

Multidimensional Screening with Rich Consumer Data*

Mira Frick Ryota Iijima Yuhta Ishii

First posted version: February 12, 2024

This version: October 2, 2024

Abstract

A multi-product monopolist faces a buyer who is privately informed about his valuations for the goods. As is well-known, optimal mechanisms are in general complicated, while simple mechanisms—such as pure bundling or separate sales—can be far from optimal and do not admit clear-cut comparisons. We show that this changes if the monopolist observes sufficiently rich data about the buyer’s valuations: Now, pure bundling always outperforms separate sales; moreover, there is a sense in which pure bundling performs essentially as well as the optimal mechanism. To formalize this, we characterize how fast the corresponding revenues converge to the first-best revenue as the monopolist’s data grows rich: Pure bundling achieves the same convergence rate to the first-best as optimal mechanisms; in contrast, the convergence rate under separate sales is suboptimal.

*Frick: Princeton University (mfrick@princeton.edu); Iijima: Princeton University (riijima@princeton.edu); Ishii: Pennsylvania State University (yxi5014@psu.edu). Acknowledgments to be added.

1 Introduction

A classic problem in economic theory is how a multi-product monopolist should sell its goods to a buyer whose valuations for the goods are unknown to the monopolist. While of clear economic relevance from both a positive and normative perspective, this problem—and related multi-dimensional screening problems—is well-known to be quite intractable. Even in simple subcases (e.g., two goods for which the buyer has independently distributed, additive values), the seller’s optimal mechanism can be difficult to characterize and look complicated: For instance, it may require the buyer to choose between a continuum of differently priced lotteries over product bundles (Manelli and Vincent, 2007; Daskalakis, Deckelbaum, and Tzamos, 2013).¹ At the same time, the multi-product selling mechanisms used in practice are often quite simple: For example, many firms only present buyers with some limited number of deterministic product bundles, or even engage in *pure bundling*, i.e., offer only the grand bundle of all products at a take-it-or-leave-it price.

In this paper, we provide a novel perspective on the use of such simple multi-product selling mechanisms. We consider a revenue-maximizing monopolistic seller who is endowed with a finite set G of indivisible goods. There is one potential buyer, whose valuations for the goods are summarized by a type vector $\theta \in \mathbb{R}_{++}^{|G|}$ (drawn from an arbitrary prior distribution) and whose payoffs are additive across goods.² The realization of θ is only known to the buyer. However, as in the literature on price discrimination, we assume that the seller observes some information about the buyer’s type that she can use in designing a selling mechanism. For example, the seller might observe various buyer characteristics (or noisy signals thereof) that are correlated with θ within the population from which the buyer is drawn.

The question we ask is: Are there simple mechanisms that perform well at exploiting such information, and, if so, which ones? If the seller’s information is arbitrary, the comparison across different simple mechanisms is not clear-cut: For example, it is well-known that, depending on the seller’s posterior, pure bundling may yield higher or lower revenue than *separate sales* (i.e., setting a separate price for each individual

¹Finding the optimal mechanism can also be computationally intractable (e.g., Daskalakis, Deckelbaum, and Tzamos, 2014), and it can exhibit counterintuitive features, e.g., nonmonotonic revenues with respect to first-order stochastic dominance increases in buyer values (Hart and Reny, 2015).

²Section 6 considers nonadditive buyer utilities, negative valuations, and seller production costs.

good). Moreover, both these simple mechanisms are in general far from optimal.³ However, the key insight of our paper is that there is a sharp answer if the seller’s information is sufficiently *precise*: In this case, we show that the seller is always better off using pure bundling than separate sales; what is more, we formalize a sense in which pure bundling performs essentially as well as the optimal mechanism.

The assumption of precise seller information may, for instance, reflect the proliferation of consumer data to which retailers have access in digital marketplaces.⁴ To tractably formalize such information, our baseline model assumes that the seller observes n independent signals from some distribution P_θ that depends on θ .⁵ Thus, the seller’s information is more precise the richer the amount of data n , and we are interested in the case where n is large.

Crucially, to evaluate and compare the performance of different classes of mechanisms, we do not explicitly derive the seller’s revenues at any particular n . Instead, we take a *convergence rate* approach: We analyze how fast the gap between the seller’s expected revenue and the first-best (i.e., known type) revenue vanishes as n grows large. This provides a parsimonious way of comparing revenues across different classes of mechanisms under rich consumer data: Mechanisms with a faster convergence rate to the first-best yield higher expected revenues than mechanisms with a slower convergence rate at *all* large enough n . As a one-dimensional statistic, convergence rates also succinctly characterize how revenues at large n depend on features of the environment, such as the seller’s signal distribution P_θ . Most importantly, while optimal mechanisms by definition approximate the first-best revenue at the fastest rate, this does not rule out that some simple (but suboptimal) mechanisms may achieve this same optimal convergence rate. If this is the case, this suggests that a seller who observes rich consumer data loses very little from using such simple mechanisms rather than the optimal mechanism.

In particular, our main result (Theorem 1) shows that, regardless of whether the seller optimizes over general selling mechanisms or is restricted to pure bundling, her expected revenue converges to the first-best equally fast as n grows large: In

³See Hart and Nisan (2019) in Section 1.1.

⁴Retailers increasingly use sophisticated tools to draw inferences about consumers’ preferences from data such as their browsing and search history, geolocation, and operating system. Retailers have also started to personalize selling mechanisms based on such data, not only via outright price discrimination but also more subtle channels such as personalized discounts, add-on offerings, or product displays (e.g., Hannak, Soeller, Lazer, Mislove, and Wilson, 2014; Gonzaga, 2018).

⁵Section 6 discusses more general settings.

both cases, the revenue gap relative to the first-best vanishes as fast as $\lambda^G \sqrt{\frac{\ln n}{n}}$. Here, the coefficient λ^G captures how fast the seller’s posterior standard deviation about the value $\sum_{g \in G} \theta_g$ of the grand bundle decays. This coefficient (and hence the corresponding revenues at large n) depends on the signal distribution P_θ only through a common statistical informativeness measure, the Fisher information.

To interpret, note that for any given n , the seller’s revenue under pure bundling is in general lower than under the optimal mechanism, which may screen the buyer via (possibly elaborate) menus of bundles or lotteries over bundles. However, Theorem 1 implies that at large n , any additional revenue gain from using such more complicated mechanisms vs. pure bundling becomes negligible: For example, as we discuss, this revenue gain is smaller than the gain from having access to even an arbitrarily small fraction of additional signals. As such, Theorem 1 suggests a novel “approximate optimality” rationale for using pure bundling.

In contrast, Theorem 1 also shows that if the seller instead optimizes over separate sales mechanisms, then her expected revenue converges to the first-best more slowly: Now, the revenue gap relative to the first-best vanishes as fast as $\sum_{g \in G} \lambda^g \sqrt{\frac{\ln n}{n}}$, where the coefficients λ^g capture how fast the seller’s posterior standard deviation about the values θ_g of each individual good decays. The coefficient $\sum_{g \in G} \lambda^g$ is greater than the aforementioned optimal coefficient λ^G : By subadditivity of standard deviations, the sum of the standard deviations of all θ_g always exceeds the standard deviation of the sum $\sum_{g \in G} \theta_g$. Thus, whereas at small n , the seller may be better off using separate sales or pure bundling depending on parameters, pure bundling always yields a higher expected revenue than separate sales when n is large enough.

A potential concern in assessing the economic relevance of Theorem 1 is whether it requires such unrealistically rich data n that the seller can almost perfectly discriminate between different buyer types: At such n , many mechanisms yield revenues that are very close to the first-best, so the differences between mechanisms may not be very interesting. This concern can be mitigated by analyzing numerical examples. For instance, in a Gaussian environment, Section 3.2 finds that pure bundling starts to approximate the optimal mechanism and to significantly outperform separate sales at n where the corresponding revenues are still only a moderate fraction of the first-best.

The proof of Theorem 1 first exploits the Bernstein-von Mises theorem to reduce the analysis to a setting in which the seller’s posterior is deterministic and Gaussian. The analysis is then based on three main insights that we illustrate in Section 4:

First, focusing on the single-good case, we shed light on how rich data affects the tradeoff between two forms of revenue loss that the seller incurs: Losses on the *intensive margin* (due to having to shade the price relative to the known type case) and on the *extensive margin* (due to the probability that the buyer is unwilling to pay the seller’s chosen price). We show that, as the seller’s posterior grows more and more precise, these two margins become strikingly imbalanced: The seller optimally prices the good in such a way that her revenue losses are driven almost entirely by the intensive margin; in contrast, extensive-margin losses become negligible at large n . This dominance of the intensive margin lies at the heart of why, in the multi-good setting, the convergence to the first-best under pure bundling cannot be improved upon by general deterministic mechanisms. Under the latter, the seller can engage in *mixed bundling*, i.e., screen the buyer by offering a menu of product bundles. Relative to pure bundling, this benefits the seller by extracting revenue from buyers who are unwilling to purchase the grand bundle. However, since this benefit only affects the extensive margin, it has a negligible impact on the revenue gap relative to the first-best at large n .⁶

Second, we show that intensive-margin losses under single-good monopoly vanish as fast as the (scaled) *standard deviation* of the seller’s posterior about the buyer’s type. In the multi-good case, this implies that comparing the revenue gap to the first-best at large n under pure bundling vs. separate sales boils down to comparing the seller’s posterior standard deviation about the value $\sum_{g \in G} \theta_g$ of the grand bundle vs. the sum of her posterior standard deviations about each θ_g . As noted, the former is smaller than the latter, as standard deviation is subadditive. At a high level, this relates to the classic intuition (e.g., Adams and Yellen, 1976; Armstrong, 1999) that bundling reduces the seller’s uncertainty relative to separate sales, as overestimating the valuations of some goods but underestimating those of others can cancel out when estimating the value of the grand bundle. However, for this intuition to be valid in our setting, it is crucial to show that the relevant measure of the seller’s uncertainty is standard deviation, as other measures of uncertainty (e.g., variance) are not subadditive.

Finally, we show that pure bundling achieves the same convergence rate as the

⁶Relative to pure bundling, mixed bundling also in principle allows the seller to reduce the magnitude of intensive-margin losses by appropriately raising the price of the grand bundle. However, we show that this increases extensive-margin losses by an order of magnitude that more than outweighs the intensive-margin gains.

optimal mechanism, which may additionally involve randomization. To get around the aforementioned challenge that optimal mechanisms are difficult to characterize, we introduce a more tractable *relaxed problem* that yields an upper bound on the seller’s optimal revenue. We then prove that even this upper bound converges to the first-best no faster than the pure bundling revenue. The relaxed problem partitions the type space into line segments and imposes incentive compatibility only within each segment. As Section 5 explains, optimal mechanisms in this problem are tractable to characterize and take an “almost deterministic” form. Moreover, we show that by carefully choosing the line segments, we can ensure that the pure bundling revenue in the relaxed problem is the same as in the original problem. Based on this, similar arguments to the comparison of pure vs. mixed bundling yield that pure bundling achieves the optimal convergence rate.

1.1 Related Literature

Our paper relates to the large classical literature on multi-dimensional screening (Adams and Yellen, 1976; McAfee, McMillan, and Whinston, 1989; Armstrong, 1996; Rochet and Choné, 1998; Manelli and Vincent, 2006, 2007; Daskalakis, Deckelbaum, and Tzamos, 2017, among many others). Several of these papers develop general methodologies to characterize optimal mechanisms. However, as discussed in the Introduction, optimal mechanisms are complicated unless specific conditions are imposed on type distributions. Some papers derive conditions under which pure bundling is optimal (e.g., Haghpanah and Hartline, 2021; Ghili, 2023).⁷ Our setting allows for general type distributions, and pure bundling is in general suboptimal at any n . Another recent strand of this literature characterizes optimal mechanisms in settings where a seller maximizes her worst-case expected revenue across a *set* of possible type distributions (e.g., Carroll, 2017; Deb and Roesler, 2023; Che and Zhong, 2021). Depending on the structure of this set, optimal mechanisms can take simple forms such as separate sales or pure bundling.

In contrast to the above papers, we depart from the criterion of exact optimality. Instead, we provide a rationale for pure bundling based on the idea that this simple mechanism allows sellers who observe rich data about buyers’ types to approximate the first-best revenue at the optimal rate. This contrasts with a notion of approximate

⁷Other recent work (e.g., Yang, 2023; Bergemann, Bonatti, Haupt, and Smolin, 2021) provides conditions under which more general menus of bundles (e.g., nested bundling) are optimal.

optimality that is often studied in computer science: worst-case guarantees, i.e., lower bounds on the performance ratio of simple vs. optimal mechanisms that are uniform with respect to type distributions (or other features of the environment; see the survey by [Roughgarden and Talgam-Cohen, 2019](#)). Neither pure bundling, separate sales, nor more generally the class of all deterministic mechanisms admit a non-zero worst-case guarantee ([Hart and Nisan, 2019](#)).⁸

While we fix a set of goods G , [Armstrong \(1999\)](#) and [Bakos and Brynjolfsson \(1999\)](#) study the many-good limit. Under independent and additive valuations, they show that pure bundling approximates the first-best as $|G| \rightarrow \infty$, because the value of the grand bundle becomes deterministic by a law of large numbers; in contrast, separate sales does not, because the value of individual goods is random.⁹ In our setting, bundling also reduces uncertainty relative to separate sales, but in a different sense: Regardless of the type distribution, both bundling and separate sales approximate the first-best as the amount of consumer data grows large, because this data allows the seller to learn the valuations for all goods; however, we show that the rate of convergence under bundling is always faster than under separate sales. More importantly, we prove that the convergence rate under bundling is the same as under the optimal mechanism.

As noted, a key step of our proof is to upper-bound the optimal revenue via a relaxed problem that partitions buyer types into line segments. Exploiting the linear structure of this problem, [Proposition 3](#) shows that optimal mechanisms involve at most one random allocation (that is consumed only by types with binding IR). This almost deterministic structure makes optimal mechanisms tractable to analyze and is notable, as even under one-dimensional types, optimal mechanisms with multi-dimensional allocations in general rely on more extensive randomization. As we discuss, this result may be useful in economic settings beyond multi-good monopoly. [Haghpanah and Hartline \(2021\)](#) impose a stochastic monotonicity assumption on type distributions and use this to construct a different relaxed problem: There, types are decomposed into (not necessarily linear) one-dimensional paths and the optimal mechanism is pure bundling.¹⁰

⁸However, positive worst-case guarantees can be derived under more restrictive assumptions on the environment ([Hart and Nisan, 2017](#); [Babaioff, Immorlica, Lucier, and Weinberg, 2020](#)).

⁹[Fang and Norman \(2006\)](#) provide joint non-asymptotic conditions on the number of goods and type distribution under which pure bundling outperforms separate sales and vice versa.

¹⁰For related decomposition approaches, see also [Armstrong \(1996\)](#); [Wilson \(1993\)](#).

More broadly, we relate to a resurgent literature on price discrimination, which studies the implications of a seller’s ability to condition selling mechanisms on additional information about buyers (e.g., Bergemann, Brooks, and Morris, 2015; Haghpanah and Siegel, 2023). Motivated by the availability of rich consumer data in digital marketplaces, some papers focus on sellers with very precise information (e.g., Rhodes and Zhou, 2024). In the context of multi-good monopoly, we highlight that sufficiently precise information can serve as a rationale for using simple selling mechanisms, such as pure bundling. An important step in our analysis is the insight that, under single-good monopoly, such information leads to a stark imbalance between intensive and extensive-margin considerations.¹¹

Finally, convergence rates to a perfect information benchmark have been used as a performance measure in other contexts: In single-agent decision problems, Moscarini and Smith (2002) study how fast an agent who observes many i.i.d. signals about the state and follows an optimal strategy can approximate the perfect information payoff; they use this to define a ranking over signal structures. Subsequent work conducts related exercises in other learning settings.¹² Aside from our different economic environment and research question, a technical departure from these papers is that our setting features continuous actions and types. This necessitates different mathematical techniques from the large-deviation theory tools that apply in the finite state and action settings of these papers. This technical difference also applies relative to our previous work, Frick, Iijima, and Ishii (2024a), which uses convergence rates to evaluate the performance of simple but suboptimal contracts in a moral hazard problem where a principal observes rich monitoring data about an agent’s action. The “hidden action” nature of that problem makes the analysis quite different from the current “hidden type” setting in other respects as well.¹³

¹¹Xie, Zhu, and Shishkin (2024) consider a single-good monopolist who observes samples from an unknown distribution of both consumer data and types. They study how fast the monopolist can approximate the second-best (i.e., known distribution) revenue.

¹²This includes multi-agent learning and misspecified learning (e.g., Harel, Mossel, Strack, and Tamuz, 2021; Frick, Iijima, and Ishii, 2023, 2024b).

¹³More broadly, convergence rates to full efficiency have been used to study the performance of simple mechanisms or strategies in the context of trade in large markets (e.g., Rustichini, Satterthwaite, and Williams, 1994) and repeated games with patient players (Sugaya and Wolitzky, 2024).

2 Setting

A monopolistic seller (“she”) is endowed with a finite set G of indivisible consumption goods. The seller is risk-neutral, attaching zero value to the goods herself and seeking to maximize her expected revenue from selling them. There is one potential buyer (“he”), described by a type vector $\theta \in \mathbb{R}_{++}^{|G|}$ whose g th entry θ_g represents his valuation of good g . The type θ is drawn from a prior density h whose support is some compact set $\Theta \subseteq \mathbb{R}_{++}^{|G|}$ with non-empty interior. The buyer’s utility from receiving the bundle of goods $B \subseteq G$ and paying a monetary transfer t to the seller takes the additively separable form

$$\sum_{g \in B} \theta_g - t = \mathbf{1}^B \cdot \theta - t,$$

where $\mathbf{1}^B \in \mathbb{R}^{|G|}$ denotes the indicator vector on bundle B (i.e., $\mathbf{1}_g^B := \mathbf{1}_{\{g \in B\}}$ for each $g \in G$). Section 6 incorporates nonadditive buyer utilities, negative valuations, and seller production costs into the analysis.

While the realization of θ is only known to the buyer, the seller has access to fairly precise (but imperfect) information about the buyer’s type. As discussed in the Introduction, such information may represent various observable consumer data (e.g., the examples in footnote 4) that is correlated with θ within the population from which the buyer is drawn. For tractability, our main focus is on the following formulation of seller information (Section 6 discusses more general formulations): Let X be a measurable space of signals that is endowed with some σ -finite measure. The seller observes n signals, $x^n = (x_1, \dots, x_n)$, that are drawn i.i.d. conditional on θ from a distribution $P_\theta \in \Delta(X)$ with density $f(\cdot, \theta)$. Thus, the seller’s information is more precise the richer her amount of data n , and we will be interested in settings where n is large.

Upon observing the signals x^n , the seller updates her prior h about the buyer’s type. She then chooses a direct mechanism, which asks the buyer to report his type and, as a function of this report, specifies allocation probabilities $q : \Theta \rightarrow \Delta(2^G)$ for each bundle along with an expected transfer $t : \Theta \rightarrow \mathbb{R}$.¹⁴ Denote by $q(B; \theta)$ the allocation probability of bundle $B \subseteq G$ to type θ .

In the second-best problem, the seller chooses q and t conditional on each signal

¹⁴This formulation rules out asking the buyer to place “bets” (à la Crémer and McLean, 1988) on the realization of x^n . This is in line with the above interpretation of x^n as a collection of personal data that is known to the buyer.

sequence x^n to maximize her expected revenue

$$R_{x^n}^{\text{SB}} := \sup_{q,t} \mathbb{E}[t(\theta) \mid x^n] \quad (1)$$

subject to the incentive compatibility (IC) and individual rationality (IR) constraints

$$\sum_{B \subseteq G} q(B; \theta) (\mathbf{1}^B \cdot \theta) - t(\theta) \geq \sum_{B \subseteq G} q(B; \theta') (\mathbf{1}^B \cdot \theta) - t(\theta'), \quad \forall \theta, \theta' \in \Theta, \quad (\text{IC})$$

$$\sum_{B \subseteq G} q(B; \theta) (\mathbf{1}^B \cdot \theta) - t(\theta) \geq 0, \quad \forall \theta \in \Theta. \quad (\text{IR})$$

The seller's *second-best revenue* is then the ex-ante expectation $R_n^{\text{SB}} := \mathbb{E}[R_{x^n}^{\text{SB}}]$.

As discussed in the Introduction, the solution to the second-best problem (i.e., the optimal mechanism) is in general complicated. Thus, we are interested in understanding how well the seller can perform at large n by using simpler mechanisms. To evaluate the performance of different classes of mechanisms, we do not explicitly derive the corresponding revenues at any given n . Instead, we analyze the convergence rates of these revenues to the first-best as n grows large. Formally, denote by $R^{\text{FB}} := \mathbb{E} \left[\sum_{g \in G} \theta_g \right]$ the *first-best revenue*, i.e., the expected revenue the seller can achieve when she directly observes the buyer's type and thus can fully extract his valuation θ_g for each good g . Note that, as n grows large, the second-best revenue converges to the first-best, i.e., $\lim_{n \rightarrow \infty} R_n^{\text{SB}} = R^{\text{FB}}$, as observing infinitely many signals perfectly reveals the buyer's type to the seller given Assumption 1.5 below. Moreover, by definition, the rate at which R_n^{SB} converges to R^{FB} represents the optimal (i.e., fastest) convergence rate of the seller's revenue to the first-best across all IC-IR mechanisms.

Convergence rates to the first-best also provide a parsimonious way to compare the performance of simple but suboptimal mechanisms: Whenever n is large enough, mechanisms whose revenues converge to the first-best faster yield higher expected revenues than mechanisms with a slower convergence rate. At the same time, achieving the optimal convergence rate to the first-best is a less demanding criterion than maximizing revenue at each n . Thus, there may be simpler classes of mechanisms than the second-best that nevertheless attain the optimal convergence rate.

We will be particularly interested in the following two classes of simple mechanisms: Under *pure bundling* (which we often refer to as *bundling* for short), the

seller posts a single price $p(G)$ for the entire grand bundle G of goods and makes a sale only if the buyer is willing to purchase all goods at this price. Under *separate sales*, the seller posts a separate price $p(g)$ for each good g and sells g to any buyer willing to pay this price.¹⁵ Let $R_{x^n}^{\text{bd}}$ (resp. $R_{x^n}^{\text{sep}}$) denote the seller's expected revenue when in (1) she is restricted to optimizing over bundling (resp. separate sales) mechanisms. Let $R_n^{\text{bd}} := \mathbb{E}[R_{x^n}^{\text{bd}}]$ and $R_n^{\text{sep}} := \mathbb{E}[R_{x^n}^{\text{sep}}]$.

Throughout, we impose the following regularity conditions. For technical convenience, we extend the signal distribution P_θ and signal density $f(\cdot, \theta)$ to all types in some compact neighborhood $\hat{\Theta} \supseteq \Theta$:

Assumption 1.

