MULTIDIMENSIONAL SCREENING WITH RETURNS

ALEXANDER HABERMAN, RAVI JAGADEESAN, AND FRANK YANG

Abstract. A monopolist wants to sell multiple goods, accounting for the possibility of returns. For any bundle purchased, a buyer can return any part of the bundle for a partial refund specified by a *return policy*. The seller is constrained to offer *additive* return policies: the total price of a bundle must be divided into return values for each of the goods in the bundle. We show that if the buyer's values are additive and independent across goods, then selling each good separately is optimal. The result applies even if the seller can use stochastic mechanisms, and holds for general multidimensional screening problems. We show that selling separately remains approximately optimal if the buyer experiences small return costs, as well as if the return policy only has to be close to additive or the buyer's values are weakly correlated. Applying our analysis to multiproduct pricing without returns, we introduce a new concept of bundle discounts and show that obtaining revenue gains from using any complex mechanism requires discounting bundles in our sense.

Date: January 9, 2025.
We thank Ben Brooks, Eric Budish, Gabriel Carroll, Piotr Dworczak, Mira Frick, Yannai Gonczarowski, Ben Golub, Nima Haghpanah, Sergiu Hart, Jason Hartline, Ryota Iijima, Jacob Leshno, Hao Li, Shengwu Li, Xiao Lin, Debraj Ray, Phil Reny, Ilya Segal, Ludvig Sinander, Alex Teytelboym, Daniel Waldinger, and several seminar audiences for helpful discussions.

Haberman: Department of Economics, Stanford University. Email: ahabermn@stanford.edu.
Jagadeesan: Department of Economics, Stanford University. Email: ravi.jagadeesan@gmail.com.
Yang: Becker Friedman Institute, University of Chicago. Email: franky@uchicago.edu.

1. Introduction

How should a seller of multiple products set prices? With a single product, the optimal mechanism simply sets a price that equates marginal revenue with marginal cost (Myerson 1981; Riley and Zeckhauser 1983; Bulow and Roberts 1989).

With multiple products, the pricing problem is much less well-behaved. For example, suppose that buyers' values are additive and statistically independent across two products. Despite the independence of demand for the products, it is generically suboptimal to sell the products separately (McAfee, McMillan, and Whinston 1989), and the optimal mechanism is generally unknown. The optimal mechanism can require allocating the products randomly (Manelli and Vincent 2006) and may require infinitely many such lotteries (Daskalakis, Deckelbaum, and Tzamos 2013).

This paper studies optimal multiproduct pricing when buyers can *return* any of the goods that they purchased. Returns have become a common feature of retail markets (National Retail Federation 2023),[1] and sellers are sometimes mandated by regulators to accept returns.[2] Our results show that accounting for the possibility of returns dramatically simplifies optimal multiproduct pricing by substantially lowering the revenue that complex selling strategies obtain.

The force we identify is that if any part of a bundle can be returned for a *partial refund*, then the possibility of returns substantively disciplines the pricing problem. To illustrate why returns matter, suppose that apples and bananas are sold for \$20 each, and a bundle of an apple and a banana is sold for \$26. If the apple in the bundle can be returned for a partial refund of more than \$6 (while keeping the banana), then buyers have an incentive to buy the bundle and return the apple, rather than pay \$20 for the banana. By contrast, if the partial refunds are \$6 or less, then returns are not profitable. Thus, buyers' incentives to strategically exploit returns depend on the return policy. To isolate this strategic effect from other possible reasons for returns, we assume buyers know their values, as in standard multiproduct pricing models.[3]

While return policies are set by the seller, a common feature is that they are *additive* across products: the total price of the bundle is divided into *return values* for each item in the bundle, which determines the partial refunds. For example, Target's

[1]For example, retail returns in the United States totaled \$743 billion in 2023.

[2]For example, the European Union's Consumer Rights Directive 2011 guarantees that "[Consumers] can return any purchase within 14 days without justification." See `https://europa.eu/youreurope/citizens/consumers/shopping/guarantees-returns`.

[3]See Krähmer and Strausz (2015), Bergemann, Castro, and Weintraub (2020), and Heumann (2020) for analyses of models with a single good in which buyers dynamically learn about their values, and Haberman and Jagadeesan (2024) for an analysis of an auction problem with interdependent values.

return policy determines return values from item list prices by allocating discounts for bundles *proportionally* to item prices: "the promotional discount is [...] applied to each item proportionately. If you choose to return any of the items, your refund will be reduced to the discounted amount for each item."[4] A similar "consumer-friendly" return policy is shared by Macy's, Whole Foods, Amazon Devices, DSW, and Lowe's.[5]

To understand the implications of additive return policies for the revenue obtained from bundling, let us return to our example with an apple and a banana. For simplicity, suppose for the moment that the return values for the apple and banana in the bundle are the same. For the return policy to be additive, the return value for each product must be \$13. Then, buyers can profit from buying the bundle and returning one product to acquire the other for \$13. Hence, the possibility of returns can make bundle pricing lead to a loss of revenue from the sales of separate products.

How do such *strategic returns* affect optimal pricing? Ignoring the possibility of returns, the pricing scheme described above is nearly optimal if buyers' values are additive across products, and values for apples and bananas are each independently distributed uniformly between 0 and 30. (The optimal bundle price is $40 - 10\sqrt{2} \approx$ \$25.86 instead of \$26.)[6] However, if buyers can obtain partial refunds of \$13 for each product in the bundle, then taking returns into account, it is as if the products are sold separately for \$13 each. But revenue would be higher from simply selling each product at its single-product monopoly price of \$15—in fact, as we show, this is the optimal pricing scheme once the seller is constrained to offer additive return policies. As a consequence, due to the seller's reoptimization in pricing, total consumer surplus in this example is actually lowered by imposing consumer-friendly return policies.[7]

We introduce the joint design problem for pricing schemes and return policies. The key constraint that we impose for much of our analysis is that the return policy for each bundle must be additive. Our main result (Theorem 1) shows that when buyers' values are additive and statistically independent across goods, it is optimal to sell goods separately at their single-product monopoly prices. This result applies even if the seller can monitor buyers' purchases and use arbitrary stochastic mechanisms.

[4]See the answer to "Will I get a refund for an item purchased as part of a BOGO promotion?" in Target's Return FAQs at `http://help.target.com/help/subcategoryarticle?childcat=Returns&parentcat=Returns+%26+Exchanges`.

[5]Another common additive return policy applies when there is a discount of the type of "free shipping if spend more than \$$X$", where returning any goods does not make a buyer forfeit the shipping discount (which is used, e.g., by Walmart, Bloomingdale's, Nike, H&M, and Uniqlo).

[6]This pricing scheme is optimal even if randomization is allowed (Manelli and Vincent 2006).

[7]In this example, the consumer surplus under separate sales is \$7.5, while consumer surplus under the optimal selling mechanism without returns is ≈ \$7.65.

Moreover, as we will discuss, our framework uncovers a general separation result for multidimensional screening problems beyond monopoly pricing (see Theorem 3).

Methodologically, our analysis shows how to incorporate a form of moral hazard to gain tractability in multidimensional screening models. Our results also have more literal interpretations that we discuss in the conclusion. For example, our main result implies that if sellers take the moral hazard problem of returns seriously, then they must decide between potentially increasing profit by using various bundling mechanisms and offering consumer-friendly return policies such as pro-rata refunds.

A simplistic intuition behind our main result can be understood as follows. Consider the running example of selling an apple and a banana, and restrict attention to deterministic mechanisms. Let p_B be the price of the bundle, and let r_1 and r_2 be the return values for the apple and the banana. By buying the bundle and returning the banana, buyers can obtain the apple for $\tilde{p}_1 = p_B - r_2$. Similarly, buyers can obtain the banana for $\tilde{p}_2 = p_B - r_1$. Depending on the pricing scheme, it is possible that the buyers can also obtain the individual products more cheaply. Letting p_1, p_2 denote the *effective prices* of the apple and banana obtained using optimal purchase and return strategies, if the return policy is additive, then we have that

$$p_1 + p_2 \leq \tilde{p}_1 + \tilde{p}_2 = (p_B - r_2) + (p_B - r_1) = 2p_B - (r_1 + r_2) = p_B.$$

Hence, under additive return policies, the effective prices that buyers face must be *superadditive* regardless of what the original prices might be. Intuitively, mixed bundling can increase profits relative to separate sales because marginally decreasing the bundle price generates a revenue gain from marginal buyers that outweighs the revenue loss from inframarginal buyers (McAfee, McMillan, and Whinston 1989). One might then argue that effectively superadditive pricing would instead require raising the bundle price, which would then lead to a net revenue loss.

However, this intuition is incomplete. With three goods, superadditive price schedules can yield strictly higher revenue than separate sales even with independent values (Example 1). Indeed, as we show, beyond the case of two goods, the concept of a bundle discount is not appropriately captured by the subadditivity/superaddivity of the price schedule. And even if we restrict attention the case of two goods, the argument of McAfee, McMillan, and Whinston (1989) only shows that decreasing the bundle price is *sufficient* to improve on separate sales. As we will discuss, one of our contributions is to formulate a general concept of bundle discounts that is applicable with arbitrary numbers of goods, and to show, conversely, that bundle discounts in

our sense are always *necessary* to improve on separate sales. The simplistic intuition described in the previous paragraph does become correct if one considers mechanisms without bundle discounts instead of superadditive pricing. But the necessity of bundle discounts to raise revenue is a consequence of our main results, not part of the proof.

Our proof instead works with direct revelation mechanisms that specify additive return policies and take into account buyers' optimal return strategies. First, we show that even under the restriction to additive return policies, it is without loss of generality to focus on mechanisms that are *return-proof*—i.e., for which the buyers optimally do not execute returns under known values. For any return-proof mechanisms with additive return policies, the seller's expected revenue is simply the sum of the expected return values of all goods. To complete the argument, the key idea is to show that the expected return value of each good i, averaged over all buyers, must be bounded above by the monopoly revenue of selling good i alone. Roughly speaking, we show that if the return value of good i is ever too high for some buyer type, then there would have to be a profitable "double deviation" by another type that involves *(i)* misreporting a higher value for good i and then *(ii)* returning that good.

Our analysis also yields new insights about multiproduct pricing without returns based on a concept of bundle discounts that we introduce. We say that a price mechanism *discounts* a bundle if the marginal prices of the items in the bundle sum to less than the price of the bundle. For example, in a "Buy One Get One Free" sale, the marginal price of the second item in a bundle is always zero, so the bundle is discounted. The absence of bundle discounts is logically independent of the superadditivity of the price schedule, but is implied by the supermodularity of the price schedule. We can apply our analysis of the model with returns to show that separate sales is optimal among all price mechanisms that have no bundle discounts (Theorem 2).[8] Moreover, we show that unlike the superadditivity or supermodularity of the price schedule, the structure of bundle discounts sharply delineates the forms of bundle pricing that can improve on the revenue of separate sales (Examples 1 and 2).

Our main result focuses on linear monopoly pricing, but the key logic behind it does not depend on a linear preference or cost structure. Consider a general screening problem (for example, multiproduct nonlinear pricing) where the principal and agent have quasilinear preferences that are additively separable across several components

[8]To relate the model without returns to our main analysis, we consider a "discount forfeited on return" policy under which the buyer loses any discounts if they engage in partial returns. While this return policy is typically not additive, we can still apply the logic behind Theorem 1 to this return policy, which allows us to derive Theorem 2.

(Carroll 2017). In this general screening problem, we show that if the agent's types are independent across components and return policies are required to be additive across components, then it is optimal to screen each component separately (Theorem 3). This result does not require any assumptions (such as single crossing) on the structure of preferences within each component. Thus, similar to Carroll (2017), our framework uncovers a general separation result for multidimensional screening. It also allows us to show that in the context of linear monopoly pricing without returns, bundle discounts defined in our sense are necessary to improve on separate sales even for randomized bundling mechanisms (Theorem 2′).

One may be concerned that our analysis relies on it being costless for buyers to execute returns. To address this concern, we incorporate return costs into our main model. Importantly, we allow the return costs to be private information to buyers, and to be positively correlated with their values; thus, the seller could, in principle, screen buyers using returns. We show that, nevertheless, the per-good profit gap between the optimal mechanism and separate sales is at most the average return cost of the buyers, regardless of the heterogeneity in return costs (Theorem 4). Thus, our main result is robust to the possibility of small return costs even if buyers' return costs, and hence their strategic use of returns, can be informative about their values.

We conclude by illustrating the robustness of our main result in three further extensions. First, we relax the assumption of additive return policies. Second, we relax the assumption that returning a full bundle results in a full refund. Last, we consider the possibility that buyers' values are correlated across products. In all three cases, we show that separate sales remains approximately optimal.

Related Literature. There is a large literature on multidimensional screening and optimal bundling without returns (Stigler 1963; Adams and Yellen 1976; McAfee, McMillan, and Whinston 1989; Armstrong 1996). It is now understood that the problem lacks tractability, even in the special case of two items with additive and independent values (Manelli and Vincent 2006; Daskalakis, Deckelbaum, and Tzamos 2013; Hart and Reny 2015). However, because of its importance, there has been a surge of recent interest in this problem, which has focused on identifying tractable special cases (Manelli and Vincent 2007; Pavlov 2011; Daskalakis, Deckelbaum, and Tzamos 2017; Bergemann, Bonatti, Haupt, and Smolin 2021; Haghpanah and Hartline 2021; Ghili 2023; Yang 2022, 2023), introducing non-Bayesian models (Carroll 2017; Deb and Roesler 2024; Che and Zhong 2025), and studying approximately optimal mechanisms (Chawla, Hartline, and Kleinberg 2007; Babaioff, Immorlica, Lucier, and Weinberg

2014; Cai, Devanur, and Weinberg 2016; Babaioff, Gonczarowski, and Nisan 2017; Hart and Nisan 2017; Frick, Iijima, and Ishii 2024).

Our paper offers a new perspective by identifying a set of return-proofness constraints on the seller's problem, motivated by observed return policies. These constraints turn out to eliminate the profitability of bundling in the benchmark model with additive and independent values. The closest papers to ours are Carroll (2017) and McAfee, McMillan, and Whinston (1989). Carroll (2017) introduces a non-Bayesian multidimensional screening model in which the principal knows the marginal distribution of the value of each good but not the joint distribution, and shows that separate sales is then max-min optimal. McAfee, McMillan, and Whinston (1989) show that separate sales is never optimal under independent values via a perturbation argument: marginally decreasing the price of the grand bundle always raises revenue. The economic insights of our application to the model without returns complement McAfee, McMillan, and Whinston (1989): our general concept of bundle discounts, which we show are necessary for any complex mechanism to outperform separate sales, includes their perturbation of separate sales as a special case.

Our model of strategic returns connects to the literature on mechanism design with withdrawal rights. To isolate the effect of strategic returns on pricing, we abstract from buyer learning and dynamic screening, which are studied by Krähmer and Strausz (2015), Bergemann, Castro, and Weintraub (2020), and Heumann (2020). Similar to several other models with withdrawal rights (Matthews and Postlewaite 1989; Forges 1999; Compte and Jehiel 2007, 2009; Heumann 2020; Battaglini and Palfrey 2024; Haberman and Jagadeesan 2024), return-proofness implies a set of incentive constraints for double deviations that play a central role in our analysis.

