

Screening with an Approximate Type Space*

Kristóf Madarász and Andrea Prat[†]
London School of Economics

December 17, 2010

Abstract

When designing institutions, we sometimes operate on the basis of approximate models of the underlying environment. Can we still achieve near-optimal solutions? We explore this issue for single-agent mechanism design problem with quasilinear preferences, where the principal knowingly uses a discrete model of the true type space. 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 model converges to the true type space. Instead, we show that solving the problem as if the model was correct is not a valid approximation.

1 Introduction

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

*We thank Mark Armstrong, Vincent Conitzer, Philippe Jehiel, Thomas Mariotti, Jim Minifie, Ilya Segal, Ennio Stachetti, Balázs Szentes, and seminar participants at Cambridge, the CRETA-IAS Workshop, LSE, UCL, NYU, Caltech, Gerzensee, Harvard-MIT, Yale, Stanford, Berkeley and Helsinki 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.

“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).

When applied to a specific organizational problem, their views spur economists to ask two related questions. Given the cognitive limits we face, can we find institutions that yield a near-optimal payoff? If so, do the features of the approximately optimal solution differ systematically from those of the optimal solution?

This paper attempts to answer these questions in the context of one of the best known problems in microeconomics: single-agent mechanism design with quasi-linear preferences. This model – commonly referred to as the ‘screening problem’ – has found a variety of important applications from taxation and regulation to insurance and labor markets. In its classic interpretation of nonlinear pricing, a multi-product monopolist offers a menu of product-price specifications to a buyer or a continuum of buyers (e.g. Wilson 1993).

In the standard formulation of the screening problem, the principal knows the true distribution of the agent’s type. We re-visit the general screening problem, but instead assume that the principal may not know or use the true type space. For reasons that we will discuss in detail below, our principal instead operates on the basis of an approximate type space. The principal is aware that her model is potentially misspecified and has a sense of the quality of her model. Can she guarantee herself an expected payoff that is not much below what she could expect if she optimized based on true space type?

To discuss type approximation in a meaningful way, we assume that the agent’s types live in a Euclidean space and there his payoff is Lipschitz-continuous in his type. The quality of the model is then how far apart the preferences of a true type and its closest type in the approximate model might be. Our principal knows an upper-bound on this distance which we call *the approximation index*.¹ Other than this, we make no further assumptions on the agent’s payoff, the principal’s cost function, or the agent’s type distribution.

To illustrate the above, consider a simple example where the agent’s preferences are determined by the geographical coordinates of his location. While

¹This assumption can be further relaxed by assuming that the principal only knows that the approximation index is satisfied with a high enough probability.

geography changes continuously the principal may use a simple model where all types in a certain area are assumed to live at some arbitrary fixed location therein (this is often the case with actual data). The principal does not know how agents are distributed within a subdivision. A measure of the quality of the approximation is the maximum distance within the largest area. One interpretation of our asymptotic analysis is that the geographical subdivision on which the principal's data is based becomes finer and finer (state, county, city, 5-digit zip code, 9-digit zip code).

Finding a near-optimal solution in our strategic setting poses a challenge that is, to the best of our knowledge, absent in non-strategic environments. Even when all primitives are well-behaved, the fact that the agent best-responds to the principal's menu choice creates room for discontinuity in the principal's expected payoff as the menu offered changes. The discontinuity is heightened by two elements. 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, (Mussa and Rosen 1978), there typically are some binding non-local incentive constraints (Wilson 1993, Armstrong 1996, Rochet and Choné 1998, Rochet and Stole 2003). This makes approximation difficult: a small perturbation of a type might lead to large changes in equilibrium choice behavior.

To address such discontinuities one could proceed in two different ways. By imposing enough assumptions one could guarantee that only local incentive constraints are binding. The resulting approximation would be valid, however, only if the principal were certain that these assumptions are satisfied. The set of known environments that satisfy this condition is extremely limited. An alternative and 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. 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. Our goal is not to find a mechanism that works very well in one specific environment, but rather to find one that produces an acceptable outcome for a large class of screening problems – one that uses generic preferences, cost functions, and type distributions.

It is essential to bear in mind that, in March and Simon's spirit, the mechanism we are looking for will only rely on information used by the principal. The

menu that is offered to the agent will only depend on the principal’s model and not on the true type space and distribution. The only link between model and truth that can enter the mechanism is the approximation index.

The core of the paper proposes a mechanism for finding near-optimal solutions to screening problems. We call our solution concept a *profit-participation mechanism*. 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 our profit-participation mechanism. Such upper bound is a smooth function of the Lipschitz constant and the approximation index. Hence 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 ways of achieving a valid approximation in our class of problems. We restrict attention to model-based mechanisms, namely approximation schemes that begin with step (i) of the profit-participation mechanism but which then can 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 which 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 to the one we propose, in that they contain an element of profit participation, or radically different, because they do not start 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 or 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 serious profit losses, the principal must act ‘magnanimously’. She needs to make the agent the residual claimant on part of her profit that she would make if her model was true. Such apparent generosity takes the form of a discount that is greater for more lucrative products.