1. The prior density h is strictly positive and locally Lipschitz continuous for all $\theta \in \Theta$.
2. For all $x \in X$, the signal densities $f(x, \theta)$ are strictly positive and C^2 in $\theta \in \text{int } \hat{\Theta}$, and there is $L > 0$ such that, for all $g, g' \in G$, $\theta, \theta' \in \text{int } \hat{\Theta}$, and $x \in X$, $\left| \frac{\partial^2 \ln f(x, \theta)}{\partial \theta_g \partial \theta_{g'}} \right| \leq L$ and $\left| \frac{\partial^2 \ln f(x, \theta)}{\partial \theta_g \partial \theta_{g'}} - \frac{\partial^2 \ln f(x, \theta')}{\partial \theta_g \partial \theta_{g'}} \right| \leq L \|\theta' - \theta\|$.¹⁶
3. The signal distributions P_θ are continuous in $\theta \in \hat{\Theta}$ with respect to the total variation distance.
4. We have $\sup_{\theta \in \Theta} \int (\sup_{\theta' \in \Theta} \ln f(x, \theta'))^2 dP_\theta(x) < \infty$.
5. The **Fisher information** matrix $I(\theta) \in \mathbb{R}^{|G| \times |G|}$, given by

$$I(\theta) := \left(- \int \frac{\partial^2}{\partial \theta_g \partial \theta_{g'}} \ln f(x, \theta) dP_\theta(x) \right)_{g, g' \in G}, \quad (2)$$

is well-defined and positive definite for each $\theta \in \hat{\Theta}$.

The main content of Assumption 1 is that the log-signal densities $\ln f(x, \theta)$ are sufficiently well-behaved. Most importantly, Assumption 1.5 requires that, conditional on θ , their expected curvatures $\frac{\partial^2}{\partial \theta_g \partial \theta_{g'}} \ln f(x, \theta)$ with respect to changes in θ are well-defined. The Fisher information matrix $I(\theta)$, which summarizes these curvatures, is a

¹⁵Under bundling, $q(G; \theta) = \mathbf{1}_{\{\sum_{g \in G} \theta_g \geq p(G)\}}$, $q(\emptyset; \theta) = \mathbf{1}_{\{\sum_{g \in G} \theta_g < p(G)\}}$, $t(\theta) = \mathbf{1}_{\{\sum_{g \in G} \theta_g \geq p(G)\}} p(G)$ for all θ . Under separate sales, $q(B; \theta) = \mathbf{1}_{\{\theta_g \geq p(g) \forall g \in B \text{ and } \theta_{g'} < p(d') \forall g' \notin B\}}$, $t(\theta) = \sum_{g \in G} \mathbf{1}_{\{\theta_g \geq p(g)\}} p(g)$ for all θ .

¹⁶Here, $\|\cdot\|$ denotes the L^1 -norm on $\mathbb{R}^{|G|}$, i.e., $\|y\| := \sum_{g \in G} |y_g|$.

standard statistical measure of how informative signals are about type θ and will feature prominently in our analysis. Fisher information provides a local approximation of the Kullback-Leibler divergence between the signal distributions at type θ vs. nearby types θ' , as $\text{KL}(P_\theta, P_{\theta'}) := \int \ln \frac{f(x, \theta)}{f(x, \theta')} dP_\theta(x) = (\theta - \theta') \cdot I(\theta)(\theta - \theta') + o(\|\theta - \theta'\|^2)$. Imposing positive definiteness on each $I(\theta)$ implies that $\text{KL}(P_\theta, P_{\theta'}) > 0$ for all $\theta \neq \theta'$, so signal distributions differ across different types. For example, if signals conditional on each θ are distributed Gaussian with mean θ and a fixed positive definite covariance matrix Σ , then $I(\theta) = \Sigma^{-1}$ for all θ , and positive definiteness rules out perfect (positive or negative) correlation across signal coordinates.

3 Main Result

3.1 Optimal Convergence Rate

Our main result, Theorem 1 below, shows that bundling allows the seller to achieve the optimal convergence rate to the first-best, whereas the convergence under separate sales is slower. Naturally, how fast the seller can approximate the perfect information benchmark depends on how informative her signals are about the buyer's type. By explicitly characterizing these convergence rates in terms of the underlying signal distributions, Theorem 1 also sheds light on the relevant measure of informativeness in this setting.

To formalize this, the Fisher information matrix $I(\theta)$ introduced in Assumption 1.5 turns out to play a key role. By the Bernstein-von Mises theorem, this measure is closely tied to the covariance of the seller's posterior at large n : Conditional on (almost) any true type θ^* , the seller's posterior as $n \rightarrow \infty$ is approximated by a normal distribution $\mathcal{N}(\theta^*, \frac{1}{n}I(\theta^*)^{-1})$ whose covariance matrix is the scaled down inverse Fisher information $\frac{1}{n}I(\theta^*)^{-1}$. Consequently, for any bundle B , the standard deviation of the seller's posterior about the buyer's valuation $\mathbf{1}^B \cdot \theta$ of B is approximated by $\frac{\lambda^B(\theta^*)}{\sqrt{n}}$, where

$$\lambda^B(\theta^*) := \sqrt{\mathbf{1}^B \cdot I(\theta^*)^{-1} \mathbf{1}^B}.$$

An important property of standard deviation is subadditivity. In particular, if $|G| \geq$

2, then letting $\lambda^g(\theta^*) := \lambda^{\{g\}}(\theta^*)$ for each $g \in G$, we have

$$\lambda^G(\theta^*) < \sum_{g \in G} \lambda^g(\theta^*), \quad (3)$$

where the inequality is strict by the positive definiteness of $I(\theta^*)$.¹⁷

Let $\lambda^B := \mathbb{E}[\lambda^B(\theta)]$ for each bundle B . For any sequences of real numbers x_n, y_n with $\lim_{n \rightarrow \infty} x_n = \lim_{n \rightarrow \infty} y_n = 0$, write $x_n \sim y_n$ if $\lim_{n \rightarrow \infty} \frac{x_n}{y_n} = 1$, i.e., if x_n and y_n vanish equally fast as n grows large.

Theorem 1. *Under the second-best mechanism and optimal bundling mechanism, the revenue gap to the first-best vanishes equally fast:*

$$R^{\text{FB}} - R_n^{\text{SB}} \sim R^{\text{FB}} - R_n^{\text{bd}} \sim \lambda^G \sqrt{\frac{\ln n}{n}}. \quad (4)$$

Under separate sales, the revenue gap to the first-best vanishes more slowly:

$$R^{\text{FB}} - R_n^{\text{sep}} \sim \sum_{g \in G} \lambda^g \sqrt{\frac{\ln n}{n}}. \quad (5)$$

By (4), when the seller optimizes over all mechanisms, the gap between her revenue R_n^{SB} and the first-best revenue R^{FB} vanishes as fast as $\lambda^G \sqrt{\frac{\ln n}{n}}$ as the amount of data n grows large. Crucially, (4) also shows that the seller can achieve this same optimal convergence rate to the first-best under pure bundling, because $R^{\text{FB}} - R_n^{\text{bd}}$ likewise vanishes as fast as $\lambda^G \sqrt{\frac{\ln n}{n}}$. In both cases, how fast the seller can approximate the first-best depends on the signal distribution only through the coefficient λ^G , which, as noted above, captures how fast the seller's posterior standard deviation about the value $\mathbf{1}^G \cdot \theta$ of the grand bundle vanishes.

To interpret this result, note that at each n , pure bundling is in general suboptimal, as optimal mechanisms may involve menus of bundles and/or lotteries over bundles. However, the result shows that bundling is an effective way for the seller to exploit rich consumer data, providing a rationale for this simple class of mechanisms. As n grows large, any revenue gain from optimizing over general mechanisms vs. pure bundling is second-order: By (4), the only potential difference between the revenue

¹⁷Since $I(\theta^*)$ is symmetric and positive definite, $\|y\|_{\theta^*} := \sqrt{y \cdot I(\theta^*)^{-1} y}$ yields a norm on $\mathbb{R}^{|G|}$ induced by an inner product. Thus, for any non-empty and disjoint $B, B' \subseteq G$, the triangle inequality implies $\lambda^{B \cup B'}(\theta^*) = \|\mathbf{1}^{B \cup B'}\|_{\theta^*} < \|\mathbf{1}^B\|_{\theta^*} + \|\mathbf{1}^{B'}\|_{\theta^*} = \lambda^B(\theta^*) + \lambda^{B'}(\theta^*)$.

gaps $R^{\text{FB}} - R_n^{\text{SB}}$ and $R^{\text{FB}} - R_n^{\text{bd}}$ is a term that is $o\left(\sqrt{\frac{\ln n}{n}}\right)$, but such a term is negligible relative to $\lambda^G \sqrt{\frac{\ln n}{n}}$ at large n .¹⁸ One way to quantify this is by saying that at large n , the benefit of optimizing over general vs. bundling mechanisms is smaller than the benefit of having access to even an arbitrarily small fraction of additional signals: For any $\varepsilon > 0$, no matter how small, we have

$$R_{\lceil(1+\varepsilon)n\rceil}^{\text{bd}} > R_n^{\text{SB}} \text{ for all large enough } n,$$

i.e., the seller is better off under pure bundling with $(1 + \varepsilon)n$ signals than under the optimal mechanism with n signals.¹⁹

In contrast, (5) shows that under separate sales the seller's revenue approximates the first-best more slowly: $R^{\text{FB}} - R_n^{\text{sep}}$ only vanishes as fast as $\sum_{g \in G} \lambda^g \sqrt{\frac{\ln n}{n}}$, and by (3), $\sum_{g \in G} \lambda^g > \lambda^G$ for $|G| \geq 2$. Thus, whereas for small n , the seller's revenue may be higher under separate sales or bundling depending on parameters, Theorem 1 implies that bundling always outperforms separate sales when the seller has access to rich enough consumer data. To understand this performance gap, note that at large n , (4)-(5) imply that to attain the same revenue as under bundling (or under the optimal mechanism) with n signals, a seller who uses separate sales would need to have access to at least $\left(\frac{\sum_{g \in G} \lambda^g}{\lambda^G}\right)^2 n$ signals.²⁰ Depending on parameters, the ratio $\frac{\sum_{g \in G} \lambda^g}{\lambda^G}$ can be arbitrarily large:

Example 1. Suppose $G = \{1, 2\}$ and signals $x_i \in \mathbb{R}^2$ are distributed Gaussian with mean θ and covariance matrix $\Sigma = \begin{pmatrix} \sigma^2 & \rho\sigma^2 \\ \rho\sigma^2 & \sigma^2 \end{pmatrix}$. Then $\frac{\lambda^1(\theta) + \lambda^2(\theta)}{\lambda^G(\theta)} = \frac{2\sigma}{\sigma\sqrt{2(1+\rho)}} > 1$ for all θ . This ratio, and hence the performance gap of bundling vs. separate sales, is decreasing in the correlation ρ of signals across goods and approaches ∞ as $\rho \rightarrow -1$ (while it approaches 1 in the perfect correlation limit $\rho \rightarrow 1$).²¹ ▲

¹⁸For any $f, g : \mathbb{N} \rightarrow \mathbb{R}_{++}$, write $f(n) = o(g(n))$ if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$.

¹⁹This has a similar flavor to Bulow and Klemperer's (1996) seminal (non-asymptotic) result that, in independent private value auctions, an auctioneer's benefit of using optimal vs. second-price auctions is smaller than the benefit of having just one additional bidder.

²⁰Another way to compare these mechanisms is in terms of their performance ratios relative to the second-best. Theorem 1 implies that $\frac{R_n^{\text{bd}}}{R_n^{\text{SB}}} = 1 - o\left(\sqrt{\frac{\ln n}{n}}\right)$ while $\frac{R_n^{\text{sep}}}{R_n^{\text{SB}}} = 1 - \sqrt{\frac{\ln n}{n}} \frac{(\sum_g \lambda^g - \lambda^G)}{R^{\text{FB}}} - o\left(\sqrt{\frac{\ln n}{n}}\right)$, so the ratio approaches 1 faster under bundling than under separate sales.

²¹This is broadly reminiscent of classical findings (e.g., Adams and Yellen, 1976; Schmalensee,

We prove Theorem 1 in Appendix A.2. Section 4 illustrates the main ideas. Notably, the proof implies that under bundling, the seller can achieve the optimal convergence rate using a pricing algorithm that does not require her to know the prior type distribution: At each realized signal sequence x^n , compute the maximum-likelihood estimate $\hat{\theta}_{x^n}$ of the buyer's type. Then price the grand bundle at

$$p_{x^n} = \mathbf{1}^G \cdot \hat{\theta}_{x^n} - \sqrt{\frac{\ln n}{n}} \sqrt{\mathbf{1}^G \cdot I_{x^n}(\hat{\theta}_{x^n})^{-1} \mathbf{1}^G}. \quad (6)$$

Here $I_{x^n}(\cdot)$ denotes the empirical Fisher information, which replaces the expectation with respect to P_θ in (2) with the sample average (see Appendix A.2.1). Thus, under (6), the price depends only on two statistics: The seller's best estimate $\mathbf{1}^G \cdot \hat{\theta}_{x^n}$ of the buyer's valuation of the grand bundle, shaded by (an amount that approximates) the standard deviation of this valuation scaled by $\sqrt{\ln n}$. While (6) need not coincide with the optimal pure bundling price (which does not typically admit a closed-form solution), we show that the induced expected revenue gap $R^{\text{FB}} - \mathbb{E}[\mathbf{1}_{\{p_{x^n} \leq \mathbf{1}^G \cdot \theta\}} p_{x^n}]$ again vanishes at the optimal rate $\lambda^G \sqrt{\frac{\ln n}{n}}$.

Remark 1. Convergence rates of other simple mechanisms. Theorem 1 highlights that the optimal convergence rate is achieved by a particularly simple, albeit stark, class of mechanisms—pure bundling. However, a fortiori, this implies that the optimal convergence rate is also attained by optimizing over other simple and less extreme classes of mechanisms that generalize pure bundling. This includes commonly observed mechanisms such as (i) nested bundling (the seller offers a menu of bundles that are nested by set inclusion) and (ii) two-part tariffs (the seller charges an entry fee along with a usage fee for each good).²² Like pure bundling, these mechanisms involve no randomization and only limited menus of bundles and are in general suboptimal at any fixed n . At the same time, the suboptimal convergence of separate sales extends to mechanisms where the seller chooses a more general partition B_1, \dots, B_k ($1 < k < |G|$) of G and prices $p(B_\ell)$ for each ℓ (e.g., Palfrey, 1983): The convergence rate is $\sum_{\ell=1}^k \lambda^{B_\ell} \sqrt{\frac{\ln n}{n}}$, which is in between that under separate sales and bundling and faster the coarser the partition.

1984) that negatively (resp. positively) correlated valuations across goods favor bundling (resp. separate sales) under some conditions. However, note that in this example ρ parametrizes signal correlations and that bundling outperforms separate sales for *all* ρ when n is large enough.

²²Pure bundling corresponds to the special case of (ii) where the entry fee is the price of G and the usage fee is 0 for all g .

Countervailing forces to pure bundling. In various natural settings, Theorem 1 breaks down as even the first-best may not involve all goods being sold. Section 6 considers two such cases—seller production costs and negative buyer valuations. A generalization of Theorem 1 remains valid: The optimal convergence rate to the first-best is achieved by mechanisms that, similar to pure bundling, offer just a single bundle (albeit generally a strict subset of G) at a take-it-or-leave price. This preserves the message that a seller with rich enough consumer data derives only a negligible benefit from using more complicated menus or lotteries over bundles. ▲

3.2 Gaussian Environment and Illustration of Convergence

As Theorem 1 highlights, an advantage of analyzing convergence rates to the first-best is that this allows for sharp comparisons across different mechanisms that do not depend on all the specifics of the environment (e.g., the details of the prior and signal distributions). On the flipside, convergence rate results are silent about revenues at fixed n , which would depend on the details of the environment. However, this does not mean that the predictions of Theorem 1 only apply once n is so large that revenues are very close to the first-best. This can be illustrated by considering numerical examples.

Specifically, we introduce the following deterministic Gaussian setting that will also play a central role in the proof of Theorem 1. For any type θ^* , denote by $R_{\text{gauss},n}^i(\theta^*)$ ($i \in \{\text{SB}, \text{bd}, \text{sep}\}$) the seller’s optimal expected revenue (under general mechanisms, bundling, and separate sales) when her belief about the buyer’s type is $\mathcal{N}(\theta^*, \frac{1}{n}I(\theta^*)^{-1})$.²³ Thus, whereas in our original model the seller’s posterior at each n is stochastic and depends on the realized signal sequence x^n , here the seller observes no signals and n parametrizes a deterministic sequence of Gaussian beliefs that concentrate on θ^* as n grows large. Nevertheless, by the aforementioned Bernstein-von Mises theorem, this deterministic Gaussian setting approximates the seller’s expected revenues conditional on type θ^* in our original model, where this approximation occurs faster than the convergence in Theorem 1:²⁴

²³That is, $R_{\text{gauss},n}^{\text{SB}}(\theta^*) = \sup_{q,t} \mathbb{E}_{\mathcal{N}(\theta^*, \frac{1}{n}I(\theta^*)^{-1})}[t(\theta)]$ subject to (IC)–(IR), and $R_{\text{gauss},n}^{\text{bd}}(\theta^*)$ ($R_{\text{gauss},n}^{\text{sep}}(\theta^*)$) is the seller’s value when the sup is over bundling (separate sales) mechanisms.

²⁴Lemma 1 is immediate from Propositions A.2 and A.3 in Appendix A.

Lemma 1. *For almost all $\theta^* \in \Theta$ and all $i \in \{\text{SB}, \text{bd}, \text{sep}\}$,*

$$\mathbb{E}[R_{x^n}^i | \theta^*] - R_{\text{gauss},n}^i(\theta^*) = o\left(\sqrt{\frac{\ln n}{n}}\right).$$

The revenues in this Gaussian setting can be studied numerically as a function of n . Figure 1 plots the revenues under pure bundling, separate sales, and mixed bundling (i.e., the optimal deterministic mechanism, see Section 4.4); it also plots an upper bound on the second-best revenue that we will derive in Section 4.5.

Note that in both panels of Figure 1, the predictions of Theorem 1 start to apply at values of n that are moderate, in the sense that all revenues are still far from the first-best revenue of 0.6.²⁵ First, pure bundling starts to outperform separate sales very quickly, and the revenue gap between the two remains significant throughout the domain of the figure. Second, the revenue gap of pure bundling relative to both mixed bundling and the second-best bound all but disappears at a point where these revenues are still only a moderate fraction (less than 70%) of the first-best revenue. Finally, reflecting the impact of correlation on the relative performance of separate sales vs. bundling highlighted in Example 1, the revenue gap between separate sales and bundling is more pronounced in the left-hand panel of Figure 1 where $\rho = -0.5$ than in the middle panel ($\rho = 0$) and right-hand panel ($\rho = 0.5$).

4 Main Ideas behind Theorem 1

This section illustrates the main ideas behind Theorem 1. Section 4.1 reduces the analysis to the Gaussian setting from Section 3.2. Section 4.2 analyzes single-good monopoly, establishing a key result about revenue losses on the intensive vs. extensive margins. Based on this, we derive the convergence rates under bundling vs. separate sales (Section 4.3), and show that offering menus of bundles does not improve the convergence rate relative to pure bundling (Section 4.4). Finally, Section 4.5 introduces a relaxed problem that yields an upper bound on the second-best revenue. Using this, we show that bundling achieves the same convergence rate as the optimal mechanism.

²⁵The slow convergence in Figure 1 reflects the slow convergence speed of $\sqrt{\frac{\ln n}{n}}$ in Theorem 1. Note that (unlike in our original model) revenues in this deterministic setting are initially decreasing in n , as very low n correspond to very high posterior variance, which benefits the seller by allowing her to charge very high prices and focus on extracting surplus from right-tail types.




Figure 1: Gaussian setting with $\theta^* = \begin{pmatrix} 0.3 \\ 0.3 \end{pmatrix}$, $I(\theta^*)^{-1} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$, and $\rho = -0.5$ (left), $\rho = 0$ (middle), or $\rho = 0.5$ (right). We plot revenues $R_{\text{gauss},n}^i(\theta^*)$ under pure bundling (orange), mixed bundling (purple), separate sales (blue), and the upper bound $\bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*)$ on the second-best revenue (yellow) as a function of n . All revenues converge to the first-best $R^{\text{FB}}(\theta^*) = 0.6$ (gray) as $n \rightarrow \infty$.

4.1 Reduction to Gaussian Setting

To prove Theorem 1, it suffices to establish the following analog of this result in the deterministic Gaussian setting we introduced in Section 3.2. Let $R^{\text{FB}}(\theta^*) := \sum_{g \in G} \theta_g^*$.

Proposition 1. *For all $\theta^* \in \Theta$, we have*

$$R^{\text{FB}}(\theta^*) - R_{\text{gauss},n}^{\text{SB}}(\theta^*) \sim R^{\text{FB}}(\theta^*) - R_{\text{gauss},n}^{\text{bd}}(\theta^*) \sim \lambda^G(\theta^*) \sqrt{\frac{\ln n}{n}},$$

$$\text{and } R^{\text{FB}}(\theta^*) - R_{\text{gauss},n}^{\text{sep}}(\theta^*) \sim \sum_{g \in G} \lambda^g(\theta^*) \sqrt{\frac{\ln n}{n}}.$$

That Proposition 1 implies Theorem 1 follows from the Gaussian approximation in Lemma 1, along with additional arguments that allow us to go from the conditional revenues in Proposition 1 to the ex-ante expected revenues in Theorem 1.²⁶ Henceforth, we focus on the Gaussian setting and illustrate the ideas behind Proposition 1.

4.2 Single-Good Monopoly: Intensive vs. Extensive Margin

First, we analyze the convergence rate to the first-best under single-good monopoly. Suppose $|G| = 1$ and the seller's belief is distributed $\mathcal{N}(\theta^*, \sigma^2/n)$ for some $\theta^* > 0$; the corresponding cdf is $F_n(x) := \Phi\left(\frac{(x-\theta^*)\sqrt{n}}{\sigma}\right)$, where Φ is the standard normal cdf. Under single-good monopoly, the optimal mechanism is a posted price: The seller chooses a price p_n and sells the good to buyers of type $\theta \geq p_n$. Let R_n^* denote the

²⁶Such arguments are needed because the approximation errors in Lemma 1 are not uniform in θ .

optimal revenue in this setting, which is given by

$$R_n^* = \max_p p(1 - F_n(p)).$$

Let p_n^* denote an optimal price.

To understand how fast the revenue gap to the first-best (i.e., $R^{\text{FB}}(\theta^*) - R_n^* = \theta^* - R_n^*$) vanishes as n grows large, it is helpful to decompose this gap as follows:

$$\theta^* - R_n^* = \theta^* - p_n^*(1 - F_n(p_n^*)) = \underbrace{\theta^* - p_n^*}_{\text{intensive margin}} + \underbrace{\theta^* F_n(p_n^*)}_{\text{extensive margin}} - \underbrace{(\theta^* - p_n^*) F_n(p_n^*)}_{\text{smaller order terms}}.$$

That is, at each n , the optimal price p_n^* must trade-off revenue losses on the *intensive margin* (i.e., due to price discounts relative to the first-best) and the *extensive margin* (i.e., due to some buyer types refusing to buy). To reduce intensive-margin losses, the seller wants to choose a price sequence p_n that approaches θ^* as fast as possible. But if the rate at which $p_n \rightarrow \theta^*$ is too fast, then extensive-margin losses $F_n(p_n) = \Phi\left(\frac{(p_n - \theta^*)\sqrt{n}}{\sigma}\right)$ may vanish too slowly or fail to vanish at all.