On a technical level, our proof strategy also builds on Hart and Nisan (2017), who study the approximate optimality of separate sales by extracting a collection of one-dimensional mechanisms from a given multidimensional mechanism. Our proofs combine revenue bounds for versions of these one-dimensional mechanisms with return-proofness constraints to bound the revenue of multidimensional mechanisms.

The remainder of the paper proceeds as follows. Section 2 introduces the model. Section 3 presents and proves the main result. Section 4 presents the implications for multidimensional screening without returns. Section 5 provides a generalization of the main result. Section 6 incorporates return costs. Section 7 studies the robustness of the main result. Section 8 is a conclusion. Appendix A contains omitted proofs. Appendix B provides approximate versions of the results of Section 4.

2. Model

A seller has a set G of n goods available for sale. There is a single buyer whose preferences are additive across the goods. Both the seller and the buyer are risk-neutral. We normalize the seller's cost for each good to be 0, and denote the buyer's private value for each good i to be $\theta_i \in \Theta_i \subset \mathbb{R}$, where Θ_i is assumed to be compact. The buyer's type $\theta := (\theta_1, \ldots, \theta_n) \in \Theta := \bigtimes_i \Theta_i$ is drawn from some joint distribution.

A *return policy* for bundle $B \in \mathcal{P}(G)$ is a function $r : \mathcal{P}(B) \to \mathbb{R}$ that specifies the partial refund $r(Z)$ for any subset $Z \subseteq B$. We always require that $r(\varnothing) = 0$, so there is no refund if nothing is returned. A *contract* (B, p, r) specifies a bundle B to purchase, a price p to pay, and a return policy r for bundle B such that $r(B) = p$. Thus, we assume that the buyer gets a full refund when returning the full bundle B.

The game proceeds as follows. The seller offers a menu of lotteries over contracts. The buyer decides whether to accept the menu, and if so, selects an option from the menu. The lottery is realized. The buyer pays the price p, receives the bundle B, and then decides whether to return any subset $Z \subseteq B$ for a partial refund of $r(Z)$.

By the extended revelation principle (Myerson 1982), it is without loss of generality to restrict attention to the direct revelation mechanisms as follows. Denote (B, Z, p, r) as a *pure outcome* x, consisting of a contract (B, p, r) and a subset $Z \subseteq B$ to return. Let $\omega \in [0, 1]$ be a randomization device. If the buyer reports type θ, then depending on the realization of ω, the seller assigns outcome $\big(B_\omega(\theta), Z_\omega(\theta), p_\omega(\theta), r_\omega(\theta, \cdot)\big)$.

Let $\mathcal{X}$ denote the space of pure outcomes. A *stochastic mechanism* M then consists of measurable functions

$$(B, Z, p, r) : \Theta \to \Delta(\mathcal{X})$$

that satisfy the *incentive compatibility* constraints: for all types $\theta, \hat{\theta}$, we have

$$\mathbb{E}_\omega\left[\sum_{i \in B_\omega(\theta) \setminus Z_\omega(\theta)} \theta_i - p_\omega(\theta) + r_\omega\big(\theta, Z_\omega(\theta)\big)\right] \geq \mathbb{E}_\omega\left[\max_{\hat{Z} \subseteq B_\omega(\hat{\theta})} \left\{\sum_{i \in B_\omega(\hat{\theta}) \setminus \hat{Z}} \theta_i - p_\omega(\hat{\theta}) + r_\omega\big(\hat{\theta}, \hat{Z}\big)\right\}\right].$$

Incentive compatibility requires that double deviations never be profitable: no type θ would want to pretend to be *any* type $\hat{\theta}$ and then return *any* subset $\hat{Z}$. Note that because we assume the buyer gets a full refund when returning a full bundle, our incentive compatibility constraints imply the usual (*ex post*) individual rationality constraints, as the buyer can always return the assigned bundle $B_\omega(\theta)$ for a full refund.

Throughout, when we refer to optimal revenue over a given set of mechanisms, we mean the supremum of expected revenue of mechanisms in the set. In our analysis, we do not need to assume that the supremum is attained as a maximum.

3. The Main Result

A return policy r for bundle B is *additive* if $r : \mathcal{P}(B) \to \mathbb{R}$ is an additive set function. Equivalently, the bundle price must be divided into return values for the items such that the partial refund for each sub-bundle is the sum of the return values of its items.

The seller wants to maximize expected revenue (net of refunds) among all stochastic mechanisms with additive return policies. Our main result shows that, under independent values, separate sales is optimal. Formally, let $\pi_i = \sup_p\{p \cdot \mathbb{P}(\theta_i \geq p)\}$ be the usual optimal monopoly revenue for good i, when sold separately. Since the buyer's values are additive across goods, separate sales achieves revenue $\sum_i \pi_i$.

Theorem 1. *Suppose that the seller is restricted to additive return policies, and that the buyer's values are statistically independent across goods. Then, the seller's optimal revenue is the separate sales revenue* $\sum_i \pi_i$.

Theorem 1 allows the seller to monitor the buyer's purchases (preventing the buyer from purchasing multiple bundles) and to use arbitrary stochastic mechanisms. However, it turns out that neither monitoring nor randomization has any value—separate sales does not require monitoring or randomization by the seller.

In the remainder of this section, we first sketch the proof of Theorem 1 under the additional assumption that the seller can only use deterministic mechanisms (Section 3.1), and then give the full proof of the theorem (Section 3.2).

3.1. Proof Sketch for Deterministic Mechanisms.

First, note that for any mechanism that involves returns in equilibrium, we can construct a new mechanism that is *return-proof* (i.e., $Z(\theta) = \varnothing$) without affecting revenue. The intuition is straightforward and is similar to the logic underlying the revelation principle: the seller can perform returns on the behalf of the buyer. Importantly, this procedure preserves the additivity of the return policies.

Now, consider any return-proof mechanism with additive return policies. Under these hypotheses, the revenue that the seller collects from type θ must be

$$p(\theta) = \sum_i r(\theta, \{i\}),$$

where we write $r(\theta, \{i\}) = 0$ for $i \notin B(\theta)$. The total revenue is then

$$\mathbb{E}_\theta\big[p(\theta)\big] = \sum_i \mathbb{E}_\theta\big[r(\theta, \{i\})\big].$$

Thus, the optimality of separate sales follows if we can show the following key claim.

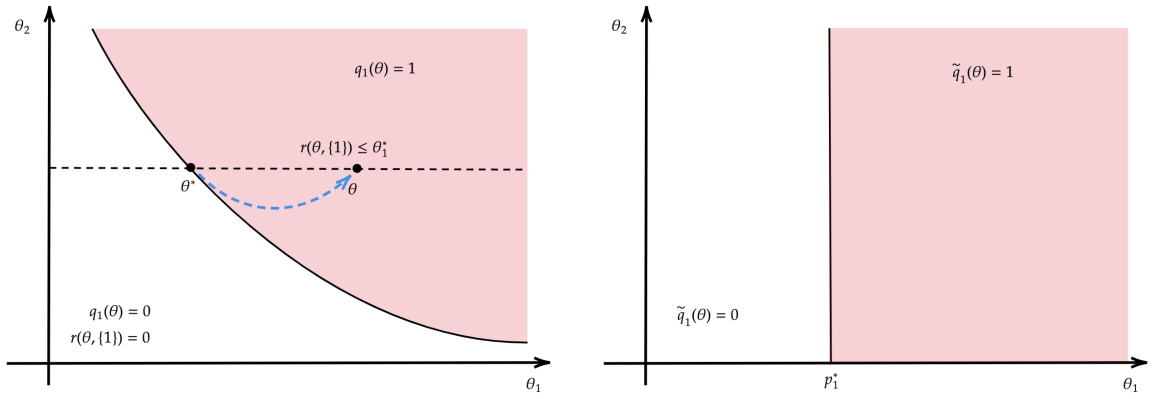

(A) hypothetical allocation rule for good 1; blue arrow indicates the deviation we use

(B) improved allocation rule for good 1; p_1^* denotes the monopoly price for good 1

FIGURE 1. Graphical illustration of intuition behind Lemma 1.

Lemma 1. *Suppose that the buyer's values are statistically independent across goods. If a deterministic mechanism is incentive-compatible and return-proof, then for all i*

$$\mathbb{E}_\theta\big[r(\theta, \{i\})\big] \le \pi_i \,.$$

We next sketch a proof of Lemma 1 assuming discrete types.[9] Fix good i. Let $q_i(\theta) \in \{0, 1\}$ denote the allocation rule of good i. Let $U(\theta)$ denote the equilibrium payoff of type θ. For any $\theta = (\theta_i, \theta_{-i})$ such that $q_i(\theta) = 1$, define $\theta^* := (\theta_i^*, \theta_{-i})$, where

$$\theta_i^* := \min\big\{s : q_i(s, \theta_{-i}) = 1\big\}$$

is the threshold value for good i for which good i starts to be sold (depending on θ_{-i}). Figure 1a illustrates with two goods.

We argue that

$$r(\theta, \{i\}) \le \theta_i^* \text{ for all } \theta \text{ such that } q_i(\theta) = 1 \,. \tag{1}$$

Fix any type θ such that $q_i(\theta) = 1$. Consider the following double deviation by type θ^*:

- Misreport to be type θ.
- Return good i to get refund $r(\theta, \{i\})$.

Figure 1a also illustrates. In order for this not to be profitable, we must have that

$$U(\theta^*) \ge \sum_{j \ne i} \theta_j q_j(\theta) - p(\theta) + r(\theta, \{i\}) = U(\theta) - \theta_i + r(\theta, \{i\}) \,.$$

[9]The sequel contains a full proof of an extension of Lemma 1 that allows for stochastic mechanisms.

The incentive constraint for the "horizontal" deviation from type θ to type θ^* is

$$U(\theta) \geq \theta_i + \sum_{j \neq i} \theta_j q_j(\theta^*) - p(\theta^*) = U(\theta^*) + \theta_i - \theta_i^*.$$

Combining the inequalities then yields

$$U(\theta) - \theta_i + \theta_i^* \geq U(\theta^*) \geq U(\theta) - \theta_i + r(\theta, \{i\})$$

which proves (1).

Since $r(\theta, \{i\}) = 0$ whenever $q_i(\theta) = 0$, (1) implies that for all θ,

$$r(\theta, \{i\}) \leq \mathbf{1}_{\theta_i \geq \theta_i^*} \cdot \theta_i^* .$$

Note that $\mathbf{1}_{\theta_i \geq \theta_i^*} \cdot \theta_i^*$ is exactly the payment rule when selling good i at a posted price θ_i^*. In particular, taking expectations conditional on θ_{-i} yields that

$$\mathbb{E}_{\theta_i}\big[r(\theta, \{i\}) \,\big|\, \theta_{-i}\big] \leq \theta_i^* \cdot \mathbb{P}(\theta_i \geq \theta_i^* \,|\, \theta_{-i}) \,,$$

where $\theta_i^* \cdot \mathbb{P}(\theta_i \geq \theta_i^* \,|\, \theta_{-i})$ is the revenue from selling good i at a posted price of θ_i^* conditional on θ_{-i}. Hence, for all θ_{-i}, we have that

$$\mathbb{E}_{\theta_i}\big[r(\theta, \{i\}) \,\big|\, \theta_{-i}\big] \leq \sup_p \Big\{ p \cdot \mathbb{P}(\theta_i \geq p \,|\, \theta_{-i}) \Big\} = \pi_i \,,$$

where the equality holds since θ_i and θ_{-i} are statistically independent. Lemma 1 follows immediately.

Figure 1b illustrates the logic geometrically in the case of two goods. Indeed, for $i = 1$, using statistical independence between θ_1 and θ_2, we can apply the above argument for each θ_2 to show that it raises revenue to change the allocation $q_1(\,\cdot\,, \theta_2)$ to $\tilde{q}_1(\,\cdot\,, \theta_2) := \mathbf{1}_{\theta_1 \geq p_1^*}$ where p_1^* is the monopoly price for good 1.

Remark 1. The incentive constraints for double deviations play an essential role in the proof of Lemma 1. In fact, if we relaxed the incentive constraints to be *(i)* truth-telling constraints plus *(ii)* obedience constraints for return recommendations without taking into account double deviations, then separate sales would *not* be optimal. Indeed, given any truthful deterministic mechanism without returns $(B, p) : \Theta \to \mathcal{P}(G) \times \mathbb{R}$ that is individually rational, the additive return policy defined by

$$r(\theta, Z) = \frac{\sum_{i \in Z} \theta_i}{\sum_{i \in B(\theta)} \theta_i} p(\theta) \quad \text{for all } Z \subseteq B(\theta)$$

would ensure that no type would want to execute any returns assuming that they reported truthfully. (The truth-telling constraints would then be automatic.) Hence, in such a relaxation, the incentive constraints would have no bite beyond the truth-

telling constraints. In particular, by McAfee, McMillan, and Whinston (1989), separate sales would be suboptimal.

Remark 2. An alternative approach to proving Theorem 1 would be to characterize the effective price schedules that are implementable under additive return policies, and then directly show that separate sales maximizes revenue within this class. Restricting attention to deterministic mechanisms, one can show that a price schedule $\{p(B)\}_{B\subseteq G}$ is implementable if and only if, when regarded as the characteristic function of a transferable utility cooperative game, it is *totally balanced*.[10] However, optimizing prices subject to this constraint is still not tractable: e.g., there are 2^n prices to optimize, and revenue is generally non-concave in bundle prices. Our mechanism design approach addresses these issues by working with allocations instead, and, as we show in Section 4, enables us to derive new results about what price schedules can outperform separate sales.

3.2. **Proof of Theorem 1.** The actual proof closely follows the argument sketched above, but also deals with stochastic mechanisms. For stochastic mechanisms, we do not have access to a threshold type θ_i^* as in the deterministic case. Therefore, we instead compare the return values $r(\theta, \{i\})$ for good i to the payment rule of a potentially stochastic mechanism for selling good i that we recover from the original multidimensional mechanism by adapting a construction of Hart and Nisan (2017).

Separate sales achieves revenue $\sum_i \pi_i$, so it suffices to prove that all other mechanisms with additive return policies achieve revenue at most $\sum_i \pi_i$.

Step 1. We first show that under additive return policies, it is without loss of generality to restrict attention to return-proof mechanisms. Formally, a mechanism M is *return-proof* if for all types θ and realizations ω, we have $Z_\omega(\theta) = \varnothing$.