Finally, let us ask why the principal uses an approximate model to begin with. We provide four answers and our results have different interpretations in each of these four cases. The most immediate interpretation – also the preferred one by the authors – 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 a model of the agent’s preferences however, and she 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. By discounting 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 smaller is the loss.²

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 preferences as in our setting (Conitzer and Sandholm 2004). To reduce this 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 computational costs are non-binding, but the principal does not know the agent's preferences only the structure of the type space. She can however sample the type space. For a fixed marketing fee, she can 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 fourth interpretation, the sole difficulty arises from the burden of communication faced by the agent. To execute a mechanism the agent needs to send information to the principal. But the communicational complexity of the optimal mechanism, the number of bits the agents needs to transmit to implement this, might be prohibitively high. Our solution method allows the principal to reduce the communication burden to any desired size in a very efficient manner. By limiting the size of the menu that is offered to the agent the principal can reduce the communication burden at the cost of some loss in profit. Our results show that using the PPM the communication burden of the problem is polynomial in the parameters of the model and in the inverse of the size of the profit loss. In fact, the communication burden grows only logarithmically in the reduction of the upper-bound on the profit loss.

²In section 5 we discuss the relation between our model uncertainty interpretation and ambiguity aversion as well as quantal response equilibria.

In the last three interpretations we can perform a comparative statics exercise on the cognitive limits of the principal, inspired by March and Simon’s goal of studying how near-optimal responses to bounded rationality shape organizational outcomes. 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 1). In section 5 we discuss the four 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 2). 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.

There is of course 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.³ There are only a small number of papers that study approximation in mechanism design, which we attempt to summarize here.

The number of contributions dealing with near-optimal mechanism design with a single agent is limited. Wilson (1993, section 8.3) discusses the approximate optimality of multi-part tariffs (with a one-dimensional type). 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 prod-

³Note 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)

uct 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). There are a number of key 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.⁴

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 theory and use it to bound the loss that the principal incurs from having to use coarser contracts. Again 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 literature at the intersection of computer science and economics on near-optimal mechanisms with multiple agents. Those papers may be concerned with computational complexity (e.g., Conitzer and Sandholm 2004) or communication complexity (e.g., Blumrosen, Nisan, Segal 2007). There, the main challenge has to do with the complexity in the allocation space created by the presence of multiple agents and multiple objects to be allocated. The preferences of individual agents for individual objects are usually relatively simple. Here instead, there is only one agent and the allocation space is relatively simple, while the difficulty lies in describing the type of the agent.

An exception is Chawla, Hartline, and Kleinberg (2007), who 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.⁵

The goal of our paper is closely related to the goal of Gabaix (2010). In that decision-theoretic framework, the agent operates on the basis of a “sparse” rep-

⁴See also Chu, Leslie, and Sorensen (2010) for a theoretical and empirical analysis of this problem.

⁵A few papers in management science study numerical heuristics in the context of monopoly pricing, Green and Krieger (1985) and Dobson and Kalish (1993).

resentation of the world, i.e. one that – like in our case – is much less complex than reality. The objective is to find the optimal sparse model compatible with the limited cognitive skills of the decision-maker. Gabaix’s set-up applies to a large number of economic problems. Our paper instead does not endogenize the approximate type space (except, partly, when looking for a lower bound to complexity). We focus on identifying and tackling the challenges to approximation that arise in strategic environments.

2 Setup

From a compact 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 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:

$$v(t, y, p) = u(t, y) - p$$

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

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

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.

2.1 Assumptions

As in most problems, approximation is only possible if the underlying environment is sufficiently smooth. If the principal does not have enough information to put some topology on the type space and to guarantee that nearby types

have similar preferences, it is difficult to see how she can find a near-optimal solution unless she has direct information on most types.⁶

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) $T \subset \mathfrak{R}^m$ is an uncountable, compact and connected set. Let D be the maximal Euclidean distance between any two points in T .

Condition 2 (Lipschitz Continuity) There exists a number k such that, for any $y \in Y$ and $t, t' \in T$,

$$\left| \frac{u(t, y) - u(t', y)}{d(t, t')} \right| \leq k$$

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.

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 parameters. Specifically, define

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

⁶There are two ways of viewing the type space in mechanism design. At a general level, it is an artificial construct, a labelling of the possible realizations of the agent's preferences which is devoid of direct economic meaning. However, sometimes the mechanism designer knows that the agent's preferences are determined by some underlying characteristics, like income, age, or location, and she builds a model where the type corresponds to actual characteristics. This second interpretation of the type space is more restrictive, but provides a more natural background for our asymptotic results.

We thus normalize $D = 1$ and $\Pi_{\max} = 1$, and refer to a problem by its normalized Lipschitz constant K . The class of problems defined by K is all the problems (each of which is described by some T, f, u, c, Y) such that the normalized Lipschitz constant is not larger than 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 latter 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 the approximate type space,

Analogously to the way the principal's true expected profit was defined, we can define what her expected profit would be if the principal's model was true. Again, this is a hypothetical object because 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 $M = ((y', p(y')), \dots, (y^k, p(y^k)))$, the principal's expected profit under her model of S and f_S is given by:

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

where $(y(t), p(y(t)))$ is the product selected by type t , namely, for all $t \in S$,

$$u(t, y(t)) - p(y(t)) \geq u(t, y) - p(y) \text{ for all } (y, p(y)) \in M,$$

with the proviso that, whenever the agent is indifferent between two or more allocations, he chooses the one that yields the highest profit to the principal. Of course, this definition can be extended to the true space as $\Pi(T, M)$.⁷

⁷At this stage, there are a number of equivalent ways to express the menu, the agent's choice, and the principal's expected profit. Perhaps, the most standard one is based on the use of a direct mechanism. For reasons that will become clear later, we prefer to use an indirect mechanism formulation in which allocations are indexed by the product.

