}(\frac{3}{2}) = \frac{1}{4} \end{array} \right. \\
& \quad \vdots \\
& \left\{ \begin{array}{l} S_n = \{0, \frac{1}{2^n}, \dots, 1, 1 + \frac{1}{2^n}, \dots, 2\} \\ f_{S_n}(0) = f_{S_n}(2) = \frac{1}{2^{n+2}}; f_{S_n}(s) = \frac{1}{2^{n+1}} \text{ for all other } s \end{array} \right. \\
& \quad \vdots
\end{aligned}$$

Given the prior f , the true approximation index for stereotype set S_n is $\varepsilon_{\text{true}}^n = \frac{1}{2} \frac{1}{2^n} = \frac{1}{2^{n+1}}$. We set $\varepsilon_n = \frac{1}{2^{n+1}}$.

Hold n fixed. The exact solution of the screening problem for S_n involves offering \bar{y} at price $p(\bar{y}) = h$ as well as a vector of alternatives identical to the vector of types $\{1 + \frac{1}{2^n}, \dots, 2 - \frac{1}{2^n}, 2\}$, each of them priced at $p(\hat{y}(s)) = h + q(t - 1)$. The minimum price is h , while the maximum price is $h + q$. Hence, by our definition of h and q , all prices are between p' and p'' .

The mechanism returns the following prices:

$$\begin{aligned}
\tilde{p}(\bar{y}) &= \Psi(p(\bar{y}), 0; K, \varepsilon_n) = \Psi(h, 0; K, \varepsilon_n) \\
\tilde{p}(\hat{y}(s)) &= \Psi(p(\hat{y}(s)), 0; K, \varepsilon_n) = \Psi(h + qs, 0; K, \varepsilon_n)
\end{aligned}$$

Now recall that by definition $\Psi(p, c)$ violates profit participation. Hence, for any $t \in [1, 2]$,

$$h + q(s - 1) - \Psi(h + q(s - 1), 0; K, \varepsilon_n) \leq h - \Psi(h, 0; K, \varepsilon_n) \quad (8)$$

Now take any type $t \in [1, 2]$ which is not a stereotype (a set of measure

1 for every S_n) and consider his choice between the allocation meant for any stereotype $s \in [1, 2]$ modified by $\Psi(\hat{y}(s))$ at price $\Psi(h + q(s - 1), 0; K, \varepsilon_n)$ and the allocation meant for stereotypes below $t = 1$ (\bar{y} at price $\Psi(h, 0)$). If he buys $\hat{y}(s)$ he gets payoff

$$h + q(t - 1 - 2|s - t|) - \Psi(h + q(s - 1), 0; K, \varepsilon_n)$$

If he buys \bar{y} he gets utility

$$h - \Psi(h, 0; K, \varepsilon_n)$$

He chooses $\hat{y}(s)$ only if

$$q(t - 1 - 2|s - t|) - \Psi(h + qs, 0; K, \varepsilon_n) \geq -\Psi(h, 0; K, \varepsilon_n)$$

which, if one subtracts (8) from it, implies:

$$q(t - 1 - 2|s - t|) - q(s - 1) \geq 0,$$

which can be re-written as

$$t - s \geq 2|s - t|,$$

which is always false. Hence, all types that are not stereotypes choose \bar{y} rather than a nearby personalized alternative. For any S_n , the expected profit of the principal if she uses Ψ is h .

Hence, the limit of the profit as $n \rightarrow \infty$ ($\varepsilon_n \rightarrow 0$) is still h , which is strictly lower than the profit with the maximal profit with the true type, which, as we saw above, is $h + \frac{q}{2}$. ■

The intuition behind Theorem 5 has to do with the knife-edge nature of mechanisms that do not include an element of profit participation. In the exact solution of the problem with the approximate type space, there is a binding constraint (IC or PC) for every alternative that is offered. The profit-participation mechanism – and analogous schemes – manage to relax these constraints in the right direction. By adding slack to those constraints

that ensure that the agent does not choose alternatives with a lower profit. Mechanisms without the profit-participation feature return a price vector that still displays binding constraints – or, even, no longer satisfies those constraints.

This means that in mechanisms where profit participation is violated types near a stereotype might choose different allocations. If only local constraints are binding, the magnitude of such misallocations vanishes as $\varepsilon \rightarrow 0$. But in multi-dimensional screening problems non-local constraints will typically be binding. In this case, the magnitude of the misallocation does not vanish even as $\varepsilon \rightarrow 0$.

The proof above proceeds by constructing a relatively straightforward class of problems where in the optimal solution non-local constraints are binding for a set of true types with positive measure. Namely, we assume that the alternative space includes a generic inferior alternative and a continuum of type-specific alternatives. In the optimal solution, a portion of types face a binding incentive-compatibility constraint between a personalized alternative and the generic alternative. Hence, all types nearby a stereotype strictly prefer the generic alternative to this stereotype's optimal allocation. This creates a non-vanishing loss for the principal.