Lemma 2. *For any incentive-compatible mechanism M with additive return policies, there exists a return-proof, incentive-compatible mechanism $\tilde{M}$ with additive return policies that yields the same expected revenue.*

[10]Recall that a transferable utility cooperative game is *totally balanced* if the core of the game, and of every subgame, is non-empty. (To understand the connection, note that the return values of items correspond to core imputations.) Applying the Bondareva–Shapley Theorem then implies that a price mechanism $\{p(B)\}_{B\subseteq G}$ is implementable under an additive return policy if and only if for all bundles B and all functions $\alpha : \mathcal{P}(B) \to [0,1]$ with

$$\sum_{Z\subseteq B \,|\, g\in Z} \alpha(Z) = 1 \quad \text{for all } g \in B,$$

we have that

$$p(B) \geq \sum_{Z\subseteq B} \alpha(Z) p(Z).$$

Proof. Construct the return-proof mechanism $\tilde{M}$ by

$$\begin{aligned}
\tilde{B}_\omega(\theta) &= B_\omega(\theta) \setminus Z_\omega(\theta) \\
\tilde{Z}_\omega(\theta) &= \varnothing \\
\tilde{p}_\omega(\theta) &= p_\omega(\theta) - r_\omega(\theta, Z_\omega(\theta)) \\
\tilde{r}_\omega(\theta, Z') &= r_\omega(\theta, Z' \cup Z_\omega(\theta)) - r_\omega(\theta, Z_\omega(\theta)) \qquad \text{for } Z' \subseteq B_\omega(\theta) \setminus Z_\omega(\theta).
\end{aligned}$$

To show that $\tilde{M}$ is incentive compatible, note that reporting truthfully and following the recommended action in $\tilde{M}$ gives the same payoff as in M for each type θ. Now, consider the deviation of misreporting to $\hat{\theta}$ and then returning $\hat{Z}_\omega \subseteq \tilde{B}_\omega(\hat{\theta})$ for each ω. If such a deviation were strictly profitable for type θ under $\tilde{M}$, then deviating to $\hat{\theta}$ and returning $\hat{Z}_\omega \cup Z_\omega(\hat{\theta}) \subseteq B_\omega(\hat{\theta})$ would be strictly profitable under M because these deviations yield the same payoff for θ by construction. Indeed, we have that

$$\sum_{i \in \tilde{B}_\omega(\hat{\theta}) \setminus \hat{Z}_\omega} \theta_i - \tilde{p}_\omega(\hat{\theta}) + \tilde{r}_\omega(\hat{\theta}, \hat{Z}_\omega) = \sum_{i \in B_\omega(\hat{\theta}) \setminus (\hat{Z}_\omega \cup Z_\omega(\hat{\theta}))} \theta_i - p_\omega(\hat{\theta}) + r_\omega(\hat{\theta}, \hat{Z}_\omega \cup Z_\omega(\hat{\theta})).$$

Since $r_\omega(\theta, B_\omega(\theta)) = p_\omega(\theta)$, we have that $\tilde{r}_\omega(\theta, \tilde{B}_\omega(\theta)) = \tilde{p}_\omega(\theta)$. Moreover, since M has additive return policies, $\tilde{M}$ also has additive return policies. □

Step 2. By Lemma 2, we can restrict attention to return-proof mechanisms. We now bound expected partial refunds for individual items in terms of monopoly revenues for selling one item conditional on knowing the buyer's values for the other items.

Lemma 3. *If $(B, Z, p, r) : \Theta \to \Delta(\mathcal{X})$ is incentive-compatible and return-proof, then for all i and all θ_{-i}, writing $r_\omega(\theta, \{i\}) = 0$ if $i \notin B_\omega(\theta)$, we have that*

$$\mathbb{E}_{\theta_i}\Big[\mathbb{E}_\omega\big[r_\omega(\theta, \{i\})\big] \,\Big|\, \theta_{-i}\Big] \leq \sup_p \Big\{p \cdot \mathbb{P}(\theta_i \geq p \,|\, \theta_{-i})\Big\}.$$

Proof. Let $q_j(\theta) := \mathbb{P}_\omega\big(j \in B_\omega(\theta)\big)$ denote the allocation rule for any good j, and let $U(\theta)$ denote the equilibrium payoff of type θ. Now, fix some good i. Let $\underline{\theta}_i := \min \Theta_i$. For any type θ, consider the following double deviation by type $\theta^* := (\underline{\theta}_i, \theta_{-i})$:

- Misreport to be type θ.
- Return good i (if purchased) to get refund $r_\omega(\theta, \{i\})$.

This deviation yields type θ^* a payoff of

$$\sum_{j \neq i} \theta_j q_j(\theta) - \mathbb{E}_\omega\big[p_\omega(\theta)\big] + \mathbb{E}_\omega\big[r_\omega(\theta, \{i\})\big] = U(\theta) - \theta_i q_i(\theta) + \mathbb{E}_\omega\big[r_\omega(\theta, \{i\})\big].$$

Incentive compatibility implies that this payoff cannot be higher than $U(\theta^*)$, and hence

$$\mathbb{E}_\omega\big[r_\omega(\theta,\{i\})\big] \le \theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*)\big).$$

Therefore, it suffices to prove that

$$\mathbb{E}_{\theta_i}\Big[\theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*)\big)\Big|\theta_{-i}\Big] \le \sup_p \Big\{p\cdot\mathbb{P}(\theta_i \ge p\,|\,\theta_{-i})\Big\}. \tag{2}$$

Now, fix θ_{-i} and suppose that the seller observes θ_{-i}. Consider an auxiliary mechanism for selling good i with allocation rule $q_i(\theta)$ and payment rule

$$\tilde{t}(\theta) := \theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*)\big).$$

This auxiliary mechanism is incentive-compatible because the payoff that type θ_i obtains from deviating to any type $\hat{\theta}_i$ is

$$\theta_i q_i(\hat{\theta}_i,\theta_{-i}) - \tilde{t}(\hat{\theta}_i,\theta_{-i}) = \sum_j \theta_j q_j(\hat{\theta}_i,\theta_{-i}) - \mathbb{E}_\omega\big[p_\omega(\hat{\theta}_i,\theta_{-i})\big] - U(\theta^*),$$

which is identical to the payoff from the same deviation in the original mechanism up to the constant $-U(\theta^*)$.[11] Moreover, the auxiliary mechanism is individually rational because it delivers utility $U(\theta)-U(\theta^*)$ to type θ_i, which is nonnegative as the incentive compatibility of the original mechanism implies that $U(\cdot,\theta_{-i})$ is monotone in θ_i. As the expected revenue of this auxiliary mechanism is

$$\mathbb{E}_{\theta_i}\Big[\theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*)\big)\Big|\theta_{-i}\Big],$$

the desired inequality (2) follows from the optimality of a posted price for selling one good without returns (Myerson 1981; Riley and Zeckhauser 1983). □

As noted in Section 3.1 for the deterministic mechanisms, by return-proofness and additivity of the return policies, the *ex post* revenue that the seller collects from type θ must be

$$p_\omega(\theta) = \sum_i r_\omega(\theta,\{i\})$$

where we write $r_\omega(\theta,\{i\}) = 0$ for $i \notin B_\omega(\theta)$. The total revenue is then

$$\mathbb{E}_\theta\big[\mathbb{E}_\omega[p_\omega(\theta)]\big] = \sum_i \mathbb{E}_\theta\big[\mathbb{E}_\omega[r_\omega(\theta,\{i\})]\big].$$

[11] The proof of Hart and Nisan (2017, Theorem A) constructs a similar mechanism but without the addition of $U(\theta^*)$ to payments, so their one-good mechanism delivers the same utility to all types as the original multidimensional mechanism. Adding $U(\theta^*)$ is crucial to obtain an exact optimality result, as it ensures that if we start out with separate sales, the resulting one-good mechanism is an optimal single-product monopoly pricing mechanism. Hart and Nisan's (2017) construction instead allows them to derive approximate optimality results without a return-proofness constraint.

Applying Lemma 3 yields that total revenue is at most

$$\sum_i \mathbb{E}_{\theta_{-i}}\Big[\sup_p \Big\{p \cdot \mathbb{P}(\theta_i \geq p \,|\, \theta_{-i})\Big\}\Big]. \tag{3}$$

To complete the proof of the theorem, note that since θ_i and θ_{-i} are statistically independent, for all θ_{-i}, we have that

$$\sup_p \Big\{p \cdot \mathbb{P}(\theta_i \geq p \,|\, \theta_{-i})\Big\} = \pi_i\,.$$

4. Application: Necessity of Bundle Discounts

Let us now assume that the seller does not offer any return policy. In this section, we introduce a new concept of bundle discounts. As an application of our main analysis, we show that to benefit from using any complex selling mechanism, a multiproduct monopolist must necessarily discount some bundle. This necessity result holds with any number of items and complements the results of McAfee, McMillan, and Whinston (1989): our concept of bundle discounts includes their revenue-improving perturbation of separate sales as a special case.

To introduce a concept of bundle discounts, we first focus on price mechanisms $\{p(B)\}_{B\subseteq G}$, where $p(\varnothing) = 0$. (We extend to stochastic mechanisms in the sequel.)

Definition 1. The *discount* that a price mechanism $\{p(B)\}_{B\subseteq G}$ offers for a bundle B is

$$\text{Discount}(B) := p(B) - \sum_{i\in B}\big(p(B) - p(B\setminus\{i\})\big).$$

We say that the mechanism *discounts* bundle B if $\text{Discount}(B) > 0$.

This definition compares the sum of *marginal* prices for the items in a bundle to the bundle price. For example, in a "Buy One Get One 50% Off" sale for two items priced at \$10 each, the discount offered for the bundle is $15-2\cdot(15-10) = \$5$—i.e., half the list price of one item. In particular, the bundle is discounted according to our definition. As a second example, consider a perturbation of separate sales in which the price of the grand bundle is decreased by ε (McAfee, McMillan, and Whinston 1989). This perturbation lowers the marginal price of each item in the grand bundle by ε too, and therefore introduces a discount for the grand bundle of $(n-1)\cdot\varepsilon$.

It turns out that under independent and additive values, bundle discounts in our sense are necessary to improve on the revenue of separate sales.

Theorem 2 (Necessity of Bundle Discounts)**.** *Consider the model without returns. If the buyer's values are independent across goods, then every price mechanism yields revenue greater than $\sum_i \pi_i$ must discount some bundle.*

The proof follows from a simple application of our main analysis (see Section 4.1). To illustrate why our concept of bundle discounts is essential to formulating Theorem 2, we next show that a similar optimality result would not hold if one instead considered *superadditive* price schedules. Indeed, such mechanisms can discount bundles in our sense and outperform separate sales.

Example 1 (Superadditive prices can discount bundles and outperform separate sales). Suppose that there are three goods, and that θ_i is distributed uniformly on $[1,2]$ for each good i, independently across goods. Under separate sales, the seller would post prices of 1 for each good. Instead, consider charging 1 for 1 item, $2+\varepsilon$ for 2 items, and $3+\varepsilon$ for 3 items, where $\varepsilon > 0$. This is a superadditive price schedule. But as the marginal price for the third item in the grand bundle is 1, the discount offered for the grand bundle is $(3+\varepsilon) - 3 \cdot 1 = \varepsilon > 0$. So, the grand bundle is discounted.

We claim that for $\varepsilon > 0$ small enough, our price mechanism outperforms separate sales. Under separate sales, all types buy all goods, leading to a revenue of 3. On changing to our superadditive price schedule, the revenue gain from the inframarginal types who continue to buy 3 items is $\varepsilon + O(\varepsilon^2)$. By contrast, the revenue loss from marginal types who switch to buying only 1 item is $O(\varepsilon^2)$, since that happens only for types who value two goods below $1+\varepsilon$. Thus, the seller obtains a strictly higher revenue for $\varepsilon > 0$ small enough; e.g., with $\varepsilon = 0.1$, the seller obtains revenue $\approx 3.07 > 3$.

With more than two items, the absence of bundle discounts in our sense does not imply superadditivity either,[12] so the two concepts are logically independent.[13]

In contrast, it is easy to verify that *supermodular* price schedules never discount any bundles. The following corollary is then immediate from Theorem 2.

Corollary 1. *In the model without returns, if the buyer's values are independent across goods, then separate sales is optimal among all supermodular price mechanisms.*

To the best of our knowledge, all of these results are new even in the well-studied model without returns. In fact, it is also known that, with correlated values, separate

[12]For example, with three goods $G = \{1,2,3\}$, consider the price schedule defined by

$$\begin{array}{lll} p(\{1\}) = 2 & p(\{1,2\}) = 2 & \\ p(\{2\}) = 0 & p(\{1,3\}) = 2 & p(\{1,2,3\}) = 3\,. \\ p(\{3\}) = 0 & p(\{2,3\}) = 2 & \end{array}$$

No bundle is discounted: the marginal prices of items in the grand bundle are all 1, so sum to 3; and each two-item bundle includes at least one item with a marginal price of 2. However, the price schedule is not superadditive as $p(\{1\}) + p(\{2,3\}) = 4 > 3 = p(\{1,2,3\})$.

[13]With two items, for price mechanisms, superaddivity and supermodularity are both equivalent to the bundle not being discounted.

sales cannot even attain a constant fraction of the optimal revenue among supermodular price mechanisms as the number of items grows (Ben-Moshe, Hart, and Nisan 2024). Perhaps surprisingly, Corollary 1 shows that, with independent values, separate sales is exactly optimal among supermodular price mechanisms.[14]

The following example illustrates the distinction between supermodularity and the absence of bundle discounts with more than two items. It also shows why our concept of discounts (unlike the modularity of the price schedule) sharply delineates the forms of bundle pricing that can raise revenue under independent and additive values.

Example 2 (Supermodularity versus no bundle discounts). In the setting of Example 1, consider instead charging 1 for 1 item, $2+(1-c)\cdot\varepsilon$ for 2 items, and $3+\varepsilon$ for 3 items, where $0<c<1$. Supermodularity requires that $(3+\varepsilon)+1 \geq 2\cdot(2+(1-c)\cdot\varepsilon)$, which holds if and only if $2c \geq 1$. But since the marginal price for the third item in the grand bundle is $1+c\cdot\varepsilon$, the discount offered for the grand bundle is $3+\varepsilon-3\cdot(1+c\cdot\varepsilon) = (1-3c)\varepsilon$. So, the grand bundle is discounted if and only if $3c<1$. Thus, for $\frac{1}{3} \leq c < \frac{1}{2}$, the mechanism is not supermodular, but nevertheless does not discount any bundles.

We next characterize the values of c for which our mechanism outperforms separate sales. The revenue gain from the inframarginal types remains $\varepsilon+O(\varepsilon^2)$. By contrast, the revenue loss from marginal types is $3c\cdot\varepsilon+O(\varepsilon^2)$, since the types who value any of the three goods below $1+c\cdot\varepsilon$ will not buy 3 items and, to the first order in ε, switch to buying 2 items. Thus, the seller obtains a first-order net revenue gain if and only if $3c<1$—that is, if and only if the grand bundle is discounted.

In particular, for $\frac{1}{3} \leq c < \frac{1}{2}$, the price mechanism is not supermodular but does not improve on separate sales. In fact, by Theorem 2, for such c, the price mechanism cannot improve on separate sales for any independent distribution of values.

Remark 3. To understand the connection between discounts and our main model with returns, note that the effective price schedules under additive return policies can never discount any bundles.[15] However, the converse is not true: there are price mechanisms

[14]In particular, with two goods and independent values, Corollary 1 implies that optimal price mechanisms are *submodular* and hence *revenue-monotone* in the sense of Hart and Reny (2015)—this special case itself recovers a recent result by Babaioff, Nisan, and Rubinstein (2018).

[15]Indeed, suppose that $\{p(B)\}_{B\subseteq G}$ are the effective prices under some deterministic mechanism and additive return policies. To show that no bundle B can be discounted, let $r:\mathcal{P}(B)\to\mathbb{R}$ be the return policy faced by a buyer who obtains B. Now given any good $i\in B$, as the buyer can obtain $B\setminus\{i\}$ by buying B and returning $\{i\}$, we have that $p(B\setminus\{i\}) \leq p(B)-r(\{i\})$. Summing over all goods $i\in B$ and using the additivity of r yields that

$$\sum_{i\in B}\big(p(B)-p(B\setminus\{i\})\big) \geq \sum_{i\in B} r(\{i\}) = p(B)\,.$$

So no bundle B can be discounted.

that do not discount any bundles but are not effective price schedules. For example, there are price mechanisms that do not offer discounts but are not superadditive (see, e.g., Footnote 12), while every effective price schedule under additive return policies is superadditive.[16] By Theorem 2, superadditivity (and other additional constraints implied by the effective price schedules under additive return policies) is not what drives the optimality of separate sales: instead, the key is that the effective prices under additive return policies can never offer discounts, which Theorem 2 shows are necessary to improve on the revenue of separate sales.