2.3 Model Quality

For our asymptotic results, we need a measure of the approximation quality of the model. Such measure should satisfy two conditions. First, it should be a scalar that reflects some ‘distance’ between the model and the truth and it should go to zero as the model tends to the truth. Second, it should have a worst-case element, which will allow us to find upper bounds to the profit loss. When the measure goes to zero, its pessimistic nature will guarantee that any other non-worst case measure would go to zero too.

Given any true type space $T \subset \mathfrak{R}^m$ with distribution f and any stereotype set $S \subset \mathfrak{R}^m$ with distribution f_S , the true approximation index $\varepsilon_{\text{true}}$ is defined as follows:

1. An approximation partition \mathcal{P} is a partition of T with $\#S$ (possibly non-connected cells), such each cell contains exactly one stereotype, and for each cell the mass of true types belonging to that cell (computed according to density f) equals the probability (according to f_S) of that cell’s stereotype. Let Γ be the set of all approximation partitions.⁸
2. For each approximation partition \mathcal{P} in Γ , define $d(\mathcal{P})$ as the maximal value in any cell of the maximal distance between the stereotype and any true type in that cell.
3. The *true approximation index* of $(T, f; S, f_S)$ is $\varepsilon_{\text{true}} = \inf_{\mathcal{P} \in \Gamma} d(\mathcal{P})$.

In step 1, the true type space is partitioned into cells that correspond to stereotypes. The mass of types contained in each cell equals the probability assigned to that stereotype. For each of these partitions \mathcal{P} , there will be a true type that is furthest away from his corresponding stereotype: this is the distance $d(\mathcal{P})$. If we minimize $d(\mathcal{P})$ over all possible partitions, we obtain $\varepsilon_{\text{true}}$. Given the assumptions on T is in, it is easy to see that there exists a partition that achieves $\varepsilon_{\text{true}}$, which we call the *best approximation partition*.

It is important to note that the best approximation partition and the true approximation index do not have to be computed in order to apply our algo-

⁸For any model (S, f_S) the corresponding set Γ is nonempty. We can always construct an approximation partition as follows. Build a ball of increasing radius around every stereotype set. Balls around stereotypes increase simultaneously at the same speed. Stop expanding the ball around stereotype \hat{t} when the mass of the ball equals $f_S(\hat{t})$. Expanding balls can only incorporate types that do not already belong to other balls. It is easy to see that the outcome of this process is an approximation partition as defined above. Of course, some of the resulting cells may be unconnected.

rithm. All the agent needs is an upper bound to the approximation index, which we call ε (and from now we refer to simply as the approximation index).

The meaning of the approximation index ε depends on the interpretation of the model. In the first interpretation – model uncertainty – the principal does not know the true T and f . In the sense that such uncertainty is radical, the principal does not formulate fully-specified Bayesian beliefs about these objects. She knows however, that her representation, S and f_S , 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, while the approximation index ε does narrow down the set of models the principal might have, it does not pin down the truth given the principal’s model uniquely. 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.

Note that, in the 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 .

In the last three interpretations – computational complexity, search cost, and communication burden – the principal knows T and f and chooses S and f_S . Hence, she will first partition T in some way that is suitable to her goals and easy to compute (for instance, in section 5 we will make use of grids). Then she will pick S and f_S to fit the partition. This will determine a particular maximal distance ε (in the case of the grid, one can pick the center of the cube and ε is half the diagonal). The resulting partition may not be the best approximation partition and hence it may be that $\varepsilon > \varepsilon_{\text{true}}$. But, given that the approximation loss that we find is an increasing function of ε , our results are valid a fortiori.⁹

3 Profit-Participation Pricing

In this section we introduce the key component of our solution method, we call this component Profit-Participation Pricing, and prove an intermediate result.

⁹In the zip code example mentioned in the introduction, an upper bound to the approximation index is given by the maximal physical distance between any two US residents who share the same zip code.

We state this result in a somewhat general form, because we will need to apply it twice in slightly different specifications in the proof of our main theorem. The definition of Profit-Participation Pricing is as follows:

Definition 1 For any menu $M = ((y^l, p(y)), \dots, (y^k, p(y^k)))$, let the menu derived by Profit-Participation Pricing be $\tilde{M} = ((y^l, \tilde{p}(y)), \dots, (y^k, \tilde{p}(y^k)))$ where the product vector is unchanged and the new price vector $\tilde{p}(y)$ is given by

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

such that

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

In words, Profit-Participation Pricing leaves the product component of a menu fixed, but gives profit-based discounts with a constant fraction. To see this, note that the above transformation can be equivalently expressed as:

$$\overbrace{\tilde{p}(y) - c(y)}^{\text{new profit}} = (1 - \tau) \overbrace{(p(y) - c(y))}^{\text{old profit}}$$

In the rest of the analysis below, we 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 .

Before proceeding to our first result, an additional definition is needed.

Definition 2 For any partition $\tilde{\mathcal{P}}$ of T , let $\mathcal{S}(\tilde{\mathcal{P}})$ be the class of stereotype sets such that each cell in the partition contains exactly one stereotype.¹⁰

The next lemma puts a bound on the profit loss when the principal replaces a particular menu with its profit-participation discounted version and offers it to any stereotype set that can be induced by a partition which is at least as fine as the best approximation partition.