The key observation is that this tradeoff dictates that p_n^* is set in such a way that revenue losses at large n are driven almost entirely by the intensive margin:

Proposition 2. *In the Gaussian environment with $|G| = 1$ and θ distributed $\mathcal{N}(\theta^*, \sigma^2/n)$ for some $\theta^* > 0$, any optimal price sequence p_n^* satisfies*

$$R^{\text{FB}}(\theta^*) - R_n^* \sim \theta^* - p_n^* \sim \sqrt{\ln n} \frac{\sigma}{\sqrt{n}} \quad \text{and} \quad \theta^* F_n(p_n^*) = o\left(\sqrt{\frac{\ln n}{n}}\right). \quad (7)$$

Moreover, for any price sequence p_n with $\theta^* - p_n \sim \delta(\theta^* - p_n^*)$ for some $\delta \in [0, 1)$, we have $\lim_n \frac{\theta^* F_n(p_n)}{\sqrt{\frac{\ln n}{n}}} = \infty$.

Proposition 2 shows that losses on the extensive margin $\theta^* F_n(p_n^*)$ become negligible relative to the intensive margin $\theta^* - p_n^*$ at large n . Hence, the revenue gap relative to the first-best vanishes at the same rate as the intensive margin. Moreover, this rate takes a simple form, given by the standard deviation $\frac{\sigma}{\sqrt{n}}$ of the seller's posterior scaled by $\sqrt{\ln n}$.²⁷ While there is no closed-form solution for the optimal prices p_n^* , the result

²⁷To understand this rate, note that if the intensive margin is proportional to the standard deviation ($\theta^* - p_n = \kappa \frac{\sigma}{\sqrt{n}}$ for some $\kappa > 0$), extensive-margin losses $F_n(p_n) = \Phi\left(\frac{(p_n - \theta^*)\sqrt{n}}{\sigma}\right) = \Phi(-\kappa)$ do

implies that the same optimal convergence rate is achieved by a simple pricing rule,

$$p_n = \theta^* - \sqrt{\ln n} \frac{\sigma}{\sqrt{n}}, \quad (8)$$

which shades the price by the amount $\sqrt{\ln n} \frac{\sigma}{\sqrt{n}}$ relative to the posterior mean θ^* .

It may appear counterintuitive that the intensive and extensive margins are highly imbalanced at the optimum. Why can't the seller reduce the intensive margin a little by correspondingly increasing the extensive margin? The "moreover" part of Proposition 2 sheds light on this point: If the intensive margin is scaled down relative to the optimum $\theta^* - p_n^*$ by some factor $\delta < 1$, then the extensive margin explodes relative to the magnitude $\sqrt{\frac{\ln n}{n}}$ of the intensive margin, which is clearly suboptimal. Notably, this occurs even when δ is arbitrarily close to 1. This reflects the relatively thin tails of the Gaussian distribution, which make the tail-event probability $F_n(p_n)$ very sensitive to the choice of p_n (Appendix C.3 extends this to more general distributions).

4.3 Bundling vs. Separate Sales: Role of Standard Deviation

We return to the multi-good setting, $|G| \geq 2$, with seller belief $\mathcal{N}(\theta^*, \frac{1}{n}I(\theta^*)^{-1})$. Proposition 2 immediately yields the convergence rates of $R_{\text{gauss},n}^{\text{bd}}(\theta^*)$ and $R_{\text{gauss},n}^{\text{sep}}(\theta^*)$, as under both bundling and separate sales, the seller's problem reduces to single-good monopoly.

Specifically, bundling is an instance of single-good monopoly, where the good is the grand bundle and the buyer's type is $\sum_{g \in G} \theta_g$. Since the standard deviation of $\sum_{g \in G} \theta_g$ is $\sqrt{\frac{\mathbf{1}^G \cdot I(\theta^*)^{-1} \mathbf{1}^G}{n}} = \frac{\lambda^G(\theta^*)}{\sqrt{n}}$, (7) then yields that the revenue gap relative to the first-best vanishes as fast as $\sqrt{\ln n} \frac{\lambda^G(\theta^*)}{\sqrt{n}}$. In contrast, under separate sales, the seller solves a collection of single-good monopoly problems, one for each good g , with corresponding type θ_g . Since the standard deviation of θ_g is $\sqrt{\frac{\mathbf{1}^g \cdot I(\theta^*)^{-1} \mathbf{1}^g}{n}} = \frac{\lambda^g(\theta^*)}{\sqrt{n}}$, (7) yields that the revenue gap relative to the first-best vanishes as fast as $\sqrt{\ln n} \frac{\sum_{g \in G} \lambda^g(\theta^*)}{\sqrt{n}}$.

Thus, the reason that convergence under bundling is faster than under separate sales boils down to the fact that the standard deviation of the sum of all valuations $\sum_{g \in G} \theta_g$ is less than the sum of the standard deviations of each individual valuation

not vanish with n . For the extensive margin to vanish, $\theta^* - p_n$ must be scaled up, where the role of the scaling factor $\sqrt{\ln n}$ comes from the exponential form of the standard normal distribution.

θ_g . By Proposition 2, this allows the seller to charge a bundling price $p_n^{\text{bd}}(G)$ that at large n exceeds the total price $\sum_{g \in G} p_n^{\text{sep}}(g)$ under separate sales. By the dominance of the intensive margin, this ensures a higher revenue under bundling.

At a high level, this reflects the intuition that bundling reduces the seller’s uncertainty about the buyer’s type relative to separate sales, as overestimating the valuations of some goods but underestimating those of others can cancel out when estimating the valuation of the grand bundle. However, for this intuition to be valid in our setting, it is crucial that Proposition 2 shows that the relevant measure of uncertainty at large n is *standard deviation*, which is subadditive. The same logic would not have been valid under other natural measures of uncertainty: For example, the variance of $\sum_{g \in G} \theta_g$ can be greater or smaller than the sum of the variances of all θ_g , depending on the correlation across dimensions.²⁸

4.4 Bundling vs. Mixed Bundling

It remains to show that general mechanisms, which may involve menus of bundles and randomization, cannot achieve a faster convergence rate than pure bundling. As a first step, we focus on the effect of menus and ignore randomization. IC-IR deterministic mechanisms correspond to ***mixed bundling***: The seller sets prices $p(B)$ for all bundles B , and the buyer chooses which bundle to purchase.²⁹ Let $(p_n^{\text{mix}}(B))_{B \subseteq G}$ denote the optimal mixed bundling prices under the seller belief $\mathcal{N}(\theta^*, \frac{1}{n}I(\theta^*)^{-1})$.

Relative to pure bundling, mixed bundling allows the seller to extract revenue from buyers who are unwilling to purchase the grand bundle G , i.e., those types with $\sum_{g \in G} \theta_g < p_n^{\text{bd}}(G)$. However, while this improves the seller’s revenue at each n , we show that the benefit in terms of the convergence rate to the first-best is negligible.

The logic relies on the intensive vs. extensive margin analysis in Proposition 2. Indeed, by (7), the extensive-margin losses under pure bundling from types $\sum_{g \in G} \theta_g < p_n^{\text{bd}}(G)$ become negligible relative to the intensive-margin losses $\sum_{g \in G} \theta_g^* - p_n^{\text{bd}}(G)$ at large n . Thus, the only way for mixed bundling to achieve a first-order improvement

²⁸To illustrate a setting where variance would have been the appropriate measure, consider a generalization of the monopolist’s problem in Section 4.2, where θ is distributed $\mathcal{N}(\theta^*, \frac{\sigma^2}{n})$ and the price choice p induces a payoff $u(p, \theta)$. If $u(p, \theta)$ is smooth, Taylor approximation arguments imply that the gap to the first-best payoff is proportional to the variance $\frac{\sigma^2}{n}$. However, in our setting, $u(p, \theta) = p\mathbf{1}_{\{p \leq \theta\}}$ is discontinuous at $p = \theta$. This makes it more difficult to approximate the first-best payoff, leading to the slower convergence rate $\sqrt{\ln n} \frac{\sigma}{\sqrt{n}}$.

²⁹We implicitly assume that a buyer cannot purchase multiple bundles at once.




Figure 2: *Illustration of segmentation when $|G| = 2$.*

over pure bundling would be if mixed bundling allowed the seller to raise the price of the grand bundle in a significant enough way that the intensive-margin losses $\sum_{g \in G} \theta_g^* - p_n^{\text{mix}}(G)$ vanish strictly faster than $\sum_{g \in G} \theta_g^* - p_n^{\text{bd}}(G)$. However, this is impossible by the “moreover” part of Proposition 2: Under any such price sequence $p_n^{\text{mix}}(G)$, extensive-margin losses (i.e., the rejection probability of the grand bundle) would increase so substantially relative to pure bundling that this would more than offset the improvements on the intensive margin.

4.5 Bounding the Optimal Revenue via a Relaxed Problem

Finally, we show that bundling achieves the same convergence rate as the second-best mechanism. Relative to mixed bundling, the second-best mechanism may involve randomization and, as noted, can be difficult to characterize. To get around this, we instead derive an upper bound on the second-best revenue $R_{\text{gauss},n}^{\text{SB}}(\theta^*)$ by considering the seller’s optimal revenue $\bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*)$ in a more tractable relaxed problem. We then show that even $R^{\text{FB}}(\theta^*) - \bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*)$ does not vanish at a faster rate than $R^{\text{FB}}(\theta^*) - R_{\text{gauss},n}^{\text{bd}}(\theta^*)$.

To define the relaxed problem, we split the type space into one-dimensional segments: Fix a full-rank matrix $A \in \mathbb{R}^{(|G|-1) \times |G|}$ and partition types into lines $(\Theta_y)_{y \in \mathbb{R}^{|G|-1}}$, where

$$\Theta_y := \{\theta \in \mathbb{R}^{|G|} : A\theta = y\} \quad \text{for each } y \in \mathbb{R}^{|G|-1}.$$

See Figure 2 for an illustration when $|G| = 2$, in which case A corresponds to a vector in \mathbb{R}^2 and each line segment is parametrized by a real number y . Given this segmentation, denote by $\bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*)$ the second-best revenue when the IC constraint governing $R_{\text{gauss},n}^{\text{SB}}(\theta^*)$ is relaxed to only hold across types in the same line segment, i.e., we only require that

$$\sum_{B \subseteq G} q(B; \theta) (\mathbf{1}^B \cdot \theta) - t(\theta) \geq \sum_{B \subseteq G} q(B; \theta') (\mathbf{1}^B \cdot \theta) - t(\theta'), \quad \forall y \in \mathbb{R}^{|G|-1}, \theta, \theta' \in \Theta_y.$$

Thus, it is as if the seller observes each type θ 's segment (i.e., the vector $A\theta$) before choosing a mechanism. Clearly, $\bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*) \geq R_{\text{gauss},n}^{\text{SB}}(\theta^*)$.

An advantage of the relaxed problem is that solving for $\bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*)$ reduces to a collection of problems with one-dimensional type spaces. That is, $\bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*)$ can be written as the average

$$\bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*) = \mathbb{E} [R_{\text{gauss},n}^{\text{SB}}(\theta^* | y)],$$

where $R_{\text{gauss},n}^{\text{SB}}(\theta^* | y)$ denotes the second-best revenue when the seller's belief over types is the conditional distribution on segment Θ_y induced by $\mathcal{N}(\theta^*, \frac{1}{n}I(\theta^*)^{-1})$, and the expectation on the right is with respect to y being distributed $\mathcal{N}(A\theta^*, \frac{1}{n}AI(\theta^*)^{-1}A^\top)$.

Since each Θ_y is one-dimensional, analyzing $R_{\text{gauss},n}^{\text{SB}}(\theta^* | y)$ is quite tractable. Section 5 below characterizes optimal mechanisms in a more general environment with one-dimensional types (dropping the Gaussian assumption). The key observation is that optimal mechanisms at each Θ_y are “almost deterministic:” In addition to deterministic bundles, they offer (at most) a *single* lottery over bundles. Thus, ruling out more complicated mechanisms with infinitely many lotteries, the optimal mechanism at each Θ_y is a simple modification of mixed bundling. Based on this, a similar logic to the comparison between mixed and pure bundling in Section 4.4 implies that, at each Θ_y , the seller's second-best revenue $R_{\text{gauss},n}^{\text{SB}}(\theta^* | y)$ does not converge to the first-best faster than does her conditional pure bundling revenue $R_{\text{gauss},n}^{\text{bd}}(\theta^* | y)$. That is, the seller's revenue gain from randomization becomes negligible at large n .

A crucial remaining issue is that the pure bundling price at each Θ_y may in general depend on y , so the bundling revenue $\mathbb{E} [R_{\text{gauss},n}^{\text{bd}}(\theta^* | y)]$ in the relaxed problem may be greater than the revenue $R_{\text{gauss},n}^{\text{bd}}(\theta^*)$ in the original problem. However, this issue can be avoided by carefully choosing how to segment the type space in the relaxed

problem. Specifically, we choose the matrix A in such a way that the conditional distribution of $\sum_{g \in G} \theta_g$ at every segment Θ_y is the same as the unconditional distribution of $\sum_{g \in G} \theta_g$. That is, observing the buyer's line segment (i.e., the vector $A\theta$) provides the seller with no additional information about the buyer's valuation of the grand bundle. Picking an A with this property is always possible, since in the current Gaussian setting this is equivalent to the requirement that $A\theta$ and $\mathbf{1}^G \cdot \theta$ are uncorrelated.³⁰ Under this choice of A , the optimal pure bundling price at each Θ_y is the same as the optimal pure bundling price in the original problem. Hence, the seller's pure bundling revenue in the original problem can be written as

$$R_{\text{gauss},n}^{\text{bd}}(\theta^*) = \mathbb{E} [R_{\text{gauss},n}^{\text{bd}}(\theta^* | y)].$$

Combined with the previous paragraph, this yields the desired conclusion that $R^{\text{FB}}(\theta^*) - \bar{R}_{\text{gauss},n}^{\text{SB}}(\theta^*)$ does not vanish at a faster rate than $R^{\text{FB}}(\theta^*) - R_{\text{gauss},n}^{\text{bd}}(\theta^*)$.

5 Analysis of One-Dimensional Type Setting

This section analyzes (a generalization of) the one-dimensional type environment at the center of the argument in Section 4.5. As claimed in Section 4.5, we show that optimal mechanisms take an almost deterministic form. This result may be of independent interest beyond its application in Section 4.5.

Let \mathcal{B} denote a finite set of allocations, which we enumerate as $\{1, 2, \dots, m\}$. An agent's type is described by an L^1 random variable τ that admits a strictly positive density on some interval $T \subseteq \mathbb{R}$. We normalize types so that $0 \in T$. Each type $\tau \in T$ is associated with a vector $\alpha(\tau) \in \mathbb{R}^m$, where $\alpha_\ell(\tau)$ represents type τ 's utility from allocation ℓ . Our key assumption is that α is linear, i.e., there is $\beta \in \mathbb{R}^m$ such that

$$\alpha(\tau) = \alpha(0) + \tau\beta.$$

Let $\bar{\beta} := \max\{\beta_1, \dots, \beta_m\}$, $\underline{\beta} := \min\{\beta_1, \dots, \beta_m\}$, and assume there exists some allocation $\bar{\ell}$ for which $\alpha_{\bar{\ell}}(0) = \beta_{\bar{\ell}} = 0$. Section 4.5 is a special case of this environment where \mathcal{B} corresponds to the set of bundles (with $\bar{\ell}$ the empty bundle), but the current formulation does not require Gaussian type distributions or additive valuations.³¹

³⁰For example, if $|G| = 2$ and $\text{Var}(\theta_1) = \text{Var}(\theta_2)$, then we can pick A to be the vector $(1, -1)$.

³¹More specifically, the current setting nests Section 4.5 by letting $T = \mathbb{R}$ and choosing α with

The current setting also encompasses some other economic applications.³²

A direct mechanism consists of a pair of measurable functions $q : T \rightarrow \Delta(\mathcal{B})$ and $t : T \rightarrow \mathbb{R}$. Let \mathcal{M} denote the set of all direct mechanisms (q, t) that satisfy both IC (i.e., $\alpha(\tau) \cdot q(\tau) - t(\tau) \geq \alpha(\tau) \cdot q(\tau') - t(\tau')$ for all $\tau, \tau' \in T$) and IR (i.e., $\alpha(\tau) \cdot q(\tau) - t(\tau) \geq 0$ for all $\tau \in T$). The designer's objective is to maximize the expected transfer $\mathbb{E}[t(\tau)]$ subject to IC and IR.

Since $\beta_{\bar{\ell}} = 0 \in [\underline{\beta}, \bar{\beta}]$, we can pick a lottery

$$b_0 \in \operatorname{argmax}_{q \in \Delta(\mathcal{B})} \alpha(0) \cdot q \text{ such that } \beta \cdot q = 0;$$

that is, b_0 yields maximal utility among all lotteries that give the same utility to all types. Let $\mathcal{D} := \{\delta_k : k \in \mathcal{B}\}$ denote the set of deterministic allocations. Let

$$\mathcal{M}^s := \{(q, t) \in \mathcal{M} : q(T) \subseteq \mathcal{D} \cup \{b_0\}\}$$

denote the subset of IC and IR mechanisms that involve only allocations that are either deterministic or b_0 . The following result shows that it is without loss of optimality for the designer to restrict attention to such mechanisms. Note that the choice of the lottery b_0 is independent of the type distribution.

Proposition 3. *We have $\sup_{(q,t) \in \mathcal{M}} \mathbb{E}[t(\tau)] = \sup_{(q,t) \in \mathcal{M}^s} \mathbb{E}[t(\tau)]$.*

We prove Proposition 3 in Appendix B. As in Manelli and Vincent (2007) and Daskalakis, Deckelbaum, and Tzamos (2017), we reformulate the designer's problem as an optimization over agent value functions $V : T \rightarrow \mathbb{R}$ that are induced by IC-IR mechanisms. However, in contrast to these papers, a challenge is that V does not uniquely pin down the corresponding mechanism (and hence the designer's payoff), because the allocation space has a higher dimension than the type space. Thus, we solve the problem in two steps: First, for each V , we optimize over the set of IC-IR mechanisms that induce V ; second, we optimize over V . We show that, while the objective in the second step is non-linear in V , it has a piecewise linear structure.

$\alpha_{\ell}(0) = \mathbb{E}[\theta \cdot \mathbf{1}^{B_{\ell}} | A\theta = y]$ and $\beta_{\ell} = \mathbf{1}^{B_{\ell}} \cdot z$ for each bundle B_{ℓ} , where $z \in \mathbb{R}^G$ is such that $Az = \mathbf{0}$.

³²For example, the linearity assumption is satisfied in settings where allocations correspond to multi-dimensional consumptions (e.g., state-contingent payoffs) and types represent weights over dimensions (e.g., beliefs). Also, Boleslavsky and Kolb (2024) consider a setting where an agent purchases allocations ("influence bundles"), whose payoff depends linearly on his one-dimensional type.

Based on this, we verify that it suffices for optimality to consider V functions that are induced by mechanisms in \mathcal{M}^s .

Relative to deterministic mechanisms, the role of offering lottery b_0 is to extract more surplus from types for whom IR binds. To illustrate, consider the case where $\alpha(0) \cdot b_0 > 0$ and b_0 is a non-degenerate lottery. Suppose that under an optimal deterministic mechanism, there is a non-trivial interval $\hat{T} \subseteq T$ of types who receive the degenerate lottery $\delta_{\bar{\ell}}$. Clearly, IR binds for these types and they pay nothing to the designer. But then the designer can strictly improve her payoff by offering b_0 at price $\alpha(0) \cdot b_0 > 0$ to these types instead of $\delta_{\bar{\ell}}$. Under this modification, IC remains valid, because $\alpha(\tau) \cdot b_0 = \alpha(0) \cdot b_0$ for all $\tau \in T$ by construction of b_0 , i.e., all types are indifferent between receiving b_0 at price $b_0 \cdot \alpha(0)$ or receiving $\delta_{\bar{\ell}}$ for free.

Remark 2 (Optimality of deterministic mechanisms). If instead $\alpha(0) \cdot b_0 = 0$, we can set $b_0 = \delta_{\bar{\ell}}$, so Proposition 3 implies the optimality of a deterministic mechanism. A sufficient condition for $\alpha(0) \cdot b_0 = 0$ is if $\beta_\ell \geq 0$ for all ℓ , i.e., the utility to each allocation is nondecreasing in types. This observation contrasts with Ghili (2023) and Yang (2023), who consider settings with one-dimensional types but without our linear utility assumption. In their settings, optimal mechanisms in general require randomization, even though all bundles' values are nondecreasing in types. \blacktriangle

6 Extensions

Production costs/negative valuations. In our main model, the first-best $R^{\text{FB}} = \mathbb{E}[\mathbf{1}^G \cdot \theta]$ involves supplying the grand bundle to all buyer types. This can fail if the seller faces a production cost or some types have negative valuations for some goods. Appendix C.1 extends the analysis to allow for both possibilities: First, to produce each bundle B , the seller incurs a cost $c(B)$, so that $R^{\text{FB}} = \mathbb{E}[\mathbf{1}^{B_\theta} \cdot \theta - c(B_\theta)]$ involves supplying a possibly different bundle $B_\theta \in \operatorname{argmax}_{B \subseteq G} (\mathbf{1}^B \cdot \theta - c(B))$ to each type θ . Second, we drop the assumption that $\Theta \subseteq \mathbb{R}_{++}^{|G|}$ and instead assume $\mathbf{1}^{B_\theta} \cdot \theta - c(B_\theta) > 0$ for all $\theta \in \Theta$, which ensures positive gains from trade.³³ In this setting, a simple generalization of pure bundling—*single-bundle mechanisms*—achieves the optimal convergence rate to the first-best: Let R_n^{1bd} denote the expected revenue when, following each signal sequence x^n , the seller optimally chooses a *single*

³³If $c \equiv 0$, this simply requires each type to have a positive valuation for at least one good.

bundle B and fixed price $p(B)$ at which to offer this bundle. Theorem C.1 shows that $R^{\text{FB}} - R_n^{\text{1bd}}$ vanishes at the same rate as $R^{\text{FB}} - R_n^{\text{SB}}$, namely as fast as $\mathbb{E}[\lambda^{B_\theta}(\theta)]\sqrt{\frac{\ln n}{n}}$.

