MULTI-DIMENSIONAL SCREENING: BUYER-OPTIMAL LEARNING AND INFORMATIONAL ROBUSTNESS

RAHUL DEB[♁] AND ANNE-KATRIN ROESLER[♔]

MAY 25, 2021

ABSTRACT: A monopolist seller of multiple goods screens a buyer whose type is initially unknown to both but drawn from a commonly known distribution. The buyer privately learns about his type via a signal. We derive the seller's optimal mechanism in two different information environments. We begin by deriving the buyer-optimal outcome. Here, an information designer first selects a signal, and then the seller chooses an optimal mechanism in response; the designer's objective is to maximize consumer surplus. Then, we derive the optimal informationally robust mechanism. In this case, the seller first chooses the mechanism, and then nature picks the signal that minimizes the seller's profits. We derive the relation between both problems and show that the optimal mechanism in both cases takes the form of pure bundling.

1. INTRODUCTION

What is the optimal mechanism that a monopolist should use to sell multiple goods to a single buyer? Despite being a classic economic problem, multi-dimensional screening is notoriously intractable. Even if the seller has just two goods and the buyer's values are additive, independent, and identically distributed, the optimal mechanism is hard to characterize generally. In this paper, we study a general version (with arbitrarily many goods and non-additive values) of this problem but with the novel feature of buyer learning. As it turns out, introducing this new feature makes the model tractable and in certain environments—including the one with independent and additive values—makes pure bundling an optimal mechanism.

The buyer in our model initially has an unknown type $(\theta_1, \ldots, \theta_n)$ that is drawn from a commonly known exchangeable[1] distribution, where each $\theta_i \in [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}_+$. The buyer's type determines his value $\kappa_b \sum_{i \in b} \theta_i$ for any bundle $b \subseteq \{1 \ldots, n\}$ of goods where $\kappa_b \geq 0$ is a non-negative constant. We assume a weak free-disposal property, which requires that the value of the grand bundle (that is, the bundle of all n goods) is greater than any other bundle for every type. This

[♁]DEPARTMENT OF ECONOMICS, UNIVERSITY OF TORONTO, RAHUL.DEB@UTORONTO.CA
[♔]DEPARTMENT OF ECONOMICS, UNIVERSITY OF TORONTO, AK.ROESLER@UTORONTO.CA

We are grateful for comments from Gabriel Carroll, Andrea Prat, Xianwen Shi, Gabor Virag, anonymous reviewers from the ACM-EC conference, and numerous seminar participants. We thank the Social Sciences and Humanities Research Council for their continued and generous financial support.

[1]This is an assumption of symmetry that requires every permutation of the type vector to have the same joint distribution. It does not rule out positive or negative correlations.

class of value functions allows for goods to be complements or substitutes and, importantly, includes additive values ($\kappa_b = 1$ for all bundles b) as a special case. The buyer learns about his type via a signal. Upon privately observing the signal realization, the buyer forms a posterior estimate of his value for different bundles.

Our aim is to derive the seller's optimal mechanism under two different informational environments. We first characterize the *buyer-optimal outcome*: this is the signal and the corresponding optimal mechanism for the seller that generate the *maximal consumer surplus*. Specifically, an information designer[2] first publicly picks the signal (the signal realization remains private to the buyer) to maximize consumer surplus anticipating that the seller will choose an optimal mechanism in response. We show that the buyer-optimal outcome is generated by a signal that makes "pure bundling" (selling the grand bundle at a given price) an optimal mechanism for the seller. Additionally, we show that the seller's profit is minimized: there is no other signal and corresponding optimal mechanism that yield a lower profit. Thus, we show that the seller's profit in the buyer-optimal outcome is the solution to a min-max problem where an adversarial nature picks a signal with the aim of minimizing the profits of a seller who best-responds with an optimal mechanism.

We then derive the *optimal informationally robust mechanism* for the seller: this is an optimal mechanism for a seller who does not know how the buyer learns and who evaluates profits according to a worst-case criterion. Here, the timing is reversed: the seller first chooses the mechanism, following which nature picks the signal to minimize the seller's profit. Therefore, in this case, the seller's profit from the optimal informationally robust mechanism is the solution to a max-min problem. Once again, we show that pure bundling is optimal for the seller, but, in this case, she randomizes over the price for the grand bundle. Moreover, we derive this result by showing that the seller's profit, in this case, is exactly equal to her profit from the buyer-optimal outcome; that is, the optimal value of the objective function in the max-min and the min-max problems coincide.

In our view, the solutions to both problems are individually economically interesting and have distinct implications. At a high level though, both demonstrate different important properties of pure bundling. The buyer-optimal outcome is a natural theoretical benchmark. The seller of a single good always finds it optimal to screen by simply posting a price. By contrast, optimal multidimensional screening can, and frequently does, involve complex menus and randomization even when values are additive and each θ_i is independently and identically distributed. Such elaborate screening helps sellers maximize profits, but the effect on the consumer is unclear. For instance, complex screening might lead to Pareto improvements where both the seller and the buyer are better off because the efficiency of trade increases. The buyer-optimal outcome is a natural benchmark to study the tradeoff between mechanism complexity and the efficiency of trade because the seller best responds to the most advantageous information structure for the buyer. Here, the optimal mechanism takes the very simple form of pure bundling, and we show that trade is efficient.

The selling practices of multi-product retailers are scrutinized by regulators who specifically express concerns about and pursue litigation against practices like tying and bundling by large

[2]As we discuss below, this could be interpreted either as a theoretical benchmark or alternatively we could think of the designer as a regulator having the best interest of consumers in mind.

firms.[3] Under the buyer-optimal signal, not only does pure bundling not cause consumer harm, it leads to the highest possible consumer surplus and efficient trade. This suggests that the information available to buyers is an important factor that should determine whether or not bundling needs to be scrutinized. Of course, this also raises the question of whether and, if so, which advertising practices should be regulated in the interest of consumers (see Bagwell, 2007 and Woodcock, 2017 for a recent case in favor of greater regulation).

Conversely, the optimal informationally robust mechanism provides a positive explanation for why we should expect to observe pure bundling in practice. Despite having historical data from different markets, sellers are unlikely to have very precise estimates of a buyer's value distribution. In particular, it is impossible for a seller to predict what information the buyer has or will acquire in any particular period. Our results show that pure bundling (albeit with a random price) provides the highest revenue guarantee. This is perhaps one reason why, in practice, we do not observe very complex screening that depends on fine details of the type distribution (as is possible in multidimensional screening). Instead, pure bundling is the common way that digital goods such as streaming services are sold. This setting is a good fit for the model: sellers such as Netflix and Spotify have considerable market power, and disposal is free.

Before moving on to the related literature, we provide some high-level intuition for the main insight driving our results. The tractability of our model, given its generality and the simple form of the seller's optimal mechanisms might come across as surprising to some readers. After all, we add buyer learning to the already complex (and unsolved at this level of generality) screening problem. Learning itself can also be significantly more complex than its one-dimensional counterpart. This is because, in this multidimensional environment, the set of possible signals is very rich and does not have a simple characterization akin to the one-dimensional case (where the prior type distribution is a mean-preserving spread of the distribution of posterior estimates that can arise from any signal). Our central insight is that it is precisely the buyer learning that makes the problem tractable. Specifically, we identify a signal that describes the buyer's information in the buyer-optimal outcome. For *this* particular signal, it is possible to derive the seller optimal mechanism. Moreover, this signal has the property that it minimizes the seller's revenue when she chooses the optimal informationally robust mechanism and plays an important role in its derivation.

An important feature of our environment is that generically there are informative signals such that the posterior estimates of each θ_i have maximal positive correlation across i. This can be done even when the prior distribution is such that different dimensions of the type vector are distributed independently or are negatively correlated. As an example, suppose the buyer observed a signal (about the sum of his type vector) of the form $\theta_1 + \cdots + \theta_n + \varepsilon$ where ε is some noise. Maximal positive correlation in the distribution of the posterior estimate of the type vector then follows from the fact that the prior distribution is exchangeable.

[3]Often, the concern is that such practices are used to prevent the entry of competitors. Importantly, however, a selling practice is not illegal if it maximizes profits even if it deters entry as a side effect. Instead, regulators typically scrutinize firms that incur deliberate losses to maintain market power.

In addition to generating such correlation, the signal in the buyer-optimal outcome adds further noise that determines the shape of the distribution of the posterior estimate of each θ_i. It might seem counterintuitive to some that the signal in the buyer-optimal outcome contains so much noise. After all, imperfect information about his type affects the buyer's ability to make good purchasing choices (in an ex-post sense) when faced with a given mechanism. However, the signal also determines the seller's optimal mechanism and, therefore, the share of surplus that the seller can appropriate. Specifically, the correlation in the signal limits the seller's ability to screen across dimensions. Further, the induced shape of the distribution of the posterior type estimate results in efficiency and gives the buyer the maximal surplus share possible (by making it optimal for the seller to offer a pure bundling mechanism at a price that is favorable for the buyer). From a technical perspective, we show that such signals allow us to map our multi-dimensional screening problem onto its one-dimensional counterpart, which in turn allows us to leverage existing results from the literature.

Related Literature

This paper lies at the intersection of two different literatures. The first literature examines the classic question of how a monopolist should jointly sell multiple goods. Despite being a mature literature (dating back till at least Adams and Yellen (1976)), the complexity of the problem is such that there are surprisingly few general insights even for the special case of additive values. A seminal result is due to McAfee, McMillan, and Whinston (1989) who show that when values for each good are independent, selling goods individually (separate sales) is never optimal for the monopolist.

In general, the optimal mechanism can be extremely complex even when values are independent and additive. Pavlov (2011) shows that optimal screening can involve randomization when values are identically and uniformly distributed. Daskalakis, Deckelbaum, and Tzamos (2017) show that the optimal mechanism for two goods features an infinite menu of lotteries when the values are drawn independently from the beta distribution. In fact, the seller might get a negligible fraction of the optimal revenue if she is restricted to using "simple mechanisms" like pure bundling or separate sales (Hart and Nisan, 2019). Unlike these results, we show that, even under general exchangeable prior type distributions, simple pure bundling is optimal in two different environments with buyer learning.

Perhaps the closest paper in this literature is Carroll (2017). He considers a seller who knows the marginal distribution of the buyer's value for each good but not the joint distribution. The monopolist chooses a mechanism that maximizes the worst-case profit computed over all joint distributions which have the given marginals. He shows that separate sales is seller optimal for this criterion. In addition to being a distinct economic problem, our setting is fundamentally different: the buyer's posterior value distribution must be obtained by Bayesian updating from the signal. In particular, we cannot introduce arbitrary correlations across goods, and the marginal posterior distribution for the value of each good can differ from the prior.

Additionally, unlike Carroll (2017), we allow for non-additive values. There is substantially less work that analyzes these environments. The recent survey of Armstrong (2016) describes a strand of this literature that does not aim to derive the optimal mechanism but instead characterize conditions under which the seller can profit from offering bundle discounts. A notable exception is Haghpanah and Hartline (2020) who characterize environments where pure bundling is seller optimal, and we employ one of their results in the proof of our main result. Specifically, we show that, while pure bundling is not an optimal mechanism for the prior type distribution, it emerges as optimal due to buyer learning.