Lemma 1 Fix T , f , S , f_S , as well as an approximation partition \mathcal{P} with an associated approximation index ε . For any feasible menu M , let \tilde{M} be the menu

¹⁰For instance, if $\tilde{\mathcal{P}}$ is a regular square grid where individual squares have diagonal length l and $\varepsilon = \frac{1}{2}l$, $\mathcal{S}(\tilde{\mathcal{P}}, \varepsilon)$ is a singleton: the only stereotype set that satisfies this condition consists of all the types at the center of each square. If instead $\varepsilon > \frac{1}{2}l$, $\mathcal{S}(\tilde{\mathcal{P}}, \varepsilon)$ contains multiple elements.

derived through Profit Participation Pricing. Take any set $S' \in \mathcal{S}(\mathcal{P}')$ where \mathcal{P}' is a partition at least as fine as \mathcal{P} . Then:

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

The above result bounds the profit loss when comparing the expected profit given the principal's model and a finer model. The profit loss used in lemma 1 is somewhat peculiar. The benchmark $\Pi(S, M)$ assumes that the stereotype space is correct. Hence, the lemma puts a bound on the difference between the profit when using a discounted menu (in various scenarios corresponding to the various S') and the profit that the principal would get with the original menu in her imaginary model world.

This intermediate result will be used twice in the next section, to derive the main positive result of the paper. Hence two special cases of the lemma will be of particular interest to us. If $\mathcal{P}' = \mathcal{P}$, the lemma compares the stereotype set S with another stereotype set induced by the same partition. If $\mathcal{P}' = T$, then there is no stereotype set at least as fine as T , and the lemma compares an original menu given to stereotype set with the discounted menu given to the true type space. Here our result claims that offering the discounted menu to the true type-space leads to a profit loss which is bounded again by $2\sqrt{2K}\varepsilon$ when compared to the profit generated by the original menu in the principal's model.

Proof. Take any menu M and compute the discounted menu \tilde{M} . Consider any two types \hat{t} and t such that they belong to the same cell of \mathcal{P} and $\hat{t} \in S$ and $t \in S'$. Hence, the distance between the two types is bounded above by ε .

Note that if $t \in S' \in \mathcal{S}(\mathcal{P}')$ and \mathcal{P}' is finer than \mathcal{P} , then there exists an element $S \in \mathcal{S}(\mathcal{P})$ such that $S' \subseteq S$. hence, the distance between a type $t \in S'$ and stereotype $\hat{t} \in S$ is a fortiori bounded above by ε .

There are two possible cases: (i) t and \hat{t} choose the same product; (ii) t and \hat{t} choose different products. Case (i) is straightforward. Suppose when \tilde{M} is offered, t chooses the allocation \hat{t} chooses from M and let's denote this by $\{\hat{y}, p(\hat{y})\}$. Here, the only loss for the principal is due to the price discount determined by τ :

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

Hence, we focus on case (ii). Suppose when \tilde{M} is offered, t chooses an allocation y' different from \hat{y} . By the Lipschitz condition and the ε distance

limit, we know that

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

implying that utility differentials for t and \hat{t} cannot be too different:

$$u(t, \hat{y}) - u(t, y') \geq u(\hat{t}, \hat{y}) - u(\hat{t}, y') - 2K\varepsilon \quad (1)$$

This does not preclude, however, that the choices of the two types are different:

Next, consider a revealed preference argument. With the *original* price vector p , the stereotype \hat{t} prefers \hat{y} to y' :

$$u(\hat{t}, \hat{y}) - p(\hat{y}) \geq u(\hat{t}, y') - p(y') \quad (2)$$

With the *discounted* price vector, type t prefers y' to \hat{y} :

$$u(t, \hat{y}) - \tilde{p}(\hat{y}) \leq u(t, y') - \tilde{p}(y') \quad (3)$$

By subtracting (3) from (2), we get that

$$\begin{aligned} &p(y') - \tilde{p}(y') - (p(\hat{y}) - \tilde{p}(\hat{y})) \\ &\geq u(t, \hat{y}) - u(t, y') - (u(\hat{t}, \hat{y}) - u(\hat{t}, y')) \end{aligned} \quad (4)$$

By (1), the right-hand side of (4) is bounded below by $2K\varepsilon$, this follows from . Given the definition of \tilde{p} , the left-hand side of (4) can also be written as:

$$\overbrace{\tau(p(y') - c(y'))}^{\text{discount for } y'} - \overbrace{\tau(p(\hat{y}) - c(\hat{y}))}^{\text{discount for } \hat{y}}.$$

Summing up,

$$\tau(p(y') - c(y') - (p(\hat{y}) - c(\hat{y}))) \geq -2K\varepsilon \quad (5)$$

There are two potential sources of loss, one due to the deviation from \hat{y} to y' , the latter due to the price discount. The loss caused by the deviation given the above inequality is

$$p(y') - c(y') - (p(\hat{y}) - c(\hat{y})) \geq -\frac{2K\varepsilon}{\tau} \quad (6)$$

The loss due to the price discount is (recalling that profit is bounded above by Π_{\max} , which was normalized to 1),

$$\tilde{p}(y') - c(y') - (p(y') - c(y')) = -\tau(p(y') - c(y')) \geq -\tau \quad (7)$$

Adding these two together we get that

$$\tilde{p}(y') - c(y') - (p(\hat{y}) - c(\hat{y})) \geq -\tau - \frac{2K\varepsilon}{\tau} \quad (8)$$