Non-additive utilities. While we focused on buyer utilities that are additive across goods, this is not essential for our main result. Appendix C.2 assumes that the buyer’s utility from receiving bundle B and paying transfer t takes the form $\omega_B - t$, where the only restriction on the valuations ω_B is that $\omega_G > \max\{\omega_B, 0\}$ for all $B \subsetneq G$. This allows for fairly general complementarity and substitutability patterns across goods, provided the buyer finds the grand bundle G most attractive.³⁴ Indeed, this abstract formulation can capture other multi-dimensional screening settings beyond multi-good monopoly; for example, bundles B might be reinterpreted as discrete quality levels, where G represents the highest quality. As in our main model, prior to choosing a mechanism, the seller observes n i.i.d. signals about the buyer’s type, which in this setting corresponds to the vector $\omega := (\omega_B)_{B \subseteq G} \in \mathbb{R}^{2^G}$. Applying analogous arguments as in Theorem 1, Theorem C.2 shows that the revenue gap to the first-best continues to vanish at the same rate under the second-best and optimal bundling mechanisms, viz., as fast as $\mathbb{E}[\lambda(\omega)]\sqrt{\frac{\ln n}{n}}$, where $\lambda(\omega) := \sqrt{\mathbf{1}^G \cdot I(\omega)^{-1} \mathbf{1}^G}$.

More general seller information. To formalize precise seller information, we assumed that the seller observes n i.i.d. signals about the buyer’s type. This formulation allowed us to apply the Bernstein-von Mises theorem to reduce the analysis to a setting with a deterministic sequence of Gaussian seller beliefs. While we do not pursue this direction in the current paper, we conjecture that our results can be generalized to certain classes of correlated signals, by applying non-i.i.d. versions of the Bernstein-von Mises theorem to extend this Gaussian approximation approach.

As an alternative modeling approach, Appendix C.3 takes as a primitive a general (not necessarily Gaussian) deterministic sequence of seller beliefs indexed by n that converge to a point-mass on the true type θ^* . Such beliefs can be interpreted as the outcome of the seller observing increasingly precise information that need not take the form of n (i.i.d. or correlated) signal draws. Under a condition that rules out heavy-tailed seller beliefs, we extend the key insight (Proposition 2) that intensive-margin losses dominate the extensive margin at large n under single-good monopoly. Based on this, Proposition C.1 shows that pure bundling continues to achieve the same convergence rate to the first-best as the optimal deterministic mechanism.³⁵

³⁴This can be further relaxed, similar to the negative type case in the previous paragraph.

³⁵We do not analyze non-deterministic mechanisms, but conjecture that the approach via the re-

A Appendix: Main Proofs

This appendix presents the proofs of Propositions 1–2 (Appendix A.1) and uses these to prove Theorem 1 (Appendix A.2). Online Appendices B–C present proofs and other omitted details for Sections 5–6.

A.1 Gaussian Environment

We analyze the Gaussian environment from Section 3.2. Appendix A.1.2 proves Proposition 2. Appendix A.1.3 uses this to prove (a strengthening of) Proposition 1.

Throughout Appendix A.1, we denote by $R^i(\mu, \Sigma)$ ($i \in \{\text{SB}, \text{bd}, \text{sep}\}$) the seller’s optimal expected revenue (under general mechanisms, bundling, and separate sales) when her belief is distributed $\mathcal{N}(\mu, \Sigma)$; when $|G| = 1$, we drop the superscripts i . Let Φ and φ be the cdf and pdf of the one-dimensional standard normal distribution. Let $\gamma := \sup_{x \in \mathbb{R}} \varphi(x)x$. We will often use the following Gaussian tail bound (Wainwright, 2019):

$$\Phi(-z) \leq \exp\left[-\frac{z^2}{2}\right], \quad \text{for all } z \geq 0. \quad (9)$$

Moreover, the tail probability admits the approximation $\lim_{z \rightarrow \infty} \frac{\Phi(-z)}{\frac{1}{\sqrt{2\pi}|z|} \exp[-\frac{z^2}{2}]} = 1$.

A.1.1 Preliminaries for Single-Good Case

Assume $|G| = 1$ and fix some $\mu, \sigma \in \mathbb{R}$. Then $R(\mu, \sigma^2) = \max_{p \in \mathbb{R}} p(1 - \Phi(\frac{p-\mu}{\sigma}))$. Let $x^*(\mu, \sigma^2) := \operatorname{argmax}_{x \in \mathbb{R}} (\mu + x\sigma)(1 - \Phi(x))$ be the optimal normalized price.

Lemma A.1. *We have $|x^*(\mu, \sigma^2)| \geq \sqrt{2 \left| \ln \frac{\mu}{\sigma\sqrt{2\pi}(1+\gamma)} \right|}$. If $\mu \geq 0$, then $x^*(\mu, \sigma^2) \leq 1$.*

Proof. Let $x^* := x^*(\mu, \sigma^2)$. To prove the first claim, assume $\mu \neq 0$, as the claim is trivial if $\mu = 0$. The first-order condition for revenue maximization implies

$$\sigma \geq \sigma(1 - \Phi(x^*)) = (\mu + \sigma x^*)\varphi(x^*),$$

which yields $1 - \varphi(x^*)x^* \geq \frac{\mu}{\sigma}\varphi(x^*)$. Thus we have

$$\frac{1}{\sqrt{2\pi}}e^{-x^{*2}/2} = \varphi(x^*) \leq (1 - \varphi(x^*)x^*) \frac{\sigma}{\mu} \leq (1 + \gamma) \frac{\sigma}{\mu}.$$

laxed problem in Section 4.5 might be extended by exploiting the non-Gaussian analysis in Section 5.

After manipulation, we get $e^{x^{*2}} \geq \left(\frac{\mu}{\sigma}\right)^2 \frac{1}{2\pi(1+\gamma)^2}$, which ensures $|x^*| \geq \sqrt{2 \left| \ln \frac{\mu}{\sigma\sqrt{2\pi(1+\gamma)}} \right|}$.

If $\mu \geq 0$, the derivative of revenue with respect to $x > 1$ satisfies

$$\sigma(1 - \Phi(x)) - \varphi(x)(\mu + \sigma x) \leq \sigma \left(\frac{\varphi(x)}{x} - \varphi(x)x \right) < 0.$$

Therefore, $x^* \leq 1$ is needed to satisfy the first-order condition. \square

Lemma A.2. *If $\mu \geq 0$, then $R(\mu, \sigma^2) \leq \max \left\{ \mu - \sigma \sqrt{2 \left| \ln \frac{\mu}{\sigma\sqrt{2\pi(1+\gamma)}} \right|}, \frac{\mu + \sigma}{2} \right\}$.*

Proof. As before, let $x^* = x^*(\mu, \sigma^2)$ denote the optimal normalized price. If $x^* \leq 0$, then by Lemma A.1, $x^* \leq -\sqrt{2 \left| \ln \frac{\mu}{\sigma\sqrt{2\pi(1+\gamma)}} \right|}$. Thus,

$$R(\mu, \sigma^2) = (\mu + x^*\sigma)(1 - \Phi(x^*)) \leq (\mu + x^*\sigma) \leq \mu - \sigma \sqrt{2 \left| \ln \frac{\mu}{\sigma\sqrt{2\pi(1+\gamma)}} \right|}.$$

If $x^* > 0$, then by Lemma A.1, $x^* \leq 1$. Thus,

$$R(\mu, \sigma^2) \leq (\mu + \sigma x^*)(1 - \Phi(x^*)) \leq \frac{\mu + \sigma}{2}.$$

Combining the above, we obtain the desired bound. \square

Proposition A.1. *If $\mu \geq 0$, then for every $n \in \mathbb{N}$,*

$$\begin{aligned} \max \left\{ \mu - \sigma \sqrt{\frac{\ln n}{n}} + \frac{\sigma}{\sqrt{n}} \sqrt{2 \left| \ln \frac{\mu}{\sigma\sqrt{2\pi(1+\gamma)}} \right|}, \frac{\mu + \sigma/\sqrt{n}}{2} \right\} &\geq R(\mu, \sigma^2/n) \\ &\geq \mu - \sigma \sqrt{\frac{\ln n}{n}} - \frac{\mu}{\sqrt{2\pi n \ln n}}. \end{aligned}$$

If $\mu \leq 0$, then for every $n \in \mathbb{N}$, $R(\mu, \sigma^2/n) \leq \sqrt{\frac{2\pi}{n}}\sigma$.

Proof. First, consider $\mu \geq 0$. By Lemma A.2,

$$R(\mu, \sigma^2/n) \leq \max \left\{ \mu - \frac{\sigma}{\sqrt{n}} \sqrt{2 \left| \ln \frac{\mu\sqrt{n}}{\sigma\sqrt{2\pi(1+\gamma)}} \right|}, \frac{\mu + \sigma/\sqrt{n}}{2} \right\}.$$

Thus, the desired upper bound follows from the observation that

$$\sqrt{2 \left| \ln \frac{\mu \sqrt{n}}{\sigma \sqrt{2\pi}(1+\gamma)} \right|} = \sqrt{\left| \ln n + 2 \ln \frac{\mu}{\sigma \sqrt{2\pi}(1+\gamma)} \right|} \geq \sqrt{\ln n} - \sqrt{2 \left| \ln \frac{\mu}{\sigma \sqrt{2\pi}(1+\gamma)} \right|}.$$

To derive the lower bound, observe that the revenue from offering price $\mu - \sigma \sqrt{\frac{\ln n}{n}}$ is

$$\begin{aligned} & \left(\mu - \sigma \sqrt{\frac{\ln n}{n}} \right) \left(1 - \Phi \left(-\sqrt{\ln n} \right) \right) \geq \mu - \sigma \sqrt{\frac{\ln n}{n}} - \mu \Phi \left(-\sqrt{\ln n} \right) \\ & = \mu - \sigma \sqrt{\frac{\ln n}{n}} - \mu \left(1 - \Phi \left(\sqrt{\ln n} \right) \right) \geq \mu - \sigma \sqrt{\frac{\ln n}{n}} - \mu \frac{\varphi(\sqrt{\ln n})}{\sqrt{\ln n}} \\ & = \mu - \sigma \sqrt{\frac{\ln n}{n}} - \frac{\mu}{\sqrt{2\pi n \ln n}}, \end{aligned}$$

where the second inequality uses the bound on the Mills ratio (Wainwright, 2019).

Next, consider $\mu \leq 0$, and let $x^* := x^*(\mu, \sigma^2/n)$. Then

$$R(\mu, \sigma^2/n) \leq \mu + x^* \frac{\sigma}{\sqrt{n}} = \frac{\sigma}{\sqrt{n}} \frac{1 - \Phi(x^*)}{\varphi(x^*)} = \frac{\sigma}{\sqrt{n}} \frac{\Phi(-x^*)}{\varphi(x^*)} \leq \frac{\sigma}{\sqrt{n}} \frac{e^{-\frac{(x^*)^2}{2}}}{\varphi(x^*)} = \sqrt{2\pi} \frac{\sigma}{\sqrt{n}},$$

where the first equality uses the first-order condition $\frac{\sigma}{\sqrt{n}}(1 - \Phi(x^*)) = (\mu + x^* \frac{\sigma}{\sqrt{n}})\varphi(x^*)$, and the second inequality uses the tail bound (9) at $x^* \geq |\mu|/(\sigma/\sqrt{n}) \geq 0$. \square

A.1.2 Proof of Proposition 2

Take any $\delta \geq 0$ and sequence of prices (p_n) with $\theta^* - p_n \sim \delta \sqrt{\ln n} \frac{\sigma}{\sqrt{n}}$. Then

$$F_n(p_n) = \Phi \left(\frac{(p_n - \theta^*)\sqrt{n}}{\sigma} \right) = \exp \left[-\frac{\delta^2}{2} \ln n + o(\ln n) \right] = n^{-\frac{\delta^2}{2} + o(1)},$$

and thus

$$\lim_{n \rightarrow \infty} \frac{\theta^* F_n(p_n)}{(\ln n)^{1/2} n^{-1/2}} = \begin{cases} 0 & \text{if } \delta > 1, \\ \infty & \text{if } \delta < 1. \end{cases} \quad (10)$$

Additionally, if $\theta^* - p_n = \sqrt{\ln n} \frac{\sigma}{\sqrt{n}}$, then $\lim_n \frac{\theta^* F_n(p_n)}{(\ln n)^{1/2} n^{-1/2}} = 0$, and hence $R^{\text{FB}}(\theta^*) - R_n(p_n) \sim \sqrt{\ln n} \frac{\sigma}{\sqrt{n}}$, where $R_n(p_n) := p_n(1 - F_n(p_n))$. This implies that, up to taking

a convergent subsequence, any sequence of optimal prices (p_n^*) must satisfy

$$1 \geq \lim_{k \rightarrow \infty} \frac{R^{\text{FB}}(\theta^*) - R_{n_k}(p_{n_k}^*)}{\sqrt{\ln n_k} \frac{\sigma}{\sqrt{n_k}}} \geq \lim_{k \rightarrow \infty} \frac{\theta^* - p_{n_k}^*}{\sqrt{\ln n_k} \frac{\sigma}{\sqrt{n_k}}}. \quad (11)$$

If either inequality is strict, then $\theta^* - p_{n_k}^* \sim \delta \sqrt{\ln n_k} \frac{\sigma}{\sqrt{n_k}}$ for some $\delta \in (0, 1)$, so (10) implies the contradiction that

$$\infty = \lim_{k \rightarrow \infty} \frac{\theta^* F_{n_k}(p_{n_k}^*)}{\sqrt{\ln n_k} \frac{\sigma}{\sqrt{n_k}}} \leq \lim_{k \rightarrow \infty} \frac{R^{\text{FB}}(\theta^*) - R_{n_k}^*(p_{n_k}^*)}{\sqrt{\ln n_k} \frac{\sigma}{\sqrt{n_k}}}.$$

Thus, both inequalities in (11) hold with equality, which implies $R^{\text{FB}}(\theta^*) - R_n(p_n^*) \sim \sqrt{\ln n} \frac{\sigma}{\sqrt{n}} \sim \theta^* - p_n^*$. This in turn implies that $\theta^* F_n(p_n^*) = o\left(\sqrt{\frac{\ln n}{n}}\right)$, as claimed.

A.1.3 Proof of Proposition 1

We prove the following stronger result than Proposition 1, which we will use in the proof of Theorem 1. Assumption 1 ensures that $I(\theta)$ is invertible for each $\theta \in \hat{\Theta}$, and that $\{I(\theta)^{-1} : \theta \in \hat{\Theta}\}$ is a compact set consisting of positive definite matrices. Let \mathcal{J} be a compact neighborhood of $\{I(\theta)^{-1} : \theta \in \hat{\Theta}\}$. By choosing \mathcal{J} to be sufficiently small, we can assume that every $J \in \mathcal{J}$ is positive definite.

Proposition A.2. *There is $K > 0$ such that for all $\theta^* \in \hat{\Theta}$, $J \in \mathcal{J}$, and $n \in \mathbb{N}$,*

$$\left| R^i \left(\theta^*, \frac{1}{n} J \right) - \theta^* \cdot \mathbf{1}^G + \left(\frac{\ln n}{n} \mathbf{1}^G \cdot J \mathbf{1}^G \right)^{1/2} \right| \leq K n^{-1/2} \quad \text{for } i = \text{bd, SB},$$

$$\text{and} \quad \left| R^{\text{sep}} \left(\theta^*, \frac{1}{n} J \right) - \theta^* \cdot \mathbf{1}^G + \sum_{g \in G} \left(\frac{\ln n}{n} \mathbf{1}^g \cdot J \mathbf{1}^g \right)^{1/2} \right| \leq K n^{-1/2}.$$

A.1.4 Proof of Proposition A.2

Since $R^{\text{bd}}(\theta^*, \frac{1}{n} J) = R\left(\sum_{g \in G} \theta_g^*, \frac{\mathbf{1}^G \cdot J \mathbf{1}^G}{n}\right)$ and $R^{\text{sep}}(\theta^*, \frac{1}{n} J) = \sum_{g \in G} R\left(\theta_g^*, \frac{\mathbf{1}^g \cdot J \mathbf{1}^g}{n}\right)$, the desired inequalities for bundling and separate sales follow from the single-good case (Proposition A.1) combined with the compactness of $\hat{\Theta}$ and \mathcal{J} . Since $R^{\text{bd}}(\theta^*, \frac{1}{n} J) \leq R^{\text{SB}}(\theta^*, \frac{1}{n} J)$, the lower bound for $R^{\text{SB}}(\theta^*, \frac{1}{n} J)$ is immediate from the bundling case, and it suffices to establish the desired upper bound for $R^{\text{SB}}(\theta^*, \frac{1}{n} J)$.

To this end, suppose types are distributed $\mathcal{N}(\theta^*, \frac{1}{n}J)$ for some $\theta^* \in \hat{\Theta}$ and $J \in \mathcal{J}$. Enumerate the set of bundles as $\mathcal{B} = \{B_\ell : \ell = 1, \dots, m\}$, where $m = 2^{|\mathcal{G}|}$ and $B_m := G$. Rewrite types as vectors $\omega \in \mathbb{R}^m$, where each coordinate ω_ℓ represents the valuation of bundle B_ℓ .³⁶ Note that ω is distributed $\mathcal{N}(y, \frac{1}{n}\Sigma)$, where $y \in \mathbb{R}^m$ satisfies $y_\ell = \sum_{g \in B_\ell} \theta_g^*$ for all ℓ , and $\Sigma \in \mathbb{R}^{m \times m}$ satisfies $\Sigma_{\ell\ell'} = \sum_{g \in B_\ell, d' \in B_{\ell'}} J_{dd'}$ for all ℓ, ℓ' . Observe that y is strictly positive and $y_m \geq y_\ell + \kappa$ for all $\ell \neq m$, where $\kappa := \min_{g \in G, \theta \in \Theta} \theta_g > 0$.

Consider the following relaxed problem. Pick a full-rank matrix $A \in \mathbb{R}^{(m-1) \times m}$ with $A\Sigma\mathbf{1}^m = \mathbf{0}$.³⁷ Conditional on $A\omega = z \in \mathbb{R}^{m-1}$, ω is distributed $\mathcal{N}(\hat{y}^z, \frac{1}{n}\hat{\Sigma})$, where $\hat{y}^z := y + B(z - Ay)$, $B := \Sigma A^\top (A\Sigma A^\top)^{-1}$, $\hat{\Sigma} := \Sigma (\mathbb{I} - A^\top (A\Sigma A^\top)^{-1} A\Sigma)$, and $\mathbb{I} \in \mathbb{R}^{m \times m}$ is the identity matrix. Define

$$\bar{R}^{\text{SB}} \left(y, \frac{1}{n}\Sigma \right) := \mathbb{E} \left[R^{\text{SB}} \left(\hat{y}^z, \frac{1}{n}\hat{\Sigma} \right) \right],$$

where the expectation is with respect to $z = A\omega$, which is distributed $\mathcal{N}(Ay, \frac{1}{n}A\Sigma A^\top)$. The value $\bar{R}^{\text{SB}} \left(y, \frac{1}{n}\Sigma \right)$ represents the seller's revenue in the relaxed problem where each type ω is restricted to only report types in $\{\omega' \in \mathbb{R}^m : A\omega' = A\omega\}$. Clearly, $R^{\text{SB}} \left(\theta^*, \frac{1}{n}J \right) \leq \bar{R}^{\text{SB}} \left(y, \frac{1}{n}\Sigma \right)$.

Conditional on $A\omega = z$, the buyer's type belongs to $\{\omega' \in \mathbb{R}^m : A\omega' = z\}$, which is a line of the form $\{\hat{y}^z + \tau\Sigma\mathbf{1}^m : \tau \in \mathbb{R}\}$. Thus, the current setting is a special case of the environment in Section 5, where $\alpha(0) := \hat{y}^z$ and $\beta := \Sigma\mathbf{1}^m$. Moreover, conditional on $A\omega = z$, the valuation ω_ℓ of each bundle B_ℓ is distributed $\mathcal{N}(\hat{y}_\ell^z, \frac{1}{n}\hat{\sigma}_\ell^2)$, where $\hat{\sigma}_\ell^2 := \hat{\Sigma}_{\ell\ell}$. Crucially, $\hat{y}_m^z = y_m$ and $\hat{\sigma}_m^2 = \sigma_m^2 := \Sigma_{mm}$, i.e., the distribution of the valuation ω_m of the grand bundle is unchanged by observing z . This is because $A\Sigma\mathbf{1}^m = \mathbf{0}$ ensures that ω_m and $A\omega$ are uncorrelated.

To apply the analysis in Section 5, we define the lottery

$$b_0^z \in \operatorname{argmax}_{q \in \Delta(\mathcal{B})} \hat{y}^z \cdot q \text{ such that } \beta \cdot q = 0.$$

Conditional on $A\omega = z$, the value of b_0^z is thus distributed $\mathcal{N}(\hat{y}_0, \frac{1}{n}\hat{\sigma}_0^2)$, where $\hat{y}_0 :=$

³⁶This more general formulation is useful for analyzing non-additive utilities (Appendix C.2).

³⁷In particular, if $\Sigma\mathbf{1}^m = \mathbf{0}$, then we can choose any full-rank matrix A . If $\Sigma\mathbf{1}^m \neq \mathbf{0}$, then $\{a : a \cdot \Sigma\mathbf{1}^m = 0\}$ is an $(m-1)$ -dimensional hyperplane in \mathbb{R}^m , and hence includes $m-1$ linearly independent vectors a_1, \dots, a_{m-1} . Then the matrix A with row vectors a_1, \dots, a_{m-1} has row rank $m-1$ and satisfies $A\Sigma\mathbf{1}^m = \mathbf{0}$.

$b_0^z \cdot \hat{y}^z$ and $\hat{\sigma}_0^2 := b_0^z \cdot \hat{\Sigma} b_0^z$. Let $R^s \left(\hat{y}^z, \frac{1}{n} \hat{\Sigma} \right)$ denote the optimal revenue when restricting attention to mechanisms that allocate to each type either a deterministic bundle $\{\delta_{B_\ell} : \ell = 1, \dots, m\}$ or b_0^z . By Proposition 3, this restriction is without loss.