This paper is also related to the growing literature on information design: Bergemann and Morris (2019) and Kamenica (2019) are recent surveys. Within this literature, we are most closely related to the recent work studying how the information environment affects the selling mechanism, efficiency, and the resulting surplus division in bilateral trade settings.[4] While most of this work considers different information environments to this paper, the key distinction is that we study multi-dimensional screening. Bergemann, Brooks, and Morris (2015) study a standard single-good monopoly pricing problem and analyze which buyer-seller surplus pairs are achievable when the seller (instead of the buyer) receives additional information which she can use to price discriminate. Ravid, Roesler, and Szentes (2020) study an environment where buyer-learning is unobservable but costly. Their main result shows that there is a distinction between free and arbitrarily cheap learning.

The closest related papers are Roesler and Szentes (2017) and Du (2018); they respectively analyze the one-dimensional versions of the two problems we study. Roesler and Szentes (2017) derive the buyer-optimal outcome for a single good, and their main insight is to show that, even if information is free, the buyer prefers not to perfectly learn his value for the product. Du (2018) derives the optimal informationally robust mechanism (a random posted price) for a single good[5] and uncovers the relation to the buyer-optimal outcome. The richness of the multi-dimensional screening environment that we consider opens the door to questions that cannot be addressed in the one-dimensional context. Namely, our main contribution is to derive the *qualitative* properties of the seller's optimal mechanism (it takes the form of pure bundling) in both information environments; for the seller of a single good, it is always just a posted price (either deterministic or random).

2. The Model

We consider a mechanism design problem with one buyer and one seller, the latter of whom has one unit of each of $n \geq 2$ goods for sale. We denote the set of goods by $N = \{1, \dots, n\}$. We assume that the seller's cost of producing these goods is 0.

Type Space: The buyer has a *type* $\theta = (\theta_1, \dots, \theta_n)$ that lies in a set $\Theta = [\theta_\ell, \theta_h]^n$ (endowed with the Borel σ-algebra $\mathscr{F}$) with $\theta_h > \theta_\ell \geq 0$. The type is initially unknown to both the buyer and seller,

[4]Similar ideas can also be found in the literature on information acquisition and disclosure in mechanism design settings such as Persico (2000), Bergemann and Välimäki (2002) or Shi (2012). In contrast to the information design literature, these papers usually consider a restricted domain of feasible information structures.

[5]He additionally constructs an informationally robust auction to sell a common-value good which has the property that, as the number of bidders gets large, its revenue guarantee converges to the full surplus.

and is drawn from a commonly known (cumulative) distribution F.[6] We assume F has a positive density for all $\theta \in \Theta$.

Additionally, we assume that F is *exchangeable*: for any permutation $\sigma : N \to N$, the joint distribution of $(\theta_1, \ldots, \theta_n)$ is the same as the joint distribution of $(\theta_{\sigma(1)}, \ldots, \theta_{\sigma(n)})$ (both of which are F).[7]

Exchangeability requires the marginal distribution of each θ_i to be the same and is clearly satisfied when each dimension of the type, θ_i, is independent and identically distributed (henceforth iid). Exchangeability allows both positive and negative correlations between dimensions of the type vector, although it does imply a lower bound on the correlation coefficient. This bound is increasing in n but, for the most widely examined case of two goods, exchangeability imposes no restriction on the degree of correlation.

Given a type θ, we use $\overline{\theta} \in \overline{\Theta} := [n\theta_\ell, n\theta_h]$ to denote the sum $\overline{\theta} := \theta_1 + \cdots + \theta_n$ and $\overline{F}$ denotes the distribution of the sum $\overline{\theta}$ induced by the type distribution F.

Value Function: Given a type θ, the buyer's value for a bundle $b \subseteq N$ is

$$u(\theta, b) = \kappa_b \sum_{i \in b} \theta_i, \tag{1}$$

where $\kappa_b \geq 0$ and we normalize the constant for the *grand bundle* N to $\kappa_N = 1$. Because of this normalization, we also refer to the sum $\overline{\theta}$ as the *grand bundle value*.

We assume that the buyer's value of not receiving a good is $u(\theta, \emptyset) = 0$ and that $u(\theta, N) \geq u(\theta, b)$ for all $b \subseteq N$ and all $\theta \in \Theta$. The latter is a *weak free-disposal* property and ensures that the greatest surplus is generated by trading the grand-bundle. In terms of the κ_b-s, this requires that $\kappa_b \leq 1 + \frac{N-|b|}{|b|} \frac{\theta_\ell}{\theta_h}$. Note that we do not require $\kappa_{b'} \geq \kappa_b$ for proper subsets $b \subset b' \subset N$. Indeed, when $\theta_\ell > 0$, we can have $\kappa_b > \kappa_N$. We assume that preferences are quasilinear and players are risk-neutral expected utility maximizers.

This framework generalizes *additive values* ($\kappa_b = 1$ for all $b \subseteq N$) and allows for goods to be either complements or substitutes[8]. This generality is not merely cosmetic and, as we will argue, has implications for the form that the seller's optimal mechanism takes. Geng, Stinchcombe, and Whinston (2005) also study multidimensional screening with such a value function, but they require that κ_b is decreasing in the number of goods in the bundle b.

Maximal Total Surplus: We use $\overline{\mu} := \int_\Theta u(\theta, N) dF(\theta) = \int_{\overline{\Theta}} \overline{\theta} d\overline{F}(\overline{\theta})$ to denote the *maximal surplus* that can be achieved from trading. Note that the weak free-disposal assumption implies that the maximal surplus is achieved by trading the grand bundle with probability one.

Signals: The buyer learns about his type via a *signal*. Given the linearity of our model and risk-neutrality of players, without loss, we will restrict attention to *unbiased* signals $(S, G_{S \times \Theta})$. The set of signal realizations $S = \Theta$ is just the type space. $G_{S \times \Theta} \in \Delta(S \times \Theta)$ is a joint distribution over

[6] We abuse notation and interchangeably refer to F as a cumulative distribution (henceforth, cdf) and a probability measure. The meaning will be clear from the argument (element vs. set).

[7] Formally, for any $X \in \mathscr{F}$, we have $F(X) = F(X_\sigma)$ where $X_\sigma = \{(\theta_{\sigma(1)}, \ldots, \theta_{\sigma(n)}) \mid (\theta_1, \ldots, \theta_n) \in X\}$.

[8] For example, our value function can capture the case of add-on items

$S \times \Theta$ such that the marginal distribution of $G_{S\times\Theta}$ over Θ is F. We denote the marginal distribution of $G_{S\times\Theta}$ over the set of signal realizations S by G.

The buyer learns about his type by observing a signal realization $s \in S$. We assume the *posterior estimate* of the type is just the signal realization $s = (s_1 \ldots, s_n)$ itself (hence, the "unbiased" terminology) so

$$s = \mathbb{E}_{G_{S\times\Theta}}[\theta \,|\, s]$$

for all s that lie in the support of G. We will refer to both the joint distribution $G_{S\times\Theta}$ and the marginal distribution G as signals since we can always convert one to the other. The buyer privately observes the signal realization.

Restricting attention to the above class of signals is without loss because the value function is linear in the type and therefore both the buyer and the seller only care about the posterior estimate of the type. To see this, suppose instead that signal realizations are drawn from an arbitrary set S' and $G_{S'\times\Theta} \in \Delta(S' \times \Theta)$ is a joint distribution that has marginal distribution F over Θ.[9] If the buyer observes a signal realization s', his value is determined by

$$\mathbb{E}[u(\theta, b) \,|\, s'] = u(\mathbb{E}_{G_{S'\times\Theta}}[\theta \,|\, s'], b) \text{ for all } b \subseteq N.$$

So we can just transform any $(S', G_{S'\times\Theta})$ into the above unbiased form by relabeling s' to $s = \mathbb{E}_{G_{S'\times\Theta}}[\theta \,|\, s']$ with G being the distribution of $\mathbb{E}_{G_{S'\times\Theta}}[\theta \,|\, s']$.

We denote the *set of signals*, by

$$\mathcal{G} := \{G \in \Delta(S) \mid G \text{ is the marginal distribution over } S \text{ induced by some signal } (S, G_{S\times\Theta})\}.$$

Note that this is the set of possible distributions over posterior estimates.

Mechanism: The seller chooses a mechanism. Formally, a *mechanism* $\mathcal{M} = (M, q, t)$ consists of a set of messages M, a (possibly random) allocation $q : M \to \Delta(2^N)$ and a transfer $t : M \to \mathbb{R}$. The allocation determines the likelihood of receiving the various bundles and of not being allocated any good; we sometimes use $q(m, b)$ to denote the probability that the buyer is allocated bundle b when he reports message $m \in M$.

If the buyer with posterior estimate s reports $m \in M$, his expected utility is

$$\mathbb{E}_{q(m)}[u(s, b)] - t(m),$$

where the expectation is taken with respect to the random allocation. Because of this structure of payoffs, it is without loss to restrict the seller to deterministic transfers.

We further assume that every mechanism is such that

$$\max_{m\in M} \Big[\mathbb{E}_{q(m)}[u(s, b)] - t(m) \Big] \geq 0,$$

for all $s \in S$. This is an individual rationality requirement. It ensures that the buyer is not forced to participate in a mechanism that gives him negative utility. Implicitly, we are also assuming that a solution to the above maximization problem exists; that is, every mechanism $\mathcal{M}$ has the property that, for every s, there is a message that maximizes the buyer's utility.

[9] We are implicitly assuming that S' is a Polish space endowed with the Borel σ-algebra. This is a technical assumption which is necessary to guarantee that the conditional distributions are well defined.

The buyer chooses a reporting strategy $\sigma : S \to \Delta(M)$ that maximizes his utility. We say that his strategy σ is a *best response* if his expected utility satisfies

$$U(s,\mathcal{M}) := \mathbb{E}_{\sigma(s)}\left[\mathbb{E}_{q(m)}[u(s,b)] - t(m)\right] \geq \mathbb{E}_{\sigma'(s)}\left[\mathbb{E}_{q(m)}[u(s,b)] - t(m)\right] \tag{2}$$

for all $s \in S$ and all other strategies σ'. For a mechanism $\mathcal{M}$, $U(s,\mathcal{M})$ denotes the buyer's *utility from best responding* and the *set of best responses* is denoted by $\Sigma(\mathcal{M})$.

Given a signal G, a mechanism $\mathcal{M}$ and a best response σ, the seller's profit is given by

$$\Pi(G,\mathcal{M},\sigma) := \mathbb{E}_G\left[\mathbb{E}_{\sigma(s)}\left[t(m)\right]\right].$$

The outer expectation is taken with respect to distribution over signal realizations and the inner with respect to the buyer's strategy.

Pure Bundling and Separate Sales: We will refer to two special classes of mechanisms repeatedly. The first is a *pure bundling mechanism* at price $\overline{p}$ which we denote by $\mathcal{M}^{PB}_{\overline{p}} = (M^{PB}, q^{PB}_{\overline{p}}, t^{PB}_{\overline{p}})$. This is the mechanism in which, whenever the buyer purchases, he is only allowed to purchase the grand bundle N at a price $\overline{p}$. Formally, we can implement this mechanism with a message space $M^{PB} = [n\theta_\ell, n\theta_h]$ and an allocation and transfer given by

$$q^{PB}_{\overline{p}}(m,b) = \begin{cases} 1 & \text{if } m \geq \overline{p} \text{ and } b = N, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad t^{PB}_{\overline{p}}(s) = \begin{cases} \overline{p} & \text{if } m \geq \overline{p}, \\ 0 & \text{otherwise.} \end{cases} \tag{PB}$$

In words, the buyer reports his value m for the grand bundle and is allocated the grand bundle at a price of $\overline{p}$ if the report is higher than the price. Clearly, it is a best response for the buyer to truthfully report his value.