We can now see the explicit trade-off between the two sources of loss. By optimizing on this, we can bound their sum. In particular, if we set τ equal to

$$\arg \min_{\tau} \tau + \frac{2K\varepsilon}{\tau} = \sqrt{2K\varepsilon}$$

we get that

$$\tilde{p}(y') - c(y') - (p(\hat{y}) - c(\hat{y})) \geq -2\sqrt{2K\varepsilon}.$$

Taking expectations appropriately, we get the statement of the lemma. ■

The lemma contains the main intuition for why this type of approximation scheme works. Profit participation puts a bound on the loss that the principal suffers if the type space is not what she thought it was. By offering profit-based price discounts, the principal ensures that allocations that generate higher profit to her become relatively more attractive to the agent. Profit-participation pricing is in effect a system of local incentives. The agent becomes a residual claimant on the principal's profit, and now types near stereotypes are encouraged to choose similarly high-margin allocations as the stereotypes.

A key feature of profit-participation pricing is that there is no guarantee that types close to a stereotype will choose in the same way as their respective stereotypes. The principal still does not know how often different allocations will be chosen by the agent. In fact, the principal cannot even guarantee that, when offered the discounted menu, stereotypes will choose the allocation they were choosing previously. However, the principal knows that whichever allocation types choose with the discounted menu, the deviation from the allocation chosen by stereotypes in the undiscouted menu cannot be very damaging to profit.

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 . Second, 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. The set of alternatives and the prices were not chosen with expected profit in mind; rather we considered simply any feasible menu. We now introduce the full version of our solution concept: we combine finding the optimal menu given the principal's model with modifying such a menu via profit-participation pricing.

Definition 3 *The profit-participation mechanism (PPM) consists of the following steps:*

- (i) Find an optimal menu \hat{M} for the screening problem defined by S, f_S, Y, u, c ;
- (ii) Apply profit-participation pricing to \hat{M} and obtain a discounted menu \tilde{M} .

PPM takes the pricing problem described in Section 2 as its input and outputs a menu \tilde{M} . Our focus now is on the profit difference comparing two scenarios: the principal's expected profit given the true optimal solution and the principal's expected profit if she offers \tilde{M} to the true type space. This comparison captures the approximation loss.

Definition 4 *Let the PPM loss be 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).*

Note that if the principal does not know the true model, she will typically not know either the true expected optimal profit or the true expected profit given menu \tilde{M} . We can now state the main result of the paper in terms of the known parameters of our setup.

Theorem 1 *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 which contains the outside option and maximizes the principal's expected profit subject to the IC constraints. 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 defined for (S, f_S, T, f) . The true approximation index is, by definition, 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,

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

The principal's profit when the agent's type is restricted to S_{\max} must clearly be at least as good as the optimal profit since fewer constraint will be binding. Hence we have that:

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

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 $S \in \mathcal{S}(\mathcal{P})$, which is the stereotype set that we are using. 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. However, the menu M' is not optimal for the stereotype set S . Pick a menu \hat{M} that is optimal for that stereotype set:

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

By definition,

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

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 proof of the theorem constructs the bound to the PPM loss by applying Lemma 1 twice. In the first application the lemma bounds the difference between the true optimal profit and the stereotype profit given any stereotype set satisfying the approximation index ε . The second application bounds the difference between the maximal stereotype profit and the true profit given the discounted menu identified by profit-participation pricing given the principal's mis-specified model. Taken together, the two steps bound the difference between the maximal profit and the profit obtained with the discounted version of the optimal stereotype menu.

Again, the bound is valid without requiring the principal to know anything beyond her model and two upper bounds to the inaccuracy of the model: the approximation index ε and the Lipschitz constant K .

5 Interpretation and Comparative Statics

As we mentioned in the introduction, we offer four different reasons why the principal might operate only on the basis of an approximate model. The first one, model uncertainty, is our main motivation. It is based on the fact that typically principals operate under weaker epistemic conditions than those that

the standard model assumes. The other three reasons are, respectively, bounds to computational capacity, sampling capacity, and communication capacity. The section ends with a comparative statics result on how the near-optimal solution changes as the cognitive limits of the principal become more binding.

5.1 Model Uncertainty

The model uncertainty is straightforward. In many situations, it is unrealistic to assume that the principal knows the true model. Rather she operates under weaker epistemic conditions, like some non-probabilistic uncertainty about the model (Walley 1991).¹¹ She knows that her model may be wrong but she does not know in which direction. She knows, however, an upper bound on how misspecified her model might be, e.g., she knows an approximation index ε . In this case, Theorem 1 allows the principal to achieve a minimum expected profit and to bound the difference between this profit and the optimal expected profit. Furthermore as her confidence in her model increases, as ε vanishes, this lower bound converges to the expected profit she would have if she knew the true type space.

Under this interpretation of model uncertainty, it is important to emphasize 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 that she knows an approximation index which will be consistent with a relatively large set of prior distributions on the true type space. Thus the upper bound in our main theorem applies to all of these distributions.

The fact that theorem 1 puts an upper bound on the inevitable loss due to model misspecification has a worst-case feel, akin to strong ambiguity aversion or maximin expected utility, e.g., Gilboa and Schmeidler (1989). Hence one may wonder why the principal should be preoccupied with worst outcomes rather than average outcomes. We would like to make two observations, however. First, as we discussed, the approximation index can be interpreted probabilistically. The principal might adopt a smaller approximation index that she believes is true only with probability $\delta < 1$. The upper bound on the profit loss then can be interpreted probabilistically in the exactly analogous manner. Second, when considering the asymptotic case where the discretized model tends to the truth,

¹¹Related 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 approach guarantees that PPM is still a valid approximation even if the principal has more prior information and uses it to compute some average case.