The following result establishes an upper bound on $R^s \left(\hat{y}^z, \frac{1}{n} \hat{\Sigma} \right)$:

Lemma A.3. *Suppose that $y_m > \hat{y}_\ell^z + \kappa/2 > 0$ for all $\ell = 0, 1, \dots, m-1$. Then $R^s \left(\hat{y}^z, \frac{1}{n} \hat{\Sigma} \right)$ is bounded above by*

$$\max \left\{ y_m - \sigma_m \sqrt{\frac{\ln n}{n}} + \frac{\sigma_m}{\sqrt{n}} \sqrt{2 \left| \ln \frac{\kappa/2}{\sigma_m \sqrt{2\pi}(1+\gamma)} \right|}, y_m - \frac{\kappa}{4} + \frac{\sigma_m}{2\sqrt{n}} \right\} + \sum_{\ell \neq m} \sqrt{\frac{2\pi}{n}} \hat{\sigma}_\ell.$$

Proof. Any IC-IR mechanism that allocates only deterministic bundles and b_0^z can be represented as a collection of prices p_ℓ for each deterministic bundle B_ℓ ($\ell = 1, \dots, m$) and price p_0 for lottery b_0^z . For the empty bundle $B_{\bar{\ell}}$, where $\omega_{\bar{\ell}} = 0$ for all types, it is without loss to set $p_{\bar{\ell}} = 0$. Denote by \mathbb{P}_z the probability measure conditional on $z = A\omega$. The revenue $\sum_{\ell=0}^m \mathbb{P}_z(\omega_\ell - p_\ell \geq \omega_{\ell'} - p_{\ell'} \forall \ell')$ is bounded above by

$$\begin{aligned} & y_m - \kappa/2 + \sum_{\ell=0}^m \mathbb{P}_z(\omega_\ell - p_\ell \geq \omega_{\ell'} - p_{\ell'} \forall \ell') \max\{p_\ell - (y_m - \kappa/2), 0\} \\ & \leq y_m - \kappa/2 + \sum_{\ell=0}^m \mathbb{P}_z(\omega_\ell \geq p_\ell) \max\{p_\ell - (y_m - \kappa/2), 0\} \\ & \leq y_m - \kappa/2 + \sum_{\ell=0}^m R(\hat{y}_\ell^z - y_m + \kappa/2, \hat{\sigma}_\ell^2/n) \\ & \leq y_m - \kappa/2 + R(\kappa/2, \sigma_m^2/n) + \sum_{\ell \neq m} \sqrt{\frac{2\pi}{n}} \hat{\sigma}_\ell, \end{aligned}$$

where the last inequality uses part 2 of Proposition A.1. Then part 1 of Proposition A.1 yields the desired bound. \square

To complete the proof of Proposition A.2, recall that $\hat{y}^z = y + B(z - Ay)$ and $\hat{y}_m^z = y_m > y_\ell + \kappa$ for all $z \in \mathbb{R}^{m-1}$ and $\ell = 0, 1, \dots, m-1$. Pick $\varepsilon > 0$ small enough that all z with $\|z - Ay\| \leq \varepsilon$ satisfy $y_m > \hat{y}_\ell^z + \kappa/2 > 0$ for all $\ell = 0, 1, \dots, m-1$. Since z is distributed $\mathcal{N}(Ay, \frac{1}{n} A \Sigma A^\top)$, there is $K_1 > 0$ such that, for all n ,

$$\left| \overline{R}^{\text{SB}} \left(y, \frac{1}{n} \Sigma \right) - \mathbb{E} \left[R^{\text{SB}} \left(\hat{y}^z, \frac{1}{n} \hat{\Sigma} \right) : \|z - Ay\| \leq \varepsilon \right] \right| \leq K_1 n^{-1}.$$

By Lemma A.3, there is $K_2 > 0$ such that, for all n ,

$$\mathbb{E} \left[R^{\text{SB}} \left(\hat{y}^z, \frac{1}{n} \hat{\Sigma} \right) : \|z - Ay\| \leq \varepsilon \right] \leq y_m - \left(\frac{\ln n}{n} \right)^{1/2} \sigma_m + K_2 n^{-1/2}.$$

Since $y_m = \sum_{g \in G} \theta_g^*$ and $\sigma_m = (\mathbf{1}^G \cdot J\mathbf{1}^G)^{1/2}$, this implies the existence of $K > 0$ such that $\bar{R}^{\text{SB}} \left(y, \frac{1}{n} \Sigma \right) \leq \sum_{g \in G} \theta_g^* - \left(\frac{\ln n}{n} \right)^{1/2} \sigma_m + Kn^{-1/2}$ for all $n \in \mathbb{N}$. Finally, by compactness of $\hat{\Theta}$ and \mathcal{J} , the coefficient K can be chosen uniformly in y and Σ .

A.2 Proof of Theorem 1

A.2.1 Preliminaries

Denote by $R^i(\lambda)$ ($i \in \{\text{SB}, \text{bd}, \text{sep}\}$) the optimal expected revenue (under general mechanisms, bundling, and separate sales) when the buyer's type is drawn from $\lambda \in \Delta(\Theta)$. Denote by $\lambda_{x^n} \in \Delta(\Theta)$ the seller's posterior after observing signals x^n . For all $i \in \{\text{SB}, \text{bd}, \text{sep}\}$, $R_n^i = \mathbb{E}[R^i(\lambda_{x^n})]$. For all signals x^n , the compactness of $\hat{\Theta}$ and continuity of $\ln f$ implies the existence of a maximum likelihood estimate (MLE)

$$\hat{\theta}_{x^n} \in \operatorname{argmax}_{\theta \in \hat{\Theta}} \sum_{i=1}^n \ln f(x_i, \theta).$$

(Note that the maximum is taken over $\hat{\Theta}$ instead of Θ). The *empirical Fisher information* at θ is the matrix

$$I_{x^n}(\theta) := \left(-\frac{1}{n} \sum_{i=1}^n \frac{\partial^2}{\partial \theta_g \partial \theta_{g'}} \ln f(x_i, \theta) \right)_{g, g' \in G}.$$

For each x^n , denote by $\hat{\lambda}_{x^n}$ the Gaussian distribution with mean $\hat{\theta}_{x^n}$ and covariance matrix $\frac{1}{n} I_{x^n}(\hat{\theta}_{x^n})^{-1}$ if $I_{x^n}(\hat{\theta}_{x^n})$ is invertible; if $I_{x^n}(\hat{\theta}_{x^n})$ is not invertible, we set the covariance matrix to be the identity matrix. Denote by TV the total-variation distance on $\Delta(\mathbb{R}^{|G|})$, i.e., $\text{TV}(\lambda, \lambda') := \sup_{\Theta'} (\lambda(\Theta') - \lambda'(\Theta'))$, where the supremum is taken over all measurable subsets $\Theta' \subseteq \mathbb{R}^{|G|}$.

By the Bernstein-von Mises theorem, the seller's posterior λ_{x^n} can be approximated (with respect to the total-variation metric) by the Gaussian belief $\hat{\lambda}_{x^n}$. Moreover, Hipp and Michel (1976) bound the approximation errors. Denote by \mathbb{P}_θ and \mathbb{E}_θ

probabilities and expectations conditional on the true type being θ :

Lemma A.4 (Hipp and Michel (1976)). *For each $\theta \in \text{int } \Theta$, there is $c_\theta \geq 0$ with*

$$\mathbb{P}_\theta(\text{TV}(\lambda_{x^n}, \hat{\lambda}_{x^n}) > c_\theta n^{-1/2}) = o(n^{-1}).$$

It is also known that the MLE $\hat{\theta}_{x^n}$ is asymptotically normal conditional on each θ . Specifically, the distribution of $n^{1/2}(\hat{\theta}_{x^n} - \theta)$ is approximated by the Gaussian distribution with mean $\mathbf{0}$ and covariance matrix $I(\theta)^{-1}$, denoted by $N_{I(\theta)^{-1}}$. Moreover, Pfanzagl (1973) bounds the approximation errors, uniformly in θ :

Lemma A.5 (Pfanzagl (1973)). *There is $c' > 0$ such that for every $\theta \in \Theta$, convex set $E \subseteq \mathbb{R}^G$, and $n \in \mathbb{N}$, we have $\left| \mathbb{P}_\theta \left(n^{1/2}(\hat{\theta}_{x^n} - \theta) \in E \right) - N_{I(\theta)^{-1}}(E) \right| \leq c'n^{-1/2}$.*

Let $T := \sup_{\theta \in \Theta, B \subseteq G} (\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B)^{1/2}$, which is finite by compactness of Θ and Assumption 1. For every $B \subseteq G$ and $n \in \mathbb{N}$, let $E_n^B := \{z \in \mathbb{R}^{|G|} : z \cdot \mathbf{1}^B > (\ln n)^{1/2} T\}$. Then for all $n \in \mathbb{N}$,

$$\begin{aligned} \mathbb{P}_\theta \left[|(\hat{\theta}_{x^n} - \theta) \cdot \mathbf{1}^B| > \left(\frac{\ln n}{n} \right)^{1/2} T \right] &\leq N_{I(\theta)^{-1}}(E_n^B) + c'n^{-1/2} \\ &\leq 2n^{-\frac{T}{2(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B)^{1/2}}} + c'n^{-1/2} \leq (2 + c')n^{-1/2}, \end{aligned} \quad (12)$$

where the first inequality uses Lemma A.5 (with corresponding coefficient c'). The second inequality uses the Gaussian tail bound (9).

We first prove some preliminary lemmas. Endow $|G| \times |G|$ -matrices with the norm $\|A\| := \sup_{y \neq \mathbf{0}} \frac{\|Ay\|}{\|y\|} = \sum_{g \in G} \max_{g' \in G} |A_{gg'}|$ induced by the L^1 -norm on $\mathbb{R}^{|G|}$:

Lemma A.6. *There is $C > 0$ such that for all $\theta \in \Theta$ and $n \in \mathbb{N}$,*

$$\mathbb{P}_\theta [A_{\theta,n}] \geq 1 - Cn^{-1/2}, \quad (13)$$

where $A_{\theta,n}$ is the event that $\|I_{x^n}(\hat{\theta}_{x^n}) - I(\theta)\| \leq \left(\frac{\ln n}{n}\right)^{1/2} C$. Moreover, there is \bar{n} such that for all $\theta \in \Theta$ and $n \geq \bar{n}$, $I_{x^n}(\hat{\theta}_{x^n})$ is invertible with $\|I_{x^n}(\hat{\theta}_{x^n})^{-1} - I(\theta)^{-1}\| \leq C\|I_{x^n}(\hat{\theta}_{x^n}) - I(\theta)\|$ for any x^n in $A_{\theta,n}$.

Proof. Inequality (13): By Assumption 1.2, there is a Lipschitz coefficient L of

$\frac{\partial^2 \ln f(x, \cdot)}{\partial \theta_g \partial \theta_{g'}}$ that is uniform in $g, g' \in G$ and $x \in X$. Then for every $n \in \mathbb{N}$ and $\theta \in \Theta$,

$$\begin{aligned} \mathbb{P}_\theta \left[\left\| I_{x^n}(\theta) - I_{x^n}(\hat{\theta}_{x^n}) \right\| > \left(\frac{\ln n}{n} \right)^{1/2} LT \right] &\leq \mathbb{P}_\theta \left[\left\| \theta - \hat{\theta}_{x^n} \right\| > \left(\frac{\ln n}{n} \right)^{1/2} T \right] \\ &\leq \sum_{g \in G} \mathbb{P}_\theta \left[\left| (\theta - \hat{\theta}_{x^n}) \cdot \mathbf{1}^g \right| > \left(\frac{\ln n}{n} \right)^{1/2} T \right] \leq |G|(2+c')n^{-1/2} \end{aligned}$$

where the last inequality uses (12). By Assumption 1.2, $L \geq \sup_{x \in X, \theta \in \Theta, g, g' \in G} \left| \frac{\partial^2 \ln f(x, \theta)}{\partial \theta_g \partial \theta_{g'}} \right|$. Then for every $n \in \mathbb{N}$ and $\theta \in \Theta$, Hoeffding's inequality implies that

$$\begin{aligned} \mathbb{P}_\theta \left[\left\| I(\theta) - I_{x^n}(\theta) \right\| > \left(\frac{\ln n}{n} \right)^{1/2} \frac{L}{2} \right] \\ \leq \sum_{g, g' \in G} \mathbb{P}_\theta \left[\left| \int \frac{\partial^2 \ln f(x, \theta)}{\partial \theta_g \partial \theta_{g'}} dP_\theta(x) - \frac{1}{n} \sum_{i=1}^n \frac{\partial^2 \ln f(x_i, \theta)}{\partial \theta_g \partial \theta_{g'}} \right| > \left(\frac{\ln n}{n} \right)^{1/2} \frac{L}{2} \right] \leq |G|^2 2n^{-1/2}. \end{aligned}$$

Combining the above observations and the triangle inequality, (13) holds for all $C \geq \max\{LT, |G|(2+c'), \frac{L}{2}, 2|G|^2\}$, $n \in \mathbb{N}$, and $\theta \in \Theta$.

Invertibility: By Assumption 1.5, $\|I(\theta)^{-1}\| > 0$ for each θ . Thus, there is $\varepsilon > 0$ such that for any $\theta \in \Theta$, $n \in \mathbb{N}$, and x^n with $\|I(\theta) - I_{x^n}(\hat{\theta}_{x^n})\| \leq \varepsilon$, the matrix $I_{x^n}(\hat{\theta}_{x^n})$ is invertible with $\|I_{x^n}(\hat{\theta}_{x^n})^{-1}\| \leq 2\|I(\theta)^{-1}\|$. Since Θ is compact and $I(\theta)$ is continuous in θ by Assumption 1, ε can be chosen uniformly in θ . Under such realizations of x^n ,

$$\begin{aligned} \|I(\theta)^{-1} - I_{x^n}(\hat{\theta}_{x^n})^{-1}\| &= \|I(\theta)^{-1}(I_{x^n}(\hat{\theta}_{x^n}) - I(\theta))I_{x^n}(\hat{\theta}_{x^n})^{-1}\| \\ &\leq \|I(\theta)^{-1}\| \|I_{x^n}(\hat{\theta}_{x^n}) - I(\theta)\| \|I_{x^n}(\hat{\theta}_{x^n})^{-1}\| \leq 2 \sup_{\theta \in \Theta} \|I(\theta)^{-1}\|^2 \|I_{x^n}(\hat{\theta}_{x^n}) - I(\theta)\|. \end{aligned}$$

By the proof of (13), the conclusion follows by taking $C \geq \max\{LT, |G|(2+c'), \frac{L}{2}, 2|G|^2, 2 \sup_{\theta \in \Theta} \|I(\theta)^{-1}\|^2\}$ and \bar{n} large enough that $(\frac{\ln \bar{n}}{\bar{n}})^{1/2} C \leq \varepsilon$. \square

Henceforth, we will often use the constant \bar{n} and C and events $(A_{\theta, n})$ obtained above without explicitly referencing Lemma A.6.

Lemma A.7. *There is $C' > 0$ such that for all $\theta \in \Theta$, $n \geq \bar{n}$, and $B \subseteq G$,*

$$\mathbb{E}_\theta \left[\left\| (\hat{\theta}_{x^n} - \theta) \right\| \mid A_{\theta, n} \right] \leq C' n^{-1/2}, \quad (14)$$

$$\text{and } \left| \mathbb{E}_\theta \left[\left(\mathbf{1}^B \cdot I_{x^n}(\hat{\theta}_{x^n})^{-1} \mathbf{1}^B \right)^{1/2} \mid A_{\theta,n} \right] - \left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} \right| \leq C' n^{-1/2}. \quad (15)$$

Proof. Inequality (14): Denote by $\hat{F}_{\theta,n}^g$ the cdf of $(\hat{\theta}_{x^n} - \theta) \cdot \mathbf{1}^g$ at θ , and by $F_{\theta,n}^g$ the cdf of the Gaussian distribution with mean 0 and variance $\frac{1}{n} \mathbf{1}^g \cdot I(\theta)^{-1} \mathbf{1}^g$. Then for every $n \in \mathbb{N}$, $g \in G$, and $\theta \in \Theta$,

$$\begin{aligned} & \left| \int_{-\kappa_n}^{\kappa_n} z d\hat{F}_{\theta,n}^g(z) \right| = \left| \int_{-\kappa_n}^{\kappa_n} z d\hat{F}_{\theta,n}^g(z) - \int_{-\kappa_n}^{\kappa_n} z dF_{\theta,n}^g(z) \right| \\ &= \left| \kappa_n \left(\hat{F}_{\theta,n}^g(\kappa_n) - F_{\theta,n}^g(\kappa_n) + \hat{F}_{\theta,n}^g(-\kappa_n) - F_{\theta,n}^g(-\kappa_n) \right) - \int_{-\kappa_n}^{\kappa_n} \left(\hat{F}_{\theta,n}^g(z) - F_{\theta,n}^g(z) \right) dz \right| \\ &\leq 4\kappa_n c' n^{-1/2} \leq 4Tc' n^{-1/2}, \end{aligned}$$

where $\kappa_n := \left(\frac{\ln n}{n}\right)^{1/2} T$, the second equality uses integration by parts, and the first inequality uses Lemma A.5. Let $M := \sup_{\theta \in \Theta} \|\theta\|$. For all $n \in \mathbb{N}$, $\theta \in \Theta$, the above inequality and (12) yield

$$\mathbb{E}_\theta \left[\|\hat{\theta}_{x^n} - \theta\| \right] = \sum_{g \in G} \left| \int z d\hat{F}_{\theta,n}^g(z) \right| \leq |G|(2Tc' + M(2 + c'))n^{-1/2}. \quad (16)$$

Then by Lemma A.6 and the fact that $\|\hat{\theta}_{x^n} - \theta\| \leq M$ for all x^n , (14) holds for some large enough C' .

Inequality (15): By part 2 of Lemma A.6, it suffices to show that $\sup_{\theta \in \Theta} \mathbb{E}_\theta[\|I_{x^n}(\hat{\theta}_{x^n}) - I(\theta)\| \mid A_{\theta,n}]$ is bounded by a constant times $n^{-1/2}$ for all $n \geq \bar{n}$. By the part 1 of Lemma A.6 and since $\sup_{\theta \in \Theta, n \in \mathbb{N}, x^n \in X^n} \|I_{x^n}(\hat{\theta}_{x^n}) - I(\theta)\| < \infty$ (by Assumption 1), it then suffices to show that $\mathbb{E}_\theta[\|I_{x^n}(\hat{\theta}_{x^n}) - I(\theta)\|]$ is bounded by a constant times $n^{-1/2}$ for all $n \geq \bar{n}$.

By Assumption 1.2, there is a Lipschitz coefficient L of $\frac{\partial^2 \ln f(x, \cdot)}{\partial \theta_g \partial \theta_{g'}}$ that is uniform in $g, g' \in G$ and $x \in X$. For all $n \in \mathbb{N}$ and $\theta \in \Theta$,

$$\mathbb{E}_\theta[\|I_{x^n}(\hat{\theta}_{x^n}) - I_{x^n}(\theta)\|] \leq |G|L\mathbb{E}_\theta[\|\hat{\theta}_{x^n} - \theta\|] \leq L|G|^2(2Tc' + M(2 + c'))n^{-1/2},$$

where the second inequality uses (16). Let $\ell := \sup_{x \in X, \theta \in \Theta, g, g' \in G} \left| \frac{\partial^2 \ln f(x, \theta)}{\partial \theta_g \partial \theta_{g'}} \right|$, which is finite by Assumption 1.2. For all $n \in \mathbb{N}$ and $\theta \in \Theta$,

$$\mathbb{E}_\theta[\|I_{x^n}(\theta) - I(\theta)\|] \leq |G| \max_{g, g' \in G} \mathbb{E}_\theta \left[\left| \sum_{i=1}^n \frac{1}{n} \frac{\partial^2}{\partial \theta_g \partial \theta_{g'}} \ln f(x_i, \theta) - \mathbb{E}_\theta \left[\frac{\partial^2}{\partial \theta_g \partial \theta_{g'}} \ln f(x, \theta) \right] \right| \right]$$

$$\leq |G| \max_{g, g' \in G} \left(\text{Var}_\theta \left[\sum_{i=1}^n \frac{1}{n} \frac{\partial^2}{\partial \theta_g \partial \theta_{g'}} \ln f(x_i, \theta) \right] \right)^{1/2} \leq |G| \ell n^{-1/2},$$

where the second (resp. third) inequality uses Hölder's (resp. Popoviciu's) inequality. Combining these observations with the triangle inequality yields the desired claim. \square

A.2.2 Conditional Payoffs

Next, we prove a conditional analog of Theorem 1:

Proposition A.3. *For all $\theta \in \text{int } \Theta$,*

$$\theta \cdot \mathbf{1}^G - \mathbb{E}_\theta[R^{\text{SB}}(\lambda_{x^n})] \sim \theta \cdot \mathbf{1}^G - \mathbb{E}_\theta[R^{\text{bd}}(\lambda_{x^n})] \sim \left(\frac{\ln n}{n} \mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G \right)^{1/2},$$

$$\text{and } \theta \cdot \mathbf{1}^G - \mathbb{E}_\theta[R^{\text{sep}}(\lambda_{x^n})] \sim \sum_{g \in G} \left(\frac{\ln n}{n} \mathbf{1}^g \cdot I(\theta)^{-1} \mathbf{1}^g \right)^{1/2}.$$

The proof is immediate from the following two lemmas:

Lemma A.8. *For all $\theta \in \text{int } \Theta$, there is $K_\theta > 0$ such that, for all $i \in \{\text{SB}, \text{bd}, \text{sep}\}$,*

$$\left| \mathbb{E}_\theta [R^i(\lambda_{x^n})] - \mathbb{E}_\theta [R^i(\hat{\lambda}_{x^n}) \mid A_{\theta, n}] \right| \leq K_\theta n^{-1/2} + o(n^{-1/2}).$$

Proof. Let $c_\theta > 0$ denote the coefficients obtained in Lemma A.4. Since $R^i(\lambda_{x^n})$ and $R^i(\hat{\lambda}_{x^n})$ are uniformly bounded under realizations in $(A_{\theta, n})$, it suffices by Lemmas A.4 and A.6 to find K_θ such that

$$\mathbb{E}_\theta \left[\left| R^i(\lambda_{x^n}) - R^i(\hat{\lambda}_{x^n}) \right| \mid A_{\theta, n} \cap \{\text{TV}(\lambda_{x^n}, \hat{\lambda}_{x^n}) \leq c_\theta n^{-1/2}\} \right] \leq K_\theta n^{-1/2} + o(n^{-1/2}). \quad (17)$$

Let $\hat{\Theta}' \subseteq \mathbb{R}^{|G|}$ be a compact neighborhood of $\hat{\Theta}$. Let $\hat{\lambda}'_{x^n}$ denote the conditional probability measure on $\hat{\Theta}'$ induced by $\hat{\lambda}_{x^n}$. By Lemma A.6 and compactness of $\{I(\theta) : \theta \in \Theta\}$, there is $M > 0$ such that $\|\frac{1}{n} I_{x^n}(\hat{\theta}_{x^n})^{-1}\| \leq n^{-1} M$ for every $n \geq \bar{n}$ and any realization x^n in $A_{\theta, n}$. Thus, since $\hat{\theta}_{x^n}$ is bounded away from $\mathbb{R}^{|G|} \setminus \hat{\Theta}'$, there is $M' > 0$ such that $\hat{\lambda}_{x^n}(\mathbb{R}^{|G|} \setminus \hat{\Theta}') \leq M n^{-1}$ and $\int_{\mathbb{R}^{|G|} \setminus \hat{\Theta}'} \sum_{g \in G} \max\{\theta_g, 0\} d\hat{\lambda}_{x^n}(\theta) \leq M' n^{-1}$ for any sequence of such realizations.

The seller's expected revenue under belief $\hat{\lambda}_{x^n}$ can be bounded as follows:

$$\hat{\lambda}_{x^n}(\hat{\Theta}')R^i(\hat{\lambda}'_{x^n}) + \int_{\mathbb{R}^{|G|} \setminus \hat{\Theta}'} \sum_{g \in G} \max\{\theta_g, 0\} d\hat{\lambda}_{x^n}(\theta) \geq R^i(\hat{\lambda}_{x^n}) \geq \hat{\lambda}_{x^n}(\hat{\Theta}')R^i(\hat{\lambda}'_{x^n}).$$

The left-hand side is the seller's expected payoff when all types outside $\hat{\Theta}'$ are perfectly observable. The right-hand side is the seller's expected payoff when all types outside $\hat{\Theta}'$ do not purchase anything. Thus, there is $M'' > 0$ such that $\left| R^i(\hat{\lambda}'_{x^n}) - R^i(\hat{\lambda}_{x^n}) \right| \leq M''n^{-1}$ for all x^n in $(A_{\theta,n})$.