Of course, the class of problems used in the proof is a set of measure zero, within the set of all possible problems. However, all we need to proceed is a class of problems where non-local constraints may bind, and is often a generic feature of the studied multi-dimensional screening problem (Rochet and Choné 1998 or Rochet and Stole 2003).

Theorem 5 only applies to the class of model-based mechanisms. Thus, we can interpret it as saying that, if there are mechanisms, beside PPM, that constitute valid approximation schemes for screening problems, such mechanisms are: (i) either similar to PPM in that they are model-based and contain an element of profit participation; (ii) or radically different because they are not model-based. We leave the exploration of mechanisms in (ii) to future research.

7 Conclusion

We consider a principal who faces a screening problem but is constrained to operate on the basis of an approximate type-space. We characterize the upper bound to the expected loss that the principal incurs if she uses the profit-participation mechanism. We show that the loss vanishes as the approximate type space tends to the true one. We prove that this is not true for any similar mechanisms that do not contain a profit participation element.

The economic insight of this paper is that a principal who operates on the basis of an approximate type space cannot just ignore the mis-specification error. She can, however, find a simple way to limit the damage. It would be interesting to know whether this insight holds beyond our set-up. Our analysis has a number of limitations that future research could address. First, we assume that the principal's cost depend only on the product characteristics but not on the type of the agent (as in insurance problems). Second, we assume that there is only one agent (or a continuum thereof). It would be interesting to extend the analysis to multiple agents. Third, we restrict attention to quasilinear mechanisms.

References

- [1] Mark Armstrong. Multiproduct Nonlinear Pricing. *Econometrica* 64(1): 51–75, January 1996.
- [2] Mark Armstrong. Price Discrimination by a Many-Product Firm. *Review of Economic Studies* 66, 151–168, 1999.
- [3] Dirk Bergemann and Karl H. Schlag. Robust Monopoly Pricing. Cowles Foundation Discussion Paper No. 1527RR, October 2008.
- [4] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC 2007)*, pp. 243–251, 2007.

- [5] Chenghuan Sean Chu, Phillip Leslie, and Alan Sorensen. Bundle-Size Pricing as an Approximation to Mixed Bundling. Working paper, Stanford University, 2009.
- [6] Vincent Conitzer and Thomas Sandholm. Automated Mechanism Design: Complexity Results Stemming from the Single-Agent Setting. In *Proceedings of the 5th International Conference on Electronic Commerce (ICEC-03)*, pp. 17-24, Pittsburgh, PA, USA, 2003.
- [7] Gregory Dobson and Shlomo Kalish. Heuristics for Pricing and Positioning a Product Line Using Conjoint and Cost Data. *Management Science* 39(2), 160–175, February 1993.
- [8] M. R. Garey and D. S. Johnson. “Strong” NP-Completeness Results: Motivation, Examples, and Implications. *Journal of the ACM* 25(3), 499–508, July 1978.
- [9] Robert Gordon. Short-Period Price Determination in Theory and Practice. *American Economic Review* 3, 265-280 1948.
- [10] Paul E. Green and Abba M. Krieger. Models and Heuristics for Product Line Selection. *Marketing Science* 4(1): 1–19, Winter 1985.
- [11] Lars Peter Hansen and Thomas J. Sargent. Wanting Robustness in Macroeconomics. Working paper, NYU. March 2010.
- [12] Jason D. Hartline and Anna Karlin. Profit Maximization in Mechanism Design. In *Algorithmic Game Theory*, Editors: Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vizarani. October 2007.
- [13] McKelvey, Richard and Palfrey, Thomas. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* 10: 6–38, 1995.
- [14] James March and Herbert Simon. *Organizations*. New York: Wiley, 1958.

- [15] Preston McAfee and John McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory* 46: 335–354, 1988.
- [16] Michael Mussa and Sherwin Rosen. Monopoly and Product Quality. *Journal of Economic Theory* 18: 301–317, August 1978.
- [17] Noam Nisan and Ilya Segal. The Communication Requirements of Efficient Allocations and Supporting Prices. *Journal of Economic Theory* 129(1): 192-224, July 2006.
- [18] Jean-Charles Rochet and Philippe Choné. Ironing, Sweeping, and Multidimensional Screening. *Econometrica* 66(4), 783–826, July 1998.
- [19] Jean-Charles Rochet and Lars Stole. The Economics of Multidimensional Screening. in *Advances in Economics and Econometrics*, Vol 1, eds. M. Dewatripont, L.P. Hansen, and S. Turnovsky, Cambridge, 2003.
- [20] Walley, Peter. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.
- [21] Robert B. Wilson. *Nonlinear Pricing*. Oxford University Press, 1992.
- [22] Yun Xu, Ji Shen, Dirk Bergemann, Edmund Yeh. Menu Pricing with Limited Information. Working paper, Yale University, 2010.