The remainder of this section proves Theorem 2 (Section 4.1), and then extends our analysis of the model without returns to allow for stochastic mechanisms (Section 4.2).

4.1. **Discounts Forfeited on Return and the Proof of Theorem 2.** To prove Theorem 2, we map the model without returns to our main model. As price mechanisms without discounts may not be effective price schedules under additive return policies (Remark 3), we cannot apply Theorem 1 directly to bound revenue. However, we can still apply the logic behind Theorem 1 to a non-additive return policy.

More precisely, consider the "discount forfeited on return" policy (used by Best Buy), which for any price mechanism $\{p(B)\}_{B\subseteq G}$ and any bundle B, is defined as

$$r(Z) = p(B) - p(B \setminus Z)\,.$$

This return policy is not additive in general. However, if the price mechanism does not discount bundle B, then we have that

$$\sum_{i\in B} r(\{i\}) = \sum_{i\in B} \big(p(B) - p(B\setminus\{i\})\big) = p(B) - \text{Discount}(B) \geq p(B)\,. \tag{4}$$

Moreover, the logic behind Theorem 1 shows that for return-proof mechanisms, the return values of items having to sum to at least the bundle price is sufficient to ensure the optimality of separate sales; Theorem 2 then follows from this conclusion.

[16]To see that the effective prices must be superadditive, suppose that $\{p(B)\}_{B\subseteq G}$ are the effective prices under some deterministic mechanism and additive return policies. Consider any bundle B, and let $r : \mathcal{P}(B) \to \mathbb{R}$ be the return policy faced by a buyer who obtains B. Suppose that B_1, B_2 are disjoint bundles with $B_1 \cup B_2 = B$. As the buyer can obtain B_1 by buying B and returning B_2 we have $p(B_1) \leq p(B) - r(B_2)$. Similarly, we have $p(B_2) \leq p(B) - r(B_1)$. As r is additive, we then have

$$p(B_1) + p(B_2) \leq 2p(B) - \big(r(B_1) + r(B_2)\big) = p(B)\,.$$

Proposition 1. *Suppose that buyer's values are statistically independent across goods. If an deterministic, incentive-compatible, return-proof mechanism* (B, Z, p, r) *satisfies*

$$\sum_{i \in B(\theta)} r(\theta, \{i\}) \geq p(\theta)$$

for all types θ*, then it yields revenue at most* $\sum_i \pi_i$*.*

Proof. The *ex post* revenue that the seller collects from type θ satisfies

$$p(\theta) \leq \sum_i r(\theta, \{i\}),$$

where we write $r(\theta, \{i\}) = 0$ for $i \notin B(\theta)$. As in Section 3.2, we can use Lemma 3 and the statistical independence of values to conclude that revenue is most $\sum_i \pi_i$. □

Proof of Theorem 2. We prove the contrapositive. Consider any price mechanism M that does not discount any bundles; we prove that the revenue of M must be at most $\sum_i \pi_i$. Under the "discount forfeited on return" policy, each type θ has an optimal choice of selecting the same bundle as under M and not returning anything, so that return policy leads to a return-proof mechanism. Since M has no bundle discounts, (4) holds. Proposition 1 then implies that M yields revenue at most $\sum_i \pi_i$. □

4.2. **Bundle Discounts for Stochastic Mechanisms.** Our concept of bundle discounts extends naturally to stochastic mechanisms. In the model without returns, a stochastic mechanism can be without loss of generality represented as a menu $\{p(q)\}_{q \in [0,1]^n}$, where each option in the menu describes the price $p(q)$ for a randomized bundle q in which each item i is obtained with probability q_i, and $p(0) = 0$.[17]

Definition 1′. The *discount* that a menu $\{p(q)\}_{q \in [0,1]^n}$ offers for a bundle q is

$$\text{Discount}(q) := p(q) - \sum_i \left(p(q_i, q_{-i}) - p(0_i, q_{-i}) \right).$$

We say that the menu *discounts* bundle q if $\text{Discount}(q) > 0$.

If the mechanism is deterministic, i.e., $q \in \{0,1\}^n$, then this definition reduces to Definition 1.[18] With this concept of discounts, Theorem 2 holds without change.

Theorem 2′ (Necessity of Bundle Discounts)**.** *Consider the model without returns. If the buyer's values are independent across goods, then every menu* $\{p(q)\}_{q \in [0,1]^n}$ *that represents a mechanism that yields revenue greater than* $\sum_i \pi_i$ *must discount a bundle.*

[17] This menu is not generally unique, but Hart and Reny (2015, Appendix A.2) shows how to construct a canonical menu for any mechanism. For our purposes, we can work with any menu.

[18] Analogous logic to the deterministic case (Remark 3 and Footnote 15) shows that effective menus under additive return policies never discount any bundles.

The proof is in the appendix. It follows a similar argument to the proof of Theorem 2. As menu descriptions of stochastic mechanisms for the model without returns do not directly involve randomization, we cannot map them to stochastic mechanisms for our main model with returns. Instead, we map such menus to deterministic mechanisms in the general screening model that we will introduce in Section 5, which allows us to view the probabilities q_i as quantities; we then use a version of Proposition 1.

Corollary 1 also extends to stochastic mechanisms. A mechanism M is *supermodular* if it can be represented by a menu $\{p(q)\}_{q\in[0,1]^n}$ where $p(\cdot)$ is supermodular. It is easy to see that if $p(\cdot)$ is supermodular, then $\{p(q)\}_{q\in[0,1]^n}$ has no bundle discounts.

Corollary 2. *In the model without returns, if the buyer's values are independent across goods, then separate sales is optimal among all supermodular mechanisms.*

5. General Multidimensional Screening

Our main result focuses on linear monopoly pricing, but our key logic does not depend on a linear preference or cost structure. In this section, we show that separate screening remains optimal under additive return policies in a general screening model.

As in Carroll (2017), consider a screening model with a set G of *additive components*, where for each component $i \in G$, there is a compact allocation space $\mathcal{X}_i$ and a compact type space Θ_i, both of which are allowed to be multidimensional. The agent's utility is additively separable across these components and quasilinear in the transfers: for any $B \subseteq G$, joint allocation $(x_i)_{i\in B}$, and payment p, the agent's payoff is

$$\sum_{i\in B} u_i(x_i,\theta_i) + \sum_{i\in G\setminus B} u_i(\underline{x}_i,\theta_i) - p\,,$$

where $\underline{x}_i \in \mathcal{X}_i$ represents an outside option in the component i problem for the agent. The principal's utility is normalized to care only about the payment p. We assume that the component utility functions u_i are bounded and measurable.

A return policy for a joint allocation $(x_i)_{i\in G}$, is a function $r : \mathcal{P}(G) \to \mathbb{R}$ that specifies the refund that an agent gets when returning part of the allocation $(x_i)_{i\in Z}$. For example, G may represent a set of goods, each of which can have some quality level x_i, where quality level $\underline{x}_i$ represents not purchasing i; a buyer can return a subset Z for refunds. As in Section 2, we require that $r(G) = p$, so returning all components leads to a full refund. A return policy r is *additive* if it is an additive set function.

The principal has access to all deterministic mechanisms. For example, if each component is a nonlinear monopoly pricing problem à la Mussa and Rosen (1978), the principal can jointly price various combinations of different qualities of different goods.

We now show that under additive return policies and independence of types across components, the optimal mechanism screens the components separately. Formally, let π_i denote the principal's expected payoff if she were to only screen component i, i.e., her optimal payoff if the agent's payoff were given by $u_i(x_i, \theta_i) - p$ (with outside option $\underline{x}_i$). In the nonlinear monopoly pricing example, π_i is the optimal profit for a menu of qualities of good i.

Theorem 3. *In the general screening model, suppose that the principal is restricted to additive return policies, and that the agent's types are independent across components. Then, the principal's optimal payoff is the separate screening payoff* $\sum_i \pi_i$.

The proof is in the appendix, and uses a similar argument to the proof of Theorem 1.

Note that Theorem 3 does not require any assumptions about the structure of preferences within each component (such as single-crossing), similar to Carroll's (2017) non-Bayesian analysis of multidimensional screening. However, in our setting, separate screening becomes optimal under type distributions that are independent across components—which are never the worst-case joint distributions in Carroll (2017).

In addition to illustrating the generality of our approach, as we discussed in Section 4.2, we also use our analysis of the general screening model to prove Theorem 2′ (the necessity of bundle discounts to improve on separate sales in the stochastic case).

6. Costly Returns

We next relax the assumption that it is costless for buyers to execute returns. More precisely, we allow the buyer to incur a cost θ_0 for returning any bundle. Importantly, we allow the return costs θ_0 to be heterogeneous and correlated with the buyer's values. Then, the seller may both deter returns more effectively, and use returns as a costly screening device. Nevertheless, we show that under additive return policies, the seller cannot do much better than separate sales as long as the average return cost across all types of the buyer is not too large—even if some types can experience high return costs.

More precisely, we enlarge the type space to be $\Theta = \Theta_0 \times \bigtimes_i \Theta_i$, where $\Theta_0 \subseteq \mathbb{R}_{\geq 0}$ is a compact space of return costs. Type $\theta = (\theta_0, \theta_1, \ldots, \theta_n) \in \Theta$ has cost θ_0 of executing any returns (i.e., choosing $Z \neq \varnothing$). The incentive compatibility constraint for a mechanism $(B, Z, p, r) : \Theta \to \Delta(\mathcal{X})$ then requires that for all types $\theta, \hat{\theta}$, we have

$$\mathbb{E}_\omega \left[\sum_{i \in B_\omega(\theta) \setminus Z_\omega(\theta)} \theta_i - p_\omega(\theta) + r_\omega\big(\theta, Z_\omega(\theta)\big) - \mathbf{1}_{Z_\omega(\theta) \neq \varnothing} \theta_0 \right]$$

$$\geq \mathbb{E}_\omega\left[\max_{\hat{Z}\subseteq B_\omega(\hat{\theta})}\left\{\sum_{i\in B_\omega(\hat{\theta})\setminus\hat{Z}}\theta_i - p_\omega(\hat{\theta}) + r_\omega\big(\hat{\theta},\hat{Z}\big) - \mathbf{1}_{\hat{Z}\neq\varnothing}\theta_0\right\}\right].$$

With return costs, this incentive compatibility condition for double deviations does not imply individual rationality, so we also impose individual rationality.[19]

We focus on the economically natural case in which types with higher values tend to have higher return costs. Formally, we assume that return costs are *stochastically monotone* in the sense that for all $\theta_{-0} \leq \theta'_{-0}$, the distribution of θ_0 conditional on θ'_{-0} first-order stochastically dominates the distribution of θ_0 conditional on θ_{-0}.[20] A special case is that return costs are given by any increasing function of values.

One might think that the seller could gain revenue by screening for high values via return costs. However, we can show that under additive return policies and independent values, the additional revenue per good that the seller can obtain from using bundling to screen via return costs is at most the average return cost $\mathbb{E}[\theta_0]$.

Theorem 4. *With stochastically monotone return costs, suppose that the seller is restricted to additive return policies, and that the buyer's values are statistically independent across goods. Then, the seller's optimal revenue is at most* $\sum_i \pi_i + n \cdot \mathbb{E}[\theta_0]$.

The proof is in the appendix, but here, we sketch the argument for the case of deterministic mechanisms under the further hypothesis that return costs are a deterministic function of values. With return costs, note that it is no longer without loss of generality to restrict attention to return-proof mechanisms, as returns can effectively serve a costly screening device for the seller. Nevertheless, we can apply a similar approach to the proof of Theorem 1 by explicitly tracking returns.

More precisely, given any deterministic mechanism, let $q_i(\theta) = \mathbf{1}_{i\in B(\theta)\setminus Z(\theta)} \in \{0,1\}$ denote the allocation rule of good i accounting for equilibrium returns. Under additive return policies, the revenue that the seller collects from type θ must be

$$p(\theta) - r(\theta, Z(\theta)) = \sum_i q_i(\theta)\cdot r(\theta,\{i\}).$$

The total revenue is then

$$\mathbb{E}_\theta\big[p(\theta) - r(\theta, Z(\theta))\big] = \sum_i \mathbb{E}_\theta\big[q_i(\theta)\cdot r(\theta,\{i\})\big].$$

Thus, the result follows if we can show the following extension of Lemma 1.

[19] That is, we require that $\mathbb{E}_\omega\Big[\sum_{i\in B_\omega(\theta)\setminus Z_\omega(\theta)}\theta_i - p_\omega(\theta) + r_\omega\big(\theta, Z_\omega(\theta)\big) - \mathbf{1}_{Z_\omega(\theta)\neq\varnothing}\theta_0\Big] \geq 0$ for all types θ.

[20] In addition to this case being economically natural, if higher types had lower return costs, then one would expect the seller's ability to screen using return costs to be more limited (Yang 2022).

Lemma 4. *With deterministic, monotone return costs, suppose that the buyer's values are statistically independent across goods. If a deterministic mechanism is incentive-compatible, then for all i*

$$\mathbb{E}\big[q_i(\theta)\cdot r(\theta,\{i\})\big] \le \pi_i + \mathbb{E}[\theta_0]\,.$$

We next sketch a proof of Lemma 4 assuming discrete types.[21] Let $\theta_0(\theta_{-0})$ be the return cost of a buyer with values θ_{-0}. Fix a good i. As in Section 3, for any $\theta = (\theta_0, \theta_i, \theta_{-i})$ such that $q_i(\theta) = 1$, define $\theta^* := (\theta_0(\theta_i^*, \theta_{-i}), \theta_i^*, \theta_{-i})$, where

$$\theta_i^* := \min\big\{s : q_i(\theta_0(s,\theta_{-i}), s, \theta_{-i}) = 1\big\}$$

is the threshold value for good i for which good i starts to be sold (depending on θ_{-i}).

We argue that

$$r(\theta,\{i\}) \le \theta_i^* + \theta_0 \text{ for all } \theta \text{ such that } q_i(\theta) = 1\,. \tag{5}$$

The argument is similar to the derivation of (1) from Section 3, but carefully accounting for equilibrium returns and for return costs, which we highlight in red. Fix any type θ such that $q_i(\theta) = 1$. In order for it not to be profitable for type θ^* to execute a double deviation of misreporting to be type θ and returning good i, taking into account the return costs of type θ^*, we must have that

$$U(\theta^*) \ge \sum_{j\neq i}\theta_j q_j(\theta) - p(\theta) + r(\theta,\{i\}) - \theta_0^* \ge U(\theta) - \theta_i + r(\theta,\{i\}) - \theta_0^*\,.$$

By casework on whether type θ^* is recommended to execute returns,[22] we can show that the incentive constraint for the "horizontal" deviation from type θ to type θ^* implies that

$$U(\theta) \ge U(\theta^*) + \theta_i - \theta_i^* + \theta_0^* - \theta_0\,.$$

[21] Appendix A contains a full proof of an extension of Lemma 4 that allows for stochastic mechanisms.