The second is a *separate sales mechanism* at prices $p = (p_1, \ldots, p_n)$ which we denote by $\mathcal{M}^{Sep}_p = (M^{Sep}, q^{Sep}_p, t^{Sep}_p)$. Here, the seller offers a price p_i for each individual good and the buyer can choose whichever bundle he likes and just pay the total price. Formally, we can implement such a mechanism with a message space $M^{Sep} = 2^N$ and an allocation, transfer given by

$$q^{Sep}_p(m,b) = \begin{cases} 1 & \text{if } m = b, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad t^{Sep}_p(m) = \begin{cases} \sum_{i\in m} p_i & \text{if } m \neq \emptyset, \\ 0 & \text{if } m = \emptyset. \end{cases} \tag{Sep}$$

Our aim is to derive the qualitative properties of the seller's optimal mechanism under two different information environments and to relate the solutions of each. We informally describe both here in words, the formal descriptions are in Section 3 and Section 4 respectively.

We first derive the buyer-optimal outcome. This is a signal for the buyer and an optimal mechanism for the seller (in response to this signal) that maximize consumer surplus. The timing here is that an information designer first publicly chooses the signal. That is, the signal structure is observed by the buyer and the seller, the signal realization is private to the buyer. The seller then chooses an optimal mechanism in response.

We then derive the seller's optimal informationally robust mechanism. This is a mechanism that maximizes the seller's profit against the worst possible signal realization. In this case, the timing is the exact opposite. The seller first chooses her mechanism. In response, nature chooses a signal and a best response for the buyer that minimizes the seller's profit.

2.1. *Discussion of the model*

We make several modeling assumptions that are worth discussing before we proceed to the analysis. The fact that we allow for an arbitrary number of goods and for goods to be either complements or substitutes makes our environment more general than the bulk of the multidimensional screening literature (especially the subset that aims to derive optimal mechanisms), which either assumes two goods, additive values or both. We chose the particular class of value functions given by (1), because it has the feature that only the posterior estimate is relevant for the mechanism design problem. For more general value functions, the posterior estimate is no longer sufficient to determine the buyer's utility and, for every signal realization s, the entire posterior distribution over Θ induced by the signal G would become relevant. This significantly complicates the information design part of our problem because we would need to optimize over a significantly larger set of signals. It is for precisely this reason that the bulk of the information design literature also restricts attention to linear environments.

We assume that the distribution F of the buyer's type is exchangeable. This assumption is primarily for tractability; we briefly discuss the complexities introduced by more general type distributions in our concluding remarks in Section 5. Nonetheless, the fact that this assumption does not rule out positive or negative correlations makes our environment significantly more general than the bulk of the existing literature. To the best of our knowledge, apart from the aforementioned Haghpanah and Hartline (2020), there are very few results that characterize the optimal mechanism when values for goods are correlated even with additional functional form assumptions about the marginal distribution of each dimension. Recent work (see, for instance, Chen and Riordan, 2013; Chen and Ni, 2017) has instead sought to understand when pure bundling can dominate separate sales (even though the former may not be the optimal mechanism).

3. THE BUYER-OPTIMAL OUTCOME

In this section, we formally define and characterize a buyer-optimal outcome. We then apply this characterization to derive a comparative statics result that relates the consumer surplus to the number of goods offered for sale. We begin with a few definitions.

Optimal Mechanism: For a signal $G \in \mathcal{G}$, an *optimal mechanism* maximizes the seller's profit. Formally, a mechanism $\mathcal{M}$ is optimal if it has a buyer best response $\sigma \in \Sigma(\mathcal{M})$ such that

$$\Pi(G, \mathcal{M}, \sigma) \geq \Pi(G, \mathcal{M}', \sigma')$$

for any other mechanism $\mathcal{M}'$ and buyer best response $\sigma' \in \Sigma(\mathcal{M}')$. Here we make the implicit assumption that the buyer chooses the best response that maximizes the seller's revenue (this could have a material impact on profits since, in particular, G may have atoms). This is the standard assumption in mechanism design. In other words, for a given signal G, the optimal mechanism is the solution to a standard multidimensional screening problem. Despite being standard, we make this tie-breaking assumption explicit to provide a clear contrast with the optimal informationally robust mechanism that we define in Section 4 (in which the buyer breaks ties against the seller).

Outcome: We will refer to a pair $(G,\ \mathcal{M})$ as an *outcome*, whenever $G \in \mathcal{G}$ is a signal and $\mathcal{M}$ is an *optimal mechanism* for the seller in response to distribution G.

Buyer-Optimal Outcome: Our goal in this section is to characterize the *buyer-optimal outcome* $(G^*,\ \mathcal{M}^*)$ that maximizes the buyer's surplus across all outcomes. Formally, we solve

$$\begin{aligned} &\max_{\mathcal{M}, G \in \mathcal{G}} \mathbb{E}_G \left[U(s, \mathcal{M}) \right] \\ &\text{such that} \\ &\mathcal{M} \text{ is an optimal mechanism for the signal } G. \end{aligned}$$

The constraint in this maximization problem is well defined because a seller optimal mechanism always exists for every signal (see, for instance, Balder (1996)). Our proof explicitly constructs the signal G^* in the buyer-optimal outcome, so we show that the maximum for the objective function is obtained.

Before stating the result, we highlight a specific subclass of signals. These will play an important role in the derivations of both the buyer-optimal outcome and the optimal informationally robust mechanism.

Perfectly correlated signal: We say that a signal $G \in \mathcal{G}$ is maximally positively correlated across dimensions, or simply *perfectly correlated*, if it is distributed along the diagonal $\{(s_1, \ldots, s_n) \in S \mid s_1 = \cdots = s_n\}$.

We are now in a position to present the first of our two main results.

THEOREM 1. *There exists a buyer-optimal outcome $(G^*,\ \mathcal{M}^*)$ which has the following properties.*
(1) Seller Best-Response: *$\mathcal{M}^*$ is a pure bundling mechanism so $\mathcal{M}^* = \mathcal{M}^{PB}_{\overline{p}^*}$ for some $\overline{p}^*$.*
(2) Signal*: G^* is perfectly correlated.*
(3) Total Surplus: *The buyer is allocated the grand bundle with probability one so trade is efficient.*

Before proceeding to the proof, it is worth briefly discussing this result. We begin by providing some intuition for why pure bundling emerges as the optimal mechanism. It turns out that perfect information about his type is not optimal for the buyer because this allows the seller to screen effectively. To prevent this, the signal in the buyer-optimal outcome injects two different kinds of noise into learning. First, it introduces perfect correlation. This effectively makes the type space one-dimensional and reduces the seller's ability to screen across dimensions. It is generically possible to construct informative perfectly correlated signals even though the θ_i's might be independently distributed or even negatively correlated.[10] In our setting, suppose that the buyer perfectly learned his grand bundle value $\theta_1 + \cdots + \theta_n$ but nothing further. Then, because the distribution F is exchangeable, the buyer's posterior estimate for each dimension will be identical and will simply be $(\theta_1 + \cdots + \theta_n)/n$.

[10] Of course, a completely uninformative signal always induces perfect correlation.

Reducing the seller's ability to screen by introducing correlation could still harm the buyer as it might simultaneously lead to a reduction of total surplus. This can be prevented by injecting further noise into the signal: instead of telling the buyer his exact grand bundle value, the signal provides a noisy estimate. Loosely speaking, perfect correlation effectively reduces the seller's problem to its one-dimensional counterpart, and hence we can adapt the methods from Roesler and Szentes (2017) (who study the sale of a single good) to show that it is possible to construct such a signal so that trade of the grand bundle always happens. The latter is the efficient outcome because of the weak free-disposal assumption. Finally, if trade always happens, the seller must be best-responding by offering a pure bundling mechanism where the price for the grand bundle is the minimum of the support of the distribution of the posterior estimate of the grand bundle value induced by G^*.

It might seem surprising that the signal in the buyer-optimal outcome inserts so much noise into learning; after all, as a result, the buyer might frequently purchase the grand bundle even when his true value is lower than the price. Both types of noise are necessary because, compared to the one-dimensional case, the larger set of mechanisms at the seller's disposal implies that a greater amount of noise is necessary to prevent effective screening. The buyer-optimal outcome exactly balances these two countervailing forces: making a good purchasing decision and preventing effective screening.

In fact, the buyer-optimal outcome is the worst outcome for the seller.

COROLLARY 1. *The seller's profit in any outcome is weakly greater than her profit π^* in any buyer-optimal outcome. As a consequence, trade happens with probability one in every buyer-optimal outcome.*

In words, this corollary (proved in Appendix A) says that every buyer-optimal outcome not only maximizes consumer surplus, it also minimizes producer surplus. The first part of the above corollary implies that the seller's profit π^* must be the same in every buyer-optimal outcome. Statement (3) in Theorem 1 then implies that trade must be efficient in every buyer-optimal outcome.

This corollary also implies that the seller's revenue in the buyer-optimal outcome is the solution to the min-max problem where an adversarial nature first picks the signal and the seller then chooses an optimal mechanism in response. In other words,

$$\pi^* = \min_{G \in \mathcal{G}} \max_{\mathcal{M}, \sigma \in \Sigma(\mathcal{M})} \Pi(G, \mathcal{M}, \sigma). \tag{3}$$

This fact will play an important role in the derivation of the optimal informationally robust mechanism.

We make one last observation before presenting the proof of Theorem 1. Note that the buyer-optimal outcome is not unique and, in particular, pure bundling need not be the unique seller best response. Specifically, if $\kappa_b \leq \kappa_N = 1$ for all $b \subset N$, then it will also be optimal for the seller to offer a separate sales mechanism where the price vector for the goods is simply the minimum of the support of G^*. This is because the buyer will still prefer to always buy the grand bundle when faced with this mechanism. However, this is no longer the case when $\kappa_b > 1$. We illustrate this in the following example.

Example. Suppose there are two goods and that $\kappa_{\{i\}} > 1$ for both $i \in \{1, 2\}$. We now argue that separate sales cannot be the seller's optimal mechanism in *any* buyer-optimal outcome. As a contradiction, suppose that there is a buyer-optimal outcome $\left(G^\star, \mathcal{M}_p^{Sep}\right)$ where $\mathcal{M}_p^{Sep}$ is a separate sales mechanism at prices $p = (p_1, p_2)$. Corollary 1 implies that trade must be efficient and so the buyer must always purchase the grand bundle. This in turn implies that, for every $\varepsilon > 0$, we must have $G^\star(\{(s_1, s_2) \mid (s_1 + s_2) - (p_1 + p_2) \leq \varepsilon\}) > 0$. In words, there must be a positive mass of buyer types whose grand bundle value is just above the total price for the grand bundle as, otherwise, the seller could earn a greater profit by instead offering a pure bundling mechanism at the higher price $p_1 + p_2 + \varepsilon$.

