

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

RAHUL DEB[∅] AND ANNE-KATRIN ROESLER[✉]

SEPTEMBER 2, 2023

ABSTRACT: A monopolist seller of multiple goods screens a buyer whose type vector is initially unknown to both but drawn from a commonly known prior distribution. The seller chooses a mechanism to maximize her worst-case profits against all possible signals from which the buyer can learn about his values for the goods. We show that it is robustly optimal for the seller to bundle goods with identical demands (these are goods that can be permuted without changing the buyer’s prior type distribution). Consequently, pure bundling is robustly optimal for exchangeable prior distributions. For exchangeable priors, pure bundling is also optimal for the seller in the information environment (with the reverse timing) where an information designer, with the objective of maximizing consumer surplus, first selects a signal for the buyer, and then the seller chooses an optimal mechanism in response. We derive a formal relationship between the seller’s problem in both information environments.

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. Moreover, in environments where the optimal mechanism can be characterized, it often takes a complex form involving (possibly uncountable) menus of lotteries. But, in practice, multiproduct sellers often use simple mechanisms. For instance, online retailers (with significant market power) such as movie and music streaming services typically offer a single price for their entire catalog—a practice economists refer to as “pure bundling”—instead of, for instance, selling separate subscription bundles for movies and music of different genres. In this paper, we provide a rationale for this practice: we show pure bundling is optimal for a seller who does not precisely know the buyer’s type distribution and wants their chosen mechanism to be robustly optimal against all possible buyer information.

[∅]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 Ben Brooks, Gabriel Carroll, Songzi Du, Nima Haghpanah, Greg Pavlov, Andrea Prat, Xianwen Shi, Gabor Virag, the editor, the reviewers and numerous seminar participants. We thank the Social Sciences and Humanities Research Council for their continued and generous financial support.

In our benchmark model, the buyer initially has an unknown type $(\theta_1, \dots, \theta_n)$ that is drawn from a commonly known, exchangeable¹ distribution, where each $\theta_i \in [\theta_\ell, \theta_h] \subset \mathbb{R}_+$. The buyer's type determines his value $\sum_{i \in b} \theta_i$ for any bundle $b \subseteq \{1 \dots, n\} =: N$ of goods. The buyer learns about his type via a signal. Upon privately observing the signal realization, the buyer forms a posterior estimate of his type. The seller chooses a *robustly optimal mechanism*: this mechanism provides the highest profit guarantee against all possible buyer signals. In other words, the seller chooses a mechanism and nature picks a signal to minimize seller profits; the robustly optimal mechanism maximizes this worst-case profit. We show that pure bundling (albeit with a random price) is robustly optimal.

In our view, this information environment is a natural way to relax the standard assumption of perfect demand knowledge for sellers. There are many scenarios where a seller is unlikely to have a precise estimate of a buyer's value distribution. This would certainly be true for a seller determining how to price new goods that she is introducing to a market. Conversely, even armed with demand estimates from historical data, it is impossible for a seller to precisely predict what information the buyer has or will acquire to learn about his value for the goods. Our result provides one explanation for why, in practice, we do not observe very complex screening that depends on fine details of the type distribution (as is possible in multi-dimensional screening even with independent and identically distributed values). Instead, pure bundling is the common way that digital goods such as streaming services from retailers like Netflix and Spotify with considerable market power are sold.

For intuition on the robust optimality of pure bundling, first note that every mechanism yields a weakly higher profit guarantee when the buyer is restricted to learning via a smaller set of signals. In particular, if only signals that lead to perfectly (positively) correlated distributions of posterior estimates for the goods were permitted, the buyer's posterior type would effectively be one-dimensional and hence a mechanism that only allocates the grand bundle (that is, the bundle of all n goods) achieves the highest worst-case profit. Conversely, suppose the seller chose to pure bundle (with any random price) and nature could pick any signal. Since the seller only sells the grand bundle, her profits only depend on the distribution over grand bundle estimates (the expected value of $\theta_1 + \dots + \theta_n$). We show that, when the prior distribution is exchangeable, the set of possible distributions over grand bundle estimates induced by perfectly correlated signals is the same as the set of possible distributions over grand bundle estimates induced by unrestricted signals. This, in turn, implies that pure bundling is robustly optimal.

Having characterized the robustly optimal mechanism, we study the information environment in which the timing is reversed: nature, with the objective of maximizing consumer surplus, first selects a signal, and then the seller chooses an optimal mechanism in response. We refer to this signal and the corresponding seller best response as a *buyer-optimal outcome*. We show that a buyer-optimal outcome is also the solution to the problem where nature, with the goal of minimizing seller profits (as opposed to maximizing consumer surplus), first picks a signal and then the seller

¹This 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. As we will discuss, we can weaken exchangeability to allow for some specific forms of asymmetry. But since this assumption is ubiquitously made and easy to state, we choose to impose it for our benchmark model to simplify the presentation.

best responds. In other words, we show that the buyer-optimal outcome is the solution to the min-max problem corresponding to the max-min problem that yields the robustly optimal mechanism. We show that there is a buyer-optimal outcome in which the seller’s optimal mechanism is pure bundling and, moreover, that the seller gets the same profit in both problems. Moreover, pure bundling remains seller-optimal in both problems even when the buyer’s value for a bundle b takes the non-additive form $\kappa_b \sum_{i \in b} \theta_i$ where $\kappa_b \geq 0$ are bundle specific constants that satisfy a weak free-disposal assumption ($\kappa_b \sum_{i \in b} \theta_i \leq \kappa_N \sum_{i \in N} \theta_i$ for all $(\theta_1, \dots, \theta_n) \in [\theta_\ell, \theta_h]^n$ and all $b \subset N$). This value function allows for goods to be either substitutes or complements (in particular, it captures add-on items) and introduces asymmetry into the model.

Finally, we examine robust optimality for non-exchangeable priors. We first show that it is always robustly optimal to bundle subsets of goods with identical demands. Here, bundling a subset $B \subseteq N$ of goods refers to a mechanism that never allocates any proper subset of B with positive probability. A subset of goods is said to have identical demands if the prior joint distribution of the type vector remains unchanged when these goods are permuted (exchangeability corresponds to the case where all goods have identical demands). We view this result—bundling goods with identical demands is robustly optimal—to be the main economic insight of the paper in that it provides a simple and general principle for multi-dimensional screening. Some readers might find this result surprising because the structure of the optimal mechanism in the standard multi-dimensional screening environment typically depends finely on the details of the entire joint distribution of values.

Our last result shows that it is not always robustly optimal to bundle goods with non-identical demands. We demonstrate this by considering the case of two goods and a prior distribution such that each dimension of the buyer’s type is distributed independently and is identical up to a shift; an example of such a prior distribution is $\mathbb{U}[0, 1] \times \mathbb{U}[a, a + 1]$ where $a \geq 0$ and \mathbb{U} refers to the uniform distribution. For any distribution \tilde{F} of θ_1 , we show that there exists a bound such that, when θ_2 ’s independent distribution is also \tilde{F} but shifted by more than the bound, there is a separate sales mechanism (that is, a mechanism in which goods are priced individually) that yields strictly higher profit guarantee than all pure bundling mechanisms.

Related Literature

This paper lies at the intersection of a few 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 \(1976\)](#)), due to the complexity of the problem, there are surprisingly few general insights. A seminal result by [McAfee, McMillan, and Whinston \(1989\)](#) is that when values are additive and each dimension of the type is distributed independently, 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. [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 profit if she is restricted to using “simple mechanisms” like pure bundling or separate sales (Hart and Nisan, 2019). However, in contrast to these (theoretical) results, very complex mechanisms are not employed in practice and our results provide an explanation for the ubiquity of pure bundling.

Moreover, all these aforementioned papers assume additive values. Our result (the robust optimality of pure bundling) is relatively unusual in that it extends to the non-additive value setting we described above and we view this to be a strength of our framework. The recent survey of Armstrong (2016) describes a strand of the screening literature (with non-additive values) that does not aim to characterize the optimal mechanism but instead derives conditions under which the seller can profit from offering bundle discounts. A notable exception is Haghpanah and Hartline (2021) who characterize environments where pure bundling is seller optimal, and we employ one of their results in our proofs.

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 mechanism chosen by the seller, the efficiency of trade, and the resulting surplus division in bilateral trade settings.² Ravid, Roesler, and Szentes (2022) study a single good environment where buyer-learning is unobservable but costly. Their main result shows that there is a distinction between free and arbitrarily cheap learning. Bergemann, Brooks, and Morris (2015) and Haghpanah and Siegel (2022) study a standard single and multiple good monopoly pricing problem, respectively, 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.

Within this literature, the closest papers are Du (2018) and Roesler and Szentes (2017); they respectively analyze the one-dimensional versions of the two problems we study. Du (2018) derives the informationally robust optimal mechanism (a random posted price) for a single good³ and uncovers the relation to the buyer-optimal outcome. 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 is better off by not learning his value for the product perfectly. 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; the seller of a single good can only ever post a price (either random or deterministic).

The (informationally) robust optimal mechanism that we derive is also related to the broader literature on robust mechanism design. Within this literature, the closest paper is Carroll (2017). He considers a seller who knows the marginal distribution of the buyer’s value for each good

²Similar 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.

³He 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 profit guarantee converges to the full surplus.

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. The main difference between his and our setting is the set over which the seller evaluates worst-case profits; these sets are not nested. In our case, the distribution of the buyer's posterior type estimate must be obtained by Bayesian updating. As a result, any signal jointly determines both the marginal distribution of each dimension of the type and the correlation across dimensions.

Finally, there are two closely related contemporaneous papers [Brooks and Du \(2021\)](#) and [Che and Zhong \(2022\)](#); we will discuss the relation to them after we present our first result ([Theorem 1](#)).

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 $N = \{1, \dots, n\}$.

This section describes our benchmark model (which is a canonical version of the multi-dimensional screening model) that has additive values, an exchangeable type distribution and zero seller costs. We discuss generalizations of first two assumptions in [Sections 3.3](#) and [4.1](#) respectively. In the concluding remarks, we discuss how our main insight is unaffected by positive seller costs.

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 \mathcal{F}) with $\theta_h > \theta_\ell \geq 0$. The type is initially unknown to both the buyer and seller, and is drawn from a commonly known (cumulative) distribution F .⁴ Throughout the paper, we assume F has a positive density for all $\theta \in \Theta$.

In the benchmark model, we assume that F is *exchangeable*: for any permutation $\sigma : N \rightarrow N$, the joint distribution of $(\theta_1, \dots, \theta_n)$ is the same as the joint distribution of $(\theta_{\sigma(1)}, \dots, \theta_{\sigma(n)})$ (both of which are F).⁵

This requires the marginal distribution of every θ_i to be the same and is clearly satisfied when each θ_i is independent and identically distributed (henceforth iid). Importantly, note that exchangeability allows for both positive and negative correlations between dimensions of the type vector. It is worth flagging that we make this assumption to simplify the presentation and, as we will discuss, our results only require a weaker assumption implied by exchangeability.

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

Value Function: A buyer of type θ has an *additive value* for each bundle $b \subseteq N$ which is given by

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

so, in particular, the buyer's value of not receiving a good is $u(\theta, \emptyset) = 0$. We assume that preferences are quasilinear in the transfers and that both the seller and buyer are risk-neutral expected utility maximizers.

⁴We will abuse notation and interchangeably refer to F as a cumulative distribution (henceforth, cdf) and a probability measure where convenient. The meaning will be clear from the argument (element vs. set) of F .

⁵Formally, for any $X \in \mathcal{F}$, we have $F(X) = F(X_\sigma)$ where $X_\sigma = \{(\theta_{\sigma(1)}, \dots, \theta_{\sigma(n)}) \mid (\theta_1, \dots, \theta_n) \in X\}$.

Signals: The buyer learns about his type via a *signal*. Given the linearity of our model and risk-neutrality of players, we will (as is common in the literature), without loss, 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 $S \times \Theta$ 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 $s = (s_1 \dots, s_n)$ 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 . The buyer privately observes the signal realization, $s \in S$.