There might appear to be a relation between our model uncertainty interpretation and the quantal response equilibrium of McKelvey and Palfrey (1995). We can imagine that the model used by our principal excludes some dimensions of the agent’s type. This means that, as in a quantal-response equilibrium, if we take the model at face value, we still have some unexplained variability in the agent’s choice. However, there are two crucial differences. First, while quantal-response equilibria operate by perturbing the players’ strategies, our agent instead always plays best response. Second, quantal-response equilibrium postulates a particular form of perturbation, while our principal may not have such information.

5.2 Computation Cost

In our second interpretation, we return to the standard assumption where the principal knows all primitives of the model, hence T and f . She now pays some cost $l > 0$, however, for each unit of computational time used to calculate a menu from the primitives of her model. To understand our contribution given these assumptions, it is useful to begin with some background information on the computational complexity of screening. A reader who is familiar with these notions might want to skip this discussion.

To clarify the argument, let’s first assume that T is discrete (even though computational complexity notions can be extended to uncountable sets). Under the Revelation Principle we can solve the mechanism design 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. Here unless we have a ‘smart’ algorithm tailored to our screening problem, finding an exact solution in our setup requires a computation time that is exponential in the size of the input. 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.¹² This means that finding the exact solution could

¹²If 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.

take very long even for relatively modest instances. With an uncountable T , the time or resources necessary to find the optimal solution to the screening problem might well become unbounded given this brute-force algorithm.

Obviously, the time required to find the exact solution would be lower if there existed a 'smart' algorithm. However, Conitzer and Sandholm (Theorem 1, 2004) have shown that the problem of finding an exact solution to the single-agent mechanism design problem with quasilinear utility, the model we consider in our paper as well, is NP-complete. Unless $P = NP$, there does not exist a polynomial-time algorithm for screening problems.

Before trying to improve on the brute force algorithm, it is useful to note that the complexity of screening depends on two joint assumption of asymmetric information and conflict of interests. If either of these is missing, we can find an exact solution in polynomial time. Absent asymmetric information, the principal could condition contracts on t and offer the surplus-maximizing, fully rent-extracting allocation for each t :

$$\begin{aligned} y^*(t) &\in \arg \max_y u(t, y) - c(y) \\ p^*(t) &= u(t, y^*(t)) - u(t, y_0) \end{aligned}$$

This would involve $\#Y \times \#T$ steps. Absent conflicts 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. The agent would select $y^*(t) \in \arg \max_y u(t, y) - c(y)$. Solving this problem would involve just $\#Y$ steps.

Given Conitzer and Sandholm's negative result on polynomial-time exact algorithms, we are interested in knowing whether there exists at least a polynomial-time near-optimal algorithm. To the best of our knowledge, this problem has not been solved for mechanism design problems, outside specific functional and distributional specifications. As we will see, our approach can be useful.

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 $\#S$.

Theorem 1 implies that through profit-participation pricing, this very sig-

nificant reduction in computational time of the size 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 minimal hypercube which contains the type space into identical m -dimensional hypercubes with diagonal length ε . Given such a partition, the maximal number of stereotypes we need is:

$$\bar{Q}(\varepsilon) = \left(\frac{1}{2\varepsilon}\right)^m \quad (9)$$

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 $\#Y * \#T$.¹³ ■

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.

The goal of this section is to show that PPM-based algorithms may be useful in dealing with complex screening problems. The specific scheme we use in the proof is based on a crude subdivision of the type space into cubic cells. In practical instances, it can be greatly improved by partitioning the type space in ways that are tailored to the problem at hand. Note also that the profit bound identified in theorem 1 applies cell by cell: if the maximal distance in cell i is ε_i , the maximal profit loss is $4\sqrt{2K}\varepsilon_i$. In specific instances, the principal can achieve further improvements by partitioning the type space so that dense type regions are assigned to small cells.

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: $u : T \times Y \rightarrow \mathfrak{R}$ (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 guar-

¹³A 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 ε^{-1} . 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.

antees that the resulting pricing scheme is an ε -approximation of the optimal pricing scheme.

5.4 Communication Cost

Our set-up can also be interpreted in terms of communication complexity faced by the agent. (Kushilevitz and Nisan 1997). Mechanisms require the agent to send information to the principal. If implementing the optimal solution requires a large communication burden, the principal might turn to approximate solutions that can be achieved with less communication (Nisan and Segal 2006, Blumrosen, Nisan and Segal 2007).

In our case, a lower bound to the communication burden depends on the number of allocations on the menu offered by the principal. This is a best-case scenario in that it assumes that the principal and the agent have already established a common language to express the set of allocations on the menu in an efficient way, so that each alternative-price pair is now represented by a natural number, just like numbered items on certain restaurant menus. For any menu M , the communication burden is then $\log(\#M)$.

The communication burden for the optimal mechanism is extremely high. To implement the exact solution for the true type space, we may need a menu with same cardinality as T . We can use our approximation scheme to reduce this communication burden to any desired level while also putting a bound on the corresponding profit loss.