[22] Indeed, if type θ^* is not recommended to execute any returns ($Z(\theta^*) = \varnothing$), then the incentive constraint for the "horizontal" deviation from type θ to type θ^* is

$$U(\theta) \ge \theta_i + \sum_{j\neq i}\theta_j q_j(\theta^*) - p(\theta^*) = U(\theta^*) + \theta_i - \theta_i^*\,.$$

The monotonicity of return costs and the definition of θ_i^* together imply $\theta_0^* \le \theta_0$, and it follows that

$$U(\theta) \ge U(\theta^*) + \theta_i - \theta_i^* \ge U(\theta^*) + \theta_i - \theta_i^* + \theta_0^* - \theta_0\,.$$

On the other hand, if $Z(\theta^*) \neq \varnothing$, then taking into account the return costs of types θ and θ^*, the same incentive constraint is instead

$$U(\theta) \ge \theta_i + \sum_{j\neq i}\theta_j q_j(\theta^*) - p(\theta^*) - \theta_0 = U(\theta^*) + \theta_i - \theta_i^* + \theta_0^* - \theta_0\,.$$

In either case, we have that

$$U(\theta) \ge U(\theta^*) + \theta_i - \theta_i^* + \theta_0^* - \theta_0\,.$$

Combining the inequalities yields that

$$U(\theta) - \theta_i + \theta_i^* + \theta_0 \geq U(\theta^*) + \theta_0^* \geq U(\theta) - \theta_i + r(\theta, \{i\}),$$

which proves (5).

We can then complete the argument as in Section 3. (5) implies that for all θ,

$$q_i(\theta) \cdot r(\theta, \{i\}) \leq \mathbf{1}_{\theta_i \geq \theta_i^*} \cdot \theta_i^* + \theta_0 .$$

As in Section 3, the statistical independence of values across goods implies that $\mathbb{E}_\theta[\mathbf{1}_{\theta_i \geq \theta_i^*} \cdot \theta_i^* \,|\, \theta_{-i}] \leq \pi_i$, and it follows that

$$\mathbb{E}_\theta[q_i(\theta) \cdot r(\theta, \{i\})] \leq \mathbb{E}_\theta[\mathbf{1}_{\theta_i \geq \theta_i^*} \cdot \theta_i^*] + \mathbb{E}_\theta[\theta_0] \leq \pi_i + \mathbb{E}_\theta[\theta_0] .$$

The general argument that applies to stochastic mechanisms is more involved. As in the proof of Theorem 1 in Section 3.2, we do not have access to a threshold type θ_i^*, so we need to compare $r(\theta, \{i\})$ to the payment rule of a potentially stochastic mechanism for selling good i. However, this one-good stochastic mechanism may not be incentive compatible in the standard model due to heterogeneity in the return costs across various types of buyers. However, by interpreting costly returns as a costly screening device, we can introduce an auxiliary one-dimensional mechanism design problem with costly misreporting for which the one-good mechanism, taking into account returns, remains incentive-compatible.

To complete the proof of Theorem 4, we then bound the optimal revenue in the auxiliary one-dimensional mechanism design problem using Lagrange multipliers that arise from standard one-dimensional screening problems. In particular, we apply the Downward Sufficiency Theorem in Yang (2022) to restrict attention to *downward* incentive constraints (including non-local ones), which allows us to uniformly bound Lagrange multipliers under a network flow interpretation of the dual introduced by Sher and Vohra (2015), Cai, Devanur, and Weinberg (2016), and Carroll (2017).

7. Discussion

In this section, we study the robustness of our main result in three extensions. Section 7.1 allows for non-additive return policies, Section 7.2 considers incomplete refundability, and Section 7.3 adds correlation between buyer's values across goods.

7.1. **Non-Additive Return Policies.**

A key assumption in Theorem 1 is that the seller is restricted to additive return policies. We now show that our insight that returns lead to the unraveling of bundling continues to hold in an approximate sense

if return policies are only required to be close to additive. For any $\alpha \in [0,1]$, we say that a return policy r for bundle B with price p is *α-additive* if

$$\sum_{i \in B \setminus Z} \big(r(Z \cup \{i\}) - r(Z)\big) \geq \alpha\big(p - r(Z)\big) \quad \text{for all } Z \subseteq B\,.$$

In the case of two items with $r(\{1\}), r(\{2\}) \leq p_B$, the definition simply asks that $r(\{1\}) + r(\{2\}) \geq \alpha \cdot p_B$. In general, the definition asks that the return values of the items in a bundle sum to at least α fraction of the bundle price; with more than two items, the definition also asks that such inequalities continue to hold for marginal return values after returning a subset $Z \subseteq B$.[23]

The next result shows that, under α-additive return policies, separate sales always guarantees at least α fraction of the optimal revenue.

Proposition 2. *Suppose that the seller is restricted to α-additive return policies, and that the buyer's values are statistically independent across goods. Then*

$$\sum_i \pi_i \geq \alpha \cdot \big(\text{Optimal Revenue}\big)\,.$$

For example, in "Buy One Get One 50% Off" sales for equally priced items, Best Buy's "discount forfeited on return" policy would refund only 50% of the item's list price if the buyer only returns a single item. Our result implies that in such a case, separate sales still yields at least $\alpha = \frac{0.5+0.5}{1.5} = \frac{2}{3}$ fraction of the mixed bundling revenue. More generally, Proposition 2 leads to new results on the approximate optimality of separate sales under a cap on bundle discounts; see Appendix B.[24]

The proof is in the appendix. For the case of deterministic mechanisms, the idea can be understood following the argument presented in Section 3.1. In particular, Lemma 1 does not rely on the additivity of return policies. Proposition 2 follows by noting that under α-additivity of the return policy, we have

$$\sum_i \pi_i \geq \sum_i \mathbb{E}\big[r(\theta, \{i\})\big] = \mathbb{E}\Big[\sum_i r(\theta, \{i\})\Big] \geq \mathbb{E}\big[\alpha \cdot p(\theta)\big] = \alpha \cdot \mathbb{E}\big[p(\theta)\big]\,.$$

7.2. **Incomplete Refundability.** We next relax the assumption that returning a full bundle results in a full refund. For example, the buyer may forfeit a deposit of up to δ. We show that for small δ, separate sales remains approximately optimal.

[23] This last requirement ensures that we can use revelation principle logic to restrict attention to return-proof mechanisms. Taking $\alpha = 1$, 1-additivity is strictly weaker than additivity, but sufficient to ensure that the hypothesis of Proposition 1 holds after applying revelation principle logic.

[24] Without returns, even with additive and independent values, separate sales cannot guarantee a constant fraction of the optimal revenue as the number of goods grows (Hart and Nisan 2017).

More precisely, we extend the space of pure outcomes to allow tuples (B, Z, p, r) with $r(B) \le p$. A mechanism is *refunding within* δ if for every contract (B, p, r) offered by the mechanism, we have that $r(B) \ge p - \delta$. For mechanisms that have additive return policies, refundability within δ simply requires that return values for the items in a bundle sum to within δ of the bundle price. With incomplete refundability, our incentive compatibility condition for double deviations does not imply individual rationality, so we also impose individual rationality.[25]

Proposition 3. *With incomplete refundability, suppose that the seller is restricted to mechanisms with additive return policies that are refunding within δ, and that the buyer's values are statistically independent across goods. Then*

$$\text{Optimal Revenue} \le \sum_i \pi_i + \delta\,.$$

The proof is in the appendix. For the case of deterministic mechanisms, the idea can be understood following the argument presented in Section 3.1. In particular, Lemma 1 continues to hold under incomplete refundability. Proposition 3 follows by noting that under additive return policies and refundability within δ, we have that

$$\delta + \sum_i \pi_i \ge \delta + \sum_i \mathbb{E}\big[r(\theta, \{i\})\big] = \mathbb{E}\Big[\delta + \sum_i r(\theta, \{i\})\Big] \ge \mathbb{E}\big[p(\theta)\big]\,.$$

7.3. **Correlated Values.** Last, we show that if the buyer's values are weakly correlated across goods, then separate sales remains approximately optimal. Let

$$\overline{\pi}_i := \mathbb{E}_{\theta_{-i}}\Big[\sup_p \big\{p \cdot \mathbb{P}(\theta_i \ge p \,|\, \theta_{-i})\big\}\Big]$$

denote the expected monopoly revenue for selling good i when the seller observes the buyer's profiles of values θ_{-i} for the other goods. Clearly (as discussed in Section 3.2), if the buyer's values are independent across goods, then $\overline{\pi}_i = \pi_i$. In general, we show that the optimal revenue lies between the separate sales revenue $\sum_i \pi_i$ and $\sum_i \overline{\pi}_i$.

Proposition 4. *Suppose that the seller is restricted to additive return policies. Then*

$$\sum_i \pi_i \le \text{Optimal Revenue} \le \sum_i \overline{\pi}_i\,.$$

This result follows from the proof of Theorem 1, in particular, (3) in Section 3.2.

When the buyer's values are close to independent across goods, the gap $\overline{\pi}_i - \pi_i$ between the monopoly revenue for selling good i with and without the information of θ_{-i} will shrink to zero, and hence separate sales remain approximately optimal

[25] That is, we require that $\mathbb{E}_\omega\Big[\sum_{i \in B_\omega(\theta) \setminus Z_\omega(\theta)} \theta_i - p_\omega(\theta) + r_\omega\big(\theta, Z_\omega(\theta)\big)\Big] \ge 0$ for all types θ.

by Proposition 4. The following result provides a formal sense in which the buyer's values being close to independent across goods implies that $\overline{\pi}_i - \pi_i$ is small. More precisely, it shows that the gap $\overline{\pi}_i - \pi_i$ is bounded by the expectation of the largest amount that learning the profile of values θ_{-i} can influence quantiles of the value θ_i.

Proposition 5. *Letting $F_i^{-1}(\cdot)$ denote the unconditional quantile function of θ_i and $F_i^{-1}(\cdot\,|\,\theta_{-i})$ denote the conditional quantile function of θ_i conditional on θ_{-i}, we have*

$$\overline{\pi}_i - \pi_i \leq \mathbb{E}_{\theta_{-i}}\Big[\sup_q \big\{F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big\}\Big],$$

In particular, if the seller is restricted to additive return policies, then

$$\text{Optimal Revenue} \leq \sum_i \pi_i + \sum_i \mathbb{E}_{\theta_{-i}}\Big[\sup_q \big\{F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big\}\Big].$$

The proof is straightforward and left to the appendix.[26]

8. Conclusion

Multiproduct sellers often need to set both prices and return policies for bundles. Consumer-friendly return policies may be required by regulators or help attract consumers, but introduce a form of moral hazard as consumers can strategically exploit them. This paper introduces the joint design problem for pricing and return policies, taking the moral hazard concern from strategic returns seriously. We show that under a restriction to additive return policies, which are common in practice, the optimal mechanism in the benchmark case of independent and additive values is simply to sell separately. The result is in stark contrast to the standard model without returns, where the optimal selling mechanism is often very complex and generally unknown.

Our result applies even if the seller can monitor the buyer's purchases and use arbitrary stochastic mechanisms, and holds for general multidimensional screening problems. Moreover, selling separately remains approximately optimal if the buyer experiences small return costs that are correlated with their values, as well as if the return policy is close to additive or the buyer's values are weakly correlated across goods. As an application to multiproduct pricing without returns, we also introduce a new concept of bundle discounts and show the necessity of bundle discounts in our sense for a seller to benefit from any complex mechanism.

[26]When the buyer's values of the goods are affiliated, a version of Proposition 5 can be established that bounds $\overline{\pi}_i - \pi_i$ in terms of how the conditional expectation $\mathbb{E}[\theta_i\,|\,\theta_{-i}]$ varies with θ_{-i}.

From an applied perspective, our results suggest that sellers face a trade-off between potentially increasing profit by using various bundling mechanisms and offering consumer-friendly return policies. In practice, the design of return policies may need to involve many forces outside of our model, such as whether consumers or regulators perceive them to be simple or fair. For example, a recent investigation by the Competition and Markets Authority in the United Kingdom has led Microsoft to change its Xbox subscription pricing to allow for pro-rata refunds.[27] A testable implication of our theory is that bundling should be more common in settings in which returns are infeasible than in settings with consumer-friendly return policies. For example, bundling should be more common in the sale of services (such as insurance and telecommunication services) than in the sale of goods to consumers, and less common when consumer-friendly return policies are required by regulators due to consumer protection concerns.[28]

From a methodological perspective, our results show how to incorporate a form of moral hazard to gain tractability in multiproduct pricing and other multidimensional screening models. More broadly, we hope that constraints imposed on the designer by broader practical considerations (such as return rights) may lead to more tractable models of multidimensional screening, with more realistic predictions.

[27]See `https://www.gov.uk/government/news/cma-secures-changes-to-xbox-subscription-practices`.
[28]Our results also suggest that once sellers' reoptimization in selling mechanisms is taken into account, such consumer protection concerns can negatively impact consumer surplus. Indeed, as discussed in the Introduction, this issue arises, for example, if there are two goods and values are independently distributed uniformly between 0 and 30 (see Footnote 7).

Appendix A. Omitted Proofs

We present proofs in the order in which results appear in the text, except that we present the proof of Theorem 2′ after the proof of Theorem 3.

A.1. **Proof of Theorem 3.** Separate screening achieves revenue $\sum_i \pi_i$, so it suffices to prove that all mechanisms with additive return policies achieve payoff at most $\sum_i \pi_i$.

As a preliminary step, note that, without loss of generality, we can normalize the values of the outside options such that

$$u(\underline{x}_i, \theta_i) = 0\,,$$

by defining $\tilde{u}(x_i, \theta_i) := u(x_i, \theta_i) - u(\underline{x}_i, \theta_i)$ if the value of the outside option is not zero. This normalization does not affect the incentive constraints even with double deviations. We can then follow the same strategy as the proof of Theorem 1.

Step 1. We first show that it remains without loss of generality to restrict attention to return-proof mechanisms. We denote mechanisms by (x, Z, p, r), where $x = (x_i)_{i \in G}$ is a joint allocation rule for the component problems and $r(\theta, \cdot) : \mathcal{P}(G) \to \mathbb{R}$ is a return policy for each type θ (where $r(\theta, G) = p(\theta)$). As in the case of linear monopoly, we say that a mechanism is *return-proof* if $Z(\theta) = \varnothing$ for all types θ.

Lemma A.1. *For any incentive-compatible mechanism M with additive return policies, there exists a return-proof, incentive-compatible mechanism $\tilde{M}$ with additive return policies that yields the same expected payoff for the principal.*

Proof. Construct the return-proof mechanism $\tilde{M}$ by

$$\begin{aligned}
\tilde{x}(\theta) &= \left(x(\theta)_{G \setminus Z(\theta)}, \underline{x}_{Z(\theta)}\right) \\
\tilde{Z}(\theta) &= \varnothing \\
\tilde{p}(\theta) &= p(\theta) - r(\theta, Z(\theta)) \\
\tilde{r}(\theta, Z') &= r(\theta, Z' \cup Z(\theta)) - r(\theta, Z(\theta))\,.
\end{aligned}$$

To show that $\tilde{M}$ is incentive compatible, note that reporting truthfully and following the recommended action in $\tilde{M}$ gives the same payoff as in M for each type θ. Now, consider the deviation of misreporting to $\hat{\theta}$ and then returning $\hat{Z} \subseteq G$. If such a deviation were strictly profitable for type θ under $\tilde{M}$, then deviating to $\hat{\theta}$ and returning $\hat{Z} \cup Z(\hat{\theta})$ would be strictly profitable under M because these deviations yield the same payoff for θ by construction. Indeed, we have that

$$\sum_{i \in G \setminus \hat{Z}} u_i(\tilde{x}_i(\hat{\theta}), \theta_i) - \tilde{p}(\hat{\theta}) + \tilde{r}(\hat{\theta}, \hat{Z}) = \sum_{i \in G \setminus (\hat{Z} \cup Z(\hat{\theta}))} u_i(x_i(\hat{\theta}), \theta_i) - p(\hat{\theta}) + r(\hat{\theta}, \hat{Z} \cup Z(\hat{\theta}))\,.$$

Since $r(\theta, G) = p(\theta)$, we have that $\tilde{r}(\theta, G) = \tilde{p}(\theta)$. Moreover, since M has additive return policies, $\tilde{M}$ also has additive return policies. □

Step 2. By Lemma A.1, we can restrict attention to return-proof mechanisms. We now prove a version of Lemma 3 for the general screening model.