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. 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 (direct) *mechanism* $\mathcal{M} = (q, t)$ that consists of an (possibly random) allocation $q : S \rightarrow \Delta(2^N)$ and a transfer $t : S \rightarrow \mathbb{R}$. The allocation determines the likelihood of receiving the various bundles and not being allocated the good; we sometimes use $q(s, b)$ to denote the probability that the buyer is allocated bundle b when he reports s .

If the buyer with posterior estimate s reports s' , his expected utility is

$$\mathbb{E}_{q(s')} [u(s, b)] - t(s'),$$

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.

A mechanism (q, t) is *incentive compatible* if

$$\mathbb{E}_{q(s)} [u(s, b)] - t(s) \geq \mathbb{E}_{q(s')} [u(s, b)] - t(s') \tag{IC}$$

for every $s, s' \in S$. Further, the mechanism is *individually rational* if

$$\mathbb{E}_{q(s)} [u(s, b)] - t(s) \geq 0. \tag{IR}$$

for all $s \in S$.

We denote the set of **IC** and **IR** mechanisms by \mathcal{M} and henceforth, when not specified, we always implicitly assume every mechanism \mathcal{M} is incentive compatible (so $\mathcal{M} \in \mathcal{M}$). Faced with a mechanism in this class, the buyer will report his posterior estimate s truthfully.

$U(s, \mathcal{M}) := \mathbb{E}_{q(s)} [u(s, b)] - t(s)$ denotes the utility of a buyer with signal realization s facing the mechanism $\mathcal{M} \in \mathcal{M}$. Conversely, we denote the seller’s profit by $\Pi(G, \mathcal{M}) := \mathbb{E}_G [t(s)]$ where the expectation is taken with respect to distribution over signal realizations.

Pure bundling and separate sales: We will refer to two special classes of mechanisms repeatedly. The first is a *pure bundling mechanism* at price \bar{p} . This is the mechanism (q, t) in which the buyer is

only allowed to purchase the grand bundle N at a price \bar{p} . Formally, for all $b \neq \emptyset$,

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

and $q(s, \emptyset) = 1 - q(s, N)$. We denote the set of pure bundling mechanisms by $\mathcal{M}^{\text{PB}} \subset \mathcal{M}$.

The second is a *separate sales mechanism* at prices $p = (p_1, \dots, p_N)$. 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, this is a mechanism (q, t) given by

$$q(s, b) = \begin{cases} 1 & \text{if } b = \hat{b}(s, p), \\ 0 & \text{otherwise,} \end{cases} \quad \text{and } t(s) = \begin{cases} \sum_{i \in b} p_i & \text{if } b = \hat{b}(s, p), \\ 0 & \text{otherwise,} \end{cases} \quad (\text{SS})$$

where $\hat{b}(s, p)$ is a bundle (possibly the empty set) that satisfies $\hat{b}(s, p) \in \operatorname{argmax}_{b' \subseteq N} \{u(s, b') - \sum_{i \in b'} p_i\}$. In words, this is the bundle that gives the buyer the highest positive utility when individual goods are priced separately. When there is more than one such bundle, the mechanism arbitrarily allocates one of the bundles to the buyer.

We will also refer to the randomized versions of these mechanisms. A *random pure bundling mechanism* has an allocation and transfer rule that only depend on the grand bundle values and moreover, the buyer is only ever allocated the grand bundle with positive probability. 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. Formally, this is a mechanism (q, t) given by

$$\begin{aligned} q(s, N) &= q(s', N), \quad t(s) = t(s') \quad \text{whenever } \sum_{i \in N} s_i = \sum_{i \in N} s'_i \quad \text{and} \\ q(s, b) &= 0 \quad \text{for } \emptyset \subset b \subset N, \quad q(s, \emptyset) = 1 - q(s, N) \quad \text{for all } s \in S. \end{aligned} \quad (\text{rPB})$$

The set of random pure bundling mechanisms is denoted by $\mathcal{M}^{\text{rPB}} \subset \mathcal{M}$.

Finally, a *random separate sales mechanism* is simply a separate sales mechanism where the seller chooses the prices randomly; for any realized prices, the buyer is free to choose any bundle and just pay the sum of the prices of the goods in that bundle. Formally, the seller chooses a distribution over prices $\mathcal{P} \in \Delta(\mathbb{R}^n)$ which results in an allocation rule and transfer

$$q(s, b) = \mathbb{E}_{\mathcal{P}} \left[\mathbb{1}_{\{b = \hat{b}(s, p)\}} \right] \quad \text{and } t(s) = \mathbb{E}_{\mathcal{P}} \left[\sum_{b \subseteq N} \mathbb{1}_{\{b = \hat{b}(s, p)\}} \sum_{i \in b} p_i \right], \quad (\text{rSS})$$

where the expectation is taken with respect to the distribution \mathcal{P} and $\mathbb{1}_{\{b = \hat{b}(s, p)\}}$ is the indicator function that takes the value 1 when $b = \hat{b}(s, p)$, 0 otherwise.

This completes the description of our model and we now proceed to showing the optimality of pure bundling under buyer learning.

3. THE OPTIMALITY OF PURE BUNDLING

This section is organized as follows. In Section 3.1, we first show that random pure bundling is a robustly optimal mechanism and discuss an application of this result. In Section 3.2, we then show

that there is a buyer-optimal outcome in which, once again, the seller chooses pure bundling, and derive a formal relation between the two problems. Finally, in Section 3.3, we present a class of non-additive value functions for which all the results generalize. All proofs are in the appendix.

3.1. Informational robustness

The main focus of this section is to derive qualitative properties of the mechanism that provides the seller the highest profit guarantee against all signals. This setting captures a seller who is uncertain about how a buyer learns about his value for the goods and the seller wants to insure herself against the worst-case scenario. In other words, the seller does not know the buyer's information acquisition technology and evaluates each mechanism based on the worst-case profits taken with respect to all possible signals and buyer best responses.

We begin with a few definitions.

Profit Guarantee: We say that a mechanism $\mathcal{M} \in \mathcal{M}$ provides a *profit guarantee* of π if

$$\Pi(G, \mathcal{M}) \geq \pi, \text{ for all signals } G \in \mathcal{G}.$$

Robustly Optimal Mechanism: Formally, we define the *robustly optimal mechanism* \mathcal{M}^* as the mechanism that solves

$$\mathcal{M}^* \in \operatorname{argmax}_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \Pi(G, \mathcal{M}). \quad (2)$$

This is the mechanism we aim to characterize. In words, it provides the seller the *highest profit guarantee* against all possible signals. We denote this highest profit guarantee by

$$\pi^* := \max_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \Pi(G, \mathcal{M}).$$

Before proceeding, it is worth making a few comments about the above definition. First, note that, implicit in the seller optimizing over direct mechanisms, is the fact that she gets to break ties in her favor (since the mechanism picks the allocation and transfer). Since signals may have atoms, such favorable tie-breaking might have a material impact on the seller's profit. It is reasonable to demand that a notion of robust optimality should not just guard against different signals but also against the profit-minimizing best response chosen by the buyer. We show in the appendix that the robust optimality of bundling does not depend on favorable tie-breaking by the buyer; we chose not to include buyer tie-breaking in the above definition to simplify the presentation.

Second, the definition implicitly suggests that a robustly optimal mechanism exists and, needless to say, this existence will be a consequence of our proof. Finally, note that, if the seller chooses to randomize the prices of different bundles, nature picks a signal *before* the price is realized. This is a consequence of the fact that the seller is choosing a direct mechanism and so can commit to random allocations after the buyer reports his signal realization to the mechanism.

We are now in the position to present our first main result.

THEOREM 1. *There is a random pure bundling mechanism that is a robustly optimal mechanism.*

An immediate consequence of this result is that an explicit characterization of the robustly optimal mechanism (that is, the menu of prices and probabilities with which the grand bundle is allocated) follows from Du (2018). This is because Theorem 1 shows that the seller’s problem can effectively be reduced to one of a single-good monopolist (whose good is the grand bundle).

Why is it robustly optimal to only sell the grand bundle? We will provide detailed intuition below but the high-level economic reason is that it is optimal to bundle together ex-ante similar goods. Choosing a mechanism that only allocates the grand bundle effectively shrinks the set of signals that the seller has to guard against since any two distinct signals (of which there are many) that generate the same distribution over posterior estimates of the grand bundle value lead to the same profit. But is there a different mechanism that can do better against all such signals? As we will argue, the answer is no because the prior distribution F is exchangeable and consequently, for any distribution of grand bundle estimates generated by a signal, there is a “perfectly correlated” signal that generates the same distribution. Against such perfectly correlated signals, it is robustly optimal for the seller to pure bundle.

We now formalize this intuition a little more, beginning with a definition.

Perfectly correlated signals: 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 $\{(s_1, \dots, s_n) \in S \mid s_1 = \dots = s_n\}$. We denote the set of perfectly correlated signals by $\mathcal{G}^{pc} \subseteq \mathcal{G}$.

Given a signal realization $s \in S$, we use $\bar{s} = s_1 + \dots + s_n$ to denote the sum. The set of all such sums \bar{s} is denoted by \bar{S} ; note that $\bar{S} = \bar{\Theta}$ (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 $\bar{G} \in \Delta(\bar{S})$ over the posterior estimates \bar{s} of the grand bundle value. We use $\bar{\mathcal{G}}$ to denote the set of these distributions of grand bundle estimates that are induced by some signal $G \in \mathcal{G}$. Similarly, we use $\bar{\mathcal{G}}^{pc}$ to denote the set of these distributions of grand bundle estimates that are induced by perfectly correlated signals $G \in \mathcal{G}^{pc}$.

Now, first observe that

$$\max_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} \leq \max_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}^{pc}} \{\Pi(G, \mathcal{M})\} = \max_{\mathcal{M} \in \mathcal{M}^{rPB}} \inf_{G \in \mathcal{G}^{pc}} \{\Pi(G, \mathcal{M})\}.$$

The inequality follows from the fact that the infimum is taken over a smaller set of signals. Loosely, the equality follows from the fact that, when the signals are perfectly correlated, the buyer’s type is effectively one-dimensional so it suffices for the seller to employ mechanisms that only allocate the grand bundle.

Conversely, observe that

$$\max_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} \geq \max_{\mathcal{M} \in \mathcal{M}^{rPB}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} = \max_{\mathcal{M} \in \mathcal{M}^{rPB}} \inf_{G \in \mathcal{G}^{pc}} \{\Pi(G, \mathcal{M})\}.$$

The inequality follows from the fact that the maximum is taken over a smaller set of mechanisms. When the mechanism only allocates the grand bundle, the seller’s profit is determined completely by the distribution of grand bundle estimates. The equality follows from the fact that $\bar{\mathcal{G}} = \bar{\mathcal{G}}^{pc}$ or, in words, that perfectly correlated signals $G \in \mathcal{G}^{pc}$ generate the same set of distributions over grand bundle estimates as (general) signals $G \in \mathcal{G}$ (Lemma 1 in the appendix shows this formally).

This is a consequence of the property that $\mathbb{E}_F[\theta_i - \theta_j | \bar{\theta}] = 0$ for all $i, j \in N$ when the prior F is exchangeable. In words, the symmetry of the prior ensures that, if the buyer only learns about his value $\bar{\theta}$ for the grand bundle, his conditional value for each good is $\bar{\theta}/n$. Note that, exchangeability is a stronger assumption than necessary to ensure this property⁶ but we chose to impose this assumption nonetheless as it is easier to state and interpret.

Taken together, both inequalities must in fact be equalities which imply that

$$\max_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} = \max_{\mathcal{M} \in \mathcal{M}^{PB}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\}$$

or that there is random pure bundling mechanism that provides the highest profit guarantee.