Suppose we wish to use a menu with at most N allocations – generating a communication burden of $\log(N)$. We can divide the type space into $n^m = N$ identical hypercubes (assume for simplicity that N is such that $N^{\frac{1}{m}}$ is an integer). The maximum distance within each cube is $\varepsilon = \frac{1}{n}$. By theorem 1, the profit loss is bounded above by

$$\Delta = 4\sqrt{2K\varepsilon}$$

The communication burden can then be expressed as

$$\begin{aligned} \log(N) &= m \log n = m \log \left(\frac{1}{\varepsilon} \right) = m \log \left(\frac{1}{\frac{\Delta^2}{32K}} \right) = \\ &= m (2 \log(\Delta^{-1}) + \log(32K)) \end{aligned}$$

There are various notions of communication efficiency for approximation

schemes (Nisan and Segal 2006). A Polynomial Communication Approximation Scheme (PCAS) in some parameters is a mechanism that achieves a loss of at most Δ using a number of bits that is polynomial in the parameters. A Fully Polynomial Communication Approximation Scheme (FPCAS) achieves a loss of at most Δ using a number of bits that is polynomial in the parameters and in Δ^{-1} . Finally, the strongest notion of the three, a Truly Polynomial Communication Approximation Scheme (TPCAS) achieves a loss of at most Δ using a number of bits that is polynomial in the parameters and in Δ^{-1} .¹⁴

The communication burden found above is polynomial in all the parameters and in $\log(\Delta^{-1})$. Hence:

Proposition 2 *PPM is a TPCAS.*

The optimal solution to a screening problem with a complex type space may involve a large menu. If the type space is uncountable, the menu can be uncountable too – leading to an unbounded communication cost. We can reduce the communication requirements drastically by operating with a coarse type space. PPM bounds the profit loss generated by using an approximate discretization of the type space rather than true type space. The bound is favorable in the sense that the communication burden grows slowly (logarithmically) in the loss bound.

5.5 Comparative Statics on Cognitive Resources

Assume now that there is some binding constraint either on the principal’s computational time or sampling cost or on the communication burden. Let’s denote this constraint by N . Then using PPM to solve the screening 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).

¹⁴An approximation scheme which is a TPCAS is then also a FPCAS and a PCAS. See Nisan and Segal (2006) for further discussion.

In words, as the thinking, searching or communication cost rises, 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. Such simpler menus will be derived from mechanisms that employ rougher categorizations of true types into stereotypes and offer items at greater discounts relative to the prices optimal for the stereotypes.

6 Alternative Mechanisms

Profit-participation mechanism is a valid approximation scheme, but will other mechanisms ‘perform better’? To address this question, we first have to note that 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, elevated performance over one class of problems tends to be offset by performance over another class (Wolpert and Macready 1997). The more prior information the principal has, the more tailored the mechanism could 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 PPM. But a more pertinent question is whether there are other valid mechanisms for the general class of problems we consider.

Since our results apply to a large class of multi-dimensional screening problems, defined only by the Lipschitz constant K and the approximation index ε , we shall now ask whether there are other mechanisms, besides PPM, that work for this class of problems. We begin by defining the class of mechanisms that are based on the principal’s model and modify prices given the optimal solution of this model:

Definition 5 *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, modifies the price vector $p(y)$ 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 6 *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 $c \leq \bar{c}$*

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

Under profit participation the principal shares her profits and hence her price discount that is strictly increasing in her profit. Profit participation is violated when there is a price/cost region including the origin where an increase in the principal's profit does not translate into a strict increase in the absolute value of the price discount.

The profit-participation condition is strong for two reasons: profit participation fails if a *weak* inequality is violated; it fails if the inequality is violated for *some* prices and cost levels, not necessarily all prices and cost levels. As our

theorem is negative, a strong condition means that we can exclude a larger class of mechanisms.

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

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

and by construction, the PPM never violates profit participation. Given the above condition, we can now show that unlike the PPM, mechanisms that violate profit participation are not valid approximation schemes for the class of problems considered here.

Theorem 2 *For any $K > 0$, 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. 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:

$$p'' - \Psi(p'', 0; K, \varepsilon) \leq p' - \Psi(p', 0; K, \varepsilon) \quad (10)$$

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 $\lim_{\bar{t} \rightarrow t} \left| \frac{\partial}{\partial t} u(t, y) \right|$ 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:

$$\begin{cases} S_0 = \{0, 1, 2\} \\ f_{S_0}(0) = f_{S_0}(2) = \frac{1}{4}; f_{S_0}(1) = \frac{1}{2} \end{cases}$$

$$\begin{cases} S_1 = \{0, \frac{1}{2}, 1, \frac{3}{2}, 2\} \\ f_{S_1}(0) = f_{S_1}(2) = \frac{1}{8}; f_{S_1}(\frac{1}{2}) = f_{S_1}(1) = f_{S_1}(\frac{3}{2}) = \frac{1}{4} \end{cases}$$

$$\vdots$$

$$\begin{cases} S_n = \{0, \frac{1}{2^n}, \dots, 1, 1 + \frac{1}{2^n}, \dots, 2\} \\ f_{S_n}(0) = f_{S_n}(2) = \frac{1}{2^{n+2}}; f_{S_n}(s) = \frac{1}{2^{n+1}} \text{ for all other } s \end{cases}$$

$$\vdots$$

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 $t \in [1, 2]$,

$$h + q(s - 1) - \Psi(h + q(s - 1), 0; K, \varepsilon_n) \leq h - \Psi(h, 0; K, \varepsilon_n) \quad (11)$$

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

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

which, if one subtracts (11) 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 proof of the above Theorem 2 proceeds by constructing a relatively straightforward class of problems with binding non-local constraints. In particular, we assume that the product space includes a generic inferior product and a continuum of type-specific products. In the optimal solution, a non-zero measure 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 the stereotype's optimal allocation. This creates a non-vanishing loss for the principal.