Lemma A.2. *If a mechanism (x, Z, p, r) is incentive-compatible and return-proof, then for all i and all θ_{-i}, we have that*

$$\mathbb{E}_{\theta_i}\left[r(\theta, \{i\}) \,\middle|\, \theta_{-i}\right] \leq \pi_i(\theta_{-i})$$

where $\pi_i(\theta_{-i})$ is the principal's optimal payoff of screening component i conditional on observing the types of other components θ_{-i}.

Proof. Fix any component i and any type θ. Consider the following double deviations by types $\theta^*(s) := (s, \theta_{-i})$ indexed by $s \in \Theta_i$:

- Misreport to be type θ.
- Withdraw from the component i allocation to get refund $r(\theta, \{i\})$.

This deviation yields type $\theta^*(s)$ a payoff of

$$\sum_{j \neq i} u_j(x_j(\theta), \theta_j) - p(\theta) + r(\theta, \{i\}) = U(\theta) - u_i(x_i(\theta), \theta_i) + r(\theta, \{i\}),$$

where $U(\theta)$ is the indirect utility function. Incentive compatibility implies that

$$U(\theta^*(s)) \geq U(\theta) - u_i(x_i(\theta), \theta_i) + r(\theta, \{i\})$$

for all $s \in \Theta_i$. Hence, we have that

$$\inf_{s \in \Theta_i}\left\{U(\theta^*(s))\right\} \geq U(\theta) - u_i(x_i(\theta), \theta_i) + r(\theta, \{i\}),$$

which we can write as

$$r(\theta, \{i\}) \leq u_i(x_i(\theta), \theta_i) - \left(U(\theta) - \inf_{s \in \Theta_i}\left\{U(\theta^*(s))\right\}\right).$$

Hence, to prove the lemma, it suffices to prove that

$$\mathbb{E}\left[u_i(x_i(\theta), \theta_i) - \left(U(\theta) - \inf_{s \in \Theta_i}\left\{U(\theta^*(s))\right\}\right) \,\middle|\, \theta_{-i}\right] \leq \pi_i(\theta_{-i}). \tag{A.1}$$

Supposing that the seller observes θ_{-i}, consider the mechanism for component i with allocation rule $x_i(\theta)$ and payment rule

$$\tilde{t}(\theta) := u_i(x_i(\theta), \theta_i) - \left(U(\theta) - \inf_{s \in \Theta_i}\left\{U(\theta^*(s))\right\}\right).$$

This auxiliary mechanism is incentive-compatible because the payoff that θ_i obtains from deviating to any type $\hat{\theta}_i$ is

$$u_i(x_i(\hat{\theta}_i, \theta_{-i}), \theta_i) - \tilde{t}(\hat{\theta}_i, \theta_{-i}) = \sum_j u_j(x_j(\hat{\theta}_i, \theta_{-i}), \theta_j) - p(\hat{\theta}_i, \theta_{-i}) - \inf_{s \in \Theta_i} \left\{ U(\theta^*(s)) \right\},$$

which is identical to the payoff for the same deviation in the original mechanism up to the constant $-\inf_{s\in\Theta_i}\{U(\theta^*(s))\}$. Moreover, the auxiliary mechanism is individually rational because it delivers utility $U(\theta) - \inf_{s\in\Theta_i}\{U(\theta^*(s))\}$ to type θ_i, which is non-negative by construction. As the expected payoff for the principal of this auxiliary mechanism is

$$\mathbb{E}\Big[u_i(x_i(\theta), \theta_i) - \Big(U(\theta) - \inf_{s \in \Theta_i} \left\{ U(\theta^*(s)) \right\} \Big) \Big| \theta_{-i} \Big],$$

the desired inequality (A.1) then follows from the definition of $\pi_i(\theta_{-i})$. □

As in the proof of Theorem 1, by return-proofness and additivity of the return policies, the revenue that the seller collects from type θ must be

$$p(\theta) = \sum_i r(\theta, \{i\}).$$

The total revenue is then

$$\sum_i \mathbb{E}_\theta \big[r(\theta, \{i\}) \big].$$

Applying Lemma A.2 yields that total revenue is at most

$$\sum_i \mathbb{E}_{\theta_{-i}} \big[\pi_i(\theta_{-i}) \big].$$

Since θ_i and θ_{-i} are statistically independent, we have that $\pi_i(\theta_{-i}) = \pi_i$, and the theorem follows.

A.2. **Proof of Theorem 2′.** We follow the same strategy as the proof of Theorem 2, but use the following extension of Proposition 1 to the general screening model.

Proposition A.1. *In the general screening model, suppose that the agent's types are independent across components. If an incentive-compatible, return-proof mechanism (x, Z, p, r) satisfies*

$$\sum_{i \in G} r(\theta, \{i\}) \geq p(\theta)$$

for all types θ, then it yields a payoff to the principal of at most $\sum_i \pi_i$.

Proof. The *ex post* payoff to the principal from type θ satisfies

$$p(\theta) \leq \sum_{i \in G} r(\theta, \{i\}),$$

As in the proof of Theorem 3, we can use Lemma A.2 and the statistical independence of types across components to conclude that the principal's payoff is most $\sum_i \pi_i$. □

To complete the proof of Theorem 2′, we prove the contrapositive. Consider any mechanism M represented by a menu $\{p(q)\}_{q\in[0,1]^n}$ that does not discount any bundle.

We map stochastic mechanisms in the model without returns to deterministic mechanisms in the general multidimensional screening model with n additive components of Section 5: for each component i, the allocations are given by $q_i \in [0,1]$, the types are given by $\theta_i \in \Theta_i$, and the outside option is given by $\underline{q}_i = 0$. For each joint allocation $q = (q_i)_{i\in G}$, define the return policy $r : \mathcal{P}(G) \to \mathbb{R}$ as

$$r(Z) := p(q_G) - p(q_{G\setminus Z}),$$

where for any $B \subseteq G$, we write $q_B := (\mathbf{1}_{1\in B}\cdot q_1, \mathbf{1}_{2\in B}\cdot q_2, \ldots, \mathbf{1}_{n\in B}\cdot q_n)$.

As in the proof of Theorem 2, each type θ has an optimal choice of selecting the same choice as in the original mechanism and not returning anything, so the resulting mechanism is return-proof. Since M does not discount any bundle, we have that

$$\sum_i r(\theta,\{i\}) = \sum_i \big(p(q_G(\theta)) - p(q_{G\setminus\{i\}}(\theta))\big) = p(q_G(\theta)) - \text{Discount}(q_G(\theta)) \geq p(q_G(\theta))$$

for all θ. By Proposition A.1, M must have revenue at most $\sum_i \pi_i$.

A.3. **Proof of Theorem 4.** The proof proceeds in three steps. First, we provide an upper bound on an auxiliary one-dimensional screening problem with costly misreporting. Second, we use the auxiliary problem to further bound the expected return value, proving the result in the case where θ_0 is a deterministic, monotone function of θ_{-0}. Last, we extend the result to the general case of stochastic return costs.

Step 1. Consider an *auxiliary one-dimensional screening problem* with a compact type space $\Theta_i \subseteq \mathbb{R}_{\geq 0}$ and cumulative distribution function $F : \Theta_i \to [0,1]$. There is *costly misreporting* in the sense that there is a bounded, measurable *auxiliary function* $\psi : \Theta_i \to \mathbb{R}$ such that if a type θ_i misreports to be type $\theta_i' < \theta_i$, they must pay a cost of $\psi(\theta_i) - \psi(\theta_i')$. Thus, the incentive constraint is

$$\theta_i q_i(\theta_i) - t(\theta_i) \geq \theta_i q_i(\theta_i') - t(\theta_i') - \big(\psi(\theta_i) - \psi(\theta_i')\big) \qquad \text{for } \theta_i > \theta_i';$$

we do not impose incentive constraints for upward misreports. We impose individual rationality only for the lowest type: writing $\underline{\theta}_i := \min \Theta_i$, we require $\underline{\theta}_i q_i(\underline{\theta}_i) - t(\underline{\theta}_i) \geq 0$.

We now bound the optimal revenue in this auxiliary one-dimensional screening problem in terms of the optimal monopoly revenue and the auxiliary function ψ.

Proposition A.2. *In the auxiliary one-dimensional screening problem*

$$\text{Optimal Revenue} \le \sup_p \left\{p \cdot (1 - F(p))\right\} + \mathbb{E}_{\theta_i}\left[\psi(\theta_i)\right] - \psi(\underline{\theta}_i)\,.$$

To prove Proposition A.2, we first prove the result for the case of discrete types, and then extend to a general compact type space using an approximation argument.

Proof of Proposition A.2 under Discrete Types. Suppose that $\Theta_i = \{\theta_i^0, \theta_i^1, \ldots, \theta_i^K\}$. Consider the standard one-dimensional screening problem. By the optimality of posted prices (Myerson, 1981; Riley and Zeckhauser, 1983), the optimal revenue in this problem is given by $\sup_p\{p\cdot(1-F(p))\}$. By Yang (2022, Theorem 2), downward incentive constraints suffice for this problem: i.e., the optimal revenue in the relaxed problem

$$\max \sum_k f(\theta_i^k)t(\theta_i^k) \text{ subject to } \begin{cases} \theta_i^k q_i(\theta_i^k) - t(\theta_i^k) \ge \theta_i^k q_i(\theta_i^{\hat{k}}) - t(\theta_i^{\hat{k}}) & \text{for } 0 \le \hat{k} < k \le K \\ \theta_i^k q_i(\theta_i^k) - t(\theta_i^k) \ge 0 & \text{for } 0 \le k \le K \\ q_i(\theta_i^k) \le 1 & \text{for } 0 \le k \le K \end{cases}$$

equals $\sup_p\{p \cdot (1 - F(p))\}$. Note that the incentive constraint for θ_i^k to θ_i^0 and the individual rationality constraint for the lowest type θ_i^0 imply that

$$\theta_i^k q_i(\theta_i^k) - t(\theta_i^k) \ge \theta_i^k q_i(\theta_i^0) - t(\theta_i^0) \ge \theta_i^0 q_i(\theta_i^0) - t(\theta_i^0) \ge 0.$$

Hence, we can drop the individual rationality constraints for types other than the lowest without changing optimal revenue: i.e., the optimal revenue in the relaxed problem

$$\max \sum_k f(\theta_i^k)t(\theta_i^k) \text{ subject to } \begin{cases} \theta_i^k q_i(\theta_i^k) - t(\theta_i^k) \ge \theta_i^k q_i(\theta_i^{\hat{k}}) - t(\theta_i^{k'}) & \text{for } 0 \le \hat{k} < k \le K \\ \theta_i^0 q_i(\theta_i^0) - t(\theta_i^0) \ge 0 & \\ q_i(\theta_i^k) \le 1 & \text{for } 0 \le k \le K \end{cases}$$

also equals $\sup_p\{p \cdot (1 - F(p))\}$. This problem is a finite-dimensional linear programming problem with nonnegative variables $q_i(\theta_i^k)$ and unconstrained variables $t(\theta_i^k)$.

Let us consider the dual of this problem. The variables are Lagrange multipliers $\lambda_{k,\hat{k}}$ for $0 \le \hat{k} < k \le K$ on the downward incentive constraints, κ on the individual rationality constraint of θ_i^0, and ν_k for $0 \le k \le K$ on the allocation constraint, which must all be nonnegative. The dual constraint corresponding to $q_i(\theta_i^k)$ is[29]

$$\sum_{\hat{k}>k} \lambda_{\hat{k},k}\theta_i^{\hat{k}} - \sum_{\hat{k}<k} \lambda_{k,\hat{k}}\theta_i^k - \mathbf{1}_{k=0}\cdot\kappa\theta_i^0 + \nu_k \ge 0\,. \tag{A.2}$$

[29] See Carroll (2017, Equation (4.6)) for a version with individual rationality constraints for all types.

The dual constraint corresponding to $t(\theta_i^k)$ is the *flow conservation equation*[30]

$$\sum_{\hat{k}<k} \lambda_{k,\hat{k}} - \sum_{\hat{k}>k} \lambda_{\hat{k},k} + \mathbf{1}_{k=0} \cdot \kappa = f(\theta_k^i)\,. \tag{A.3}$$

Strong duality then guarantees that there is an (optimal) dual with

$$\sum_k \nu_k = \sup_p \left\{p \cdot (1 - F(p))\right\}. \tag{A.4}$$

Now consider any (incentive-compatible and individually rational) mechanism (q_i, t) for the auxiliary screening problem. (A.2) and the allocation constraint imply that

$$0 \le \sum_k \left(\sum_{\hat{k}>k} \lambda_{\hat{k},k} \theta_i^{\hat{k}} q_i(\theta_i^k) - \sum_{\hat{k}<k} \lambda_{k,\hat{k}} \theta_i^k q_i(\theta_i^k) + \nu_k \right) - \kappa \theta_i^0 q_i(\theta_i^0)\,.$$

By the flow conservation equation (A.3), we can write revenue as

$$\mathbb{E}[t(\theta_i)] = \sum_k f(\theta_i^k) t(\theta_i^k) = \sum_k \left(\sum_{\hat{k}<k} \lambda_{k,\hat{k}} t(\theta_i^k) - \sum_{\hat{k}>k} \lambda_{\hat{k},k} t(\theta_i^k) \right) + \kappa t(\theta_i^0)\,.$$

Using the previous inequality and grouping terms, it follows that revenue is at most

$$\mathbb{E}[t(\theta_i)] \le \sum_k \nu_k - \sum_k \sum_{\hat{k}<k} \lambda_{k,\hat{k}} \left(\theta_i^k q_i(\theta_i^k) - t_i(\theta_i^k) - \theta_i^k q_i(\theta_i^{\hat{k}}) + t_i(\theta_i^{\hat{k}}) \right) - \kappa \left(\theta_i^0 q_i(\theta_i^0) - t(\theta_i^0) \right).$$

Individual rationality for the lowest type then implies that

$$\mathbb{E}[t(\theta_i)] - \sum_k \nu_k \le - \sum_k \sum_{\hat{k}<k} \lambda_{k,\hat{k}} \left(\theta_i^k q_i(\theta_i^k) - t_i(\theta_i^k) - \theta_i^k q_i(\theta_i^{\hat{k}}) + t_i(\theta_i^{\hat{k}}) \right).$$

By incentive compatibility with costly misreporting, we then have that

$$\mathbb{E}[t(\theta_i)] - \sum_k \nu_k \le \sum_k \sum_{\hat{k}<k} \lambda_{k,\hat{k}} \left(\psi(\theta_i^k) - \psi(\theta_i^{\hat{k}}) \right) = \sum_k \left(\sum_{\hat{k}<k} \lambda_{k,\hat{k}} - \sum_{\hat{k}>k} \lambda_{\hat{k},k} \right) \psi(\theta_i^k)\,.$$

By the flow conservation equation (A.3), it follows that

$$\mathbb{E}[t(\theta_i)] - \sum_k \nu_k \le \sum_k f(\theta_i^k) \psi(\theta_i^k) - \kappa \psi(\theta_i^0) = \mathbb{E}_{\theta_i}[\psi(\theta_i^k)] - \kappa \psi(\underline{\theta}_i)\,.$$

Summing the flow conservation equation (A.3) over all types yields $\kappa = \sum_k f(\theta_k^i) = 1$, and by the dual optimality condition (A.4), we can conclude that

$$\mathbb{E}[t(\theta_i)] \le \sup_p \left\{p \cdot (1 - F(p))\right\} + \mathbb{E}_{\theta_i}\left[\psi(\theta_i)\right] - \psi(\underline{\theta}_i)\,.$$

This concludes the proof of the discrete type case of Proposition A.2. □

[30]See Sher and Vohra (2015, Equation (12)), Cai, Devanur, and Weinberg (2016, Lemma 4), and Carroll (2017, Equation (4.7)). For interpretation, $\lambda_{k,k'}$ is the flow from node k to node k', κ is the flow from node 0 to the sink, and $f(\theta_k^i)$ is the flow from the source to node k.