Before proceeding, it is worth contrasting our theorem with results in some recent related papers. As mentioned earlier, [Haghpanah and Hartline \(2021\)](#) provide a condition under which pure-bundling is optimal in the standard (non-robust) multidimensional screening problem. Their setting is more general than ours in that they allow for arbitrary value functions but theirs is a joint condition on the value function *and* the type distribution.⁷ The best way to highlight the difference is to discuss the restriction of their condition to additive values. When the type distribution is perfectly correlated, their condition applies and the optimal mechanism is pure bundling. However, their condition does not hold in general for exchangeable type distributions; as is well known, pure bundling is not always optimal for iid type distributions. Conversely, our result holds for all exchangeable distributions. Indeed, as the above discussion suggests, our proof exploits the fact that the robust optimality criterion allows us to restrict attention to perfectly correlated signals.

We would also like to acknowledge some contemporaneous work that generalizes [Theorem 1](#) along two distinct dimensions.⁸ [Brooks and Du \(2021\)](#) independently sketch an argument similar to the one described above and argue that [Theorem 1](#) extends to the case of an auction with multiple buyers. [Che and Zhong \(2022\)](#) build on our work by showing that [Theorem 1](#) generalizes to (i) the seller being uncertain about the exact prior distribution as well as buyer-learning and (ii) certain non-exchangeable prior distributions, requiring that the vector of valuations is built from a co-monotonic component and an idiosyncratic component that has zero ex-ante means. The main aim of their paper is to generalize the results of [Carroll \(2017\)](#) and their proof follows a distinct approach they develop towards achieving this goal. Specifically, they consider a setting where there is a partition of the set of goods and there is an exogenously given set of distributions for the value $(\sum_{i \in b} \theta_i)$ of each bundle b that is an element of the partition. Their seller evaluates worst-case profits across all possible joint distributions whose marginal distributions of the value of each partition element b lie in the given set.

As mentioned earlier, we are going to generalize [Theorem 1](#) along two different dimensions not covered by either paper. First, in [Section 3.3](#), we first describe how our result applies to a certain class of non-additive values that now allow goods to be either complements or substitutes. Then,

⁶For example, $\mathbb{E}_F[\theta_i - \theta_j | \bar{\theta}] = 0$ for the non-exchangeable prior F for three goods that assigns equal probability to the types $(1, 2, 3)$, $(3, 1, 2)$, and $(2, 3, 1)$.

⁷They show pure bundling is optimal if, for all bundles $b \subset N$, the distribution of $\frac{u(\theta, b)}{u(\theta, N)}$ conditional on $u(\theta, N) = \bar{u}$ is first-order stochastically non-decreasing in \bar{u} .

⁸Both results appeared online after the first version of our working paper was circulated.

in Section 4.1, we study non-exchangeable priors (that do not need to satisfy the co-monotonicity property of Che and Zhong, 2022) and show that it is always robustly optimal for the seller to bundle goods with identical demands. But more importantly, our goal is to develop the economic implications of Theorem 1 which we do in the remainder of this subsection and in Section 4.

The careful reader would have noted that Theorem 1 does not state that a robustly optimal mechanism *must* be random pure bundling. Indeed the intuition above, involving perfectly correlated signals, might suggest that a random separate sales mechanism (for instance, with perfectly correlated prices) can also provide the same profit guarantee. But this is not the case.

COROLLARY 1. *Suppose the value for each good is iid with distribution \check{F} (so the prior $F = \check{F} \times \cdots \times \check{F}$). No random separate sales mechanism provides the highest profit guarantee π^* .*

The proof of this result is in an online appendix but we provide some intuition here. Let $\check{\pi}^*$ denote the highest profit guarantee that a seller can achieve when she sells a single good to a buyer whose value is drawn from \check{F} . As Theorem 1 shows, π^* is the highest profit guarantee the seller can achieve by selling the grand bundle to a buyer whose value is drawn from \bar{F} (the prior distribution of the sum $\bar{\theta}$). Put differently, π^* and $\check{\pi}^*$ can both be found by deriving the highest profit guarantee in a one-dimensional setting which is a problem that has been solved in the literature. Now, if $n\check{\pi}^* < \pi^*$ (by definition, it cannot be $>$), then the seller cannot achieve a profit guarantee of π^* by random separate sales.

In words, this inequality states that a single-good monopolist can achieve a strictly higher revenue guarantee when she faces a buyer whose type is drawn from distribution \check{F} as opposed to a buyer whose type is the average of n iid draws from \check{F} . To see this, note that the distribution of the average value $\frac{\theta_1 + \cdots + \theta_n}{n}$ is a mean-preserving contraction⁹ of \check{F} . This implies that the infimum in the definition of the highest profit guarantee is taken over a larger set when the prior type distribution is \check{F} (since, in this one-dimensional setting, every signal is a mean-preserving contraction of the prior type distribution).

For intuition on why the highest profit guarantee from a separate sales mechanism is $n\check{\pi}^*$, recall that a random separate sales mechanism is a joint distribution over prices $\mathcal{P} \in \Delta(\mathbb{R}^n)$. Now, consider signals of the form $G = \check{G}_1 \times \cdots \times \check{G}_n \in \mathcal{G}$ in which the posterior estimate of the value of each good i is drawn independently from \check{G}_i . Against such signals, no distribution \mathcal{P} over prices can yield a profit guarantee higher than $n\check{\pi}^*$ because the profit guarantee from the sale of each good is determined by the marginal distribution of prices for that good derived from \mathcal{P} . This simply amounts to separately deriving the optimal distribution of prices for each of n identical one-dimensional problems in which the buyer's value is drawn from \check{F} and, by definition, the seller cannot achieve a profit guarantee higher than $\check{\pi}^*$ from the sale of each good.

As an example for two goods, when \check{F} is $\mathbb{U}[0, 1]$ and therefore, \bar{F} is the triangle distribution on $[0, 2]$ with mode 1, we have $.408 = 2\check{\pi}^* < \pi^* = .482$.

We end this subsection with an application of Theorem 1. The canonical multi-dimensional screening problem (without buyer learning), in addition to being generally intractable, has several

⁹A distribution is a mean-preserving contraction of another if both have the same mean and the latter second-order stochastically dominates the former. A formal definition can be found in the proof of Theorem 2 in the appendix.

counterintuitive features. One particularly surprising property is shown by [Hart and Reny \(2015\)](#). Unlike the case of a single good, the optimal mechanism for a seller of multiple goods might yield a lower profit for a type distribution that first-order stochastically dominates another.¹⁰ This is because the optimal mechanism need not be “monotone” in that a higher type (component by component) might end up paying strictly less than a lower type. In fact, [Ben Moshe, Hart, and Nisan \(2022\)](#) have recently shown that restricting attention to monotone mechanisms may yield a negligible fraction of the maximal profit.

The next result shows that this counterintuitive property disappears for the robustly optimal mechanism (the result is actually a little stronger). This is because a random pure bundling mechanism is effectively one-dimensional and so the highest profit guarantee only depends on the prior distribution \bar{F} of grand-bundle values.

THEOREM 2. *Suppose the type distribution F first-order stochastically dominates the type distribution F' . Then, for any random pure bundling mechanism $\mathcal{M} \in \mathcal{M}^{rPB}$, we have*

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

where \mathcal{G} and \mathcal{G}' are the sets of signals corresponding to priors F and F' respectively.

Consequently, the seller’s highest profit guarantee from F is greater than that from F' .

3.2. The buyer-optimal outcome

In this subsection, we formally define and characterize a buyer-optimal outcome and relate it to the robustly optimal mechanism. We begin by defining an outcome and what it means for an outcome to be buyer-optimal.

Outcome: An *outcome* is a pair (G, \mathcal{M}) where $G \in \mathcal{G}$ is a signal and \mathcal{M} is an optimal mechanism for the seller in response to distribution G (that is, $\Pi(G, \mathcal{M}) \geq \Pi(G, \mathcal{M}')$ for all $\mathcal{M}' \in \mathcal{M}$).

Buyer-Optimal Outcome: A *buyer-optimal outcome* maximizes the buyer’s surplus across all outcomes and is given by

$$\begin{aligned} (G^*, \mathcal{M}^*) \in \operatorname{argmax}_{\mathcal{M} \in \mathcal{M}, G \in \mathcal{G}} \mathbb{E}_G [U(s, \mathcal{M})] \\ \text{such that } \Pi(G, \mathcal{M}) \geq \Pi(G, \mathcal{M}') \text{ for all } \mathcal{M}' \in \mathcal{M}. \end{aligned} \tag{3}$$

Note that the constraint corresponds to verifying that \mathcal{M} is the solution to the standard multi-dimensional screening problem for the signal G . It is well-defined because an optimal mechanism exists for every type distribution (see [Balder, 1996](#)). Since optimal multi-dimensional screening for arbitrary type distributions is intractable, our proof approach circumvents having to evaluate the objective at every signal G and its corresponding optimal mechanism \mathcal{M} . Moreover, our proof explicitly constructs the signal G^* in the buyer-optimal outcome thereby showing that the maximum for the objective function is obtained.

¹⁰A type distribution F first-order stochastically dominates distribution F' if for all increasing functions $y : \Theta \rightarrow \mathbb{R}$, we have $\int_{\Theta} y(\theta) dF(\theta) \geq \int_{\Theta} y(\theta) dF'(\theta)$. (A function y is increasing if $y(\theta) \geq y(\theta')$ when $\theta_i \geq \theta'_i$ for all $1 \leq i \leq n$.) This is the standard definition of first-order stochastic dominance for multivariate distributions (see Section 6 in [Shaked and Shanthikumar, 2007](#)).

We are now in a position to present our second main result.

THEOREM 3. *There exists a buyer-optimal outcome (G^*, \mathcal{M}^*) which has the following properties.*

- (i) Seller Best-Response: \mathcal{M}^* is a pure bundling mechanism.
- (ii) Signal: G^* is perfectly correlated.
- (iii) Total Surplus: The buyer is allocated the grand bundle with probability one so trade is efficient.

It might be surprising to some readers that the signal in the buyer-optimal outcome does not just perfectly reveal the buyer's type. After all, this would allow him to always get the highest ex-post consumer surplus by making the optimal purchase decision. Indeed, the fact that trade is efficient implies that the buyer sometimes purchases the grand bundle when he should not. To see this, note that when $\theta_\ell = 0$, the seller's optimal mechanism never gives away the grand bundle for free since she can always guarantee herself a strictly positive profit by selling the grand bundle at a price just above zero. So the optimal mechanism \mathcal{M}^* must be providing the seller a strictly positive profit which, combined with efficient trade, implies that the buyer makes purchases with positive probability when his true value for the grand bundle is below the seller's profit.

The reason perfect information about his type is not optimal for the buyer is because this allows the seller to extract more surplus by screening effectively. To prevent this, the signal in the buyer-optimal outcome injects two different types of noise into learning. First, it only provides information about the value of the grand bundle. As discussed in the previous subsection, for exchangeable prior distributions F , this is possible without additional information about the individual dimensions trickling through, yielding a perfectly correlated distribution of posterior estimates. This effectively makes the type space one-dimensional and reduces the seller's ability to screen across dimensions.

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. Hence, we can employ the argument of [Roesler and Szentes \(2017\)](#) (who study the sale of a single good) to show that it is possible to construct a signal such that there is an optimal pure bundling mechanism for the seller in which the grand bundle is always traded. Finally, since the signal is perfectly correlated, this mechanism is also a best response of the seller in the original problem. 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^* .

Note the relation between the robustly optimal mechanism that solves (2) and the buyer-optimal outcome that solves (3). In the former, the seller chooses to pure bundle in order to protect against non perfectly correlated signals that provide the buyer differential information about the dimensions of his type. In the latter, the signal is perfectly correlated in order to prevent the seller from screening effectively by using the full power of multi-dimensional screening. In fact, these problems are formally related; the next result describes a property of the buyer-optimal outcome that will make this connection clear.