Now consider the positive mass of types that satisfy $0 \leq (s_1 + s_2) - (p_1 + p_2) \leq \varepsilon$. The buyer prefers to purchase the grand bundle instead of just good i whenever $\kappa_{\{i\}} s_i - p_i \leq (s_1 + s_2) - (p_1 + p_2)$. Adding up over the two goods, we get

$$(\kappa_b - 1)(p_1 + p_2) \leq (\kappa_b - 1)(s_1 + s_2) \leq (\kappa_{\{1\}} - 1)s_1 + (\kappa_{\{2\}} - 1)s_2 \leq (s_1 + s_2) - (p_1 + p_2) \leq \varepsilon$$

where $\kappa_b = \min\{\kappa_{\{1\}}, \kappa_{\{2\}}\}$. For small enough $\varepsilon > 0$, the above inequality cannot be true since we must have $p_1 + p_2 > 0$. This provides the requisite contradiction because it implies that there will be a positive mass of buyer types that will strictly prefer to not buy the grand bundle thereby implying trade is not efficient in the buyer-optimal outcome $\left(G^\star,\ \mathcal{M}_p^{Sep}\right)$.

3.1. *Proof of Theorem 1*

In this subsection, we prove Theorem 1. Readers who are not interested in the details of the proof can skip to Section 3.2 without loss of continuity. However, the proof of Theorem 3 (the characterization of the optimal informationally robust mechanism) builds on the arguments that follow.

We begin by introducing some additional notation and terminology. Given $s \in S$, we use $\bar{s} = s_1 + \cdots + s_n$ to denote the sum. The set of all such $\bar{s}$ is denoted by $\overline{S}$; note that $\overline{S} = \overline{\Theta}$ (because $S = \Theta$) but we use distinct notation nonetheless to distinguish the sum of the signal realization vector from the sum of the type vector. Every signal $G \in \mathcal{G}$ induces a distribution $\overline{G} \in \Delta(\overline{S})$ over the posterior estimates $\bar{s}$ of the grand bundle value. We use $\overline{\mathcal{G}}$ to denote the set of these distributions of grand bundle estimates that are induced by some signal $G \in \mathcal{G}$.

Given two distributions $\overline{G}, \overline{G}' \in \Delta(\overline{S})$, we say that $\overline{G}$ is a *mean-preserving spread* of $\overline{G}'$ or $\overline{G} \succsim \overline{G}'$ if

$$\int_{n\theta_\ell}^{\bar{s}} \overline{G}(x)\ \mathrm{d}x \geq \int_{n\theta_\ell}^{\bar{s}} \overline{G}'(x)\ \ \mathrm{d}x \quad \text{ for all } \bar{s} \in [n\theta_\ell, n\theta_h] \text{ with equality for } \bar{s} = n\theta_h. \tag{4}$$

We now establish some properties of the set $\overline{\mathcal{G}}$.

Lemma 1. *The set of distributions over grand bundle estimates induced by signals in $\mathcal{G}$ has the following properties:*

(i) $\overline{G}$ is a distribution over grand bundle estimates iff $\overline{F}$ is a mean-preserving spread of $\overline{G}$, or equivalently,

$$\overline{\mathcal{G}} = \left\{ \overline{G} \in \Delta(\overline{S}) \mid \overline{F} \succsim \overline{G} \right\}.$$

(ii) For every signal $G \in \mathcal{G}$, there exists a signal $G' \in \mathcal{G}$ that is perfectly correlated such that G and G' induce the same distribution $\overline{G} = \overline{G}'$ over posterior grand bundle estimates.

PROOF OF LEMMA 1. We begin by defining one-dimensional signals that only provide the buyer information about his grand bundle value. These are (unbiased) signals $(\overline{S}, H_{\overline{S}\times\overline{\Theta}})$ where the set of signal realizations $\overline{S}$ is just the set of possible grand bundle values and $H_{\overline{S}\times\overline{\Theta}} \in \Delta(\overline{S} \times \overline{\Theta})$ is a joint distribution over $\overline{S} \times \overline{\Theta}$ such that the marginal distribution of $H_{\overline{S}\times\overline{\Theta}}$ over $\overline{\Theta}$ is $\overline{F}$ and

$$\bar{s} = \mathbb{E}_{H_{\overline{S}\times\overline{\Theta}}}[\bar{\theta} \mid \bar{s}]$$

for all $\bar{s}$ in the support. We refer to such one-dimensional signals as *grand-bundle signals*.

We use H to denote the marginal distribution of $H_{\overline{S}\times\overline{\Theta}}$ over the set of signal realizations $\overline{S}$ and use $\mathcal{H}$ to denote the set of all such distributions H over grand bundle estimates. As we argued when we defined signals $(S, G_{S\times\Theta})$ for the type vector, it is without loss to restrict attention to such unbiased signals.

Now note that elements of $\mathcal{H}$ are just cdfs of real-valued random variables. We defined this set because it has a well known characterization

$$\mathcal{H} = \left\{ H \in \Delta(\overline{S}) \mid \overline{F} \succsim H \right\}, \tag{5}$$

or, in words, H is the marginal distribution of posterior estimates of grand bundle values for some grand bundle signal iff $\overline{F}$ is a mean-preserving spread of H. We will use this characterization to establish the first part of the lemma.

We first argue that $\overline{\mathcal{G}} \subseteq \mathcal{H}$. To see this, observe that signal $G_{S\times\Theta}$ induces a joint distribution $\overline{G}_{\overline{S}\times\overline{\Theta}}$ over $\overline{S} \times \overline{\Theta}$ such that the marginal distribution of $\overline{G}_{\overline{S}\times\overline{\Theta}}$ over $\overline{\Theta}$ is $\overline{F}$ and the marginal distribution over $\overline{S}$ is $\overline{G}$. Formally, this is the image measure of $G_{S\times\Theta}$ generated by the mapping $a(s,\theta) = (s_1 + \cdots + s_n, \theta_1 + \cdots + \theta_n)$ which implies that, for any measurable set $A \subset \overline{S} \times \overline{\Theta}$, we have $\overline{G}_{\overline{S}\times\overline{\Theta}}(A) = G_{S\times\Theta}(a^{-1}(A))$. Moreover, observe that

$$\mathbb{E}_{\overline{G}_{\overline{S}\times\overline{\Theta}}}\left[\bar{\theta}|\bar{s}\right] = \mathbb{E}_{G_{S\times\Theta}}\left[\bar{\theta}|\bar{s}\right] = \mathbb{E}_{G_{S\times\Theta}}\left[\mathbb{E}_{G_{S\times\Theta}}[\theta_1 + \cdots + \theta_n|s]|\bar{s}\right] = \mathbb{E}_{G_{S\times\Theta}}[s_1 + \cdots + s_n|\bar{s}] = \bar{s}$$

for all $\bar{s}$ in the support. Therefore the marginal distribution $\overline{G}$ over $\overline{S}$ induced by $\overline{G}_{\overline{S}\times\overline{\Theta}}$ satisfies $\overline{G} \in \mathcal{H}$.

To complete the proof of part (i), we need to show that $\mathcal{H} \subseteq \overline{\mathcal{G}}$. We will in fact also show part (ii) by arguing that, for every $H \in \mathcal{H}$, there exists a perfectly correlated signal $G \in \mathcal{G}$ such that the distribution it induces on grand bundle estimates satisfies $\overline{G} = H$.

By definition, H is the marginal distribution over $\overline{S}$ corresponding to a joint distribution $H_{\overline{S}\times\overline{\Theta}} \in \Delta(\overline{S} \times \overline{\Theta})$ that has marginal distribution $\overline{F}$ over $\overline{\Theta}$ and that satisfies $\bar{s} = \mathbb{E}_{H_{\overline{S}\times\overline{\Theta}}}[\bar{\theta} \mid \bar{s}]$. We use $H_{\overline{S}\times\overline{\Theta}}$ to define a family of conditional distributions $\hat{H}(\cdot|\theta) \in \Delta(\overline{S})$ as follows

$$\hat{H}(\cdot|\theta) := H(\cdot|\theta_1 + \cdots + \theta_n). \tag{6}$$

This combined with the distribution F over Θ generates a joint distribution $\hat{H}_{\overline{S}\times\Theta}$ over $\overline{S}\times\Theta$ whose marginal distributions over $\overline{S}$ and Θ are $\hat{H}=H$ and F respectively.

Now observe that

$$\mathbb{E}_{\hat{H}_{\overline{S}\times\Theta}}\left[\theta_1+\cdots+\theta_n|\bar{s}\right]=\bar{s}$$

for all $\bar{s}$ in the support. This is a consequence of the definition (6) of $\hat{H}_{\overline{S}\times\Theta}$ and from the fact that $\bar{s}=\mathbb{E}_{H_{\overline{S}\times\overline{\Theta}}}[\overline{\theta}\,|\,\bar{s}]$.

Given the joint distribution $\hat{H}_{\overline{S}\times\Theta}$, we can derive the conditional distribution $\hat{H}(\cdot|\bar{s})$ over Θ. Now observe that the conditional distribution $\hat{H}(\cdot|\bar{s})$ is exchangeable. This follows from the definition of $\hat{H}_{\overline{S}\times\Theta}$ and because F is assumed to be exchangeable. This in turn implies

$$\mathbb{E}_{\hat{H}}\left[\theta_i|\bar{s}\right]=\frac{\bar{s}}{n}\quad\text{for all } i\in\{1,\ldots,n\}.$$

Now define a joint distribution $G_{S\times\Theta}$ over $S\times\Theta$ that is the image measure of $\hat{H}_{\overline{S}\times\Theta}$ generated by the mapping $\hat{a}(\bar{s},\theta)=\left(\frac{\bar{s}}{n},\ldots,\frac{\bar{s}}{n},\theta\right)$. Formally, for any measurable $\hat{A}\subseteq S\times\Theta$, we have $G_{S\times\Theta}(\hat{A})=\hat{H}_{\overline{S}\times\Theta}(\hat{a}^{-1}(\hat{A}))$. By construction, the marginal distribution of $G_{S\times\Theta}$ over Θ is F, the marginal distribution G over S is distributed along the diagonal $\{(s_1,\ldots,s_n)\in S\mid s_1=\cdots=s_n\}$ and so is perfectly correlated.

Observe that

$$\mathbb{E}_{G_{S\times\Theta}}\left[\theta_i\,\middle|\,s=\left(\frac{\bar{s}}{n},\ldots,\frac{\bar{s}}{n}\right)\right]=\mathbb{E}_{\hat{H}_{\overline{S}\times\Theta}}\left[\theta_i|\bar{s}\right]=\frac{\bar{s}}{n}\quad\text{for all } i\in\{1,\ldots,n\}$$

and so $G\in\mathcal{G}$ or, in words, that G is an unbiased signal. Finally, by construction, the distribution $\overline{G}$ over posterior grand bundle estimates induced by G satisfies $\overline{G}=\hat{H}=H$ which completes the proof. ■

We now define the special class of truncated Pareto distributions that we will employ in the proof of Theorem 1 below. These are defined as

$$H_{\alpha,\beta}(\overline{\theta})=\begin{cases}0 & \text{if } \overline{\theta}<\underline{\alpha},\\ 1-\alpha/\overline{\theta} & \text{if } \overline{\theta}\in[\alpha,\ \beta),\\ 1 & \text{if } \overline{\theta}\geq\beta,\end{cases}\tag{7}$$

where $\alpha\leq\beta$. Truncated Pareto distributions are supported on $[\alpha,\beta]$, are continuous on (α,β) and have an atom of size $\frac{\alpha}{\beta}$ at the truncation point β. When $\alpha=\beta$, $H_{\alpha,\alpha}$ is the degenerate distribution with an atom of size 1 at α.