Completion of the Proof of Proposition A.2. Consider any mechanism (q_i, t_i) for the auxiliary screening problem. It suffices to prove that for each $\varepsilon > 0$, we have that

$$\mathbb{E}[t(\theta_i)] \le \varepsilon + \sup_p \left\{p \cdot (1 - F(p))\right\} + \mathbb{E}\left[\psi(\theta_i)\right] - \psi(\underline{\theta}_i)\,.$$

Let $S \ge \max\{\varepsilon, \underline{\theta}_i + \sup_{\theta_i \in \Theta}\{|\psi(\theta_i) - \psi(\underline{\theta}_i)| + \theta_i\}\}$, which is finite as ψ is bounded. Individual rationality for the lowest type implies that $t(\underline{\theta}_i) \le \underline{\theta}_i$. Incentive compatibility for deviations to $\underline{\theta}_i$ then implies that $t(\theta_i) \le S$ for all types θ_i. Let $\gamma < \frac{\varepsilon}{2S + \psi(\underline{\theta}_i)}$.

By Lusin's Theorem, there exists a compact subset $\tilde{\Theta}_i \subset \Theta_i$ such that (t_i, ψ) are continuous on $\tilde{\Theta}_i$ and $\mathbb{P}\left[\theta \in (\Theta_i \backslash \tilde{\Theta}_i)\right] < \gamma$. Let $\tilde{F}$ be the probability distribution of θ_i conditional on $\theta_i \in \tilde{\Theta}_i$. Moreover, without loss, we can assume $\underline{\theta}_i \in \tilde{\Theta}_i$. Then, we have

$$\mathbb{E}[t(\theta_i)] - \mathbb{E}[\psi(\theta_i)] \le (1 - \gamma)\Big(\mathbb{E}_{\tilde{F}}[t_i(\theta_i)] - \mathbb{E}_{\tilde{F}}[\psi(\theta_i)]\Big) + \gamma \cdot 2S\,.$$

Applying Yang (2022, Lemma 5) to the standard one-dimensional screening problem yields a sequence of finite approximations $(\Theta_i^{(m)}, F^{(m)})$ to the type space and distribution $(\tilde{\Theta}_i, \tilde{F})$, where $\Theta_i^{(m)} \subset \Theta_i$ and $\underline{\theta}_i^{(m)} = \underline{\theta}_i$, such that $F^{(m)} \to_d \tilde{F}$ and

$$\sup_p \left\{p \cdot (1 - \tilde{F}(p))\right\} \ge \limsup_m \left\{\sup_p \left\{p \cdot (1 - F^{(m)}(p))\right\}\right\}.$$

By the case of the proposition for discrete types, for all m, we have that

$$\mathbb{E}_{F^{(m)}}\left[t_i(\theta_i)\right] \le \sup_p \left\{p \cdot (1 - F^{(m)}(p))\right\} + \mathbb{E}_{F^{(m)}}\left[\psi(\theta_i)\right] - \psi(\underline{\theta}_i)\,.$$

Taking limits superior, by construction and continuity of t and ψ on $\tilde{\Theta}_i$, we have

$$\mathbb{E}_{\tilde{F}}\left[t(\theta_i)\right] \le \sup_p \left\{p \cdot (1 - \tilde{F}(p))\right\} + \mathbb{E}_{\tilde{F}}\left[\psi(\theta_i)\right] - \psi(\underline{\theta}_i)\,,$$

It follows that

$$\mathbb{E}[t(\theta_i)] - \mathbb{E}[\psi(\theta_i)] \le (1 - \gamma) \sup_p \left\{p \cdot (1 - \tilde{F}(p))\right\} - \psi(\underline{\theta}_i) + \gamma \cdot (2S + \psi(\underline{\theta}_i))\,.$$

Note also that for all $p \ge 0$, we have that

$$(1 - \gamma)p \cdot (1 - \tilde{F}(p)) = p \cdot \mathbb{P}\left[\theta_i \in \tilde{\Theta}_i \cap (p, \infty)\right] \le p \cdot \mathbb{P}\left[\theta_i \in (p, \infty)\right] = p \cdot (1 - F(p))\,.$$

Taking suprema over p yields that

$$(1 - \gamma) \sup_p \left\{p \cdot (1 - \tilde{F}(p))\right\} \le \sup_p \left\{p \cdot (1 - F(p))\right\},$$

and it follows that

$$\mathbb{E}[t(\theta_i)] - \mathbb{E}[\psi(\theta_i)] \le \sup_p \left\{p \cdot (1 - F(p))\right\} - \psi(\underline{\theta}_i) + \gamma \cdot (2S + \psi(\underline{\theta}_i))$$

$$= \sup_p \left\{ p \cdot (1 - F(p)) \right\} - \psi(\underline{\theta}_i) + \varepsilon \,,$$

as desired. □

Step 2. Let us now suppose that θ_0 is a deterministic monotone function of θ_{-0}. We use Proposition A.2 to prove an extension of Lemma 4 to stochastic mechanisms.

Lemma A.3. *With deterministic, monotone return costs, if* $(B, Z, p, r) : \Theta \to \Delta(\mathcal{X})$ *is incentive-compatible, then for all* i *and all* θ_{-i}*, we have that*

$$\mathbb{E}_{\theta_i}\left[\mathbb{E}_\omega \left[\mathbf{1}_{i \in B_\omega(\theta) \setminus Z_\omega(\theta)} \cdot r_\omega(\theta, \{i\}) \right] \middle| \theta_{-i} \right] \leq \sup_p \left\{ p \cdot \mathbb{P}(\theta_i \geq p \,|\, \theta_{-i}) \right\} + \mathbb{E}_{\theta_i}\left[\theta_0 \middle| \theta_{-i} \right].$$

Proof. Let $q_j(\theta) := \mathbb{P}_\omega\big(j \in (B_\omega(\theta) \setminus Z_\omega(\theta))\big)$ denote the allocation rule for any good j accounting for equilibrium returns, similarly let $\tilde{p}_\omega(\theta)$ denote the net payment accounting for equilibrium refund, and let $U(\theta)$ denote the equilibrium payoff of type θ. Also, let $\theta_0(\theta_{-0})$ denote the return cost of a type with value vector θ_{-0}.

Now, fix some good i. Let $\underline{\theta}_i := \min \Theta_i$. For any type θ, consider the following double deviation by type $\theta^* := (\theta_0(\underline{\theta}_i, \theta_{-i}), \underline{\theta}_i, \theta_{-i})$:

- Misreport to be type θ.
- Follow the return recommendations for type θ, but also return good i (if purchased) to get refund $r_\omega(\theta, \{i\})$.

This deviation yields type θ^* a payoff of at least

$$\sum_{j \neq i} \theta_j q_j(\theta) - \mathbb{E}_\omega\left[\tilde{p}_\omega(\theta)\right] + \mathbb{E}_\omega\left[\mathbf{1}_{i \in B_\omega(\theta) \setminus Z_\omega(\theta)} \cdot r_\omega(\theta, \{i\}) \right] - \theta_0^*$$
$$\geq U(\theta) - \theta_i q_i(\theta) + \mathbb{E}_\omega\left[\mathbf{1}_{i \in B_\omega(\theta) \setminus Z_\omega(\theta)} \cdot r_\omega(\theta, \{i\}) \right] - \theta_0^*.$$

Incentive compatibility implies that the above cannot be higher than $U(\theta^*)$, and thus

$$\mathbb{E}_\omega\left[\mathbf{1}_{i \in B_\omega(\theta) \setminus Z_\omega(\theta)} \cdot r_\omega(\theta, \{i\}) \right] \leq \theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*) \big) + \theta_0^*.$$

Hence, it suffices to prove that

$$\mathbb{E}_{\theta_i}\left[\theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*) \big) \middle| \theta_{-i} \right] \leq \sup_p \left\{ p \cdot \mathbb{P}(\theta_i \geq p \,|\, \theta_{-i}) \right\} + \mathbb{E}_{\theta_i}\left[\theta_0 \middle| \theta_{-i} \right] - \theta_0^*. \tag{A.5}$$

Now, fix θ_{-i} and suppose that the seller observes θ_{-i}. Consider an auxiliary mechanism for selling good i with allocation rule $q_i(\theta)$ and payment rule

$$\tilde{t}(\theta) := \theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*) \big).$$

We claim that this auxiliary mechanism is feasible in the auxiliary one-dimensional screening problem with the auxiliary function $\psi(\theta_i) = \theta_0(\theta_i, \theta_{-i})$. Note that type θ_i

obtains payoff $U(\theta)-U(\theta^*)$ from reporting truthfully under the auxiliary mechanism. To show the auxiliary mechanism is incentive-compatible under costly misreporting, note that the payoff that type θ_i obtains from deviating to any type $\hat{\theta}_i < \theta_i$ is

$$\begin{aligned}
&\theta_i q_i(\hat{\theta}_i,\theta_{-i}) - \tilde{t}(\hat{\theta}_i,\theta_{-i}) - \big(\theta_0(\theta_i,\theta_{-i}) - \theta_0(\hat{\theta}_i,\theta_{-i})\big)\\
&= \sum_j \theta_j q_j(\hat{\theta}_i,\theta_{-i}) - \mathbb{E}_\omega\Big[\tilde{p}_\omega(\hat{\theta}_i,\theta_{-i}) + \mathbf{1}_{Z_\omega(\hat{\theta}_i,\theta_{-i})\neq\varnothing}\theta_0(\hat{\theta}_i,\theta_{-i})\Big]\\
&\qquad\qquad - U(\theta^*) - \big(\theta_0(\theta_i,\theta_{-i}) - \theta_0(\hat{\theta}_i,\theta_{-i})\big)\\
&= \sum_j \theta_j q_j(\hat{\theta}_i,\theta_{-i}) - \mathbb{E}_\omega\Big[\tilde{p}_\omega(\hat{\theta}_i,\theta_{-i}) + \mathbf{1}_{Z_\omega(\hat{\theta}_i,\theta_{-i})\neq\varnothing}\theta_0(\theta_i,\theta_{-i})\Big]\\
&\qquad\qquad - U(\theta^*) - \mathbb{P}_\omega\big(Z_\omega(\hat{\theta}_i,\theta_{-i})=\varnothing\big)\cdot\big(\theta_0(\theta_i,\theta_{-i}) - \theta_0(\hat{\theta}_i,\theta_{-i})\big).
\end{aligned}$$

This deviation payoff is identical to the payoff from the same deviation in the original mechanism up to the constant $-U(\theta^*)$ and the quantity

$$-\mathbb{P}_\omega\big(Z_\omega(\hat{\theta}_i,\theta_{-i})=\varnothing\big)\cdot\big(\theta_0(\theta_i,\theta_{-i}) - \theta_0(\hat{\theta}_i,\theta_{-i})\big),$$

which is nonpositive due to the monotonicity of return costs. Moreover, the auxiliary mechanism is individually rational for the lowest type because it delivers utility 0 to type θ^*. As the expected revenue of this auxiliary mechanism is

$$\mathbb{E}_{\theta_i}\Big[\theta_i q_i(\theta) - \big(U(\theta) - U(\theta^*)\big)\,\Big|\,\theta_{-i}\Big],$$

the desired inequality (A.5) follows from Proposition A.2. □

Under costly returns with deterministic return costs, consider any incentive-compatible mechanism. As noted in Section 6 for deterministic mechanisms, by the additivity of the return policies, the *ex post* revenue that the seller collects from type θ must be

$$p_\omega(\theta) = \sum_i \mathbf{1}_{i\in B_\omega(\theta)\setminus Z_\omega(\theta)}\cdot r_\omega(\theta,\{i\})$$

The total revenue is then

$$\mathbb{E}_\theta\Big[\mathbb{E}_\omega\big[p_\omega(\theta)\big]\Big] = \sum_i \mathbb{E}_\theta\Big[\mathbb{E}_\omega\big[\mathbf{1}_{i\in B_\omega(\theta)\setminus Z_\omega(\theta)}\cdot r_\omega(\theta,\{i\})\big]\Big].$$

Applying Lemma A.3 yields that total revenue is at most

$$\sum_i \mathbb{E}_{\theta_{-i}}\Big[\sup_p\big\{p\cdot\mathbb{P}(\theta_i\geq p\,|\,\theta_{-i})\big\} + \mathbb{E}_{\theta_i}[\theta_0\,|\,\theta_{-i}]\Big] = \sum_i \mathbb{E}_{\theta_{-i}}\Big[\sup_p\big\{p\cdot\mathbb{P}(\theta_i\geq p\,|\,\theta_{-i})\big\}\Big] + n\cdot\mathbb{E}[\theta_0].$$

As in the proof of Theorem 1, we then use the statistical independence of θ_i and θ_{-i} to conclude that total revenue must be at most $\sum_i \pi_i + n\cdot\mathbb{E}[\theta_0]$.

Step 3. To complete the argument, we remove the assumption that θ_0 is a deterministic function of θ_{-0} using a coupling argument.