COROLLARY 2. *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 says that every buyer-optimal outcome not only maximizes consumer surplus, it also minimizes profits. The first part of the above corollary implies that the seller’s profit π^* must be the same in every buyer-optimal outcome from which it then follows by statement (iii) of [Theorem 3](#) that trade must be efficient in *every* buyer-optimal outcome. Thus [Corollary 2](#) strengthens statement (iii) of [Theorem 3](#) as the latter simply refers to the existence of a buyer-optimal outcome with efficient trade.

This corollary also implies that the seller’s profit 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,

$$\pi^* = \min_{G \in \mathcal{G}} \max_{\mathcal{M} \in \mathcal{M}} \Pi(G, \mathcal{M}). \quad (4)$$

In fact, the seller’s profit is identical in both problems (2) and (4):

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

The fact that the max-min and min-max problems have the same value for the case of a single good was first observed by [Du \(2018\)](#). [Brooks and Du \(2021\)](#) 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 profit in both is a consequence, but not a main focus of this paper. By contrast, [Brooks and Du \(2021\)](#) 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.

We end this subsection with an application of [Theorem 3](#) to derive a comparative static relating consumer surplus in the buyer-optimal outcome to the number of goods. We begin with some context that motivates this application. 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 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 the case of two goods and a buyer whose value for each good is independently and uniformly distributed on $[0, 1]$. Here, the optimal separate sales mechanism charges a price of $\frac{1}{2}$ for each good. Now, suppose that in addition to selling the goods individually, the seller could pure bundle. She would choose to do the latter with the optimal pure bundling mechanism charging a price of $\sqrt{\frac{2}{3}} < \frac{1}{2} + \frac{1}{2}$ for the grand bundle. This 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 $n \rightarrow \infty$. 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. 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. We show that such an analysis is possible for the buyer-optimal outcome.

The next result refers to the *average consumer surplus*, which is the consumer surplus in the buyer-optimal outcome divided by the number of goods. CS_n refers to the average consumer surplus from the buyer-optimal outcome corresponding to the prior distribution F . Now consider the marginal distribution over the first $n - 1$ goods derived from F and note this distribution is also exchangeable. We use CS_{n-1} to denote the average consumer surplus in the buyer-optimal outcome corresponding to this prior distribution.

THEOREM 4. *The average consumer surplus in the buyer-optimal outcome is decreasing in the number of goods: $CS_n \leq CS_{n-1}$.*

This result highlights the nuanced interplay between information and screening but might seem counterintuitive in light of the following two facts discussed above. First, there is always a buyer-optimal outcome in which the seller chooses to pure bundle. Second, without buyer learning, the optimal pure bundling mechanism can yield a strictly *higher* average consumer surplus when the number of goods increases (for example when $n = 2$ and the value of each good is independently drawn from the uniform $[0, 1]$ distribution). These statements do not conflict because the information that the buyer receives in the buyer-optimal outcome changes as the number of goods increases. In order to prevent the seller from effectively screening across dimensions, the signal in the buyer-optimal outcome introduces 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 increases, such correlation surrenders more surplus to the seller.

3.3. Substitutes and complements

All the results from Sections 3.1 and 3.2 hold¹¹ when the buyer has the *non-additive value function*

$$u(\theta, b) = \kappa_b \sum_{i \in b} \theta_i \text{ where } \kappa_b \geq 0, \quad (5)$$

which implies the value of not receiving a good is $u(\theta, \emptyset) = 0$. We assume that $u(\theta, N) \geq u(\theta, b)$ for all $b \subseteq N$ and all $\theta \in \Theta$. This a weak free-disposal property and ensures that the greatest surplus is generated by trading the grand-bundle (as in the additive values case). In terms of the κ_b -s, this requires that $\kappa_b \leq \left(1 + \frac{N-|b|}{|b|} \frac{\theta_\ell}{\theta_n}\right) \kappa_N$. Note that we do not require $\kappa_{b'} \geq \kappa_b$ (or $u(\theta, b') \geq u(\theta, b)$) for proper subsets $b \subset b' \subset N$. Indeed, when $\theta_\ell > 0$, we can have $\kappa_b > \kappa_N$. Clearly, this value function subsumes the additive value framework of [Section 2](#).

¹¹Our older working paper [Deb and Roesler \(2021\)](#) contains the proofs of Theorems 1 and 3 for this value function.

This value function allows for goods to be either substitutes or complements, generality that is important for modeling the demand for multiple goods.¹² As an example, consider a buyer evaluating two bundles from a movie streaming service: the first bundle b only offers comedy movies and the second bundle N is the entire catalog. It would be natural for the buyer to assign a lower marginal value to any individual comedy movie when they have the entire catalog instead of when they only have the comedy bundle since the likelihood that they will ever watch the movie is lower when they have more choice. In this case, $\kappa_N < \kappa_b$.

Conversely, consider a two-partner household evaluating the same two-tiered bundles. In this case, the marginal value for a movie might be higher for the entire catalog. This is because the household may end up watching fewer comedy movies (the preference of one partner) if they do not have the option of watching movies from other genres (the preference of the second partner). In this case, $\kappa_N > \kappa_b$. Of course, in both scenarios the total value of the entire catalog will be higher ($\kappa_N \sum_{i \in N} \theta_i \geq \kappa_b \sum_{i \in b} \theta_i$) since disposal is free for a digital product.

While we situate this example in the context of a streaming service, it should be clear that goods can be substitutes or complements in a variety of different multiproduct settings. While the above value function is more general than the bulk of the literature (that focuses on additive values), it retains the feature that only the posterior type estimate is relevant for the mechanism design problem. In other words, the value of each bundle is linear in the buyer's type; this allows us to restrict attention to unbiased signals which is an essential feature for tractability. For this reason, linearity is a commonly made assumption in the information design literature.

We end this section with two additional observations. First, note that this general value function allows the buyer to assign different values to two bundles b and b' even though $\sum_{i \in b} \theta_i = \sum_{i' \in b'} \theta_{i'}$. This allows us to accommodate some limited asymmetry in the environment even with an exchangeable type distribution F . In particular, this value function can capture add-on items. Returning to the example of a streaming service, consider a buyer whose primary concern is to acquire a specific TV show for his children which is say captured by good i . He values having access to additional shows and movies but only if the bundle contains the show i . For such a buyer, $\kappa_b = 0$ for all $b \subset N$, $i \notin b$, and $0 < u(\theta, \{i\}) \leq u(\theta, b') \leq u(\theta, N)$ for all $b' \subset N$, $i \in b'$. In the next section, we maintain additive values but discuss the implications of a non-exchangeable prior.

Finally, note that [Theorem 3](#) did not claim that the buyer-optimal outcome is unique and, indeed, pure bundling need not be the unique seller best response to a perfectly correlated signal. Specifically (and unlike [Theorem 1](#)), with additive values, it is also 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, when values can take the more general form of this subsection, there are cases where separate sales may not be part of any buyer-optimal outcome. This observation, demonstrated in the next example, is the reason we highlight pure bundling (as opposed to separate sales) in [Theorem 3](#).

¹²[Geng, Stinchcombe, and Whinston \(2005\)](#) study a standard multi-dimensional screening setting with such a value function, but they require that κ_b is decreasing in the number of goods in the bundle b . Within this framework, they provide a sufficient condition for the approximate optimality of pure bundling.

Example. Suppose there are two goods and that $\kappa_{\{i\}} > \kappa_N = 1$ 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 (G^*, \mathcal{M}^*) where \mathcal{M}^* is a separate sales mechanism at prices $p = (p_1, p_2)$. **Corollary 2** 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 $G^*(\{(s_1, s_2) \mid (s_1 + s_2) - (p_1 + p_2) \leq \varepsilon\}) > 0$. 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 these inequalities over both goods, we get

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

where $\kappa_\ell = \min\{\kappa_{\{1\}}, \kappa_{\{2\}}\}$. 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 not to buy the grand bundle thereby implying trade is not efficient in the buyer-optimal outcome (G^*, \mathcal{M}^*) .

4. BEYOND EXCHANGEABILITY

In this section, we discuss properties of the robustly optimal mechanism (for additive values) when the assumption of an exchangeable prior distribution is relaxed. In Section 4.1, we first show that our main economic insight—goods with identical demands should be bundled—extends to non-exchangeable priors. In Section 4.2, we then show that this is not true for goods whose demands differ. Specifically, random pure bundling is *not* a robustly optimal mechanism even for two goods whose values are independently distributed and are identical up to a shift.

4.1. Bundling goods with identical demands

Without loss, let the bundle $B = \{1, \dots, m\}$ consist of the first $2 \leq m \leq n$ goods. We use θ_B and θ_{-B} (and a similar notation for s) to denote $(\theta_1, \dots, \theta_m)$ and $(\theta_{m+1}, \dots, \theta_n)$ respectively.

Identical demands: We say goods in bundle B have *identical demands* if, for any permutation $\sigma : B \rightarrow B$, the joint distributions of $(\theta_1, \dots, \theta_n)$ and $(\theta_{\sigma(1)}, \dots, \theta_{\sigma(m)}, \theta_{m+1}, \dots, \theta_n)$ are both F .¹³

One simple example of such a prior is when the joint distribution of bundle B is independent from the joint distribution of θ_{-B} and the marginal distribution of B is exchangeable.

We now define what it means to bundle a subset of goods.

Bundling a subset of goods: We say that a mechanism $(q, t) \in \mathcal{M}$ bundles the set $B \subseteq N$ of goods if it satisfies the following two properties:

- (i) *Proper subsets of B are never allocated:* if $q(s, b) > 0$ for $s \in S$ and $b \subseteq N$, then $B \cap b \in \{\emptyset, B\}$;
- (ii) *Only the value of bundle B matters:* for types $(s_B, s_{-B}), (s'_B, s_{-B})$ such that $\sum_{i \in B} s_i = \sum_{i \in B} s'_i$, the allocation $q(s_B, s_{-B}) = q(s'_B, s_{-B})$ and transfer $t(s_B, s_{-B}) = t(s'_B, s_{-B})$ are the same.

¹³Formally, for any $X \in \mathcal{F}$, we have $F(X) = F(X_\sigma)$ where $X_\sigma = \{(\theta_{\sigma(1)}, \dots, \theta_{\sigma(m)}, \theta_{m+1}, \dots, \theta_n) \mid (\theta_1, \dots, \theta_n) \in X\}$.

Our next result generalizes [Theorem 1](#) by showing that it is robustly optimal to bundle goods with identical demands.

THEOREM 5. *Suppose the goods in bundle B have identical demands. If a mechanism provides a profit guarantee of π then, there is a mechanism that bundles B that also provides a profit guarantee of π .*

The intuition for this result mirrors that of [Theorem 1](#). Bundling goods in B implies that only the distribution of $\bar{s}_B = s_1 + \dots + s_m$ matters for profits so there are fewer signals to guard against. Moreover, because goods in B have identical demands, for any signal $G \in \mathcal{G}$, there is another signal $G' \in \mathcal{G}$ that induces the same distribution over $(\bar{s}_B, s_{m+1}, \dots, s_n)$ but for which the marginal distribution of (s_1, \dots, s_m) is perfectly correlated. Against such signals, it is robustly optimal to bundle the goods in B .

The observant reader would have noticed that [Theorem 5](#) is phrased differently to [Theorem 1](#) in that it does not refer to a robustly optimal mechanism. We chose this particular statement since it has the same economic content and a relatively simple proof that does not require us to establish the existence of a robustly optimal mechanism.

Finally, note that, unlike [Theorem 1](#) which provides a complete characterization, [Theorem 5](#) only describes a qualitative property of the robustly optimal mechanism. Specifically, it does not characterize the probability with which bundle B is allocated as a function of the posterior estimates s_{-B} of the other goods $N \setminus B$ (and vice versa). Deriving a complete characterization is beyond the scope of this paper but we view this to be an interesting question for future work.

4.2. When bundling is not robustly optimal

The previous subsection allowed for non-exchangeable priors but focused on the allocation of goods with identical demands. What about goods whose demands differ? We now argue that bundling together such goods can be strictly suboptimal.