We use $\overline{\mathcal{G}}^P\subset\overline{\mathcal{G}}$ to denote the subset of distributions over grand bundle estimates (induced by signals in $\mathcal{G}$) that are truncated Pareto distributions. Formally,

$$\overline{\mathcal{G}}^P=\left\{\overline{G}\in\overline{\mathcal{G}}\mid\overline{G}=H_{\alpha,\beta}\quad\text{for some}\quad n\theta_\ell\leq\alpha\leq\beta\leq n\theta_h\right\}.$$

This set is non-empty because it includes the distribution $H_{\overline{\mu},\overline{\mu}}\in\overline{\mathcal{G}}$ over grand bundle estimates induced by the completely uninformative signal.

It is easy to show that there is a continuous, strictly decreasing function $\beta(\alpha)$ such that every element of $\overline{\mathcal{G}}^P$ is of the form $H_{\alpha,\beta(\alpha)}$. Hence, in what follows, we simplify notation and only use the lower bound of the support $H_\alpha\in\overline{\mathcal{G}}^P$ to denote a truncated Pareto distribution over grand bundle estimates induced by a signal, since this alone suffices to describe this distribution. In other words, whenever we refer to a distribution H_α we are implicitly assuming $H_\alpha \in \overline{\mathcal{G}}^P$ and that $H_\alpha = H_{\alpha,\beta(\alpha)}$.

Roesler and Szentes (2017) showed that the class of truncated Pareto distributions can be used to characterize the buyer-optimal outcome for a single good. The truncated Pareto distribution $H_{\alpha,\beta}$ has the property that all pure bundling mechanisms with price $\overline{p}$ in $[\alpha, \beta]$ yield the same profit α. Since their work, the properties of this class of distributions have been exploited in several information design papers.

We are now ready to prove Theorem 1.

PROOF OF THEOREM 1. We prove the theorem in two steps. The first step adapts the proof of Lemma 1 in Roesler and Szentes (2017) to our multidimensional context. While the proofs are similar, we reproduce the complete argument here so that our proof is self-contained, and we flag the key differences.

Step 1: Consider an arbitrary outcome $(G,\ \mathcal{M})$ at which the seller's profit is $\pi \in [n\theta_\ell, n\theta_h]$. There exists a perfectly correlated signal $G' \in \mathcal{G}$ that satisfies the following properties.

i $G' \in \mathcal{G}$ induces a truncated Pareto distribution $\overline{G}' = H_\pi \in \overline{\mathcal{G}}^P$ over grand bundle estimates.

ii The buyer's consumer surplus under signal G' is weakly higher (than that under outcome $(G,\ \mathcal{M})$) when the seller offers the pure bundling mechanism $\mathcal{M}_\pi^{PB}$ at price π.

iii Pure bundling mechanism $\mathcal{M}_\pi^{PB}$ yields the seller the highest profit in the set of pure bundling mechanisms.

Proof of Step 1: Let $\overline{G}$ be the distribution on the grand bundle estimates induced by G. Given this signal, if the seller offers a pure bundling mechanism $\mathcal{M}_{\overline{p}}^{PB}$ at price $\overline{p}$, the buyer purchases the grand bundle if his posterior estimate of the grand bundle value satisfies $\overline{s} \geq \overline{p}$ (where note that we assume the buyer purchases when $\overline{s} = \overline{p}$). Therefore, the seller's profit from this pure bundling mechanism satisfies

$$\overline{p}\left[1 - \overline{G}(\overline{p}) + \delta_{\overline{G}}(\overline{p})\right] \leq \pi \quad \iff \quad 1 - \frac{\pi}{\overline{p}} \leq \overline{G}(\overline{p}) - \delta_{\overline{G}}(\overline{p})$$

where $\delta_{\overline{G}}(\overline{p})$ is the mass of distribution $\overline{G}$ at $\overline{p}$ (if $\overline{G}$ has no atom at $\overline{p}$, this will be 0). The left inequality follows from the fact that $\mathcal{M}$ is an optimal mechanism in response to signal G.

Observe that since $\delta_{\overline{G}}(\overline{p}) \geq 0$, we can conclude from the right inequality that $H_{\pi,n\theta_h}(\overline{p}) \leq \overline{G}(\overline{p})$ for all $\overline{p} \in [n\theta_\ell, n\theta_h]$ because, recall that,

$$H_{\pi,n\theta_h}(\overline{s}) = \begin{cases} 0 & \text{if } \overline{s} < \pi, \\ 1 - \frac{\pi}{\overline{s}} & \text{if } \overline{s} \in [\pi, n\theta_h), \\ 1 & \text{if } \overline{s} \geq n\theta_h. \end{cases}$$

In words, $H_{\pi,n\theta_h}$ first-order stochastically dominates $\overline{G}$. This implies

$$\int_{n\theta_\ell}^{n\theta_h} \overline{s}\ \mathrm{d}H_{\pi,n\theta_h}(\overline{s}) \geq \int_{n\theta_\ell}^{n\theta_h} \overline{s}\ \mathrm{d}\overline{G}(\overline{s}) = \overline{\mu} \geq \pi = \int_{n\theta_\ell}^{n\theta_h} \overline{s}\ \mathrm{d}H_{\pi,\pi}(\overline{s}). \tag{8}$$

The left inequality follows from the first-order stochastic dominance we established above. The left equality is a consequence of the fact that G is a signal and so $\int_{n\theta_\ell}^{n\theta_h} \overline{s}\ \mathrm{d}\overline{G}(\overline{s}) = \int_{n\theta_\ell}^{n\theta_h} \overline{\theta}\ \mathrm{d}\overline{F}(\overline{\theta}) = \overline{\mu}$.

The right equality follows from the fact that the distribution $H_{\pi,\pi}$ is just the Dirac measure that has an atom of mass 1 at π.

The simple but key observation in the proof of this step (relative to the one dimensional case analyzed in Roesler and Szentes (2017)) is the right inequality. This follows from the weak free-disposal assumption on the buyer's value function, which implies that always trading the grand bundle generates the maximal surplus, and therefore we must have $\pi \leq \overline{\mu}$. This implies that there exists a pure bundling mechanism that would generate a higher consumer surplus.

Since $\int_{n\theta_\ell}^{n\theta_h} \overline{s} \ \mathrm{d}H_{\pi,z}(\overline{s})$ is continuous and strictly increasing in $z \in [\pi, n\theta_h]$, the inequality (8) combined with the intermediate value theorem imply that there is a unique $\tau \in [\pi, n\theta_h]$ such that $\int_{n\theta_\ell}^{n\theta_h} \overline{s} \ \mathrm{d}H_{\pi,\tau}(\overline{s}) = \overline{\mu}$. In words, the truncated Pareto distribution $H_{\pi,\tau}$ has the same mean as $\overline{F}$ and $\overline{G}$.

We will now argue that $\overline{G}$ is a mean-preserving spread of $H_{\pi,\tau}$ or, in our notation, that $\overline{G} \succsim H_{\pi,\tau}$. To see this, first observe that $H_{\pi,\tau}(\overline{s}) = H_{\pi,n\theta_h}(\overline{s}) \leq \overline{G}(\overline{s})$ on $\overline{s} \in [n\theta_\ell, \tau)$. This implies that for any $z \in [n\theta_\ell, \tau)$

$$\int_{n\theta_\ell}^{z} \overline{G}(\overline{s}) \ \mathrm{d}\overline{s} \geq \int_{n\theta_\ell}^{z} H_{\pi,\tau}(\overline{s}) \ \mathrm{d}\overline{s},$$

and for any $z \in [\tau, n\theta_h]$,

$$\int_{n\theta_\ell}^{z} \overline{G}(\overline{s}) \ \mathrm{d}\overline{s} = \overline{\mu} - \int_{z}^{n\theta_h} \overline{G}(\overline{s}) \ \mathrm{d}\overline{s} \geq \overline{\mu} - \int_{z}^{n\theta_h} H_{\pi,\tau}(\overline{s}) \ \mathrm{d}\overline{s} = \int_{0}^{z} H_{\pi,\tau}(\overline{s}) \ \mathrm{d}\overline{s}, \tag{9}$$

where the inequality follows from $H_{\pi,\tau}(\overline{s}) = 1 \geq \overline{G}(\overline{s})$ for all $\overline{s} \geq \tau$.

Part (i) of Lemma 1 shows that $\overline{F} \succsim \overline{G}$ and therefore $\overline{G} \succsim H_{\pi,\tau}$ implies $\overline{F} \succsim H_{\pi,\tau}$. In other words, $H_\pi \in \overline{\mathcal{G}}$ is a distribution over posterior grand bundle estimates induced by a signal.

Part (ii) of Lemma 1 implies that we can find a signal $G' \in \mathcal{G}$ that is perfectly correlated and that induces a distribution $\overline{G}'$ over posterior grand bundle estimates that satisfies $\overline{G}' = H_\pi \in \overline{\mathcal{G}}$.

Finally, observe that the pure bundling mechanism $\mathcal{M}_\pi^{PB}$ at price π is an optimal pure bundling mechanism (though not necessarily an optimal mechanism) for the seller when the buyer's posterior grand bundle estimate is distributed according to H_π. Since this price is the minimum of the support of H_π, trade is efficient and therefore the buyer's surplus is weakly higher in this case than compared to outcome $(G, \mathcal{M})$. This is because profits are the same in both cases, and the maximal surplus $\overline{\mu}$ is generated by $\mathcal{M}_\pi^{PB}$.

Step 2: There exists a buyer-optimal outcome $(G^*,\ \mathcal{M}^*)$ such that:

i Trade is efficient.
ii G^* is perfectly correlated.
iii G^* induces a truncated Pareto distribution $\overline{G}^* \in \overline{\mathcal{G}}^P$ over grand bundle estimates.
iv $\mathcal{M}^*$ is a pure bundling mechanism.

Proof of Step 2: Let $(G,\ \mathcal{M})$ be an outcome. Then by Step 1, there exists a signal $G' \in \mathcal{G}$ that is perfectly correlated such that $\overline{G}' = H_\alpha \in \overline{\mathcal{G}}$ for some $n\theta_\ell \leq \alpha \leq n\theta_h$ and the buyer is weakly better off under this signal when the seller offers a pure bundling mechanism $\mathcal{M}_\alpha^{PB}$ at price α. We now argue that $\mathcal{M}_\alpha^{PB}$ is an optimal mechanism for the seller in response to G'.

So suppose that the buyer learns via signal G'. Then, for any mechanism $\mathcal{M}' = (M', q', t')$ and any buyer best response $\sigma' \in \Sigma(\mathcal{M}')$, the seller's revenue satisfies

$$\int_{s \in S} \mathbb{E}_{\sigma'(s)} [t(m)] \; \mathrm{d}G'(s) = \int_{\overline{s} \in \overline{S}} \mathbb{E}_{\sigma'\left(\frac{\overline{s}}{n}, \ldots, \frac{\overline{s}}{n}\right)} [t(m)] \; \mathrm{d}\overline{G}'(\overline{s})$$

because G' is perfectly correlated and so the distribution is supported on the diagonal. Observe that this implies that, given the signal G', the seller is effectively solving a one-dimensional mechanism design problem. The Revelation Principle applies so the seller's problem is equivalent to choosing a direct mechanism $\overline{\mathcal{M}} = (\overline{M}, \overline{q}, \overline{t})$ with $\overline{M} = \overline{S}$, $\overline{q} : \overline{S} \to \Delta(2^N)$, $\overline{t} : \overline{S} \to \mathbb{R}$ (defined on a one-dimensional type space where the type is the grand bundle value) to solve

$$\begin{aligned}
&\max_{(\overline{q}, \overline{t})} \int_{\overline{s} \in \overline{S}} t(\overline{s})\, d\overline{G}'(\overline{s}) \\
&\text{subject to} \\
&\mathbb{E}_{\overline{q}(\overline{s})} \left[u\left(\tfrac{\overline{s}}{n}, \ldots, \tfrac{\overline{s}}{n}, b\right)\right] - \overline{t}(\overline{s}) \geq \mathbb{E}_{\overline{q}(\hat{s})} \left[u\left(\tfrac{\overline{s}}{n}, \ldots, \tfrac{\overline{s}}{n}, b\right)\right] - \overline{t}(\hat{s}) \quad \text{for all } \overline{s}, \hat{s} \in \overline{S} \text{ and} \\
&\mathbb{E}_{\overline{q}(\overline{s})} \left[u\left(\tfrac{\overline{s}}{n}, \ldots, \tfrac{\overline{s}}{n}, b\right)\right] - \overline{t}(\overline{s}) \geq 0 \quad \text{for all } \overline{s} \in \overline{S}.
\end{aligned} \tag{10}$$

The constraints are simply the standard incentive compatibility and individual rationality constraints.