Observe that, for all $\theta \in \Theta$, $n \geq 2M$, and x^n in $A_{\theta,n}$,

$$\begin{aligned} \text{TV}(\lambda_{x^n}, \hat{\lambda}'_{x^n}) &= \sup_{E \subseteq \hat{\Theta}'} |\lambda_{x^n}(E) - \hat{\lambda}'_{x^n}(E)| = \sup_{E \subseteq \hat{\Theta}'} |\lambda_{x^n}(E) - \hat{\lambda}_{x^n}(E)| + \hat{\lambda}_{x^n}(E) \frac{1 - \hat{\lambda}_{x^n}(\hat{\Theta}')}{\hat{\lambda}_{x^n}(\hat{\Theta}')} \\ &\leq \text{TV}(\lambda_{x^n}, \hat{\lambda}_{x^n}) + \frac{Mn^{-1}}{1 - Mn^{-1}} \leq \text{TV}(\lambda_{x^n}, \hat{\lambda}_{x^n}) + 2Mn^{-1}. \end{aligned}$$

Since $\hat{\Theta}'$ is bounded, there is $K > 0$ such that for all x^n in $A_{\theta,n}$,

$$\begin{aligned} \left| R^i(\lambda_{x^n}) - R^i(\hat{\lambda}_{x^n}) \right| &\leq \left| R^i(\lambda_{x^n}) - R^i(\hat{\lambda}'_{x^n}) \right| + \left| R^i(\hat{\lambda}'_{x^n}) - R^i(\hat{\lambda}_{x^n}) \right| \\ &\leq K\text{TV}(\lambda_{x^n}, \hat{\lambda}'_{x^n}) + M''n^{-1}. \end{aligned}$$

The above observations yield (17) with $K_\theta = Kc_\theta$. □

Lemma A.9. *There exist $K' > 0$ and $\bar{n}' \in \mathbb{N}$ such that for all $\theta \in \Theta$ and $n \geq \bar{n}'$,*

$$\left| \theta \cdot \mathbf{1}^G - \left(\frac{\ln n}{n} \mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G \right)^{1/2} - \mathbb{E}_\theta[R^i(\hat{\lambda}_{x^n}) \mid A_{\theta,n}] \right| \leq K'n^{-1/2} \quad \text{for } i \in \{\text{SB}, \text{bd}\},$$

$$\text{and} \quad \left| \theta \cdot \mathbf{1}^G - \sum_{g \in G} \left(\frac{\ln n}{n} \mathbf{1}^g \cdot I(\theta)^{-1} \mathbf{1}^g \right)^{1/2} - \mathbb{E}_\theta[R^{\text{sep}}(\hat{\lambda}_{x^n}) \mid A_{\theta,n}] \right| \leq K'n^{-1/2}.$$

Proof. By Proposition A.2, there is $K > 0$ such that for all $n \in \mathbb{N}$ and x^n with $I_{x^n}(\hat{\theta}_{x^n})^{-1} \in \mathcal{J}$, we have

$$\left| R^i(\hat{\lambda}_{x^n}) - \hat{\theta}_{x^n} \cdot \mathbf{1}^G + \left(\frac{\ln n}{n} \mathbf{1}^G \cdot I_{x^n}(\hat{\theta}_{x^n})^{-1} \mathbf{1}^G \right)^{1/2} \right| \leq K \left(\frac{\ln n}{n} \right)^{1/2} \quad \text{for } i \in \{\text{SB}, \text{bd}\},$$

$$\text{and } \left| R^{\text{sep}}(\hat{\lambda}_{x^n}) - \hat{\theta}_{x^n} \cdot \mathbf{1}^G + \sum_{g \in G} \left(\frac{\ln n}{n} \mathbf{1}^g \cdot I_{x^n}(\hat{\theta}_{x^n})^{-1} \mathbf{1}^g \right)^{1/2} \right| \leq K \left(\frac{\ln n}{n} \right)^{1/2}.$$

Lemma A.6 yields \bar{n}' such that for all $n \geq \bar{n}'$, $\theta \in \Theta$, and x^n in $A_{\theta, n}$, we have $I_{x^n}(\hat{\theta}_{x^n})^{-1} \in \mathcal{J}$. Thus, Lemma A.7 implies the claim by setting $K' = K + C'$. \square

A.2.3 Completing the Proof of Theorem 1

Lower bound of revenue gap: We first show that $R^{\text{FB}} - R_n^{\text{SB}}$ and $R^{\text{FB}} - R_n^{\text{bd}}$ (resp. $R^{\text{FB}} - R_n^{\text{sep}}$) vanish no faster than at rate $\left(\frac{\ln n}{n}\right)^{1/2} \mathbb{E}[(\mathbf{1} \cdot I(\theta)^{-1} \mathbf{1})^{1/2}]$ (resp. $\left(\frac{\ln n}{n}\right)^{1/2} \sum_{g \in G} \mathbb{E}[(\mathbf{1}^g \cdot I(\theta)^{-1} \mathbf{1}^g)^{1/2}]$). For $i \in \{\text{SB}, \text{bd}\}$,

$$\begin{aligned} \liminf_{n \rightarrow \infty} (R^{\text{FB}} - R_n^i) \left(\frac{\ln n}{n} \right)^{-1/2} &= \liminf_{n \rightarrow \infty} \int (\theta \cdot \mathbf{1}^G - \mathbb{E}_\theta[R^i(\lambda_{x^n})]) \left(\frac{\ln n}{n} \right)^{-1/2} g(\theta) d\theta \\ &\geq \int (\mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G)^{1/2} g(\theta) d\theta = \mathbb{E}[(\mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G)^{1/2}], \end{aligned}$$

where the inequality holds by Fatou's Lemma, as $\liminf_{n \rightarrow \infty} (\theta \cdot \mathbf{1}^G - \mathbb{E}_\theta[R^i(\lambda_{x^n})]) \left(\frac{\ln n}{n}\right)^{-1/2} = (\mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G)^{1/2}$ for almost all $\theta \in \Theta$ by Proposition A.3. The same argument implies

$$\liminf_{n \rightarrow \infty} (R^{\text{FB}} - R_n^{\text{sep}}) \left(\frac{\ln n}{n} \right)^{-1/2} \geq \sum_{g \in G} \mathbb{E}[(\mathbf{1}^g \cdot I(\theta)^{-1} \mathbf{1}^g)^{1/2}].$$

Upper bound of revenue gap: Next, we show that $R^{\text{FB}} - R_n^{\text{SB}}$ and $R^{\text{FB}} - R_n^{\text{bd}}$ (resp. $R^{\text{FB}} - R_n^{\text{sep}}$) vanish no slower than at rate $\left(\frac{\ln n}{n}\right)^{1/2} \mathbb{E}[(\mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G)^{1/2}]$ (resp. $\left(\frac{\ln n}{n}\right)^{1/2} \sum_{g \in G} \mathbb{E}[(\mathbf{1}^g \cdot I(\theta)^{-1} \mathbf{1}^g)^{1/2}]$). For each x^n , consider the mixed bundling mechanism that prices each $B \subseteq G$ at 0 if $I_{x^n}(\hat{\theta}_{x^n})$ is not invertible, and otherwise at

$$p_{x^n}(B) := \hat{\theta}_{x^n} \cdot \mathbf{1}^B - \left(\frac{\ln n}{n} \mathbf{1}^B \cdot I_{x^n}(\hat{\theta}_{x^n})^{-1} \mathbf{1}^B \right)^{1/2}.$$

We first lower-bound the probability that the buyer's willingness-to-pay for B exceeds $p_{x^n}(B)$:

Lemma A.10. *There exists $M > 0$ such that for all $B \subseteq G$, $n \geq \bar{n}$, and $\theta \in \Theta$,*

$$\mathbb{P}_\theta[p_{x^n}(B) \leq \theta \cdot \mathbf{1}^B] \geq 1 - Mn^{-1/2}.$$

Proof. Let C denote the coefficient obtained in Lemma A.6. By Lemma A.6, there is $M_1 > 0$ such that $\left| \left(\mathbf{1}^B \cdot I_{x^n}(\hat{\theta}_{x^n})^{-1} \mathbf{1}^B \right)^{1/2} - \left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} \right| \leq \left(\frac{\ln n}{n} \right)^{1/2} M_1$ for all $B \subseteq G$, $\theta \in \Theta$, $n \geq \bar{n}$ and x^n with $\|I(\theta) - I_{x^n}(\hat{\theta}_{x^n})\| \leq \left(\frac{\ln n}{n} \right)^{1/2} C$.

Letting $E_n^B := \left\{ z \in \mathbb{R}^{|G|} : z \cdot \mathbf{1}^B > (\ln n)^{1/2} \left(\left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} + M_1 \left(\frac{\ln n}{n} \right)^{1/2} \right) \right\}$ for all $B \subseteq G$, $n \geq \bar{n}$, and $\theta \in \Theta$, we have

$$\begin{aligned}
& \mathbb{P}_\theta \left[\left(\hat{\theta}_{x^n} - \theta \right) \cdot \mathbf{1}^B > \left(\frac{\ln n}{n} \right)^{1/2} \left(\left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} + M_1 \left(\frac{\ln n}{n} \right)^{1/2} \right) \right] \\
& \leq N_{I(\theta)^{-1}}(E_n^B) + c'n^{-1/2} \\
& \leq \exp \left[- \frac{(\ln n) \left(\left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right) + 2M_1 \left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \frac{\ln n}{n} \right)^{1/2} + M_1^2 \frac{\ln n}{n} \right)}{2\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B} \right] + c'n^{-1/2} \\
& = n^{-1/2} n^{\frac{2M_1 \left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \frac{\ln n}{n} \right)^{1/2} + M_1^2 \frac{\ln n}{n}}{2\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B}} + c'n^{-1/2} \\
& = n^{-1/2} \exp \left[(\ln n) \frac{2M_1 \left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \frac{\ln n}{n} \right)^{1/2} + M_1^2 \frac{\ln n}{n}}{2\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B} \right] + c'n^{-1/2} \leq M_2 n^{-1/2},
\end{aligned}$$

where $M_2 > 0$ is chosen uniformly in θ . The first inequality uses Lemma A.5, the second inequality uses the Gaussian tail bound (9), and the last inequality uses the inequality $\ln n \leq \frac{n^\alpha}{\alpha}$, which holds for all $n > 0$ and $\alpha > 0$.

For every $B \subseteq G$, $n \geq \bar{n}$ and $\theta \in \Theta$, by definition of $p_{x^n}(B)$ and choice of M_1 , $\mathbb{P}_\theta[p_{x^n}(B) \leq \theta \cdot \mathbf{1}^B]$ is weakly greater than the probability

$$\begin{aligned}
& \mathbb{P}_\theta \left[\|I(\theta) - I_{x^n}(\hat{\theta}_{x^n})\| > \left(\frac{\ln n}{n} \right)^{1/2} C \text{ and} \right. \\
& \quad \left. \left(\hat{\theta}_{x^n} - \theta \right) \cdot \mathbf{1}^B \leq \left(\frac{\ln n}{n} \right)^{1/2} \left(\left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} + M_1 \left(\frac{\ln n}{n} \right)^{1/2} \right) \right],
\end{aligned}$$

which in turn is no less than

$$\begin{aligned}
& 1 - \mathbb{P}_\theta \left[\|I(\theta) - I_{x^n}(\hat{\theta}_{x^n})\| > \left(\frac{\ln n}{n} \right)^{1/2} C \right] \\
& - \mathbb{P}_\theta \left[\left(\hat{\theta}_{x^n} - \theta \right) \cdot \mathbf{1}^B > \left(\frac{\ln n}{n} \right)^{1/2} \left(\left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} + M_1 \left(\frac{\ln n}{n} \right)^{1/2} \right) \right] \\
& \geq 1 - (C + M_2)n^{-1/2},
\end{aligned}$$

where the inequality uses (13) from Lemma A.6. This establishes the desired claim with $M := C + M_2$. \square

Next, we approximate the revenue from offering each bundle B at price $p_{x^n}(B)$:

Lemma A.11. *There exists $M' > 0$ such that for all $B \subseteq G$, $n \geq \bar{n}$, and $\theta \in \Theta$,*

$$\left| \mathbb{E}_\theta[p_{x^n}(B)\mathbf{1}_{\{p_{x^n}(B) \leq \theta \cdot \mathbf{1}^B\}}] - \theta \cdot \mathbf{1}^B + \left(\frac{\ln n}{n} \mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} \right| \leq M' n^{-1/2}.$$

Proof. Let C' be the coefficient obtained in Lemma A.7. For all $B \subseteq G$, $n \geq \bar{n}$ and $\theta \in \Theta$, (14)–(15) from Lemma A.7 imply

$$\begin{aligned} & \left| \mathbb{E}_\theta[p_{x^n}(B) \mid A_{\theta,n}] - \theta \cdot \mathbf{1}^B + \left(\frac{\ln n}{n} \mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} \right| \\ & \leq \left| \mathbb{E}_\theta[\hat{\theta}_{x^n} \cdot \mathbf{1}^B \mid A_{\theta,n}] - \theta \cdot \mathbf{1}^B \right| \\ & \quad + \left(\frac{\ln n}{n} \right)^{1/2} \left| \mathbb{E}_\theta \left[\left(\mathbf{1}^B \cdot I_{x^n}(\hat{\theta}_{x^n})^{-1} \mathbf{1}^B \right)^{1/2} \mid A_{\theta,n} \right] - \left(\mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} \right| \leq 2C' n^{-1/2}. \end{aligned}$$

Since $\sup_{x \in X, \theta \in \text{int } \hat{\Theta}, g, g' \in G} \left| \frac{\partial^2 \ln f(x, \theta)}{\partial \theta_g \partial \theta_{g'}} \right| < \infty$ by Assumption 1.2 and $\hat{\Theta}$ is bounded, $\sup_{n \in \mathbb{N}, x^n \in X^n} |p_{x^n}(B)| < \infty$. Given this and Lemmas A.6 and A.10, $|\mathbb{E}_\theta[p_{x^n}(B) \mid A_{\theta,n}] - \mathbb{E}_\theta[p_{x^n}(B)\mathbf{1}_{\{p_{x^n}(B) \leq \theta \cdot \mathbf{1}^B\}}]|$ is bounded by $n^{1/2}$ times a constant that is uniform in θ . Thus, the desired bound follows from the above inequality. \square

To complete the proof of Theorem 1, denote by $\hat{R}_n^{\text{bd}} := \mathbb{E}[p_{x^n}(G)\mathbf{1}_{\{p_{x^n}(G) \leq \theta \cdot \mathbf{1}^G\}}]$ the expected revenue from the pure bundling mechanism that prices the grand bundle at $p_{x^n}(G)$ at each x_n (as in (6)). By definition, $\hat{R}_n^{\text{bd}} \leq R_n^{\text{bd}}$ for each n . By Lemma A.11 applied to $B = G$, the difference between \hat{R}_n^{bd} and $\mathbb{E} \left[\theta \cdot \mathbf{1}^G - \left(\frac{\ln n}{n} \mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G \right)^{1/2} \right]$ is bounded by $n^{-1/2}$ times a constant. This ensures that, for $i \in \{\text{SB}, \text{bd}\}$,

$$\mathbb{E} \left[\left(\mathbf{1}^G \cdot I(\theta)^{-1} \mathbf{1}^G \right)^{1/2} \right] \geq \limsup_{n \rightarrow \infty} (R^{\text{FB}} - \hat{R}_n^{\text{bd}}) \left(\frac{\ln n}{n} \right)^{-1/2} \geq \limsup_{n \rightarrow \infty} (R^{\text{FB}} - R_n^i) \left(\frac{\ln n}{n} \right)^{-1/2}.$$

Analogously, denote by \hat{R}_n^{sep} the expected revenue from the separate sales mech-

anism that prices each good g at $p_{x^n}(g)$ at each x^n . Then, by the same arguments,

$$\begin{aligned} \sum_{g \in G} \mathbb{E} [(\mathbf{1}^g \cdot I(\theta)^{-1} \mathbf{1}^g)^{1/2}] &\geq \limsup_{n \rightarrow \infty} (R^{\text{FB}} - \hat{R}_n^{\text{sep}}) \left(\frac{\ln n}{n} \right)^{-1/2} \\ &\geq \limsup_{n \rightarrow \infty} (R^{\text{FB}} - R_n^{\text{sep}}) \left(\frac{\ln n}{n} \right)^{-1/2}. \end{aligned}$$

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. (1996): “Multiproduct nonlinear pricing,” *Econometrica*, 64(1), 51–75.
- (1999): “Price discrimination by a many-product firm,” *Review of Economic Studies*, 66(1), 151–168.
- BABAIOFF, M., N. IMMORLICA, B. LUCIER, AND S. M. WEINBERG (2020): “A simple and approximately optimal mechanism for an additive buyer,” *Journal of the ACM (JACM)*, 67(4), 1–40.
- BAKOS, Y., AND E. BRYNJOLFSSON (1999): “Bundling information goods: Pricing, profits, and efficiency,” *Management Science*, 45(12), 1613–1630.
- BERGEMANN, D., A. BONATTI, A. HAUPT, AND A. SMOLIN (2021): “The optimality of upgrade pricing,” in *International Conference on Web and Internet Economics*, pp. 41–58. Springer.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The limits of price discrimination,” *American Economic Review*, 105(3), 921–957.
- BOLES LAVSKY, R., AND A. KOLB (2024): “The Design and Price of Influence,” .
- BÖRGERS, T. (2015): *An introduction to the theory of mechanism design*. Oxford University Press.
- BULOW, J., AND P. KLEMPERER (1996): “Auctions versus negotiations,” *American Economic Review*, 86(1), 180–194.
- CARROLL, G. (2017): “Robustness and separation in multidimensional screening,” *Econometrica*, 85(2), 453–488.
- CHE, Y.-K., AND W. ZHONG (2021): “Robustly-Optimal Mechanism for Selling Multiple Goods,” in *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 314–315.
- CRÉMER, J., AND R. P. MCLEAN (1988): “Full extraction of the surplus in Bayesian and dominant strategy auctions,” *Econometrica*, 56(6), 1247–1257.
- DASKALAKIS, C., A. DECKELBAUM, AND C. TZAMOS (2013): “Mechanism design via optimal transport,” in *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pp. 269–286.
- (2014): “The Complexity of Optimal Mechanism Design,” in *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1302–1318. SIAM.
- (2017): “Strong Duality for a Multiple-Good Monopolist,” *Econometrica*, 85(3), 735–767.
- DEB, R., AND A.-K. ROESLER (2023): “Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness,” *Review of Economic Studies*, 91(5).
- FANG, H., AND P. NORMAN (2006): “To bundle or not to bundle,” *RAND Journal of Economics*, 37(4), 946–963.
- FRICK, M., R. IJIMA, AND Y. ISHII (2023): “Learning Efficiency of Multiagent Information Structures,” *Journal of Political Economy*, 131(12), 3377–3414.
- (2024a): “Monitoring with rich data,” *arXiv preprint arXiv:2312.16789*.
- (2024b): “Welfare Comparisons for Biased Learning,” *American Economic Review*, 114(6), 1612–1649.

- GHILI, S. (2023): “A Characterization for Optimal Bundling of Products with Nonadditive Values,” *American Economic Review: Insights*, 5(3), 311–326.
- GONZAGA, P. (2018): “Personalised Pricing in the Digital Era,” OECD Discussion Paper.
- HAGHPANAH, N., AND J. HARTLINE (2021): “When is pure bundling optimal?,” *Review of Economic Studies*, 88(3), 1127–1156.
- HAGHPANAH, N., AND R. SIEGEL (2023): “Pareto-improving segmentation of multiproduct markets,” *Journal of Political Economy*, 131(6), 1546–1575.
- HANNAK, A., G. SOELLER, D. LAZER, A. MISLOVE, AND C. WILSON (2014): “Measuring price discrimination and steering on e-commerce web sites,” in *Proceedings of the 2014 Internet Measurement Conference*, pp. 305–318.
- HAREL, M., E. MOSSEL, P. STRACK, AND O. TAMUZ (2021): “Rational Groupthink,” *Quarterly Journal of Economics*, 136(1), 621–668.
- HART, S., AND N. NISAN (2017): “Approximate revenue maximization with multiple items,” *Journal of Economic Theory*, 172, 313–347.
- (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.
- HIPP, C., AND R. MICHEL (1976): “On the Bernstein-v. Mises approximation of posterior distributions,” *The Annals of Statistics*, 4(5), 972–980.
- MANELLI, A. M., AND D. R. VINCENT (2006): “Bundling as an optimal selling mechanism for a multiple-good monopolist,” *Journal of Economic Theory*, 127(1), 1–35.
- (2007): “Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly,” *Journal of Economic Theory*, 137(1), 153–185.
- McAFEE, R. P., J. McMILLAN, AND M. D. WHINSTON (1989): “Multiproduct monopoly, commodity bundling, and correlation of values,” *Quarterly Journal of Economics*, 104(2), 371–383.
- MOSCARINI, G., AND L. SMITH (2002): “The law of large demand for information,” *Econometrica*, 70(6), 2351–2366.
- PALFREY, T. R. (1983): “Bundling decisions by a multiproduct monopolist with incomplete information,” *Econometrica*, 51(2), 463–483.
- PFANZAGL, J. (1973): “The accuracy of the normal approximation for estimates of vector parameters,” *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 25(3), 171–198.
- RHODES, A., AND J. ZHOU (2024): “Personalized pricing and competition,” *American Economic Review*, 114(7), 2141–2170.
- ROCHET, J.-C., AND P. CHONÉ (1998): “Ironing, sweeping, and multidimensional screening,” *Econometrica*, 66(4), 783–826.
- ROUGHGARDEN, T., AND I. TALGAM-COHEN (2019): “Approximately optimal mechanism design,” *Annual Review of Economics*, 11, 355–381.
- RUSTICHINI, A., M. A. SATTERTHWAITTE, AND S. R. WILLIAMS (1994): “Convergence to efficiency in a simple market with incomplete information,” *Econometrica*, 62(5), 1041–1063.
- SCHMALENSEE, R. (1984): “Gaussian demand and commodity bundling,” *Journal of Business*, pp. S211–S230.
- SUGAYA, T., AND A. WOLITZKY (2024): “Non-Recursive Dynamic Incentives: A Rate of Convergence Approach,” working paper.
- WAINWRIGHT, M. J. (2019): *High-dimensional statistics: A non-asymptotic viewpoint*, vol. 48. Cambridge university press.
- WILSON, R. B. (1993): *Nonlinear pricing*. Oxford University Press.
- XIE, H., Y. ZHU, AND D. SHISHKIN (2024): “On the limitations of data-based price discrimination,” *Theoretical Economics*, forthcoming.
- YANG, F. (2023): “Nested Bundling,” Working Paper.