To demonstrate this point, we focus on a class of simple non-exchangeable environments that are obtained from an iid setting by simply shifting the support of one marginal. Specifically, we consider the case of two goods (so $n = 2$) whose joint distribution is given by $F = \check{F} \times \check{F}_x$ so each dimension of the type is distributed independently with marginal distributions \check{F} and \check{F}_x respectively. Moreover, the distribution of the second dimension is an x -shift of the first where $x > 0$. Formally, $\check{F}_x(\theta_2) = \check{F}(\theta_2 - x)$ for all $\theta_2 \in [\theta_\ell + x, \theta_h + x]$.¹⁴

Our last result shows, that for sufficiently large x , there is a random separate sales mechanism that provides a profit guarantee that no random pure bundling mechanism can match. The implication, of course, is that no random pure bundling mechanism is robustly optimal.

THEOREM 6. *Suppose $n = 2$ and the prior distribution is given by $F = \check{F} \times \check{F}_x$ where \check{F} is supported on $[\theta_\ell, \theta_h]$ with $\theta_\ell = 0$ and \check{F}_x is an x -shift of \check{F} .*

Then, there is an $\bar{x} > 0$ such that for all $x \geq \bar{x}$, there exists a random separate sales mechanism that provides a profit guarantee that no random pure bundling mechanism can provide.

¹⁴The fact that the distributions of both dimensions of the type are supported on distinct intervals is not critical for the point we are trying to make. We focus on shifted distributions since this seemed like the smallest qualitative departure from exchangeability.

The following is brief intuition for this result. The distribution of grand bundle values is supported on $[2\theta_\ell + x, 2\theta_h + x]$. As x becomes larger, the highest profit guarantee achievable with a random pure bundling mechanism converges to the lower bound of the support $2\theta_\ell + x$. Loosely speaking, this is because the profit loss from not selling the good increases relative to the benefits from randomizing the price of the grand bundle. Conversely, the seller can always guarantee herself a profit of $\tilde{\pi} + \theta_\ell + x$ by setting a price of $\theta_\ell + x$ for the second good (thereby always selling it) and randomizing the price of the first good in a way that corresponds to the robustly optimal mechanism for a single good whose prior value is distributed by \check{F} . Since $\tilde{\pi} > 0$ when $\theta_\ell = 0$ (the seller can always guarantee a positive profit by pricing just above 0), we have $\tilde{\pi} + \theta_\ell + x > 2\theta_\ell + x$ as required.¹⁵

We end this section by noting that the bound \bar{x} above which random separate sales starts providing a strictly higher profit guarantee than every random pure bundling mechanism need not be very large. For instance, if \check{F} is $\mathbb{U}[0, 1]$, the bound $\bar{x} = 1.75$ reflects a relatively small shift.

5. CONCLUDING REMARKS

Before concluding, it is worth providing a brief discussion of the assumption of zero seller costs that we imposed throughout. Our most general result on robust optimality ([Theorem 5](#)) holds when the seller has a cost $c > 0$ of producing each good; indeed, the proof of this result in the appendix incorporates such costs. We did not impose this generality at the outset because [Theorem 3](#) that characterizes the buyer-optimal outcome does not similarly generalize. The proof of that result was built around the fact that trade is efficient in the buyer-optimal outcome. This continues to be the case when the cost is positive but sufficiently low but not when the cost becomes larger.

To summarize, in this paper, we study a general multi-dimensional screening problem with buyer learning. We derive qualitative features of the robustly optimal mechanism: this is the mechanism that provides the highest profit guarantee for the seller against all signals from which the buyer learns his value. We main economic insight is simple and general: it is robustly optimal to bundle goods with identical demands.

APPENDIX

This appendix contains the proofs of all the results in the text.

APPENDIX A. PROOFS OF [THEOREM 1](#) AND [THEOREM 5](#)

We first prove [Theorem 5](#) and then invoke it to prove [Theorem 1](#).

We need to introduce some additional notation and terminology. Given $s \in S$, recall that $\bar{s}_B = s_1 + \dots + s_m$ denotes the sum of posterior estimates of goods in the bundle B . The set of all (\bar{s}_B, s_{-B}) is denoted by

$$\tilde{S}_B := [m\theta_\ell, m\theta_h] \times [\theta_\ell, \theta_h]^{n-m} =: \tilde{\Theta}_B.$$

¹⁵This intuition can be generalized to arbitrary $\theta_\ell > 0$ and we actually prove a slightly more general statement of [Theorem 6](#) in the appendix ([Theorem 6'](#)). We chose the special case of $\theta_\ell = 0$ to simplify the statement.

Every signal $G \in \mathcal{G}$ induces a distribution $\tilde{G}_B \in \Delta(\tilde{S}_B)$ over the posterior estimates (\bar{s}_B, s_{-B}) . We use $\tilde{\mathcal{G}}_B$ to denote the set of these distributions that are induced by signals $G \in \mathcal{G}$. Similarly, \tilde{F}_B denotes the distribution over $\tilde{\Theta}_B$ induced by the prior F .

We say that a signal $G \in \mathcal{G}$ is *perfectly correlated for bundle B* if $G(\{s \in S \mid s_1 = \dots = s_m\}) = 1$ or in words, that all the mass of the signal is on the subset of S for which the signal realizations of goods in B are identical. We denote the set of such signals by $\mathcal{G}_B^{pc} \subseteq \mathcal{G}$.

We now present a lemma that will be useful to prove the optimality of pure bundling.

LEMMA 1. *For every signal $G \in \mathcal{G}$, there exists a signal $G' \in \mathcal{G}_B^{pc}$ that is perfectly correlated for bundle B such that G and G' induce the same distribution \tilde{G}_B over \tilde{S}_B .*

PROOF. We begin by defining the $n - m + 1$ -dimensional signals that only provide information about the sum of values for goods in B but there are no restrictions on the information provided about goods not in the bundle B . These are (unbiased) signals $(\tilde{S}_B, H_{\tilde{S}_B \times \tilde{\Theta}_B})$ where $H_{\tilde{S}_B \times \tilde{\Theta}_B} \in \Delta(\tilde{S}_B \times \tilde{\Theta}_B)$ is a joint distribution over $\tilde{S}_B \times \tilde{\Theta}_B$ such that the marginal distribution of $H_{\tilde{S}_B \times \tilde{\Theta}_B}$ over $\tilde{\Theta}_B$ is \tilde{F}_B and

$$(\bar{s}_B, s_{-B}) = \mathbb{E}_{H_{\tilde{S}_B \times \tilde{\Theta}_B}} [(\bar{\theta}_B, \theta_{-B}) \mid (\bar{s}_B, s_{-B})]$$

for all (\bar{s}_B, s_{-B}) in the support.

We use H to denote the marginal distribution of $H_{\tilde{S}_B \times \tilde{\Theta}_B}$ over the set of signal realizations \tilde{S}_B and use \mathcal{H} to denote the set of all such distributions. As with the signals $(S, G_{S \times \Theta})$ for the type vector, it is without loss to restrict attention to such unbiased signals.

We first argue that $\tilde{\mathcal{G}}_B \subseteq \mathcal{H}$. To see this, observe that signal $G_{S \times \Theta}$ induces a joint distribution $G_{\tilde{S}_B \times \tilde{\Theta}_B}$ over $\tilde{S}_B \times \tilde{\Theta}_B$ such that the marginal distribution over $\tilde{\Theta}_B$ is \tilde{F}_B . Formally, $G_{\tilde{S}_B \times \tilde{\Theta}_B}$ is the image measure of $G_{S \times \Theta}$ generated by the mapping $a(s, \theta) = (\sum_{i \in B} s_i, s_{-B}, \sum_{i \in B} \theta_i, \theta_{-B})$ which implies that, for any measurable set $A \subseteq \tilde{S}_B \times \tilde{\Theta}_B$, we have $G_{\tilde{S}_B \times \tilde{\Theta}_B}(A) = G_{S \times \Theta}(a^{-1}(A))$. Moreover, observe that

$$\begin{aligned} \mathbb{E}_{G_{\tilde{S}_B \times \tilde{\Theta}_B}} [(\bar{\theta}_B, \theta_{-B}) \mid (\bar{s}_B, s_{-B})] &= \mathbb{E}_{G_{S \times \Theta}} [(\bar{\theta}_B, \theta_{-B}) \mid (\bar{s}_B, s_{-B})] \\ &= \mathbb{E}_{G_{S \times \Theta}} [\mathbb{E}_{G_{S \times \Theta}} [(\bar{\theta}_B, \theta_{-B}) \mid s] \mid (\bar{s}_B, s_{-B})] = \mathbb{E}_{G_{S \times \Theta}} [(\bar{s}_B, s_{-B}) \mid (\bar{s}_B, s_{-B})] = (\bar{s}_B, s_{-B}) \end{aligned}$$

for all (\bar{s}_B, s_{-B}) in the support. Therefore the marginal distribution \tilde{G}_B over \tilde{S}_B induced by $G_{\tilde{S}_B \times \tilde{\Theta}_B}$ satisfies $\tilde{G}_B \in \mathcal{H}$.

We now show that, for every $H \in \mathcal{H}$, there exists a $G \in \mathcal{G}$ that is perfectly correlated for bundle B such that the distribution it induces on \tilde{S}_B satisfies $\tilde{G}_B = H$. (This also shows that $\mathcal{H} \subseteq \tilde{\mathcal{G}}_B$.)

By definition, H is the marginal distribution over \tilde{S}_B corresponding to an unbiased signal $H_{\tilde{S}_B \times \tilde{\Theta}_B} \in \Delta(\tilde{S}_B \times \tilde{\Theta}_B)$. We use the distribution $H_{\tilde{S}_B \times \tilde{\Theta}_B}$ to define a family of conditional distributions $\hat{H}(\cdot \mid \theta) \in \Delta(\tilde{S}_B)$ as follows

$$\hat{H}(\cdot \mid \theta) := H(\cdot \mid \theta_1 + \dots + \theta_m, \theta_{m+1}, \dots, \theta_n). \quad (6)$$

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

Now observe that

$$\mathbb{E}_{\hat{H}_{\bar{s}_B \times \Theta}} [(\bar{\theta}_B, \theta_{-B}) \mid (\bar{s}_B, s_{-B})] = (\bar{s}_B, s_{-B})$$

for all (\bar{s}_B, s_{-B}) in the support. This is a consequence of the definition (6) of $\hat{H}_{\bar{s}_B \times \Theta}$ and from the fact that $(\bar{s}_B, s_{-B}) = \mathbb{E}_{H_{\bar{s}_B \times \Theta_B}} [(\bar{\theta}_B, \theta_{-B}) \mid (\bar{s}_B, s_{-B})]$.

Given the joint distribution $\hat{H}_{\bar{s}_B \times \Theta}$, we can derive the conditional distribution $\hat{H}(\cdot \mid \bar{s}_B)$ over Θ . Now observe that the conditional distribution $\hat{H}(\cdot \mid \bar{s}_B)$ has the feature that goods in B have identical demands. This follows from the definition of $\hat{H}_{\bar{s}_B \times \Theta}$ and because F is such that goods in B have identical demands. This in turn implies

$$\mathbb{E}_{\hat{H}_{\bar{s}_B \times \Theta}} [\theta_i \mid (\bar{s}_B, s_{-B})] = \frac{\bar{s}_B}{m} \quad \text{for all } i \in \{1, \dots, m\}.$$

Now define a joint distribution $G_{S \times \Theta}$ over $S \times \Theta$ that is the image measure of $\hat{H}_{\bar{s}_B \times \Theta}$ generated by the mapping $\hat{a}(\bar{s}_B, s_{-B}, \theta) = \left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{b+1}, \dots, s_n \right)$. Formally, for any measurable $\hat{A} \subseteq S \times \Theta$, we have $G_{S \times \Theta}(\hat{A}) = \hat{H}_{\bar{s}_B \times \Theta}(\hat{a}^{-1}(\hat{A}))$. By construction, the marginal distribution of $G_{S \times \Theta}$ over Θ is F , the marginal distribution G over S assigns measure 1 to the set $\{s \in S \mid s_1 = \dots = s_m\}$ and so is perfectly correlated for bundle B .