Let $F(\theta_0 \,|\, \theta_{-0})$ denote the conditional cumulative distribution function of of θ_0 conditional on θ_{-0}, and let $F^{-1}(q \,|\, \theta_{-0}) = \inf\{\theta_0 \,|\, F(\theta_0 \,|\, \theta_{-0}) \geq q\}$ denote the corresponding quantile function. Letting q be a uniform $[0,1]$ random variable independent of θ_{-0}, by the inverse probability integral transform, we have that

$$(\theta_0, \theta_{-0}) \stackrel{d}{=} (F^{-1}(q \,|\, \theta_{-0}), \theta_{-0}) \,.$$

Note that stochastic monotonicity implies that $F^{-1}(q \,|\, \cdot)$ is an increasing function. Hence, for each q, applying the result for the deterministic return costs case, we have

$$\mathbb{E}_\theta[t(\theta) \,|\, q] \leq \sum_i \pi_i + n \cdot \mathbb{E}_{\theta_{-0}}\big[F^{-1}(q \,|\, \theta_{-0}) \,\big|\, q\big] \,.$$

Taking expectations over q then yields that

$$\begin{aligned} \mathbb{E}[t(\theta)] = \mathbb{E}_q\big[\mathbb{E}[t(\theta) \,|\, q]\big] &\leq \mathbb{E}_q\Big[\sum_i \pi_i + n \cdot \mathbb{E}_{\theta_{-0}}\big[F^{-1}(q \,|\, \theta_{-0})\big]\Big] \\ &= \sum_i \pi_i + n \cdot \mathbb{E}_{q,\theta_{-0}}[F^{-1}(q \,|\, \theta_{-0})] = \sum_i \pi_i + n \cdot \mathbb{E}[\theta_0] \,. \end{aligned}$$

A.4. **Proof of Proposition 2.** We use the same strategy as the proof of Theorem 1.
Step 1. We first show that it remains without loss of generality to restrict to return-proof mechanisms under α-additive return policies.

Lemma A.4. *For any incentive-compatible mechanism M with α-additive return policies, there exists a return-proof, incentive-compatible mechanism $\tilde{M}$ with α-additive return policies that yields the same revenue.*

Proof. The lemma follows from the proof of Lemma 2: if M has α-additive return policies, then the constructed mechanism $\tilde{M}$ also has α-additive return policies. □

Step 2. By Lemma A.4, we can restrict without loss of generality to the return-proof mechanisms. By the α-additivity of the return policies, the *ex post* revenue that the seller collects from type θ must then be at most

$$p_\omega(\theta) \leq \frac{1}{\alpha} \sum_i r_\omega(\theta, \{i\})$$

where we write $r_\omega(\theta, \{i\}) = 0$ for $i \notin B_\omega(\theta)$. The total revenue is then

$$\mathbb{E}_\theta\big[\mathbb{E}_\omega[p_\omega(\theta)]\big] \leq \frac{1}{\alpha} \sum_i \mathbb{E}_\theta\big[\mathbb{E}_\omega[r_\omega(\theta, \{i\})]\big] \,.$$

As in the proof of Theorem 1, we can apply Lemma 3 to conclude that total revenue must be at most $\frac{1}{\alpha}\sum_i \pi_i$.

A.5. **Proof of Proposition 3.** We use the same strategy as the proof of Theorem 1. **Step 1.** We first show that it remains without loss of generality to restrict to return-proof mechanisms with incomplete refundability under refundability within δ.

Lemma A.5. *With incomplete refundability, for any incentive-compatible mechanism M with additive return policies that is refunding within δ, there exists a return-proof, incentive-compatible mechanism $\tilde{M}$ with additive return policies that is refunding within δ and yields the same expected revenue.*

Proof. The lemma follows from the proof of Lemma 2: if M is refunding within δ, then the constructed mechanism $\tilde{M}$ is also refunding within δ. □

Step 2. By Lemma A.5, we can restrict without loss of generality to the return-proof mechanisms. By the additivity of the return policies and within-δ refundability, the *ex post* revenue that the seller collects from type θ must then be at most

$$p_\omega(\theta) \le r_\omega(B_\omega(\theta)) + \delta \le \sum_i r_\omega(\theta, \{i\}) + \delta$$

where we write $r_\omega(\theta, \{i\}) = 0$ for $i \notin B_\omega(\theta)$. The total revenue is then

$$\mathbb{E}_\theta\big[\mathbb{E}_\omega[p_\omega(\theta)]\big] \le \sum_i \mathbb{E}_\theta\big[\mathbb{E}_\omega[r_\omega(\theta, \{i\})]\big] + \delta\,.$$

As in the proof of Theorem 1, we can apply Lemma 3 to conclude that total revenue must be at most $\sum_i \pi_i + \delta$.

A.6. **Proof of Proposition 5.** Note that for all $q \in [0,1]$, we have that

$$(1-q)\cdot\big(F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big) \le \max\big\{F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q), 0\big\}.$$

Taking suprema over q yields that

$$\sup_q\Big\{(1-q)\cdot\big(F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big)\Big\} \le \max\Big\{\sup_q\big\{F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big\}, 0\Big\}.$$

As $F^{-1}(0\,|\,\theta_i) \ge F^{-1}(0)$ almost surely, it follows that

$$\sup_q\Big\{(1-q)\cdot\big(F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big)\Big\} \le \sup_q\Big\{F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\Big\}$$

almost surely. Noting that

$$\sup_p\Big\{p\cdot\mathbb{P}(\theta_i \ge p\,|\,\theta_{-i})\Big\} = \sup_q\Big\{(1-q)\cdot F_i^{-1}(q\,|\,\theta_{-i})\Big\} \quad\text{and}\quad \pi_i = \sup_q\Big\{(1-q)\cdot F_i^{-1}(q)\Big\},$$

taking expectations yields that

$$\begin{aligned}\overline{\pi}_i - \pi_i &\leq \mathbb{E}_{\theta_{-i}}\Big[\sup_q \Big\{(1-q)\cdot\big(F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big)\Big\}\Big] \\ &\leq \mathbb{E}_{\theta_{-i}}\Big[\sup_q \big\{F_i^{-1}(q\,|\,\theta_{-i}) - F_i^{-1}(q)\big\}\Big].\end{aligned}$$

Appendix B. Necessity of Bundle Discounts: Quantitative Version

Consider the model without returns. In this appendix, we show that under independent and additive values, separate sales provides a constant fraction of the revenue of any bundling mechanism with a cap on the magnitude of bundle discounts as a fraction of bundle prices. Our argument for the case of deterministic mechanisms uses our analysis of non-additive return policies (Proposition 2 from Section 7.1).

We consider mechanisms that discount bundles by at most β fraction of their prices. Formally, we say that a menu $\{p(q)\}_{q\in[0,1]^n}$ offers *at most β-bundle discounts* if

$$\text{Discount}(q) \leq \beta \cdot p(q)$$

for all (randomized) bundles q. For example, in a "Buy One Get One Free" sale for two equally priced items, the discount offered for the bundle is equals the bundle price, so the mechanism does not offer at most β-bundle discount for any $\beta < 1$. By contrast in a "Buy One Get One 50% Off" sale for two items priced at \$10 each, the discount offered for the bundle is \$5, which is $\frac{1}{3}$ of the bundle price. Thus, this second mechanism offers at most β-bundle discount for all $\beta \geq \frac{1}{3}$.

We show that under independent and additive values, separate sales yields at least $(1-\beta)$ fraction of the revenue of any mechanism that offers at most β-bundle discounts.

Proposition B.1. *Consider the model without returns and let $0 \leq \beta \leq 1$. Suppose that the buyer's values are independent across goods. For any mechanism M that is represented by a menu $\{p(q)\}_{q\in[0,1]^n}$ that offers at most β-bundle discounts, we have*

$$\sum_i \pi_i \geq (1-\beta)\cdot \text{Rev}(M)\,.$$

Without a bound on bundle discounts, separate sales cannot achieve a constant fraction of optimal revenue as the number of items grows even under independent values (Hart and Nisan 2017). In the examples constructed in the literature, the optimal mechanism is *pure bundling*, where the item prices are set to equal bundle price, and hence (as in the case of a "Buy One Get One Free" sale) has an extreme

discount for the grand bundle of the entire bundle price.[31] Conversely, Proposition B.1 shows that whenever we can bound the extent of bundle discounts, then separate sales can achieve a substantial faction of the revenue depending on the bundle discounts.

To illustrate the logic behind the proof of Proposition B.1, we sketch an argument for the case of deterministic mechanisms. As in Section 4.1, for any bundle B, consider the "discount forfeited on return" policy $r(Z) = p(B) - p(B \setminus Z)$. Since M offers at most β-bundle discounts for each bundle $B \setminus Z$, we have

$$\begin{aligned}\sum_{i \in B \setminus Z} \big(r(Z \cup \{i\}) - r(Z)\big) &= \sum_{i \in B \setminus Z} \Big(p(B \setminus Z) - p\big((B \setminus Z) \setminus \{i\}\big)\Big) \\ &= p(B \setminus Z) - \text{Discount}(B \setminus Z) \\ &\geq (1-\beta) \cdot p(B \setminus Z) = (1-\beta) \cdot \big(p(B) - r(Z)\big)\,.\end{aligned}$$

So, the return policy is $(1-\beta)$-additive. We can then use Proposition 2 to show that separate sales must achieve at least $(1-\beta)$ fraction of the revenue of M.

For the general argument, as in the proof of Theorem 2′, we instead map stochastic mechanisms in the model without returns to deterministic mechanisms in the general multidimensional screening model of Section 5. We then combine the logic behind Proposition A.1 with the approximation argument behind Proposition 2.[32]

Proof of Proposition B.1. As in the proof of Theorem 2′, defining the return policy $r : \mathcal{P}(G) \to \mathbb{R}$ for each joint allocation $q = (q_i)_{i \in G}$ as

$$r(Z) := p(q_G) - p(q_{G \setminus Z})\,,$$

the resulting mechanism is return-proof. Hence, by Lemma A.2, we have that

$$\mathbb{E}\big[r(\theta, \{i\})\big] \leq \pi_i\,,$$

where $r(\theta, \{i\})$ is the refund for component i, given the original allocation being type θ's equilibrium allocation $q(\theta)$. Since M offers at most β bundle discounts, we have

$$\sum_i r(\theta, \{i\}) = \sum_i \big(p(q_G(\theta)) - p(q_{G \setminus \{i\}}(\theta))\big) \geq (1-\beta) \cdot p(q_G(\theta))$$

for all θ. Thus, we have that

$$\sum_i \pi_i \geq \mathbb{E}\big[r(\theta, \{i\})\big] \geq (1-\beta)\mathbb{E}\big[p(q(\theta))\big] = (1-\beta)\text{Rev}(M)\,,$$

proving the result. □

[31] Indeed, under independent values, *either* separate sales *or* pure bundling (depending on the distributions) guarantees $\geq \frac{1}{6}$ of the optimal revenue (Babaioff, Immorlica, Lucier, and Weinberg 2014).
[32] Similar logic shows that Proposition 2 extends to the general screening model of Section 5.

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.

Babaioff, M., Y. A. Gonczarowski, and N. Nisan (2017). The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC '17, pp. 869–877. Association for Computing Machinery.

Babaioff, M., N. Immorlica, B. Lucier, and S. M. Weinberg (2014). A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of the IEEE 55th Annual Symposium on Foundations of Computer Science*, FOCS '14, pp. 21–30. Institute of Electrical and Electronics Engineers.

Babaioff, M., N. Nisan, and A. Rubinstein (2018). Optimal deterministic mechanisms for an additive buyer. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, New York, NY, USA, pp. 429. Association for Computing Machinery.

Battaglini, M. and T. R. Palfrey (2024). Organizing for collective action: Olson revisited. *Journal of Political Economy 132*(9), 2881–2936.

Ben-Moshe, R., S. Hart, and N. Nisan (2024). Monotonic mechanisms for selling multiple goods. Working paper, Hebrew University of Jerusalem.

Bergemann, D., A. Bonatti, A. Haupt, and A. Smolin (2021). The optimality of upgrade pricing. In *International Conference on Web and Internet Economics*, WINE '21, pp. 41–58. Springer.

Bergemann, D., F. Castro, and G. Y. Weintraub (2020). The scope of sequential screening with ex post participation constraints. *Journal of Economic Theory 188*, 105055.

Bulow, J. and J. Roberts (1989). The simple economics of optimal auctions. *Journal of Political Economy 97*(5), 1060–1090.

Cai, Y., N. R. Devanur, and S. M. Weinberg (2016). A duality-based unified approach to bayesian mechanism design. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, STOC '16, pp. 926–939. Association for Computing Machinery.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica 85*(2), 453–488.

Chawla, S., J. D. Hartline, and R. Kleinberg (2007). Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pp. 243–251. Association for Computing Machinery.

Che, Y.-K. and W. Zhong (Forthcoming, 2025). Robustly optimal mechanisms for selling multiple goods. *Review of Economic Studies*.

Compte, O. and P. Jehiel (2007). On quitting rights in mechanism design. *American Economic Review 97*(2), 137–141.

Compte, O. and P. Jehiel (2009). Veto constraint in mechanism design: Inefficiency with correlated types. *American Economic Journal: Microeconomics 1*(1), 182–206.

Daskalakis, C., A. Deckelbaum, and C. Tzamos (2013). Mechanism design via optimal transport. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pp. 269–286. Association for Computing Machinery.

Daskalakis, C., A. Deckelbaum, and C. Tzamos (2017). Strong duality for a multiple-good monopolist. *Econometrica 85*(3), 735–767.

Deb, R. and A.-K. Roesler (2024). Multi-dimensional screening: Buyer-optimal learning and informational robustness. *Review of Economic Studies 91*(5), 2744–2770.

Forges, F. (1999). Ex post individually rational trading mechanisms. In A. Alkan, C. D. Aliprantis, and N. Yannelis (Eds.), *Current Trends in Economics: Theory and Applications*, Volume 8 of *Studies in Economic Theory*, pp. 157–175. Springer.

Frick, M., R. Iijima, and Y. Ishii (2024). Multidimensional screening with rich consumer data. Working paper, Princeton University.

Ghili, S. (2023). A characterization for optimal bundling of products with nonadditive values. *American Economic Review: Insights 5*(3), 311–326.

Haberman, A. and R. Jagadeesan (2024). Auctions with withdrawal rights: A foundation for uniform price. Working paper, Stanford University.

Haghpanah, N. and J. Hartline (2021). When is pure bundling optimal? *Review of Economic Studies 88*(3), 1127–1156.

Hart, S. and N. Nisan (2017). Approximate revenue maximization with multiple items. *Journal of Economic Theory 172*, 313–347.

Hart, S. and P. J. Reny (2015). Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics 10*(3), 893–922.

Heumann, T. (2020). Information design and sequential screening with ex post participation constraint. *Theoretical Economics 15*(1), 319–359.

Krähmer, D. and R. Strausz (2015). Optimal sales contracts with withdrawal rights. *Review of Economic Studies 82*(2), 762–790.

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.

Manelli, A. M. and D. R. Vincent (2007). Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory 137*(1), 153–185.

Matthews, S. A. and A. Postlewaite (1989). Pre-play communication in two-person sealed-bid double auctions. *Journal of Economic Theory 48*(1), 238–263.

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.

Mussa, M. and S. Rosen (1978). Monopoly and product quality. *Journal of Economic Theory 18*(2), 301–317.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–73.

Myerson, R. B. (1982). Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics 10*(1), 67–81.

National Retail Federation (2023). 2023 consumer returns in the retail industry. Available at `http://nrf.com/research/2023-consumer-returns-retail-industry`.

Pavlov, G. (2011). Optimal mechanism for selling two goods. *B.E. Journal of Theoretical Economics 11*(1), 000010220219351704 1664.

Riley, J. and R. Zeckhauser (1983). Optimal selling strategies: When to haggle, when to hold firm. *Quarterly Journal of Economics 98*(2), 267–289.

Sher, I. and R. Vohra (2015). Price discrimination through communication. *Theoretical Economics 10*(2), 597–648.

Stigler, G. J. (1963). United States v. Loew's Inc.: A note on block-booking. *Supreme Court Review 1963*, 152–157.

Yang, F. (2022). Costly multidimensional screening. Working paper, Stanford University.

Yang, F. (2023). Nested bundling. Working paper, Stanford University.