Now observe that, for $\overline{s} > 0$, the ratio of the value of the grand bundle N to any other bundle b is a constant because

$$\frac{u\left(\frac{\overline{s}}{n}, \ldots, \frac{\overline{s}}{n}, N\right)}{u\left(\frac{\overline{s}}{n}, \ldots, \frac{\overline{s}}{n}, b\right)} = \frac{\overline{s}}{\kappa_b \frac{|b|}{n} \overline{s}} = \frac{n}{\kappa_b |b|}.$$

We can then directly apply Proposition 1 in Haghpanah and Hartline (2020)[11] to conclude that a pure bundling mechanism $\mathcal{M}^{PB}_{\overline{p}}$ at some price $\overline{p}$ solves (10). Because $\overline{G}' = H_\alpha$, we can then conclude that the pure bundling mechanism $\mathcal{M}^{PB}_{\alpha}$ at the price α is a best response of the seller to G'.

Now define,

$$\alpha^* := \min \left\{ \alpha \;\middle|\; H_\alpha \in \overline{\mathcal{G}} \right\}.$$

In words, α^* is the lowest minimum of the support of a truncated Pareto distribution over grand bundle estimates induced by a signal; the existence of this minimum can be established by a simple continuity argument. By Lemma 1, there exists a perfectly correlated signal G^* such that the distribution over grand bundle estimates that it induces satisfies $\overline{G}^* = H_{\alpha^*}$ and, by the above argument, the pure bundling mechanism $\mathcal{M}^{PB}_{\alpha^*}$ at the price α^* is an optimal mechanism for the seller in response to G^*.

By construction, the outcome $(G^*, \mathcal{M}^{PB}_{\alpha^*})$ satisfies all the properties (1)–(3) of the theorem. Moreover, from Step 1, it is buyer-optimal. If there were another outcome that generated strictly higher consumer surplus, Step 1 implies that we would be able to find a truncated Pareto distribution $H_\alpha \in \overline{\mathcal{G}}$ where $\alpha < \alpha^*$ which is a contradiction. This completes the proof of this step and of the theorem. ■

[11] Informally, their result states that a pure bundling mechanism is optimal when types are one-dimensional and the ratio of the value of the grand bundle N to every other bundle is non-increasing.

3.2. *Comparative statics*

In this section, we apply Theorem 1 to derive a comparative static relating consumer surplus to the number of goods. We begin with some context. It is clearly beneficial for the monopolist to have the ability to screen over all n goods as opposed to having to set a price for each good individually. This is because maximizing profits over a strictly larger set of mechanisms must achieve a weakly higher profit. However, as Salinger (1995) observes, increased profits need not be at the expense of consumer surplus. For instance, consider a buyer with additive values for $n = 2$ goods where his value for each good is independently and uniformly distributed on $[0, 1]$. Here, the optimal separate sales mechanism is to charge a price of $\frac{1}{2}$ for each good. Now, suppose that in addition to selling the goods individually, the seller was allowed to pure bundle. She would choose to do the latter, and the optimal pure bundling mechanism is to charge a price of $\sqrt{\frac{2}{3}} < \frac{1}{2} + \frac{1}{2}$ for the grand bundle. The latter mechanism (which exploits the fact that there are multiple goods) leads to *both* higher profits and consumer surplus.

By contrast, Bakos and Brynjolfsson (1999) derive a limit result that shows the seller can extract all the surplus from a buyer with additive, iid values when the number of goods $n \to \infty$. They use a law of large numbers to argue that the value of the grand bundle divided by the number of goods n converges, and so the seller can extract all the surplus by just offering a pure bundling mechanism at a price of n-times that limit. As we argued in the introduction, because it is hard to characterize the optimal mechanism, we are not aware of any general results for finitely many goods that describe whether the seller's ability to screen across multiple dimensions hurts consumers. The goal of this section is to show that such an analysis is possible for the buyer-optimal outcome.

In this section, we restrict attention to the case of additive values and we assume each θ_i is iid with distribution $\tilde{F}$ that has a positive density (so the joint distribution is $F = \tilde{F} \times \cdots \times \tilde{F}$). We denote $\tilde{\mu} = \mathbb{E}_{\tilde{F}}[\theta_i]$ to be the mean of each dimension of the agent's type. Because we will vary the number of goods we define $\tilde{S} = [\theta_\ell, \theta_h]$ and use $\tilde{S}^n := [\theta_\ell, \theta_h]^n$ (instead of just S) to denote the set of possible signal realizations. Additionally, we will use $\overline{S}_n := [n\theta_\ell, n\theta_h]$ to denote the set of possible grand bundle posterior estimates.

Theorem 1 states that there is a buyer-optimal outcome $\left(G_n^*, \mathcal{M}^{PB}_{\overline{p}_n^*}\right)$ where G_n^* is perfectly correlated and $\mathcal{M}^{PB}_{\overline{p}_n^*}$ is a pure bundling mechanism at price $\overline{p}_n^*$. As we explained in the discussion following Theorem 1, when values are additive, the separate sales mechanism $\mathcal{M}^{Sep}_{p_n^*}$ where $p_n^* = \left(\frac{\overline{p}_n^*}{n}, \ldots, \frac{\overline{p}_n^*}{n}\right)$ is also an optimal mechanism for the seller in response to G_n^*.

Finally, we define

$$\mathrm{CS}_n := \tilde{\mu} - \frac{\overline{p}_n^*}{n},$$

to be the *average consumer surplus*. This is the total consumer surplus from the buyer-optimal outcome divided by the number of goods n. The following result, the proof of which is in Appendix A, shows that this average consumer surplus decreases in the number of goods.

THEOREM 2. *Suppose valuations are additive and each dimension of the type is iid. Then the average consumer surplus in the buyer-optimal outcome is decreasing in the number of goods,*

$$CS_n \geq CS_{n+1} \quad \textit{for all } n \geq 1.$$

Moreover, as n grows large, the average consumer surplus in the buyer-optimal outcome converges to zero ($\lim_{n\to\infty} CS_n = 0$) and the seller extracts all the surplus.

This result (proved in the appendix) is a straightforward consequence of Theorem 1. For intuition, consider the buyer-optimal outcome $\left(G^*_{n+1}, \mathcal{M}^{Sep}_{p^*_{n+1}}\right)$ for the $n+1$ goods in which G^*_{n+1} is perfectly correlated and the seller chooses a separate sales mechanism. Let $\widetilde{G}^*_n$ be the marginal distribution of G^*_{n+1} over the first n dimensions $(s_1, \ldots, s_n)$. Note that $\widetilde{G}^*_n$ is a perfectly correlated signal for the type space $[\theta_\ell, \theta_h]^n$ distributed by $\tilde{F} \times \cdots \times \tilde{F}$. For the perfectly correlated signal $\widetilde{G}^*_n$, the same separate sales mechanism $\mathcal{M}^{Sep}_{p^*_{n+1}}$ at prices p^*_{n+1} is an optimal mechanism for the seller implying that buyer-optimal outcome must generate at least CS_{n+1} in average consumer surplus.

This result highlights the nuanced interplay between information and screening. In an additive values environment, there is always a buyer-optimal outcome in which the seller's optimal mechanism is separate sales. Thus, when we increase the number of goods, the seller still finds it optimal to sell each good separately or, in other words, does not strictly benefit from bundling together different goods. But yet the average consumer surplus decreases. This is because the information that the buyer receives changes as the number of goods increases. In order to prevent the seller from multi-dimensional screening, the signal in the buyer-optimal outcome must introduce correlation (by injecting noise) and does so by only providing information to the buyer about his value for the grand bundle. As the number of goods increase, such correlation surrenders more surplus to the seller until, in the limit, she can extract all the surplus.

4. THE OPTIMAL INFORMATIONALLY ROBUST MECHANISM

The buyer-optimal outcome characterizes the informational environment that is most advantageous for the buyer. Here, the timing is such that an information designer chooses the signal first, and then the seller best responds. It is equally natural to think of the alternative timing: the seller first picks her mechanism, and then nature chooses the signal in response. This has at least two interpretations. The first captures a seller who does not know the exact type distribution and is worried about model misspecification. The second interpretation is that the buyer acquires information after observing the mechanism, but the seller does not know the buyer's information acquisition technology. The seller evaluates each mechanism based on the worst-case profits taken with respect to all possible signal realizations and buyer best-responses.

We begin with a few definitions.

Revenue Guarantee: We say that a mechanism $\mathcal{M}$ provides a *revenue guarantee* of π if

$$\Pi(G, \mathcal{M}, \sigma) \geq \pi$$

for all signals $G \in \mathcal{G}$ and buyer best responses $\sigma \in \Sigma(\mathcal{M})$.

Informationally Robust Mechanism: Formally, we define the *optimal informationally robust mechanism* $\underline{\mathcal{M}}^\star$ as the mechanism that solves

$$\underline{\mathcal{M}}^\star \in \operatorname*{argmax}_{\mathcal{M}} \inf_{G \in \mathcal{G}, \sigma \in \Sigma(\mathcal{M})} \Pi(G, \mathcal{M}, \sigma).$$

This is the mechanism we aim to characterize. In words, it provides the seller the highest revenue guarantee against all possible signals and best responses by the buyer. Note the difference with the (standard) definition of an optimal mechanism in Section 3 where the buyer is assumed to break indifference in favor of the seller. As with the buyer-optimal outcome, we will explicitly construct the mechanism which will show that the maximum is obtained.

In fact, we will show that it takes the following simple form.

Random Pure Bundling Mechanism: A *random pure bundling mechanism* $\mathcal{M}^{rPB} = (M^{rPB}, q^{rPB}, t^{rPB})$ has a message space $M^{rPB} = \overline{S}$ given by the set of possible posterior grand bundle estimates and an allocation rule

$$q^{rPB}(m, b) = 0 \text{ if } b \neq N \quad \text{and} \quad q^{rPB}(m, N) + q^{rPB}(m, \emptyset) = 1 \tag{rPB}$$

for all $m \in M^{rPB}$. In words, these are mechanisms in which the buyer is only ever allocated the grand bundle; however this allocation could be random. Put differently, the buyer is effectively offered a menu of prices and probabilities where each menu item corresponds to the buyer paying a price in exchange for receiving the grand bundle with the given probability.