Observe that

$$\mathbb{E}_{G_{S \times \Theta}} \left[\theta \mid s = \left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{b+1}, \dots, s_n \right) \right] = \mathbb{E}_{\hat{H}_{\bar{s}_B \times \Theta}} [\theta \mid (\bar{s}_B, s_{-B})] = \left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{b+1}, \dots, s_n \right),$$

and so $G \in \mathcal{G}$ or, in words, that G is an unbiased signal. Finally, by construction, the distribution \tilde{G}_B over posterior estimates of $(\bar{\theta}_B, \theta_{-B})$ induced by G satisfies $\tilde{G}_B = H$ which completes the proof. \blacksquare

With this lemma in hand, we are now in a position to prove [Theorem 5](#). As we mentioned in the conclusion, we will in fact prove a more general version of the result in which the seller has a cost $c \geq 0$ of each good. In the proof, we use one additional piece of notation: $\mathcal{M}^B \subset \mathcal{M}$ denotes the set of mechanisms that bundle B .

PROOF OF THEOREM 5. We first show that

$$\sup_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} \geq \sup_{\mathcal{M} \in \mathcal{M}^B} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} = \sup_{\mathcal{M} \in \mathcal{M}^B} \inf_{G \in \mathcal{G}_B^{pc}} \{\Pi(G, \mathcal{M})\}. \quad (7)$$

Here, the inequality simply follows from the fact that the right side takes the supremum over the smaller set of mechanisms that bundle B . The equality is a consequence of [Lemma 1](#). When the seller offers a mechanism $\mathcal{M} \in \mathcal{M}^B$ that bundles goods in B , for any signal $G \in \mathcal{G}$, the profit only depends on the distribution \tilde{G}_B that the signal G induces over the vector \tilde{s}_B . But as [Lemma 1](#) shows, there is another signal $G' \in \mathcal{G}_B^{pc}$ that is perfectly correlated for bundle B that induces the same distribution \tilde{G}_B on \tilde{s}_B . Therefore, taking the infimum over \mathcal{G} or \mathcal{G}_B^{pc} leads to the same value.

Next, we show that

$$\sup_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} \leq \sup_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}_B^{pc}} \{\Pi(G, \mathcal{M})\} = \sup_{\mathcal{M} \in \mathcal{M}^B} \inf_{G \in \mathcal{G}_B^{pc}} \{\Pi(G, \mathcal{M})\}. \quad (8)$$

The inequality is a consequence of taking the infimum over the smaller set of signals. We will show the equality by the following argument: for every mechanism $(q, t) \in \mathcal{M}$, we will construct a mechanism $(\hat{q}, \hat{t}) \in \mathcal{M}^B$ that bundles B and generates the same profit for the seller at all signals in \mathcal{G}_B^{pc} . Since every signal in \mathcal{G}_B^{pc} is perfectly correlated in bundle B , the seller's profit only depends on how the mechanism is defined for types of the form $(\hat{s}, \dots, \hat{s}, s_{-B})$.

So given a mechanism (q, t) , we first define \hat{q} as follows. Consider an arbitrary $b \subset N$ such that $b \cap B = \emptyset$ (such a b can be the empty set). We set

$$\hat{q}(s, b') = 0 \text{ for all } b \subset b' \subset b \cup B \text{ and all } s \in S.$$

In words, the allocation \hat{q} assigns zero probabilities to bundles b' which contain a strict subset of the goods in B as required for a mechanism that bundles B .

Then, we set

$$\hat{q}(s, B \cup b) = \sum_{b \subset b' \subseteq b \cup B} \frac{|B \cap b'|}{m} q \left(\left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{-B} \right), b' \right)$$

and

$$\hat{q}(s, b) = \sum_{b \subset b' \subseteq b \cup B} q \left(\left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{-B} \right), b' \right) - \hat{q}(s, B \cup b),$$

for all $s \in S$. In words, the altered allocation rule \hat{q} is different from q in two ways. First, as required for a mechanism that bundles B , for any signal realization $s \in S$, the allocation \hat{q} only depends on (\bar{s}_B, s_{-B}) . Second, \hat{q} is constructed from q by moving the allocation probability from any bundle b' such that $b \subset b' \subset b \cup B$ to the bundles b and $B \cup b$.

Note that \hat{q} is a well defined allocation rule in that, for all $s \in S$, $\hat{q}(s, b) \in [0, 1]$ for all $b \subseteq N$, and $\sum_{b \subseteq N} \hat{q}(s, b) = 1$. The former follows immediately from the definition; to see the latter, observe that

$$\begin{aligned} \sum_{b \subseteq N} \hat{q}(s, b) &= \sum_{b \cap B = \emptyset} (\hat{q}(s, B \cup b) + \hat{q}(s, b)) \\ &= \sum_{b \cap B = \emptyset} \left(\hat{q}(s, B \cup b) + \sum_{b \subset b' \subseteq b \cup B} q \left(\left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{-B} \right), b' \right) - \hat{q}(s, B \cup b) \right) \\ &= \sum_{b \cap B = \emptyset} \sum_{b \subset b' \subseteq b \cup B} q \left(\left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{-B} \right), b' \right) \\ &= \sum_{b \subseteq N} q \left(\left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{-B} \right), b \right) \\ &= 1. \end{aligned}$$

The transfer for a type $s \in S$ is simply set to

$$\hat{t}(s) = t \left(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{-B} \right).$$

and observe that \hat{t} is identical to t at all types $(\hat{s}, \dots, \hat{s}, s_{-B}) \in S$. Therefore, if (\hat{q}, \hat{t}) is IC and IR, it yields the same revenue for the seller as (q, t) for every signal $G \in \mathcal{G}_B^{pc}$.

We now show that (\hat{q}, \hat{t}) is indeed **IC** and **IR**. Note that, by construction, every type $s \in S$ gets the same utility as $(\frac{\bar{s}_B}{m}, \dots, \frac{\bar{s}_B}{m}, s_{-B})$. Also note that the utility that every type $(\hat{s}, \dots, \hat{s}, s_{-B})$ gets from (mis)reporting as type $(\hat{s}', \dots, \hat{s}', s'_{-B})$ is the same in both mechanisms (q, t) and (\hat{q}, \hat{t}) because

$$\begin{aligned}
 & \mathbb{E}_{\hat{q}(\hat{s}', \dots, \hat{s}', s'_{-B})} [u((\hat{s}, \dots, \hat{s}, s_{-B}), b)] \\
 = & \sum_{b \cap B = \emptyset} [\hat{q}((\hat{s}', \dots, \hat{s}', s'_{-B}), B \cup b) u((\hat{s}, \dots, \hat{s}, s_{-B}), b \cup B) + \hat{q}((\hat{s}', \dots, \hat{s}', s'_{-B}), b) u((\hat{s}, \dots, \hat{s}, s_{-B}), b)] \\
 = & \sum_{b \cap B = \emptyset} \left[\hat{q}((\hat{s}', \dots, \hat{s}', s'_{-B}), B \cup b) \left(\sum_{i \in b} s_i + m\hat{s} \right) + \hat{q}((\hat{s}', \dots, \hat{s}', s'_{-B}), b) \left(\sum_{i \in b} s_i \right) \right] \\
 = & \sum_{b \cap B = \emptyset} \left[(\hat{q}((\hat{s}', \dots, \hat{s}', s'_{-B}), B \cup b) + \hat{q}((\hat{s}', \dots, \hat{s}', s'_{-B}), b)) \left(\sum_{i \in b} s_i \right) + \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), B \cup b) (m\hat{s}) \right] \\
 = & \sum_{b \cap B = \emptyset} \left[\sum_{b \subseteq b' \subseteq b \cup B} q((\hat{s}', \dots, \hat{s}', s'_{-B}), b') \left(\sum_{i \in b} s_i \right) + \sum_{b \subseteq b' \subseteq b \cup B} q((\hat{s}', \dots, \hat{s}', s'_{-B}), b') (|B \cap b'| \hat{s}) \right] \\
 = & \sum_{b \cap B = \emptyset} \left[\sum_{b \subseteq b' \subseteq b \cup B} q((\hat{s}', \dots, \hat{s}', s'_{-B}), b') u((\hat{s}, \dots, \hat{s}, s_{-B}), b') \right] \\
 = & \sum_{b \subseteq N} q((\hat{s}', \dots, \hat{s}', s'_{-B}), b) u((\hat{s}, \dots, \hat{s}, s_{-B}), b) \\
 = & \mathbb{E}_{q(\hat{s}', \dots, \hat{s}', s'_{-B})} [u((\hat{s}, \dots, \hat{s}, s_{-B}), b)],
 \end{aligned}$$

and the transfers under t and \hat{t} are identical for such types. This shows (\hat{q}, \hat{t}) is **IC** and **IR** because (q, t) is **IC** and **IR**.

It remains to be shown that the seller's expected cost under mechanism (\hat{q}, \hat{t}) is the same as under (q, t) . Take any type $(\hat{s}, \dots, \hat{s}, s_{-B})$. The cost of serving this type under mechanism (\hat{q}, \hat{t}) is

$$\begin{aligned}
 & \sum_{b \subseteq N} c|b| \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), b) \\
 = & \sum_{b \cap B = \emptyset} \sum_{b \subseteq b' \subseteq b \cup B} c|b'| \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), b') \\
 = & \sum_{b \cap B = \emptyset} [c|b| \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), b) + c|b \cup B| \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), b \cup B)] \\
 = & \sum_{b \cap B = \emptyset} \sum_{b \subseteq b' \subseteq b \cup B} c|b| q((\hat{s}, \dots, \hat{s}, s_{-B}), b') \\
 & - \sum_{b \cap B = \emptyset} c|b| \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), B \cup b) + \sum_{b \cap B = \emptyset} c|b \cup B| \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), b \cup B) \\
 = & \sum_{b \cap B = \emptyset} \sum_{b \subseteq b' \subseteq b \cup B} c|b| q((\hat{s}, \dots, \hat{s}, s_{-B}), b') + \sum_{b \cap B = \emptyset} cm \hat{q}((\hat{s}, \dots, \hat{s}, s_{-B}), b \cup B) \\
 = & \sum_{b \cap B = \emptyset} \sum_{b \subseteq b' \subseteq b \cup B} c|b| q((\hat{s}, \dots, \hat{s}, s_{-B}), b') + \sum_{b \cap B = \emptyset} cm \sum_{b \subseteq b' \subseteq b \cup B} \frac{|B \cap b'|}{m} q((\hat{s}, \dots, \hat{s}, s_{-B}), b') \\
 = & \sum_{b \cap B = \emptyset} \sum_{b \subseteq b' \subseteq b \cup B} c|b'| q((\hat{s}, \dots, \hat{s}, s_{-B}), b') \\
 = & \sum_{b \subseteq N} c|b| q((\hat{s}, \dots, \hat{s}, s_{-B}), b).
 \end{aligned}$$

Since the revenue and cost for the seller are the same under both mechanisms, we have shown $\Pi(G, (q, t)) = \Pi(G, (\hat{q}, \hat{t}))$ for all $G \in \mathcal{G}_B^{pc}$ as required.