Online Appendix to “Multidimensional Screening with Rich Consumer Data”

Mira Frick, Ryota Iijima, and Yuhta Ishii

B Proof of Proposition 3

B.1 Preliminaries

For each $x \in [\underline{\beta}, \overline{\beta}]$, define

$$H(x) := \max_{q \in \Delta(\mathcal{B})} \alpha(0) \cdot q \text{ such that } \beta \cdot q = x,$$

$$\mathcal{H}(x) := \operatorname{argmax}_{q \in \Delta(\mathcal{B})} \alpha(0) \cdot q \text{ such that } \beta \cdot q = x.$$

Note that the constraint set $\{q \in \Delta(\mathcal{B}) : \beta \cdot q = x\}$ is nonempty for all $x \in [\underline{\beta}, \overline{\beta}]$. Furthermore, $0 \in [\underline{\beta}, \overline{\beta}]$ because $\beta_{\bar{i}} = 0$. Henceforth, we fix some arbitrary $b_0 \in \mathcal{H}(0)$, and we order the allocations such that $\underline{\beta} = \beta_1 \leq \dots \leq \beta_m = \overline{\beta}$.

Each mechanism (q, t) induces the agent indirect utility function given by $V(\tau) = \alpha(\tau) \cdot q(\tau) - t(\tau)$ for each τ . For any convex function $V : T \rightarrow \mathbb{R}$, the set of subgradients at τ is

$$\partial V(\tau) := \{\lambda \in \mathbb{R} : V(\tau') \geq V(\tau) + \lambda(\tau' - \tau) \forall \tau'\}.$$

It is well-known that (i) $\partial V(\tau) \neq \emptyset$; (ii) ∂V is a monotone correspondence, i.e., for all $\tau' > \tau$, $\inf \partial V(\tau') \geq \sup \partial V(\tau)$; and (iii) $\partial V(\tau)$ is a singleton for almost all $\tau \in T$. The last property implies that for any two functions h_1 and h_2 with $h_1(\tau), h_2(\tau) \in \partial V(\tau)$ for all τ and any measurable function φ , we have $\mathbb{E}[\varphi(h_1(\tau))] = \mathbb{E}[\varphi(h_2(\tau))]$. Thus, we write $\mathbb{E}[\varphi(\partial V(\tau))]$ for this expectation.

Let \mathcal{V} denote the set of all convex functions defined on T such that $V(\tau) \geq 0$ and $\partial V(\tau) \subseteq [\underline{\beta}, \overline{\beta}]$ for all τ . By standard arguments, any IC-IR mechanism (q, t) induces an indirect utility function $V \in \mathcal{V}$. In particular, its subgradient takes the form

$$\beta \cdot q(\tau) \in \partial V(\tau) \tag{18}$$

for all τ . However, in contrast to the single-good monopoly problem, in our setting there are typically many mechanisms that induce the same indirect utility function but may lead to different payoffs for the designer. Let $\mathcal{M}(V)$ be the set of all IC-IR mechanisms $(q, t) \in \mathcal{M}$ that induce the indirect utility function V , and let $\mathcal{M}^s(V)$ be the subset of such mechanisms that are in \mathcal{M}^s .

The following four preliminary lemmas are proved in Appendix B.3. First, we

derive the designer's optimal payoff among all mechanisms that induce the same indirect utility function. Moreover, we explicitly construct an optimal mechanism that induces the indirect utility function V :

Lemma B.1. *Consider any $V \in \mathcal{V}$. Then $\mathcal{M}(V) \neq \emptyset$ and*

$$\max_{(q,t) \in \mathcal{M}(V)} \mathbb{E}[t(\tau)] = \mathbb{E}[H(\partial V(\tau)) + \tau \partial V(\tau) - V(\tau)].$$

Moreover, for any mechanism (q, t) such that for all τ ,

1. $q(\tau) \in \mathcal{H}(\partial V(\tau)) := \bigcup_{x \in \partial V(\tau)} \mathcal{H}(x)$,
2. $t(\tau) = \alpha(\tau) \cdot q(\tau) - V(\tau)$,

we have $(q, t) \in \arg \max_{(q', t') \in \mathcal{M}(V)} \mathbb{E}[t'(\tau)]$.

Lemma B.1 immediately yields the following reformulation of the designer's problem, in terms of optimization with respect to V :

Corollary B.1. *We have*

$$\sup_{(q,t) \in \mathcal{M}} \mathbb{E}[t(\tau)] = \sup_{V \in \mathcal{V}} \mathbb{E}[H(\partial V(\tau)) + \tau \partial V(\tau) - V(\tau)].$$

The next lemma establishes a structural property of H that is useful for simplifying the optimization problem with respect to V :

Lemma B.2. *There exists a subsequence $(\beta_{\ell_i})_{i=1}^{k+1}$ of $(\beta_{\ell})_{\ell=1}^m$ with $\beta_{\ell_1} = \underline{\beta}$ and $\beta_{\ell_{k+1}} = \overline{\beta}$ such that*

1. H is linear on $[\beta_{\ell_i}, \beta_{\ell_{i+1}}]$ for every $i = 1, \dots, k$,
2. $\mathcal{H}(\beta_{\ell_i}) \cap (\mathcal{D} \cup \{b_0\}) \neq \emptyset$ for every $i = 1, \dots, k + 1$,
3. $\beta_{\ell_i} = 0$ for some $i = 1, \dots, k + 1$.

For the remainder of this section, we fix the sequence $(\beta_{\ell_i})_{i=1, \dots, k+1}$ derived in Lemma B.2. Let \mathcal{V}^s be the set of all $V \in \mathcal{V}$ such that $\partial V(\tau) \in \{\beta_{\ell_1}, \dots, \beta_{\ell_{k+1}}\}$ for almost all τ . The following lemma shows that the designer can restrict attention to indirect utility functions in \mathcal{V}^s (as long as the type space is compact):

Lemma B.3. *Assume T is compact. For every $V \in \mathcal{V}$, there exists $\tilde{V} \in \mathcal{V}^s$ such that*

$$\mathbb{E} \left[\left(H(\partial \tilde{V}(\tau)) + \tau \partial \tilde{V}(\tau) - \tilde{V}(\tau) \right) \right] \geq \mathbb{E} [H(\partial V(\tau)) + \tau \partial V(\tau) - V(\tau)].$$

The final lemma guarantees that any indirect utility function in \mathcal{V}^s can be induced by some mechanism in \mathcal{M}^s :

Lemma B.4. *For every $\tilde{V} \in \mathcal{V}^s$, there exists $(q, t) \in \mathcal{M}^s(\tilde{V})$ such that*

$$\mathbb{E}[t(\tau)] = \mathbb{E} \left[H(\partial \tilde{V}(\tau)) + \tau \partial \tilde{V}(\tau) - \tilde{V}(\tau) \right].$$

B.2 Completing the Proof of Proposition 3

First, assume that T is compact. Take any $(q, t) \in \mathcal{M}$ and let V denote the corresponding indirect utility. By Lemma B.3, there exists $\tilde{V} \in \mathcal{V}^s$ such that

$$\mathbb{E} \left[H(\partial\tilde{V}(\tau)) + \tau\partial\tilde{V}(\tau) - \tilde{V}(\tau) \right] \geq \mathbb{E} [(H(\partial V(\tau)) + \tau\partial V(\tau) - V(\tau))] \geq \mathbb{E}[t(\tau)]$$

where the second inequality uses Lemma B.1. By Lemma B.4, there exists some $(\tilde{q}, \tilde{t}) \in \mathcal{M}^s$ with $\mathbb{E}[\tilde{t}(\tau)] = \mathbb{E} \left[H(\partial\tilde{V}(\tau)) + \tau\partial\tilde{V}(\tau) - \tilde{V}(\tau) \right]$ that delivers a weakly higher payoff than (q, t) . This proves the result when T is compact.

To prove the result for a general interval $T \subseteq \mathbb{R}$, we want to show that, for any $(q, t) \in \mathcal{M}$ and $\varepsilon > 0$, there is $(q^*, t^*) \in \mathcal{M}^s$ such that $\mathbb{E}[t(\tau)] \leq \mathbb{E}[t^*(\tau)] + \varepsilon$. To this end, note that for any $\varepsilon > 0$, there is a compact interval $\tilde{T} \subseteq \mathbb{R}$ of types with $\mathbb{E}[\mathbf{1}_{\tau \in \tilde{T}} \max_{\ell} \alpha_{\ell}(\tau)] \leq \varepsilon/2$ because τ is L^1 . For any $(q, t) \in \mathcal{M}$, IR implies $t(\tau) \leq \max_{\ell} \alpha_{\ell}(\tau)$ for every τ , and hence $\mathbb{E}[t(\tau)] \leq \sup_{(\tilde{q}, \tilde{t}) \in \mathcal{M}} \mathbb{E}[\mathbf{1}_{\tau \in \tilde{T}} \tilde{t}(\tau)] + \varepsilon/2$. By the result for the compact case, there is $(q^*, t^*) \in \mathcal{M}^s$ such that $\mathbb{E}[\mathbf{1}_{\tau \in \tilde{T}} t^*(\tau)] + \varepsilon/2 \geq \sup_{(\tilde{q}, \tilde{t}) \in \mathcal{M}} \mathbb{E}[\mathbf{1}_{\tau \in \tilde{T}} \tilde{t}(\tau)]$, where $t^*(\cdot) \geq 0$. Thus,

$$\mathbb{E}[t(\tau)] \leq \mathbb{E}[\mathbf{1}_{\tau \in \tilde{T}} t^*(\tau)] + \varepsilon \leq \mathbb{E}[t^*(\tau)] + \varepsilon.$$

B.3 Proofs of Lemmas B.1–B.4

Proof of Lemma B.1. We first show that

$$\sup_{(q, t) \in \mathcal{M}(V)} \mathbb{E}[t(\tau)] \leq \mathbb{E}[H(\partial V(\tau)) + \tau\partial V(\tau) - V(\tau)].$$

To see this, consider any $(q, t) \in \mathcal{M}(V)$. By (18), $\beta \cdot q(\tau) \in \partial V(\tau)$ for all τ . Thus,

$$\begin{aligned} \mathbb{E}[t(\tau)] &= \mathbb{E}[\alpha(0) \cdot q(\tau) + (\beta \cdot q(\tau))\tau - V(\tau)] \leq \mathbb{E}[H(\beta \cdot q(\tau)) + (\beta \cdot q(\tau))\tau - V(\tau)] \\ &= \mathbb{E}[H(\partial V(\tau)) + \tau\partial V(\tau) - V(\tau)]. \end{aligned}$$

Now, consider any mechanism (q, t) such that for all τ ,

1. $q(\tau) \in \mathcal{H}(\partial V(\tau))$,
2. $t(\tau) = \alpha(\tau) \cdot q(\tau) - V(\tau)$.

To see that (q, t) is IC, note that type τ 's utility from misreporting to be type $\hat{\tau}$ is

$$\alpha(\tau) \cdot q(\hat{\tau}) - t(\hat{\tau}) = (\tau - \hat{\tau})(\beta \cdot q(\hat{\tau})) + V(\hat{\tau}) \leq V(\tau) = \alpha(\tau) \cdot q(\tau) - t(\tau),$$

where the inequality uses $\beta \cdot q(\hat{\tau}) = \partial V(\hat{\tau})$, which follows from $q(\hat{\tau}) \in \mathcal{H}(\partial V(\hat{\tau}))$. Thus, (q, t) is IC. Moreover, by construction it induces $V \geq 0$ as the indirect utility function. As a result, $(q, t) \in \mathcal{M}(V)$.

Finally, because again $q(\tau) \in \mathcal{H}(\partial V(\tau))$ for all τ ,

$$\mathbb{E}[t(\tau)] = \mathbb{E}[\alpha(0) \cdot q(\tau) + p(\beta \cdot q(\tau)) - V(\tau)] = \mathbb{E}[H(\partial V(\tau)) + \tau \partial V(\tau) - V(\tau)],$$

as claimed. \square

Proof of Lemma B.2. The following lemma implies Lemma B.2. This is because one can add 0 to the sequence $(\beta_{\ell_i})_{i=1}^{k+1}$ constructed in Lemma B.5, since we have $[\underline{\beta}, \bar{\beta}] \ni 0$ and $\mathcal{H}(0) \ni b_0$ by construction.

Lemma B.5. *There exists a subsequence $(\beta_{\ell_i})_{i=1}^{k+1}$ of $(\beta_{\ell})_{\ell=1}^m$ with $\beta_{\ell_1} = \underline{\beta}$ and $\beta_{\ell_{k+1}} = \bar{\beta}$ such that*

1. H is linear on $[\beta_{\ell_i}, \beta_{\ell_{i+1}}]$ for every $i = 1, \dots, k$; and
2. $\mathcal{H}(\beta_{\ell_i}) \cap \mathcal{D} \neq \emptyset$ for every $i = 1, \dots, k + 1$.

Proof of Lemma B.5. Consider the linear program associated with the value $H(x)$:

$$H(x) = \max_{y \in \Delta(\mathcal{B})} \alpha(0) \cdot y \text{ such that } \beta \cdot y = x.$$

By strong duality, the above is equal to the value of the dual program:

$$\begin{aligned} H(x) &= \min_{(z_1, z_2) \in \mathbb{R}^2} z_1 x + z_2 \text{ such that } z_1 \beta + z_2 \mathbf{1} \geq \alpha(0) \\ &= \min_{z \in \mathbb{R}} z x + \left(\max_{\ell} \alpha_{\ell}(0) - z \beta_{\ell} \right) \\ &= \min_{z \in \mathbb{R}} z x + G(z), \end{aligned}$$

where $G(z) := \max_{\ell} \alpha_{\ell}(0) - z \beta_{\ell}$.

Note that $G(z)$ is a piecewise linear function. That is, there exist $z_1 > \dots > z_k$ and $S = \{\ell_1, \dots, \ell_{k+1}\} \subseteq \{1, 2, \dots, m\}$ with $1 = \ell_1 < \ell_2 < \dots < \ell_{k+1} = m$ such that

$$G(z) = \begin{cases} \alpha_{\ell_1}(0) - z\beta_{\ell_1} & \text{if } z \in [z_1, +\infty), \\ \alpha_{\ell_i}(0) - z\beta_{\ell_i} & \text{if } z \in [z_i, z_{i-1}] \text{ for } i = 2, \dots, k, \\ \alpha_{\ell_{k+1}}(0) - z\beta_{\ell_{k+1}} & \text{if } z \in (-\infty, z_k]. \end{cases} \quad (19)$$

Observe that $\partial G(z_i) = [-\beta_{\ell_{i+1}}, -\beta_{\ell_i}]$ for all i . Thus, whenever $-x \in \partial G(z_i)$, which occurs when $x \in [\beta_{\ell_i}, \beta_{\ell_{i+1}}]$, we have

$$H(x) = z_i x + G(z_i). \quad (20)$$

This proves the first part of the lemma.

To prove the second part, suppose that $i = 1, \dots, k$. Then by (19) and (20),

$$H(\beta_{\ell_i}) = z_i \beta_{\ell_i} + G(z_i) = \alpha_{\ell_i}(0).$$

Thus, $\delta_{\ell_i} \in \mathcal{H}(\beta_{\ell_i})$, whence $\mathcal{D} \cap \mathcal{H}(\beta_{\ell_i}) \neq \emptyset$. If instead $i = k + 1$, then by (19)–(20),

$$H(\beta_{\ell_{k+1}}) = z_k \beta_{\ell_{k+1}} + G(z_k) = \alpha_{\ell_{k+1}}(0).$$

Again, $\delta_{\ell_{k+1}} \in \mathcal{H}(\beta_{\ell_{k+1}})$, and hence $\mathcal{D} \cap \mathcal{H}(\beta_{\ell_{k+1}}) \neq \emptyset$. □

Proof of Lemma B.3. Write $T = [\underline{\tau}, \bar{\tau}]$, and let (I_1, \dots, I_k) be any interval partition of T .³⁸ Let $M(I_1, \dots, I_k)$ be the set of non-decreasing functions $h : T \rightarrow [\underline{\beta}, \bar{\beta}]$ for which $h(I_i) \subseteq [\beta_{\ell_i}, \beta_{\ell_{i+1}}]$ for each $i = 1, \dots, k$. We endow $M(I_1, \dots, I_k)$ with the L^1 norm, making it a metric space. Let $S(I_1, \dots, I_k) \subseteq M(I_1, \dots, I_k)$ denote the set of all functions in $M(I_1, \dots, I_k)$ that are step functions taking values in $\{\beta_1, \dots, \beta_{k+1}\}$.

Claim 1. *The set $M(I_1, \dots, I_k)$ is compact and convex, with $\text{ext}(M(I_1, \dots, I_k)) = S(I_1, \dots, I_k)$.*

Proof of Claim 1. Clearly, $M(I_1, \dots, I_k)$ is convex. To prove compactness, it suffices to show that $M(I_1, \dots, I_k)$ is sequentially compact, since $M(I_1, \dots, I_k)$ is a metric space. To this end, consider a sequence $h_n \in M(I_1, \dots, I_k)$. By Helly's selection theorem, there exist a subsequence (n_k) and a function h such that $\lim_{j \rightarrow \infty} h_{n_j}(\tau) = h(\tau)$ for

³⁸We allow $I_j = \emptyset$ for some j .

all τ . Hence, we also have $h \in M(I_1, \dots, I_k)$. Moreover, by the Lebesgue dominated convergence theorem, $\int_{\underline{\tau}}^{\bar{\tau}} |h_{n_j}(\tau) - h(\tau)| d\tau = 0$. Thus, $M(I_1, \dots, I_k)$ is sequentially compact. The proof that $\text{ext}(M(I_1, \dots, I_k)) = S(I_1, \dots, I_k)$ follows from standard arguments (cf. Lemma 2.7 in [Börger, 2015](#)). \square

To complete the proof of Lemma B.3, consider any function $V \in \mathcal{V}$. Let $p^* = \min(\arg \min_{\tau \in T} V(\tau))$. Consider any monotone function $\varphi : T \rightarrow \mathbb{R}$ such that $\varphi(\tau) \in \partial V(\tau)$ for all $\tau \in T$. Define the interval partition I_1, I_2, \dots, I_k by

$$I_1 = \varphi^{-1}([\beta_{\ell_1}, \beta_{\ell_2})), \quad \dots, \quad I_k = \varphi^{-1}([\beta_{\ell_k}, \beta_{\ell_{k+1}}]).$$

Note the following two properties of I_1, \dots, I_k :

1. $\varphi \in M(I_1, \dots, I_k)$,
2. for all $h \in M(I_1, \dots, I_k)$, $h(\tau)(\tau - p^*) \geq 0$ for all τ .

To see that the second property holds, note that $\beta_i = 0$ for some i . Thus, p^* is the lower bound of interval I_i . Then any $h \in M(I_1, \dots, I_k)$ is decreasing at $\tau < p^*$ and weakly increasing at $\tau \geq p^*$.

Next, define the functional $T : M(I_1, \dots, I_k) \rightarrow \mathbb{R}$ by

$$T(h) = \mathbb{E} \left[\left(H(h(\tau)) + \tau h(\tau) - \int_{p^*}^{\tau} h(\tau') d\tau' \right) \right].$$

By the dominated convergence theorem, T is continuous. Moreover, T is linear, since H is linear on $[\beta_i, \beta_{i+1}]$ for every $i = 1, \dots, k$ by Proposition B.2. Thus,

$$\begin{aligned} & \mathbb{E} [\mathbf{1}_{\{\tau \in I\}} (H(\partial V(\tau)) + \tau \partial V(\tau) - V(\tau))] = \mathbb{E} [\mathbf{1}_{\{\tau \in I\}} (H(\varphi(\tau)) + \tau \varphi(\tau) - V(\tau))] \\ & = T(\varphi) - V(p^*) \leq \sup_{h \in M(I_1, \dots, I_k)} T(h) - V(p^*) = \max_{h \in S(I_1, \dots, I_k)} T(h) - V(p^*), \end{aligned}$$

where the last equality follows from Claim 1 and Bauer's maximum principle.

Thus, there exists some $\tilde{h} \in S(I_1, \dots, I_k)$ for which

$$T(\tilde{h}) - V(p^*) \geq \mathbb{E} [(H(\partial V(\tau)) + \tau \partial V(\tau) - V(\tau))].$$

Define \tilde{V} by $\tilde{V}(\tau) = \int_{p^*}^{\tau} \tilde{h}(\tau') d\tau'$. Since $h \in S(I_1, \dots, I_k)$, $\partial \tilde{V}(\tau) \in \{\beta_{\ell_1}, \dots, \beta_{\ell_{k+1}}\}$ for almost all $\tau \in T$. Moreover, because $\tilde{h} \in M(I_1, \dots, I_k)$, we have $\tilde{h}(\tau)(\tau - p^*) \geq 0$

for all $\tau \in T$. Thus,

$$\min_{\tau \in T} \tilde{V}(\tau) = \tilde{V}(p^*) = 0.$$

Together, this implies $\tilde{V} \in \mathcal{V}^s$. Hence,

$$\begin{aligned} \mathbb{E} \left[(H(\partial \tilde{V}(\tau)) + \tau \partial \tilde{V}(\tau) - \tilde{V}(\tau)) \right] &= \mathbb{E} \left[(H(\tilde{h}(\tau)) + \tau \tilde{h}(\tau) - \int_{p^*}^{\tau} \tilde{h}(\tau') d\tau') \right] \\ &\geq T(\tilde{h}) - V(p^*) \geq \mathbb{E} [(H(\partial V(\tau)) + \tau \partial V(\tau) - V(\tau))], \end{aligned}$$

as claimed. \square

Proof of Lemma B.4. Let q be an allocation rule with $q(\tau) \in \mathcal{H}(\partial \tilde{V}(\tau))$ for all τ . Since by Lemma B.2 $\partial \tilde{V}(\tau) \cap \{\beta_{\ell_1}, \dots, \beta_{\ell_{k+1}}\} \neq \emptyset$ for all τ , we can choose q such that $q(\tau) \in \{b_0\} \cup \mathcal{D}$ for all τ . Define transfers $t(\tau) := \alpha(\tau) \cdot q(\tau) - \tilde{V}(\tau)$. By Lemma B.1, $(q, t) \in \arg \max_{(q', t') \in \mathcal{M}(\tilde{V})} \mathbb{E} [t'(\tau)]$ and $\mathbb{E} [t(\tau)] = \mathbb{E} [H(\partial \tilde{V}(\tau)) + \tau \partial \tilde{V}(\tau) - \tilde{V}(\tau)]$. Moreover, by construction, $q(\tau) \subseteq \mathcal{D} \cup \{b_0\}$. Hence, $(q, t) \in \mathcal{M}^s(\tilde{V})$. \square

C Details for Section 6

C.1 Production Costs and Negative Types

In our main model, the first-best $R^{\text{FB}} = \mathbb{E}[\mathbf{1}^G \cdot \theta]$ involves supplying the grand bundle to all buyer types. This can fail if the seller faces a production cost or some types have negative valuations for some goods. We now extend the analysis to allow for both possibilities: First, to produce each bundle B , the seller incurs a cost $c(B)$, so that $R^{\text{FB}} = \mathbb{E} [\mathbf{1}^{B_\theta} \cdot \theta - c(B_\theta)]$ involves supplying a possibly different bundle $B_\theta \in \arg \max_{B \subseteq G} (\mathbf{1}^B \cdot \theta - c(B))$ to each type θ . Second, we drop the assumption that $\Theta \subseteq \mathbb{R}_{++}^{|G|}$ and instead assume $\mathbf{1}^{B_\theta} \cdot \theta - c(B_\theta) > 0$ for all $\theta \in \Theta$, which ensures positive gains from trade.