Before characterizing the optimal informationally robust mechanism, it is worth relating it to the buyer-optimal outcome. Observe that we must have

$$\max_{\mathcal{M}} \inf_{G \in \mathcal{G}, \sigma \in \Sigma(\mathcal{M})} \Pi(G, \mathcal{M}, \sigma) \leq \max_{\mathcal{M}} \inf_{\sigma \in \Sigma(\mathcal{M})} \Pi(G^*, \mathcal{M}, \sigma) = \pi^* = \min_{G \in \mathcal{G}} \max_{\mathcal{M}, \sigma \in \Sigma(\mathcal{M})} \Pi(G, \mathcal{M}, \sigma) \tag{11}$$

where, recall that, G^*, π^* are respectively the signal we construct and the profit in the buyer-optimal outcome (Theorem 1). The first equality follows from the observation that $\mathcal{M}^{PB}_{\pi^*}$ is an optimal mechanism in response to the signal G^* and that, for this signal and mechanism, all buyer best responses result in the same profit. In words, this shows that the seller's worst case profit from the optimal informationally robust mechanism must be lower than that from the buyer-optimal outcome.

In the next result, we show that these are actually equal and, more importantly, we use this fact to characterize the optimal informationally robust mechanism. Note that we cannot immediately employ a minimax theorem to show the equality because the infimum on the left and the maximum on the right in (11) are taken over both signals and buyer best responses. This fact is the point of departure for Brooks and Du (2020) who study the relation between the max-min and min-max problems in general multi-agent environments (we discuss their paper below).

THEOREM 3. *There is a random pure bundling mechanism that is an optimal informationally robust mechanism. This mechanism provides a revenue guarantee of π^*, the seller's revenue in the buyer-optimal outcome.*

PROOF. The proof adapts the argument of the proof of Proposition 1 in Du (2018). We provide it here for completeness and flag the main step where we reduce our multidimensional problem to the single good case that he analyzes.

We begin by stating a property of the signal G^* from the buyer-optimal outcome we derived in the proof of Theorem 1. Recall that this signal is perfectly correlated and induces a truncated Pareto distribution $\overline{G}^* = H_{\pi^*}$ on grand bundle estimates. It is possible to show[12] that there must be a $\bar{s}^* \in (\pi^*, \beta(\pi^*))$ in the interior of the support of H_{π^*} such that

$$\int_{n\theta_\ell}^{\bar{s}^*} \overline{F}(x)\,\mathrm{d}x = \int_{n\theta_\ell}^{\bar{s}^*} H_{\pi^*}(x)\,\mathrm{d}x. \tag{12}$$

For some intuition, it is straightforward to show that the function $I_\pi(y) := \int_{n\theta_\ell}^{y} H_\pi(x)\,\mathrm{d}x$ is continuous and decreasing in π (for a given y). In other words, whenever $\pi < \pi'$, then $I_\pi(y) \geq I_{\pi'}(y)$ for all $y \in [n\theta_\ell, n\theta_h]$ with strict inequality on $(\pi, \beta(\pi))$. Moreover, recall that $H_{\pi^*} \in \overline{\mathcal{G}}$ being a distribution over posterior grand bundle estimates (induced by a signal) implies that $\int_{n\theta_\ell}^{\bar{s}} \overline{F}(x)\,\mathrm{d}x - \int_{n\theta_\ell}^{\bar{s}} H_{\pi^*}(x)\,\mathrm{d}x \geq 0$ for all $\bar{s} \in (n\theta_\ell, n\theta_h)$ or that $\overline{F}$ is a mean-preserving spread of H_{π^*}. If this inequality was always slack, it would be possible to construct another truncated Pareto distribution $H_{\pi^*-\varepsilon} \in \overline{\mathcal{G}}$ for sufficiently small $\varepsilon > 0$ which would violate the optimality of H_{π^*}.

Equation (12) has two immediate consequences. First, because $\int_{n\theta_\ell}^{\bar{s}} \overline{F}(x)\,\mathrm{d}x - \int_{n\theta_\ell}^{\bar{s}} H_{\pi^*}(x)\,\mathrm{d}x \geq 0$ for all $\bar{s} \in (n\theta_\ell, n\theta_h)$, this function must be minimized at $\bar{s} = \bar{s}^*$. Hence, from the first-order condition, we obtain

$$\overline{F}(\bar{s}^*) = H_{\pi^*}(\bar{s}^*). \tag{13}$$

Second, doing integration by parts on both sides of (12) and using (13), we get

$$\int_{n\theta_\ell}^{\bar{s}^*} x\,\mathrm{d}\overline{F}(x) = \int_{n\theta_\ell}^{\bar{s}^*} x\,\mathrm{d}H_{\pi^*}(x). \tag{14}$$

We now use the revenue π^* from the buyer optimal outcome and the value $\bar{s}^*$ defined in (12) to construct a random pure bundling mechanism. We first provide an intuitive implementation before formally describing the allocation and transfer. So suppose the seller only offers the grand bundle for sale but randomizes over the price. Specifically, suppose the price is randomly drawn from the interval $\bar{p} \in [\pi^*, \bar{s}^*]$ with the cumulative distribution $\mathcal{P}$ given by

$$\mathcal{P}(\bar{p}) = \begin{cases} 1 & \text{if } \bar{p} > \bar{s}^*, \\ \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \log\left(\frac{\bar{p}}{\pi^*}\right) & \text{if } \pi^* \leq \bar{p} \leq \bar{s}^*, \\ 0 & \text{if } \bar{p} < \pi^*. \end{cases}$$

Clearly, a type $s \in S$ purchases the grand bundle when the realized price $\bar{p} \leq \bar{s}$; since the distribution of prices has no atoms, the seller's revenue is unaffected by the decision that the buyer makes at the zero probability event $\bar{p} = \bar{s}$. Therefore, the allocation and transfer for this

[12] See the proof of Proposition 1 in Du (2018) or Lemma 3 in Ravid, Roesler, and Szentes (2020).

mechanism can be written in the form of a random pure bundling mechanism $\mathcal{M}^{rPB}$ with

$$q^{rPB}(s,b) = 0 \quad \text{if } b \neq N,$$

$$q^{rPB}(s,N) = \mathcal{P}(\bar{s}) = \begin{cases} \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \min\left\{\log\left(\frac{\bar{s}}{\pi^*}\right), \log\left(\frac{\bar{s}^*}{\pi^*}\right)\right\} & \text{if } \bar{s} > \pi^*, \\ 0 & \text{if } \bar{s} \leq \pi^*, \end{cases} \tag{15}$$

$$t^{rPB}(s) = \int_{n\theta_\ell}^{\bar{s}} \overline{p}\, \mathrm{d}\mathcal{P}(\overline{p}) = \begin{cases} \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \min\left\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\right\} & \text{if } \bar{s} > \pi^*, \\ 0 & \text{if } \bar{s} \leq \pi^*. \end{cases}$$

Note that faced with this mechanism, truth-telling or $\sigma(s) = s$ for all $s \in S$ is a best response, and all best responses give the seller the same revenue.

Now observe that, for any signal $G \in \mathcal{G}$, the seller's profit from the mechanism $\mathcal{M}^{rPB}$ satisfies

$$\begin{aligned} \int_{\Theta} t^{rPB}(s)\, \mathrm{d}G(s) &= \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \int_{\pi^*}^{n\theta_h} \left(\min\left\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\right\}\right) \mathrm{d}\overline{G}(\bar{s}) \\ &\geq \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \int_{n\theta_\ell}^{n\theta_h} \left(\min\left\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\right\}\right) \mathrm{d}\overline{G}(\bar{s}) \\ &\geq \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \int_{n\theta_\ell}^{n\theta_h} \left(\min\left\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\right\}\right) \mathrm{d}\overline{F}(\bar{s}). \end{aligned} \tag{16}$$

The first inequality follows from the fact that $\min\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\} < 0$ for $\bar{s} < \pi^*$. The second inequality is a consequence of two facts. First, G is a signal and so (from part (i) of Lemma 1) $\overline{F}$ is a mean-preserving spread of the distribution $\overline{G}$ (induced by G over grand bundle estimates). Second, the function being integrated is concave and hence the inequality follows from the fact that $\overline{F} \succsim \overline{G}$.

Observe that (16) implies

$$\inf_{G\in\mathcal{G},\sigma\in\Sigma(\mathcal{M}^{rPB})} \Pi(G, \mathcal{M}^{rPB}, \sigma) \geq \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \int_{n\theta_\ell}^{n\theta_h} \left(\min\left\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\right\}\right) \mathrm{d}\overline{F}(\bar{s})$$

or, in words, that mechanism $\mathcal{M}^{rPB}$ has a revenue guarantee given by the right side. This is the key insight that allows us to employ Du (2018)'s argument; it shows that it is possible to generate this revenue guarantee in our multidimensional environment by effectively using a one dimensional mechanism.

Recall that G^* induces a truncated Pareto distribution $\overline{G}^* = H_{\pi^*}$ over grand bundle estimates and so the seller's worst-case profit from $\mathcal{M}^{rPB}$ when the signal is G^* is equal to

$$\begin{aligned} \inf_{\sigma\in\Sigma(\mathcal{M}^{rPB})} \Pi(G^*, \mathcal{M}^{rPB}, \sigma) &= \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \int_{\pi^*}^{n\theta_h} \left(\min\left\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\right\}\right) \mathrm{d}H_{\pi^*}(\bar{s}) \\ &= \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \int_{n\theta_\ell}^{\bar{s}^*} (\bar{s}-\pi^*)\, \mathrm{d}H_{\pi^*}(\bar{s}) + \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} (1 - H_{\pi^*}(\bar{s}^*))(\bar{s}^*-\pi^*) \\ &= \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1} \int_{n\theta_\ell}^{n\theta_h} \left(\min\left\{\bar{s}-\pi^*, \bar{s}^*-\pi^*\right\}\right) \mathrm{d}\overline{F}(\bar{s}), \end{aligned}$$

where the last equality follows from (13) and (14). In other words, for the signal G^* from the buyer optimal outcome, the seller's profit precisely equals the revenue guarantee (16).

Finally, note that the seller's revenue also satisfies

$$\begin{aligned}\inf_{\sigma\in\Sigma(\mathcal{M}^{rPB})}\Pi(G^*,\mathcal{M}^{rPB},\sigma) &= \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1}\int_{\pi^*}^{n\theta_h}\left(\min\left\{\bar{s}-\pi^*,\,\bar{s}^*-\pi^*\right\}\right)\,\mathrm{d}H_{\pi^*}(\bar{s}) \\ &= \left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1}\int_{\pi^*}^{\bar{s}^*}(\bar{s}-\pi^*)\frac{\pi^*}{\bar{s}^2}\,\mathrm{d}\bar{s}+\left[\log\left(\frac{\bar{s}^*}{\pi^*}\right)\right]^{-1}(\bar{s}^*-\pi^*)\frac{\pi^*}{\bar{s}^*} \\ &=\pi^*.\end{aligned}$$

But inequality (11) shows that π^* is the highest possible revenue guarantee a mechanism can provide and this implies that

$$\max_{\mathcal{M}}\inf_{G\in\mathcal{G},\sigma\in\Sigma(\mathcal{M})}\Pi(G,\mathcal{M},\sigma)=\inf_{G\in\mathcal{G},\sigma\in\Sigma(\mathcal{M}^{rPB})}\Pi(G,\mathcal{M}^{rPB},\sigma)=\inf_{\sigma\in\Sigma(\mathcal{M}^{rPB})}\Pi(G^*,\mathcal{M}^{rPB},\sigma)=\pi^*$$

and so $\mathcal{M}^{rPB}$ is an optimal informationally robust mechanism and this completes the proof. ■

Theorem 3 provides one possible motivation for the ubiquity of pure bundling (especially for digital goods): this mechanism guarantees the best possible profit evaluated against model misspecification. Note that, unlike Theorem 1, even with additive values, separate sales cannot immediately be employed here as an alternative optimal informationally robust mechanism. To see this, suppose the seller offered a separate sales mechanism with a random but perfectly correlated price vector. Observe that there can be two separate signals G and G' that induce the same distribution over grand bundle estimates, but that generate different profits given this separate sales mechanism. Thus, we can no longer employ the simple argument that we did to generate the revenue guarantee given by (16). Loosely speaking, random pure bundling mechanisms have an advantage in providing higher revenue guarantees because they reduce the buyer's private information to a single dimension and, in doing so, the worst-case signal is effectively drawn from a smaller set.