The above argument shows that the inequalities in (7) and (8) are in fact equalities which in turn implies that

$$\sup_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} = \sup_{\mathcal{M} \in \mathcal{M}^B} \inf_{G \in \mathcal{G}_B^{pc}} \{\Pi(G, \mathcal{M})\} = \sup_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}_B^{pc}} \{\Pi(G, \mathcal{M})\}.$$

If, for the left term, the value of the supremum is not attained by any $\mathcal{M} \in \mathcal{M}$, then the theorem follows from the first equality. Conversely, if the value of the supremum is achieved by some $\mathcal{M} \in \mathcal{M}$, the above construction together with the second equality imply that there must also be a mechanism $\mathcal{M} \in \mathcal{M}^B$ that attains the value of the supremum. This completes the proof. ■

Recall that, after we defined a robustly optimal mechanism in equation (2), we had mentioned that this definition implicitly assumed that the buyer, when indifferent between reporting type s or s' , broke ties in favor of the seller. We had mentioned that adversarial tie breaking will not affect our results; this is demonstrated by the proof of [Theorem 5](#). To see this, redefine the profit function $\Pi(G, \mathcal{M})$ so that the profit for a given signal G and mechanism \mathcal{M} is computed by assuming the buyer, from his set of best responses, chooses the reporting strategy that minimizes the seller's profit. With this new definition, the proof goes through unchanged. This is because the key step of the proof that shows $\sup_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}_B^{pc}} \{\Pi(G, \mathcal{M})\} = \sup_{\mathcal{M} \in \mathcal{M}^B} \inf_{G \in \mathcal{G}_B^{pc}} \{\Pi(G, \mathcal{M})\}$ is constructive: for any $\mathcal{M} \in \mathcal{M}$, we constructed a mechanism $\widehat{\mathcal{M}} \in \mathcal{M}^B$ such that all types in the support of any signal $G \in \mathcal{G}_B^{pc}$ get the same utility, the seller gets the same transfer and the cost is the same. Therefore, even with the new definition of profits (with adversarial tie-breaking), it must be the case that $\Pi(G, \mathcal{M}) = \Pi(G, \widehat{\mathcal{M}})$ for all $G \in \mathcal{G}_B^{pc}$. Finally, also observe that this argument remains unchanged if we defined mechanisms more generally with message spaces that were not necessarily the type space S .

PROOF OF THEOREM 1. The proof of [Theorem 5](#) (taking $B = N$) implies that

$$\sup_{\mathcal{M} \in \mathcal{M}} \inf_{G \in \mathcal{G}} \{\Pi(G, \mathcal{M})\} = \sup_{\mathcal{M} \in \mathcal{M}^{rPB}} \inf_{G \in \mathcal{G}^{pc}} \{\Pi(G, \mathcal{M})\}.$$

When the seller chooses a random pure bundling mechanism, for any signal $G \in \mathcal{G}$, the profits are only determined by the distribution \bar{G} over grand bundle values induced by G . But this implies that the problem $\sup_{\mathcal{M} \in \mathcal{M}^{rPB}} \inf_{G \in \mathcal{G}^{pc}} \{\Pi(G, \mathcal{M})\}$ is equivalent to deriving the robustly optimal mechanism for a single good (whose prior value distribution is \bar{F}). [Du \(2018\)](#) shows such an optimal mechanism exists (he additionally characterizes it) and this completes the proof. ■

APPENDIX B. PROOF OF THEOREM 2

PROOF OF THEOREM 2. Analogously to \bar{F} and \bar{G} , we use \bar{F}' and \bar{G}' to respectively denote the distribution of grand bundle values and the set of distributions of posterior estimates corresponding to F' . It is well known that the set \bar{G} (respectively \bar{G}') consists of all mean-preserving contractions of \bar{F} (respectively \bar{F}'). A distribution $H' \in \Delta(\bar{\Theta})$ is a *mean-preserving contraction* of another

distribution $H \in \Delta(\bar{\Theta})$ if

$$\int_{n\theta_\ell}^z H(\bar{s}) \, d\bar{s} \geq \int_{n\theta_\ell}^z H'(\bar{s}) \, d\bar{s} \quad \text{for all } z \in [n\theta_\ell, n\theta_h] \text{ with equality for } z = n\theta_h. \quad (9)$$

With this definition in place, we prove the theorem in two steps.

Step 1: For every distribution $\bar{G} \in \bar{\mathcal{G}}$, there exists a distribution $\bar{G}' \in \bar{\mathcal{G}}'$ such that \bar{G} first-order stochastically dominates \bar{G}' .

We begin by noting that

$$\int_{n\theta_\ell}^z \bar{F}'(\bar{s}) \, d\bar{s} \geq \int_{n\theta_\ell}^z \bar{F}(\bar{s}) \, d\bar{s} \geq \int_{n\theta_\ell}^z \bar{G}(\bar{s}) \, d\bar{s} \quad \forall z \in [n\theta_\ell, n\theta_h]$$

where the first inequality follows from the fact that \bar{F} first-order stochastically dominates \bar{F}' and the second from the fact that \bar{G} is a mean-preserving contraction of \bar{F} .

First observe that, if \bar{F} and \bar{F}' have equal means, then the above inequalities hold with equality at $z = n\theta_h$ and so the claim in this step can be proved by choosing $\bar{G}' = \bar{G}$.

So suppose the mean of \bar{F} is strictly greater than \bar{F}' . Define

$$\bar{z} := \max \left\{ z \in [n\theta_\ell, n\theta_h] \mid \int_{n\theta_\ell}^z \bar{F}'(\bar{s}) \, d\bar{s} = \int_{n\theta_\ell}^z \bar{G}(\bar{s}) \, d\bar{s} \right\}$$

and note the maximum exists because both terms on either side of the equality are continuous in z and equal at $z = n\theta_\ell$. A consequence of this definition is that

$$\int_{n\theta_\ell}^z \bar{F}'(\bar{s}) \, d\bar{s} > \int_{n\theta_\ell}^z \bar{G}(\bar{s}) \, d\bar{s} \quad \text{for all } z \in (\bar{z}, n\theta_h].$$

Now, for $t \in [\bar{z}, n\theta_h]$, define the distribution

$$\bar{G}'_t(\bar{s}) := \begin{cases} \bar{G}(\bar{s}) & \text{if } \bar{s} \in [n\theta_\ell, t) \\ 1 & \text{if } \bar{s} \in [t, n\theta_h], \end{cases}$$

In words, the distribution \bar{G}'_t is the same as \bar{G} before the point t and assigns all remaining mass in \bar{G} (to the right of t) to an atom of \bar{G}'_t at t . Observe that, by construction, \bar{G} first-order stochastically dominates \bar{G}'_t for every $t \in [\bar{z}, n\theta_h]$.

When $t = n\theta_h$, we have

$$\int_{n\theta_\ell}^{n\theta_h} \bar{G}'_{n\theta_h}(\bar{s}) \, d\bar{s} = \int_{n\theta_\ell}^{n\theta_h} \bar{G}(\bar{s}) \, d\bar{s} < \int_{n\theta_\ell}^{n\theta_h} \bar{F}'(\bar{s}) \, d\bar{s},$$

and when $t = \bar{z}$, we have

$$\int_{n\theta_\ell}^{n\theta_h} \bar{G}'_{\bar{z}}(\bar{s}) \, d\bar{s} = \int_{n\theta_\ell}^{\bar{z}} \bar{G}(\bar{s}) \, d\bar{s} + \int_{\bar{z}}^{n\theta_h} 1 \, d\bar{s} = \int_{n\theta_\ell}^{\bar{z}} \bar{F}'(\bar{s}) \, d\bar{s} + \int_{\bar{z}}^{n\theta_h} 1 \, d\bar{s} \geq \int_{n\theta_\ell}^{n\theta_h} \bar{F}'(\bar{s}) \, d\bar{s}.$$

Now, since $\int_{n\theta_\ell}^{n\theta_h} \bar{G}'_t(\bar{s}) \, d\bar{s}$ is continuous in t , the intermediate value theorem implies that there exists a $\hat{z} \in [\bar{z}, n\theta_h]$ at which

$$\int_{n\theta_\ell}^{n\theta_h} \bar{G}'_{\hat{z}}(\bar{s}) \, d\bar{s} = \int_{n\theta_\ell}^{n\theta_h} \bar{F}'(\bar{s}) \, d\bar{s}.$$

We end the proof of this step by arguing that \overline{G}'_z is a mean-preserving contraction of \overline{F}' . To see this, first observe that for $z \in [n\theta_\ell, \hat{z}]$, we have

$$\int_{n\theta_\ell}^z \overline{G}'_z(\bar{s}) \, d\bar{s} = \int_{n\theta_\ell}^z \overline{G}(\bar{s}) \, d\bar{s} \leq \int_{n\theta_\ell}^z \overline{F}(\bar{s}) \, d\bar{s} \leq \int_{n\theta_\ell}^z \overline{F}'(\bar{s}) \, d\bar{s}.$$

Since, by construction, the means of \overline{G}'_z and F' are equal and $1 = \overline{G}'_z(\bar{s}) \geq F'(\bar{s})$ for $\bar{s} \in [\hat{z}, n\theta_h]$, the above inequality also implies that

$$\int_{n\theta_\ell}^z \overline{G}'_z(\bar{s}) \, d\bar{s} \leq \int_{n\theta_\ell}^z \overline{F}'(\bar{s}) \, d\bar{s} \quad \text{for all } z \in (\hat{z}, n\theta_h].$$

Thus $\overline{G}'_z \in \overline{\mathcal{G}}'$ is the distribution of posterior estimates of grand bundle values required to complete the proof of this step.

Step 2: For any random pure bundling mechanism $\mathcal{M} = (q, t) \in \mathcal{M}^{rPB}$, we have

$$\inf_{G \in \mathcal{G}} \Pi(G, \mathcal{M}) \geq \inf_{G' \in \mathcal{G}'} \Pi(G', \mathcal{M}).$$

We can always write a random pure mechanism as a one-dimensional mechanism. Simply take $\bar{q} : \overline{S} \rightarrow [0, 1]$, $\bar{t} : \overline{S} \rightarrow \mathbb{R}$ to be

$$\bar{q}(\bar{s}) = q(s, N) \text{ and } \bar{t}(\bar{s}) = t(s) \quad \text{for any } s \in S, \sum_{i \in N} s_i = \bar{s}$$

and recall this is well defined because the allocation and transfer of a random pure bundling mechanism only depends on the grand bundle value \bar{s} . Clearly, (\bar{q}, \bar{t}) is incentive compatible for the message space \overline{S} because \mathcal{M}^{rPB} is incentive compatible.

Take any signal $G \in \mathcal{G}$ and note that, by construction,

$$\int_{\overline{S}} \bar{t}(\bar{s}) \, d\overline{G}(\bar{s}) = \int_S t(s) \, dG(s).$$

From [Lemma 1](#) and Step 1, there exists a signal $G' \in \mathcal{G}'$ such that the corresponding distribution of posterior estimates of grand bundle values \overline{G}' is first order stochastically dominated by \overline{G} . Hence,

$$\int_S t(s) \, dG(s) = \int_{\overline{S}} \bar{t}(\bar{s}) \, d\overline{G}(\bar{s}) \geq \int_{\overline{S}} \bar{t}(\bar{s}) \, d\overline{G}'(\bar{s}) = \int_S t(s) \, dG'(s)$$

where the inequality follows from Proposition 2 in [Hart and Reny \(2015\)](#). (They show that for the sale of a single good, any incentive compatible mechanism yields a higher profit from a type distribution that first-order stochastically dominates another.) This completes the proof of this step and the theorem since for every signal $G \in \mathcal{G}$ we can find another signal $G' \in \mathcal{G}'$ such that the random pure bundling mechanism \mathcal{M} yields a lower profit on the latter. \blacksquare

APPENDIX C. PROOFS OF RESULTS IN SECTION 3.2

PROOF OF THEOREM 3 AND COROLLARY 2. We first show that

$$\inf_{G \in \mathcal{G}} \sup_{\mathcal{M} \in \mathcal{M}} \{\Pi(G, \mathcal{M})\} \geq \inf_{G \in \mathcal{G}} \sup_{\mathcal{M} \in \mathcal{M}^{PB}} \{\Pi(G, \mathcal{M})\} = \inf_{G \in \mathcal{G}^{pc}} \sup_{\mathcal{M} \in \mathcal{M}^{PB}} \{\Pi(G, \mathcal{M})\}. \quad (10)$$

Here, the inequality simply follows from the fact that the right side takes the supremum over the smaller set of pure bundling mechanisms. The equality is a consequence of [Lemma 1](#). When the

seller offer a pure bundling mechanism $\mathcal{M} \in \mathcal{M}^{PB}$, for any signal $G \in \mathcal{G}$, the profit only depends on the distribution \bar{G} that the signal G induces grand bundle values. But as [Lemma 1](#) shows, there is a perfectly correlated signal $G' \in \mathcal{G}^{pc}$ that induces the same distribution \bar{G} on \bar{s} . Therefore, taking the supremum over \mathcal{G} or \mathcal{G}^{pc} leads to the same value.