The intuition behind the Theorem has to do with the knife-edge nature of mechanisms that do not include profit participation. In the exact solution of the principal's model there is a binding constraint (IC or PC) for every alternative offered. The profit-participation mechanism – and analogous schemes – relaxe these constraints in the right direction. They add slack to constraints that ensures that the agent does not choose alternatives with a lower profit. Mechanisms without the profit-participation return a price vector that still displays binding constraints – or don't even satisfy those constraints. When 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 $\varepsilon \rightarrow 0$. In multi-dimensional screening problems, however, non-local constraints will typically bind in optimum. Here, the magnitude of the misallocation does not vanish even as $\varepsilon \rightarrow 0$.

Given the nature of the statement to be proven, the proof is by example. One may object that the class of problems used in the proof is a zero-measure set within the set of all possible screening problems. We chose this particular class because it can be given a direct economic interpretation. However, all we need is a class of problems where non-local constraints are binding for a positive measure of types. As binding nonlocal constraints are an endemic feature of multi-dimensional screening problems (Rochet and Choné 1998, Rochet and Stole 2003), endless other examples could be found.

Theorem 2 says that if a mechanism performs at least as well as PPM, it must be either very similar to PPM in that it relies on profit participation, or very different in that it is not even model based. Hence, the theorem points to three interesting questions that we leave for future research. Are there other profit-participation mechanisms that perform better than PPM? Are there non-model based mechanisms that perform better than PPM? Are there not-too-restrictive classes of screening problems where the Naive Mechanism is guaranteed to be

a valid approximation?

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, but she can 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. Fourth, it would be interesting to explore non-model based mechanisms.

References

- [1] Armstrong, Mark. (1996): Multiproduct Nonlinear Pricing. *Econometrica* 64(1): 51–75.
- [2] Armstrong, Mark (1999): Price Discrimination by a Many-Product Firm. *Review of Economic Studies* 66: 151–168.
- [3] Blumrosen, Liad, Noam Nisan, and Ilya Segal. (2007): Auctions with Severely Bounded Communication. *Journal of Artificial Intelligence Research*, 28: 233-266.
- [4] Chawla, Shuchi, Jason D. Hartline, and Robert D. Kleinberg. (2007): Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC 2007)*, pp. 243-251.

- [5] Chu, Chenghuan Sean, Phillip Leslie, and Alan Sorensen. (2010): Bundle-Size Pricing as an Approximation to Mixed Bundling. forthcoming, *American Economic Review*.
- [6] Conitzer, Vincent, and Thomas Sandholm.(2004): Self-interested Automated Mechanism Design and Implications for Optimal Combinatorial Auctions In Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04), pp. 132-141, New York, NY, USA.
- [7] Dobson, Gregory and Shlomo Kalish.(1993): Heuristics for Pricing and Positioning a Product Line Using Conjoint and Cost Data. *Management Science* 39(2): 160–175.
- [8] Gabaix, Xavier (2010). A Sparsity-Based Model of Bounded Rationality. Mimeo, New York University.
- [9] Garey M.R. and D. S. Johnson. “Strong” NP-Completeness Results: Motivation, Examples, and Implications. *Journal of the ACM* 25(3), 499–508, July 1978.
- [10] Gilboa, Itzhak, and David Schmeidler. (1989) Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2): 141-153.
- [11] Green, Paul E., and Abba M. Krieger.(1985) Models and Heuristics for Product Line Selection. *Marketing Science* 4(1): 1–19.
- [12] Hansen, Lars Peter, and Thomas J. Sargent. (2010): Wanting Robustness in Macroeconomics. *Working paper*, NYU.
- [13] Kushilevitz, Eyal, and Noam Nisan: *Communication Complexity*, Cambridge University Press, Cambridge. 1997.
- [14] McKelvey, Richard, and Thomas Palfrey. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* 10: 6–38, 1995.
- [15] March, James, and Herbert Simon.(1958): *Organizations*. New York: Wiley.
- [16] Mussa, Michael, and Sherwin Rosen. (1978): Monopoly and Product Quality. *Journal of Economic Theory* 18: 301–317.

- [17] Nisan, Noam and Ilya Segal. (2006): The Communication Requirements of Efficient Allocations and Supporting Prices. *Journal of Economic Theory* 129(1): 192-224.
- [18] Rochet, Jean-Charles, and Philippe Choné. (1998): Ironing, Sweeping, and Multidimensional Screening. *Econometrica* 66(4): 783–826.
- [19] Rochet, Jean-Charles, and Lars Stole. (2003): The Economics of Multidimensional Screening. in *Advances in Economics and Econometrics*, Vol 1, eds. M. Dewatripont, L.P. Hansen, and S. Turnovsky, Cambridge.
- [20] Walley, Peter. (1991): *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London.
- [21] Robert B. Wilson. (1993): *Nonlinear Pricing*. Oxford University Press.
- [22] Wolpert, David H. and William G. Macready. (1997): No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation* 1(1): 67-82
- [23] Yun Xu, Ji Shen, Dirk Bergemann, and Edmund Yeh. (2010): Menu Pricing with Limited Information. *Working paper*, Yale.