Finally, as we have already mentioned, the fact that the max-min and min-max problems have the same solution for the sale of a single good was first observed by Du (2018). In independent and contemporaneous work, Brooks and Du (2020) generalize this insight to a variety of different settings (including multiple goods with additive values) with multiple buyers who have interdependent values. They consider finite type spaces, so the fact that this equivalence arises in our model is not per se implied by any of their results. More substantively, however, the aims of our respective papers are different. Our goal is to derive qualitative properties of the seller's optimal mechanism in two different information environments; the fact that the seller gets the same revenue in both is a fact we use to prove Theorem 3 not a main focus of this paper. Moreover, we show (in Corollary 1) that the seller's profit in every buyer-optimal outcome equals the value of the objective from the solution to the min-max problem (this equivalence is not a priori apparent). By contrast, Brooks and Du (2020) aim to show the equivalence of the min-max and max-min problems very generally, but they do not derive the seller's optimal mechanism in either; instead, their results are meant to provide a means for efficient numerical simulation. In the multiple-agent environment they consider, this equivalence is harder to establish than in our single agent setting due to the possibility of multiple equilibria.

5. Concluding remarks

In this paper, we study a general multidimensional screening problem for a seller in two different information environments with buyer learning. In the first, we derive the optimal mechanism under the information structure that maximizes consumer surplus. In the second, we derive the optimal informationally robust mechanism which provides the highest revenue guarantee for the seller against all possible information structures. We show that pure bundling emerges in both and that the seller's profit is the same in both problems. Our main theoretical insight is that the introduction of buyer learning allows us to reduce the seller's problem to a one-dimensional counterpart which, in turn, considerably simplifies this typically intractable problem.

We end with a brief discussion of the key assumption driving our results: the prior type distribution F is exchangeable. As we discussed, this is a considerably more general environment than what is typically analyzed in the multidimensional screening literature, but it is worth pointing out that our proof strategy does not extend to general distributions. This can happen even if F is such that each θ_i has identical marginal distributions; for instance, when the buyer's values for two goods are positive correlated but negatively correlated for a different pair. Our proof of Lemma 1, which uses the exchangeability of the prior, does not generalize to this case. Therefore, Theorem 1 cannot be directly generalized, since it builds on this result that any distribution over posterior grand bundle estimates (induced by a signal) can be induced by a perfectly correlated signal. We view this richness to be yet another interesting feature of this multi-dimensional environment and, in future work, we hope to generalize our results to such asymmetric environments.

APPENDIX A. MISSING PROOFS FROM THE TEXT

PROOF OF COROLLARY 1. Suppose, for contradiction, that there is an outcome $(G,\ (q,t))$ in which the seller's profit π is strictly lower than the profit π^* from a buyer-optimal outcome $(G^*,\ (q^*,t^*))$. Steps 1 and 2 in the proof of Theorem 1 show that it is possible to find an outcome $(G',\ (q',t'))$ in which trade is efficient and the seller's profit from her optimal mechanism is exactly π. But this is a contradiction because the consumer surplus in outcome $(G',\ (q',t'))$ will be $\overline{\mu}-\pi > \overline{\mu}-\pi^*$ contradicting the optimality of $(G^*,\ (q^*,t^*))$.

■

PROOF OF THEOREM 2. Denote the set of signals when there are n goods by $\mathcal{G}_n$ and the set of distributions on grand bundle estimates that they induce by $\overline{\mathcal{G}}_n$. Let $\tilde{F}_n$ denote the distribution of the average value $\frac{\theta_1+\cdots+\theta_n}{n} \in [0,1]$. Then observe that any Pareto distributed grand bundle signal $H_{\alpha_n} \in \overline{\mathcal{G}}_n$ will satisfy $\tilde{F}_n \succsim H_{\frac{\alpha_n}{n},\frac{\beta(\alpha_n)}{n}}$ or, in words, that $H_{\frac{\alpha_n}{n},\frac{\overline{\alpha}_n}{n}}$ is a *signal for the average value of n goods*. The converse is also true and $\tilde{F}_n \succsim H_{\frac{\alpha_n}{n},\frac{\beta(\alpha_n)}{n}}$ implies $H_{\alpha_n} \in \overline{\mathcal{G}}_n$. Therefore, if we have a signal for the average value, there will be a corresponding grand bundle signal which implies we can use Lemma 1 to find a perfectly correlated signal for the type vector.

Additionally, note that $\tilde{F}_n \succsim \tilde{F}_{n+1}$. Clearly $\tilde{F}_n$ and $\tilde{F}_{n+1}$ have the same mean $\tilde{\mu}$ and the former will have a higher variance because each θ_i is iid with distribution $\tilde{F}$. The fact that $\tilde{F}_n$ also second order stochastically dominates $\tilde{F}_{n+1}$ is a well known property of the distribution of sample means.

From the proof of Theorem 1, we know that there are buyer-optimal outcomes $\left(G_n^*,\ \mathcal{M}^{Sep}_{p_n^*}\right)$ and $\left(G_{n+1}^*,\ \mathcal{M}^{Sep}_{p_{n+1}^*}\right)$ for n and $n+1$ goods respectively where G_n^* and G_{n+1}^* are perfectly correlated and induce truncated Pareto distributed grand bundle estimates $H_{\overline{p}_n^*} \in \overline{\mathcal{G}}_n$ and $H_{\overline{p}_{n+1}^*} \in \overline{\mathcal{G}}_{n+1}$.

We have argued above that

$$\tilde{F}_n \succsim \tilde{F}_{n+1} \succsim H_{\frac{\overline{p}_{n+1}^*}{n+1},\frac{\beta(\overline{p}_{n+1}^*)}{n+1}}. \tag{17}$$

In words, this implies that, in addition to being a signal for the average value of $n+1$ goods, $H_{\frac{\overline{p}_{n+1}^*}{n+1},\frac{\beta(\overline{p}_{n+1}^*)}{n+1}}$ is also a signal for the average value of n goods. The optimality of $\overline{p}_n^*$ then implies that $\frac{\overline{p}_{n+1}^*}{n+1} \geq \frac{\overline{p}_n^*}{n}$ or, equivalently, that $CS_n \geq CS_{n+1}$.

Finally, as $n \to \infty$, the weak law of large numbers implies that $\tilde{F}_n \xrightarrow{d} \delta_{\tilde{\mu}}$ where $\xrightarrow{d}$ denotes convergence in distributions and $\delta_{\tilde{\mu}}$ is the Dirac measure that assigns mass 1 to the point $\tilde{\mu}$. Therefore, we will also have $H_{\frac{\overline{p}_n^*}{n},\frac{\beta(\overline{p}_n^*)}{n}} \xrightarrow{d} \delta_{\tilde{\mu}}$ and so, in the limit, $\frac{\overline{p}_n^*}{n} \to \tilde{\mu}$ and the seller extracts all the surplus. ■

References

ADAMS, W. J., AND J. L. YELLEN (1976): "Commodity Bundling and the Burden of Monopoly," *Quarterly Journal of Economics*, 90(3), 475–498.

ARMSTRONG, M. (2016): "Nonlinear Pricing," *Annual Review of Economics*, 8(1), 583–614.

BAGWELL, K. (2007): "Signalling and Entry Deterrence: A Multidimensional Analysis," *RAND Journal of Economics*, 38(3), 670–697.

BAKOS, Y., AND E. BRYNJOLFSSON (1999): "Bundling Information Goods: Pricing, Profits, and Efficiency," *Management science*, 45(12), 1613–1630.

BALDER, E. J. (1996): "On the Existence of Optimal Contract Mechanisms for Incomplete Information Principal-Agent Models," *Journal of Economic Theory*, 68(1), 133–148.

BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): "The Limits of Price Discrimination," *American Economic Review*, 105(3), 921–57.

BERGEMANN, D., AND S. MORRIS (2019): "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57(1), 44–95.

BERGEMANN, D., AND J. VÄLIMÄKI (2002): "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70(3), 1007–1033.

BROOKS, B., AND S. DU (2020): "A Strong Minimax Theorem for Informationally-Robust Auction Design," *Available at SSRN 3663721*.

CARROLL, G. (2017): "Robustness and Separation in Multidimensional Screening," *Econometrica*, 85(2), 453–488.

CHEN, B., AND D. NI (2017): "Optimal Bundle Pricing under Correlated Valuations," *International Journal of Industrial Organization*, 52, 248–281.

CHEN, Y., AND M. H. RIORDAN (2013): "Profitability of Product Bundling," *International Economic Review*, 54(1), 35–57.

DASKALAKIS, C., A. DECKELBAUM, AND C. TZAMOS (2017): "Strong Duality for a Multiple-Good Monopolist," *Econometrica*, 85(3), 735–767.

DU, S. (2018): "Robust Mechanisms under Common Valuation," *Econometrica*, 86(5), 1569–1588.

GENG, X., M. B. STINCHCOMBE, AND A. B. WHINSTON (2005): "Bundling Information Goods of Decreasing Value," *Management science*, 51(4), 662–667.

HAGHPANAH, N., AND J. HARTLINE (2020): "Is Pure Bundling Optimal?," *Review of Economic Studies*, forthcoming.

HART, S., AND N. NISAN (2019): "Selling Multiple Correlated Goods: Revenue Maximization and Menu-Size Complexity," *Journal of Economic Theory*, 183, 991–1029.

KAMENICA, E. (2019): "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 11(1), 249–272.

MCAFEE, R. P., J. MCMILLAN, AND M. D. WHINSTON (1989): "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *Quarterly Journal of Economics*, 104(2), 371–383.

PAVLOV, G. (2011): "Optimal Mechanism for Selling Two Goods," *BE Journal of Theoretical Economics*, 11(1).

PERSICO, N. (2000): "Information Acquisition in Auctions," *Econometrica*, 68(1), 135–148.

RAVID, D., A.-K. ROESLER, AND B. SZENTES (2020): "Learning Before Trading: On the Inefficiency of Ignoring Free Information," Available at SSRN: https://ssrn.com/abstract=3317917.

ROESLER, A.-K., AND B. SZENTES (2017): "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107(7), 2072–80.

SALINGER, M. A. (1995): "A Graphical Analysis of Bundling," *Journal of business*, 68(1), 85–98.

SHI, X. (2012): "Optimal Auctions with Information Acquisition," *Games and Economic Behavior*, 74(2), 666–686.

WOODCOCK, R. A. (2017): "The Obsolescence of Advertising in the Information Age," *Yale Law Journal*, 127, 2270–2341.