Next, we show that

$$\inf_{G \in \mathcal{G}} \sup_{\mathcal{M} \in \mathcal{M}} \{\Pi(G, \mathcal{M})\} \leq \inf_{G \in \mathcal{G}^{pc}} \sup_{\mathcal{M} \in \mathcal{M}} \{\Pi(G, \mathcal{M})\} = \inf_{G \in \mathcal{G}^{pc}} \sup_{\mathcal{M} \in \mathcal{M}^{PB}} \{\Pi(G, \mathcal{M})\}. \quad (11)$$

The inequality is a consequence of taking the infimum over the smaller set of signals. Finally, the equality follows from the fact that, for any perfectly correlated signal $G \in \mathcal{G}^{pc}$, it is optimal for the seller to choose a pure bundling mechanism. This is easy to show directly but we do not need to because it is a consequence of [Proposition 1](#) in [Haghpanah and Hartline \(2021\)](#).¹⁶

Taken together, the inequalities (10) and (11) imply that

$$\inf_{G \in \mathcal{G}} \sup_{\mathcal{M} \in \mathcal{M}} \{\Pi(G, \mathcal{M})\} = \inf_{G \in \mathcal{G}^{pc}} \sup_{\mathcal{M} \in \mathcal{M}^{PB}} \{\Pi(G, \mathcal{M})\}.$$

But the problem on the right is equivalent to solving the one-dimensional problem

$$\inf_{\bar{G} \in \bar{\mathcal{G}}} \sup_{\bar{p} \in \bar{\Theta}} \left\{ \bar{p} \int_{\bar{p}}^{n\theta_h} d\bar{G}(\bar{s}) \right\},$$

in which \bar{p} corresponds to a price for the grand bundle. [Roesler and Szentes \(2017\)](#) show that a solution (\bar{G}^*, \bar{p}^*) to this problem is the single good buyer-optimal outcome. Importantly, this solution has the feature that trade occurs with probability one (that is $\bar{G}^*(\bar{p}^*) = 0$).

Therefore, there is a buyer-optimal outcome (G^*, \mathcal{M}^*) in which $G^* \in \mathcal{G}^{pc}$ is a perfect correlated signal that generates the distribution \bar{G}^* over grand bundle estimates (such a signal exists because of [Lemma 1](#)) and \mathcal{M}^* is the pure bundling mechanism at price \bar{p}^* . The reason this outcome is buyer-optimal is that the maximal surplus is realized and the signal G^* generates the lowest possible profit for the seller. This also implies that trade must be efficient in every buyer-optimal outcome. This completes the proof of the theorem and the corollary. \blacksquare

PROOF OF THEOREM 4. Let \check{F}_n denote the distribution of the average value $\frac{\theta_1 + \dots + \theta_n}{n}$ and similarly, let \check{F}_{n-1} denote the distribution of $\frac{\theta_1 + \dots + \theta_{n-1}}{n-1}$ where both are computed with respect to prior F . Let $\check{\mathcal{G}}_n, \check{\mathcal{G}}_{n-1}$ denote the set of signals corresponding to $\check{F}_n, \check{F}_{n-1}$ respectively; these signals provide information to the buyer about the posterior estimate of the average value.

First, consider the case with n goods and observe (from the proof of [Theorem 3](#)) that the consumer surplus from a buyer-optimal outcome for distribution F is identical to the consumer surplus from a buyer-optimal outcome for the sale of a single good whose prior value distribution is \bar{F} . Then note that for every $\bar{G} \in \bar{\mathcal{G}}$, there is a $\check{G}_n \in \check{\mathcal{G}}_n$ (and vice versa) such that $\bar{G}(n\check{s}) = \check{G}_n(\check{s})$ for all $\check{s} \in [\theta_\ell, \theta_h]$ and therefore $\int_{\bar{p}}^{n\theta_h} [\bar{s} - \bar{p}] d\bar{G}(\bar{s}) = n \int_{\bar{p}/n}^{\theta_h} [\check{s} - \bar{p}/n] d\check{G}_n(\check{s})$ for all \bar{p} . This implies that CS_n is the consumer surplus from a buyer-optimal outcome for the sale of a single good whose prior value distribution is \check{F}_n and a similar relation holds when there are $n - 1$ goods.

¹⁶Informally, 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.

Second, the distribution \check{F}_n is a mean-preserving contraction of \check{F}_{n-1} because F is exchangeable. This in turn implies $\check{G}_n \subseteq \check{G}_{n-1}$ since every mean preserving contraction of \check{F}_n is also a mean preserving contraction of \check{F}_{n-1} . Consequently, we must have $CS_{n-1} \geq CS_n$ since the consumer surplus in the latter is maximized over a smaller set of signals. ■

APPENDIX D. PROOF OF THEOREM 6

The proof of [Theorem 6](#) employs the following lemma.

LEMMA 2. *Take any distribution \bar{F} with a positive density on its support $[2\theta_\ell, 2\theta_h]$. Let \bar{F}_x be the distribution that satisfies $\bar{F}_x(\bar{\theta}) = \bar{F}(\bar{\theta} - x)$ for all $\bar{\theta} \in [2\theta_\ell + x, 2\theta_h + x]$. Let π_x^* be the highest profit guarantee (provided by the robustly optimal mechanism) for this distribution \bar{F}_x .*

Then, for every $y > 0$, there exists an $\bar{x} > 0$, such that

$$\pi_x^* \leq 2\theta_\ell + x + y \quad \text{for all } x \geq \bar{x}.$$

PROOF. Take an arbitrary $y > 0$. If there is no $x > 0$ such that $\pi_x^* > 2\theta_\ell + x + y$, we are done. When there is such an x , we have to show that there exists a bound $\bar{x} > 0$ such that $\pi_x^* > 2\theta_\ell + x + y$ implies $x < \bar{x}$. So first, observe that

$$2\theta_\ell + x + y < \pi_x^* \leq \max_{\bar{p} \in [2\theta_\ell + x, 2\theta_h + x]} \bar{p}[1 - \bar{F}_x(\bar{p})] = \max_{\bar{p} \in [2\theta_\ell, 2\theta_h]} [\bar{p} + x][1 - \bar{F}(\bar{p})] \quad (12)$$

where the right side is the monopoly profit the seller would get against the signal that perfectly reveals the buyer's value. Obviously, the highest profit guarantee must be lower than the profit that the seller can get against any fixed signal.

Then, (12) implies that every $\bar{p}^* \in \operatorname{argmax}_{\bar{p} \in [2\theta_\ell, 2\theta_h]} [\bar{p} + x][1 - \bar{F}(\bar{p})]$ must satisfy

$$\bar{p}^* > 2\theta_\ell + y$$

as otherwise even trade with probability one would not yield the necessary profit. Therefore,

$$\pi_x^* \leq \max_{\bar{p} \in [2\theta_\ell, 2\theta_h]} [\bar{p} + x][1 - \bar{F}(\bar{p})] \leq [2\theta_h + x][1 - \bar{F}(2\theta_\ell + y)]$$

because $\bar{p} \leq 2\theta_h$ and trade happens with at most probability $[1 - \bar{F}(2\theta_\ell + y)]$. Now since $\bar{F}(2\theta_\ell + y) > 0$ (as \bar{F} has positive density throughout the support), there exists an $\bar{x} > 0$ such that

$$[2\theta_h + \bar{x}][1 - \bar{F}(2\theta_\ell + y)] \leq 2\theta_\ell + \bar{x}.$$

For all $x' \geq \bar{x}$, we therefore have $\pi_{x'}^* \leq 2\theta_\ell + x'$ which in turn implies $x < \bar{x}$ as required. ■

As mentioned in the body of the paper, we prove a slightly more general result.

THEOREM 6'. *Suppose $n = 2$ and the prior distribution is given by $F = \check{F} \times \check{F}_x$ where \check{F} is supported on $[\theta_\ell, \theta_h]$ and \check{F}_x is an x -shift of \check{F} . Moreover, suppose a seller of a single good facing a buyer whose prior value distribution is \check{F} can guarantee herself a profit $\check{\pi} > \theta_\ell$.*

Then, there is an $\bar{x} > 0$ such that for all $x \geq \bar{x}$, there exists a random separate sales mechanism that provides a profit guarantee that no random pure bundling mechanism can provide.

PROOF. Let \bar{F}_x be the distribution over grand bundle values induced by the joint distribution $\check{F} \times \check{F}_x$. From [Lemma 2](#) (taking $y = (\check{\pi} - \theta_\ell)/2$), there exists an \bar{x} such that, for all $x \geq \bar{x}$, the profit guarantee provided by every random pure bundling mechanism is less than $2\theta_\ell + x + \frac{\check{\pi} - \theta_\ell}{2}$.

But the seller can always guarantee herself as least a payoff of

$$\theta_\ell + x + \check{\pi} = 2\theta_\ell + x + \check{\pi} - \theta_\ell > 2\theta_\ell + x + \frac{\check{\pi} - \theta_\ell}{2}$$

by setting the price of good 2 to $\theta_\ell + x$ (so a deterministic price equal to the lower bound of the support) and randomizing the price of good 1 using the distribution $\mathcal{P} \in \Delta(\mathbb{R})$ that provides a seller of a single good who faces a buyer whose prior value distribution is \check{F} a profit guarantee of $\check{\pi}$. The latter follows from the fact that the marginal distribution G_1 of good 1 corresponding to any signal $G \in \mathcal{G}$ of the joint distribution $\check{F} \times \check{F}_x$ is a mean preserving contraction of \check{F} . This completes the proof. ■

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.
- 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.
- BEN MOSHE, R., S. HART, AND N. NISAN (2022): “Monotonic Mechanisms for Selling Multiple Goods,” *arXiv preprint arXiv:2210.17150*.
- 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 (2021): “On the Structure of Informationally Robust Optimal Auctions,” *Working Paper*.
- CARROLL, G. (2017): “Robustness and Separation in Multidimensional Screening,” *Econometrica*, 85(2), 453–488.
- CHE, Y.-K., AND W. ZHONG (2022): “Robustly-Optimal Mechanism for Selling Multiple Goods,” *arXiv preprint arXiv:2105.02828*.
- DASKALAKIS, C., A. DECKELBAUM, AND C. TZAMOS (2017): “Strong Duality for a Multiple-Good Monopolist,” *Econometrica*, 85(3), 735–767.
- DEB, R., AND A.-K. ROESLER (2021): “Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness,” *CEPR Working Paper*.
- 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 (2021): "When is Pure Bundling Optimal?," *Review of Economic Studies*, 88(3), 1127–1156.
- HAGHPANAH, N., AND R. SIEGEL (2022): "The Limits of Multiproduct Price Discrimination," *American Economic Review: Insights*, 4(4), 443–58.
- HART, S., AND N. NISAN (2019): "Selling Multiple Correlated Goods: Revenue Maximization and Menu-Size Complexity," *Journal of Economic Theory*, 183, 991–1029.
- HART, S., AND P. J. RENY (2015): "Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations," *Theoretical Economics*, 10(3), 893–922.
- KAMENICA, E. (2019): "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 11(1), 249–272.
- MCAFFEE, 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 (2022): "Learning Before Trading: On the Inefficiency of Ignoring Free Information," *Journal of Political Economy*, 130(2), 346–387.
- 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.
- SHAKED, M., AND J. G. SHANTHIKUMAR (2007): *Stochastic Orders*, Springer Series in Statistics. Springer, New York, NY.
- SHI, X. (2012): "Optimal Auctions with Information Acquisition," *Games and Economic Behavior*, 74(2), 666–686.