In this setting, a simple generalization of pure bundling—*single-bundle mechanisms*—achieves the optimal convergence rate to the first-best: For each signal sequence x^n , let $R_{x^n}^{\text{1bd}}$ denote the expected revenue when, conditional on x^n , the seller optimally chooses a single bundle B (possibly a strict subset of G) and price $p(B)$ at which to offer this bundle. Let $R_n^{\text{1bd}} := \mathbb{E}[R_{x^n}^{\text{1bd}}]$.

Theorem C.1. *Under the second-best mechanism and optimal single-bundle mechanism, the revenue gap to the first-best vanishes equally fast: Letting $\lambda := \mathbb{E}[\lambda^{B(\theta)}(\theta)]$,*

$$R^{\text{FB}} - R_n^{\text{SB}} \sim R^{\text{FB}} - R_n^{\text{1bd}} \sim \lambda \sqrt{\frac{\ln n}{n}}.$$

C.1.1 Proof of Theorem C.1

The proof proceeds analogously to Theorem 1. Denote by Θ^* the set of types $\theta \in \text{int } \Theta$ such that $\text{argmax}_{B \subseteq G} \theta \cdot \mathbf{1}^B - c(B)$ is a singleton. Types are in Θ^* with probability 1. The analysis of the one-dimensional environment in Appendix B extends unchanged; that is, optimal mechanisms only use deterministic bundles and one specific random bundle.

We extend the analysis of the Gaussian environment in Appendix A.1. Denote by $R^i(\theta^*, \frac{1}{n}J)$ ($i \in \{\text{SB}, \text{1bd}\}$) the seller's (second-best and single-bundling) payoffs under the belief $\mathcal{N}(\theta^*, \frac{1}{n}J)$. An analogous argument as in Proposition A.2 yields a $K > 0$ such that for all $\theta^* \in \Theta^*$, $J \in \mathcal{J}$, $n \in \mathbb{N}$, and $i \in \{\text{SB}, \text{1bd}\}$,

$$\left| R^i\left(\theta^*, \frac{1}{n}J\right) - \theta^* \cdot \mathbf{1}^{B(\theta^*)} + c(B(\theta^*)) + \left(\frac{\ln n}{n} \mathbf{1}^{B(\theta^*)} \cdot J \mathbf{1}^{B(\theta^*)}\right)^{1/2} \right| \leq Kn^{-1/2}.$$

Combined with the Gaussian approximation in Appendix A.2.1, we obtain the following extension of Proposition A.3 that characterizes the convergence rate of the seller's payoff conditional on each type: For each $\theta \in \Theta^*$ and $i \in \{\text{SB}, \text{1bd}\}$,

$$\sum_{g \in B(\theta)} \theta_g - c(B(\theta)) - \mathbb{E}_\theta[R^i(\lambda_{x^n})] = \left(\frac{\ln n}{n}\right)^{1/2} \lambda^{B(\theta)}(\theta) + o\left(\left(\frac{\ln n}{n}\right)^{1/2}\right).$$

This approximation, together with Fatou's Lemma, implies that both $R^{\text{FB}} - R_n^{\text{SB}}$ and $R^{\text{FB}} - R_n^{\text{1bd}}$ vanish at least as slow as $\lambda \sqrt{\frac{\ln n}{n}}$.

To show that $R^{\text{FB}} - R_n^{\text{SB}}$ and $R^{\text{FB}} - R_n^{\text{1bd}}$ vanish at least as fast as $\lambda \sqrt{\frac{\ln n}{n}}$, we extend the argument in Appendix A.2.3. Denote by \hat{R}_n^{1bd} the expected revenue when, following each realization x^n , the seller uses a single-bundle mechanism that offers the bundle B_{x^n} at price $p_{x^n}(B_{x^n})$, where

$$B_{x^n} \in \text{argmax}_{B \subseteq G} p_{x^n}(B) - c(B)$$

and $p_{x^n}(B)$ for each $B \subseteq G$ is as defined in Appendix A.2.3. Clearly, $R_n^{\text{1bd}} \geq \hat{R}_n^{\text{1bd}}$.

By the same argument as in Lemma A.11, there exists $M' > 0$ such that for all $n \geq \bar{n}$, $\theta \in \Theta$, and $B \subseteq G$,

$$\left| \mathbb{E}_\theta[(p_{x^n}(B) - c(B))\mathbf{1}_{\{p_{x^n}(B) \leq \theta \cdot \mathbf{1}^B\}}] - \theta \cdot \mathbf{1}^B + c(B) + \left(\frac{\ln n}{n} \mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} \right| \leq M' n^{-1/2}.$$

For every $n \geq \bar{n}$ and $\theta \in \Theta$, we have

$$\begin{aligned} & \mathbb{E}_\theta[(p_{x^n}(B_{x^n}) - c(B_{x^n}))\mathbf{1}_{\{p_{x^n}(B_{x^n}) \leq \theta \cdot \mathbf{1}^{B_{x^n}}\}}] \geq \mathbb{E}_\theta[p_{x^n}(B_{x^n}) - c(B_{x^n})] - M'' n^{-1/2} \\ & \geq \max_{B \subseteq G} \mathbb{E}_\theta[p_{x^n}(B) - c(B)] - M'' n^{-1/2} \\ & \geq \max_{B \subseteq G} \mathbb{E}_\theta[(p_{x^n}(B) - c(B))\mathbf{1}_{\{p_{x^n}(B) \leq \theta \cdot \mathbf{1}^B\}}] - M''' n^{-1/2} \\ & \geq \max_{B \subseteq G} \theta \cdot \mathbf{1}^B - c(B) - \left(\frac{\ln n}{n} \mathbf{1}^B \cdot I(\theta)^{-1} \mathbf{1}^B \right)^{1/2} - (M' + M''') n^{-1/2} \end{aligned}$$

for some $M'', M''' > 0$ that are chosen uniformly in θ , where the first and third inequalities use Lemma A.10, the second uses the definition of B_{x^n} , and the fourth uses the above extension of Lemma A.11.

Since $\hat{R}_n^{\text{1bd}} = \mathbb{E}[(p_{x^n}(B_{x^n}) - c(B_{x^n}))\mathbf{1}_{\{p_{x^n}(B_{x^n}) \leq \theta \cdot \mathbf{1}^{B_{x^n}}\}}]$, we can lower-bound $\hat{R}_n^{\text{1bd}} - \mathbb{E} \left[\theta \cdot \mathbf{1}^{B(\theta)} - c(B(\theta)) - \left(\frac{\ln n}{n} \mathbf{1}^{B(\theta)} \cdot I(\theta)^{-1} \mathbf{1}^{B(\theta)} \right)^{1/2} \right]$ by $n^{-1/2}$ times a constant. Then, as claimed, we have that for $i \in \{\text{SB}, \text{1bd}\}$,

$$\begin{aligned} \mathbb{E} \left[(\mathbf{1}^{B(\theta)} \cdot I(\theta)^{-1} \mathbf{1}^{B(\theta)})^{1/2} \right] & \geq \limsup_{n \rightarrow \infty} (R^{\text{FB}} - \hat{R}_n^{\text{1bd}}) \left(\frac{\ln n}{n} \right)^{-1/2} \\ & \geq \limsup_{n \rightarrow \infty} (R^{\text{FB}} - R_n^i) \left(\frac{\ln n}{n} \right)^{-1/2}. \end{aligned}$$

C.2 Non-Additive Utilities

We extend the main result to a setting where the buyer's payoffs are not additive across goods. Let $\mathcal{B} = 2^G$ be the space of all bundles with $m = |\mathcal{B}|$. The buyer's type is represented by a vector $\omega \in \mathbb{R}^{\mathcal{B}}$, where entry ω_B represents his valuation of bundle B . His utility from receiving bundle $B \in \mathcal{B}$ and paying transfer t is $\omega_B - t$. Type ω is drawn from a density g whose support is some compact set $\Omega \subseteq \mathbb{R}^{\mathcal{B}}$ with non-empty interior. We assume that $\omega_G > \max\{\omega_B, 0\}$ for all $B \subsetneq G$ and $\omega \in \Omega$. Thus, every buyer type finds the grand bundle G most attractive.

As in our main model, conditional on buyer type ω , the seller observes a sequence of n signals, $x^n = (x_1, \dots, x_n)$, that are drawn i.i.d. from a distribution $P_\omega \in \Delta(X)$ with density $f(\cdot, \omega)$. The seller then chooses a direct mechanism, described by mappings $q : \Omega \rightarrow \Delta(2^G)$ and $t : \Omega \rightarrow \mathbb{R}$. Conditional on x^n , the second-best revenue is $R_{x^n}^{\text{SB}} := \sup_{q,t} \mathbb{E}[t(\omega)|x^n]$ subject to

$$\sum_{B \in \mathcal{B}} q(B; \omega) \omega_B - t(\omega) \geq \sum_{B \in \mathcal{B}} q(B; \omega') \omega_B - t(\omega'), \quad \forall \omega, \omega' \in \Omega, \quad (\text{IC})$$

$$\sum_{B \in \mathcal{B}} q(B; \omega) \omega_B - t(\omega) \geq 0, \quad \forall \omega \in \Omega. \quad (\text{IR})$$

Let $R_n^{\text{SB}} := \mathbb{E}[R_{x^n}^{\text{SB}}]$ and $R^{\text{FB}} := \mathbb{E}[\omega_G]$. Let $R_{x^n}^{\text{bd}}$ denote the seller's optimal expected revenue at x^n when she is restricted to using (pure) bundling. Let $R_n^{\text{bd}} := \mathbb{E}[R_{x^n}^{\text{bd}}]$.³⁹

We impose an analog of Assumption 1. As before, we extend the signal distribution P_ω and density $f(\cdot, \omega)$ to all types in some compact neighborhood $\hat{\Omega} \supseteq \Omega$. Here, we pick $\hat{\Omega}$ sufficiently close to Ω so that $\omega_G > \max\{\omega_B, 0\}$ holds for all $\omega \in \hat{\Omega}$ and $B \subsetneq G$.

Assumption C.1.

1. The prior density g is strictly positive and locally Lipschitz continuous for all $\omega \in \Omega$.
2. For all $x \in X$, the signal densities $f(x, \omega)$ are strictly positive and C^2 in $\omega \in \text{int } \hat{\Omega}$, and there is $L > 0$ such that, for all $B, B' \in \mathcal{B}$, $\omega, \omega' \in \text{int } \hat{\Omega}$, and $x \in X$, $\left| \frac{\partial^2 \ln f(x, \omega)}{\partial \omega_B \partial \omega_{B'}} \right| \leq L$ and $\left| \frac{\partial^2 \ln f(x, \omega)}{\partial \omega_B \partial \omega_{B'}} - \frac{\partial^2 \ln f(x, \omega')}{\partial \omega_B \partial \omega_{B'}} \right| \leq L \|\omega' - \omega\|$.
3. The signal distributions P_ω are continuous in $\omega \in \hat{\Omega}$ with respect to the total variation distance.
4. We have $\sup_{\omega \in \Omega} \int (\sup_{\omega' \in \Omega} \ln f(x, \omega'))^2 dP_\omega(x) < \infty$.
5. The Fisher information matrix $I(\omega) \in \mathbb{R}^{\mathcal{B} \times \mathcal{B}}$, given by

$$I(\omega) := \left(\int -\frac{\partial^2}{\partial \omega_B \partial \omega_{B'}} \ln f(x, \omega) dP_\omega(x) \right)_{B, B' \in \mathcal{B}},$$

is well-defined and positive definite for each $\omega \in \hat{\Omega}$.

³⁹We do not analyze separate sales, which in the current setting no longer reduces to single-good monopoly, as the optimal price of each good g may depend on the prices of other goods. However, Theorem 1 already shows that it cannot in general achieve the optimal rate of convergence.

The following result extends Theorem 1. Analogously to Section 3.1, conditional on each buyer type ω^* , the seller's posterior standard deviation about ω_G is approximated by $\frac{\lambda_{\omega^*}}{\sqrt{n}}$ at large n , where $\lambda_{\omega^*} := \sqrt{\mathbf{1}^G \cdot I(\omega^*)^{-1} \mathbf{1}^G}$. Let $\lambda := \mathbb{E}[\lambda(\omega)]$.

Theorem C.2. *Under the second-best mechanism and optimal bundling mechanism, the revenue gap to the first-best revenue vanishes equally fast:*

$$R^{\text{FB}} - R_n^{\text{SB}} \sim R^{\text{FB}} - R_n^{\text{bd}} \sim \lambda \sqrt{\frac{\ln n}{n}}.$$

We omit the proof, which is analogous to that of Theorem 1. In particular, as in Appendix A.2, the seller's posterior at large n can be approximated by a Gaussian distribution whose mean is $\hat{\omega}_{x^n} \in \operatorname{argmax}_{\omega \in \hat{\Omega}} \sum_{i=1}^n \ln f(x_i, \omega)$ and whose covariance matrix is the inverse of

$$I_{x^n}(\omega) := \left(-\frac{1}{n} \sum_{i=1}^n \frac{\partial^2}{\partial \omega_B \partial \omega_{B'}} \ln f(x_i, \omega) \right)_{B, B' \in \mathcal{B}}$$

evaluated at $\omega = \hat{\omega}_{x^n}$. To upper-bound R_n^{SB} , we make use of a one-dimensional relaxed problem as in Appendix A.1, which did not assume additive values.

C.3 Deterministic Setting with Non-Gaussian Beliefs

Extending the deterministic Gaussian environment studied in the main text, we consider a setting where the seller has a general deterministic sequence of beliefs indexed by n that converge to a point-mass on the true type θ^* . Such beliefs can be interpreted as the outcome of the seller observing increasingly precise information that need not take the form of n (i.i.d. or correlated) signal draws.

Formally, conditional on true type $\theta^* \in \mathbb{R}_{++}^{|G|}$, the seller's belief at n is that the buyer's type is $\theta = \theta^* + n^{-1/2}z$, where $z \in \mathbb{R}^{|G|}$ is drawn from a cdf F that is continuously differentiable and positive. For each $B \subseteq G$, let F^B denote the induced cdf of $\mathbf{1}^B \cdot z$. We impose the following tail conditions: There exist $\alpha^+(B), \alpha^-(B), \beta^+(B), \beta^-(B) > 0$ such that, as $z \rightarrow \infty$,

$$\begin{aligned} 1 - F^B(z) &= \exp[-\alpha^+(B)z^{\beta^+(B)} + o(z^{\beta^+(B)})], \\ \text{and } F^B(-z) &= \exp[-\alpha^-(B)(-z)^{\beta^-(B)} + o(z^{\beta^-(B)})]. \end{aligned}$$

These conditions rule out heavy-tailed distributions (i.e., whose moment-generating function $\int \exp[tz]dF^B(z)$ is infinite at all t). The Gaussian setting in the main text corresponds to the case $\alpha^+(B) = \alpha^-(B) = \frac{1}{2(\lambda^B(\theta^*))^2}$ and $\beta^+(B) = \beta^-(B) = 2$, based on the Gaussian tail bound in (9).

We restrict attention to deterministic mechanisms. Denote by $R_n^{\text{mix}}(\theta^*)$ the seller's optimal expected revenue under IC-IR deterministic mechanisms (which can be represented as mixed bundling). Denote by $R_n^{\text{bd}}(\theta^*)$ the optimal revenue under (pure) bundling. Denote by $R^{\text{FB}}(\theta^*) := \sum_{g \in G} \theta_g^*$ the first-best revenue.

Proposition C.1. *Under the optimal deterministic mechanism and optimal bundling mechanism, the revenue gap to the first-best vanishes equally fast:*

$$R^{\text{FB}}(\theta^*) - R_n^{\text{mix}}(\theta^*) \sim R^{\text{FB}}(\theta^*) - R_n^{\text{bd}}(\theta^*) \sim \left(\frac{\ln n}{2\alpha^-(G)} \right)^{1/\beta^-(G)} n^{-1/2}.$$

Thus, optimizing over pure bundling achieves the same convergence rate to the first-best as optimizing over all possible deterministic mechanisms. This generalizes our main insight that pure bundling is an effective mechanism under precise seller beliefs, as using menus of bundles only allows for a negligible improvement relative to pure bundling. The main qualification is that, unlike Theorem 1, Proposition C.1 does not allow for stochastic mechanisms, whose analysis we leave as an open question.

Appendix C.3.1 first analyzes the single-good case. Appendix C.3.2 then proves Proposition C.1.

C.3.1 Single-Good Case

We extend the intensive vs. extensive margin analysis from the Gaussian environment in Section 4.2 to the current setting. Assume $|G| = 1$, so we can omit the dependence of parameters $\alpha^+, \alpha^-, \beta^+, \beta^-$ on B . Denote by $R_n^*(\theta)$ the optimal revenue, which is achieved by a posted price p_n^* . Let F_n denote the cdf of the seller's belief, i.e., $F_n(\theta) = F((\theta - \theta^*)\sqrt{n})$.

Generalizing Proposition 2, the following result shows that at large n , the seller optimally sets prices in such a way that revenue losses relative to the first-best are driven by the intensive margin, while the extensive margin becomes negligible. Write $\alpha = \alpha^-, \beta = \beta^-$.

Proposition C.2. *Any optimal price sequence p_n^* satisfies*

$$R^{\text{FB}}(\theta^*) - R_n^* \sim \theta^* - p_n^* \sim \left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2} \quad \text{and} \quad \theta^* F_n(p_n^*) = o\left((\ln n)^{1/\beta} n^{-1/2}\right).$$

Moreover, under any price sequence p_n with $\theta^* - p_n \sim \delta \left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}$ for some $\delta \in [0, 1)$, we have $\lim_n \frac{\theta^* F_n(p_n)}{(\ln n)^{1/\beta} n^{-1/2}} = \infty$.

Proof. Take any $\delta \geq 0$ and sequence of prices (p_n) with $\theta^* - p_n \sim \delta \left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}$. Then by the tail conditions on F ,

$$F_n(p_n) = \exp\left[-\frac{\delta^\beta}{2} \ln n + o(\ln n)\right] = n^{-\frac{\delta^\beta}{2} + o(1)},$$

and thus

$$\lim_{n \rightarrow \infty} \frac{\theta^* F_n(p_n)}{(\ln n)^{1/\beta} n^{-1/2}} = \begin{cases} 0 & \text{if } \delta > 1, \\ \infty & \text{if } \delta < 1. \end{cases}$$

Additionally, if $\theta^* - p_n = \left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}$ for all n , then $\lim_n \frac{\theta^* F_n(p_n)}{(\ln n)^{1/\beta} n^{-1/2}} = 0$, and hence $R^{\text{FB}}(\theta^*) - R_n(p_n) \sim \left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}$, where $R_n(p_n) := p_n(1 - F_n(p_n))$.

Given these observations, the remainder of the proof is analogous to that of Proposition 2. \square

C.3.2 Proof of Proposition C.1

Let $(p_n(B))_{B \subseteq G}$ denote optimal mixed-bundling prices that yield $R_n^{\text{mix}}(\theta^*)$, where $p_n(\emptyset) = 0$. Observe first that for any $B \subseteq G$ with $\lim_{n \rightarrow \infty} p_n(B) > \mathbf{1}^B \cdot \theta^*$,

$$\begin{aligned} & \text{Prob}[\mathbf{1}^B \cdot \theta - p_n(B) \geq \mathbf{1}^{B'} \cdot \theta - p_n(B'), \forall B' \subseteq G] \leq \text{Prob}[\mathbf{1}^B \cdot \theta - p_n(B) \geq 0] \\ & = 1 - F^B(n^{1/2}(p_n(B) - \mathbf{1}^B \cdot \theta^*)) \\ & = \exp[-\alpha^+(B)n^{\frac{\beta^+(B)}{2}} + o(n^{\frac{\beta^+(B)}{2}})] = o(n^{-1}), \end{aligned} \tag{21}$$

where the second equality uses the assumption on right-tail probabilities.

Write $\alpha = \alpha^-(G)$ and $\beta = \beta^-(G)$. To prove the proposition, note that $R^{\text{FB}}(\theta^*) - R_n^{\text{bd}}(\theta^*) \sim \left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}$ follows from Proposition C.2. Suppose toward a contradiction that $R^{\text{FB}}(\theta^*) - R_n^{\text{mix}}(\theta^*) \sim \left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}$ fails. Since $R^{\text{FB}}(\theta^*) - R_n^{\text{mix}}(\theta^*) \leq$

$R^{\text{FB}}(\theta^*) - R_n^{\text{bd}}(\theta^*)$ holds for each n by definition, we have

$$\liminf_{n \rightarrow \infty} \frac{R^{\text{FB}}(\theta^*) - R_n^{\text{mix}}(\theta^*)}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} < 1. \quad (22)$$

Observe by definition of $R_n^{\text{mix}}(\theta^*)$ that

$$R^{\text{FB}}(\theta^*) - R_n^{\text{mix}}(\theta^*) = \sum_{B \subseteq G} \text{Prob}[\mathbf{1}^B \cdot \theta - p_n(B) \geq \mathbf{1}^{B'} \cdot \theta - p_n(B'), \forall B' \subseteq G] (\mathbf{1}^G \cdot \theta^* - p_n(B)). \quad (23)$$

Consider the case $\liminf_{n \rightarrow \infty} \frac{\mathbf{1}^G \cdot \theta^* - p_n(G)}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} \geq 1$. Given (22)-(23), there exists $B \subsetneq G$ such that $\liminf_{n \rightarrow \infty} \text{Prob}[\mathbf{1}^B \cdot \theta - p_n(B) \geq \mathbf{1}^{B'} \cdot \theta - p_n(B'), \forall B' \subseteq G] > 0$ and $\liminf_{n \rightarrow \infty} \frac{\mathbf{1}^G \cdot \theta^* - p_n(B)}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} < 1$. This contradicts (21).

Consider the case $\liminf_{n \rightarrow \infty} \frac{\mathbf{1}^G \cdot \theta^* - p_n(G)}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} < 1$. Then

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \frac{\sum_{B \subsetneq G} \text{Prob}[\mathbf{1}^B \cdot \theta - p_n(B) \geq \mathbf{1}^{B'} \cdot \theta - p_n(B'), \forall B' \subseteq G]}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} \\ & \geq \liminf_{n \rightarrow \infty} \frac{\text{Prob}[\mathbf{1}^G \cdot \theta - p_n(G) < 0]}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} = \infty, \end{aligned}$$

where the equality uses the second part of Proposition C.2. But given (22)-(23), there exists $B \subsetneq G$ such that $\liminf_{n \rightarrow \infty} \frac{\text{Prob}[\mathbf{1}^B \cdot \theta - p_n(B) \geq \mathbf{1}^{B'} \cdot \theta - p_n(B'), \forall B' \subseteq G]}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} > 0$ and $\liminf_{n \rightarrow \infty} \frac{\mathbf{1}^G \cdot \theta^* - p_n(B)}{\left(\frac{\ln n}{2\alpha}\right)^{1/\beta} n^{-1/2}} < 1$. This contradicts (21).