Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Sequential Mechanisms with Ex Post Individual Rationality

Itai Ashlagi, Constantinos Daskalakis, Nima Haghpanah

informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 71, No. 1, January–February 2023, pp. 245–258
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

Sequential Mechanisms with Ex Post Individual Rationality

Itai Ashlagi,[a] **Constantinos Daskalakis,**[b] **Nima Haghpanah**[c,*]

[a] Department of Management Science and Engineering, Stanford University, Stanford, California 94305; [b] Computer Science & Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; [c] Department of Economics, Pennsylvania State University, University Park, Pennsylvania 16802

*Corresponding author

Contact: iashlagi@stanford.edu (IA); costis@csail.mit.edu (CD); haghpanah@psu.edu, https://orcid.org/0000-0001-7025-4282 (NH)

Received: April 28, 2021
Revised: October 11, 2021
Accepted: May 31, 2022
Published Online in Articles in Advance: August 12, 2022

Area of Review: Market Analytics and Revenue Management

https://doi.org/10.1287/opre.2022.2332

Abstract. We study optimal mechanisms for selling multiple products to a buyer who learns her values for those products sequentially. A mechanism may use static prices or adjust them over time, and it may sell the products separately or as bundles. We study mechanisms that provide the buyer a nonnegative ex post utility. We show that there exists an optimal mechanism that determines the allocation of each product as soon as the buyer learns her value for that product. This observation allows us to solve for optimal mechanisms recursively. We use this recursive characterization to show that static mechanisms are suboptimal if the buyer first learns her values for products that are ex ante less valuable. Under this condition, the ability to bundle products is less profitable than the ability to adjust prices dynamically.

Funding: This work was supported by ONR [Grant N00014-12-1-0999] and NSF Awards CCF-0953960 (CAREER), CCF-1551875 and SES-1254768. Part of this work was done while the authors were visiting the Simons Institute for Theory of Computing.

Supplemental Material: The e-companion is available at https://doi.org/10.1287/opre.2022.2332.

Keywords: sequential mechanisms • multiproduct monopolist • ex post individual rationality

1. Introduction

Online multiproduct sellers increasingly use interactive websites to customize their offers to individual buyers. For example, a user who clicks on the first Harry Potter movie on Amazon is shown a "Bundle and save" offer to buy all eight Harry Potter movies at a discounted price. Additionally, insurance providers make personalized offers to a user, involving discounts to buy bundles of insurance products, after she fills out an online form that solicits the buyer's preferences and characteristics. What selling strategy should a multiproduct seller use to maximize profit? Should he offer products as bundles or sell them individually? Should he use static prices or adjust them dynamically based on user interaction? Should he combine these two instruments and use both bundling and dynamically adjusted prices?

We study these questions in a setting with a rich class of selling strategies. There are a number of products, and the buyer learns her values for products sequentially, one product in each period. These values are drawn independently from known distributions. The selling strategy is a mechanism that specifies a set of possible decisions for the buyer in each period as well as the eventual allocation of products and the payment as a function of all these decisions. A special case is the class of static mechanisms in which prices do not change over time but products may be sold as bundles. Another special case is when the products are sold separately but at dynamically adjusted prices. A general mechanism may combine these two instruments and sell products as bundles and at dynamically adjusted prices.

We restrict attention to mechanisms in which the buyer has an ex post nonnegative utility. That is, after the buyer learns all of her values, her utility for the allocation and prices specified by the mechanism must be nonnegative. This restriction excludes mechanisms in which the seller "sells the store in advance" to the buyer. In such a mechanism, before the buyer learns her values, the seller offers her the grand bundle of products at a price equal the buyer's expected value. The constraint that ex post utility must be nonnegative allows us to compare dynamic and static mechanisms on an equal footing by isolating the ability to adjust prices over time from the ability to charge the buyer advance payments before she learns her values. For a static mechanism, our ex post nonnegative utility constraint is equivalent to the standard notion of individual rationality. Thus, our class of mechanisms includes well-studied static multiproduct screening mechanisms (going back to Stigler 1963, Adams and Yellen 1976, McAfee et al. 1989).

Our first result is that, in order to maximize profit, it is without loss of generality to restrict attention to a class of *separable* mechanisms. A separable mechanism has two features. First, it sells the products separately. That is, the allocation of each product is specified immediately once the value for that product is revealed to the buyer. Second, in each period, the buyer simply reports the value learned in that period. A special case is when the buyer is offered a deterministic price based on previous interactions, but in general, a separable mechanism may be randomized.

A separable mechanism is handicapped because it sells the products separately. It does not have the ability to "bundle" the products by arbitrarily tying the allocation of one product to the value for another. To see this, consider selling two products, and suppose that the value for each product is either 1 or 2. A static mechanism can bundle the products. For example, it can offer the two products only as a bundle at a take it or leave it price of 3. If the buyer's value for either one of the products is 2, she buys the bundle, and otherwise, she buys nothing. In this static mechanism, the allocation of the first product depends on the buyer's value for the second product. So, no separable mechanism can implement this allocation.

Bundling is a strong instrument to screen types in static settings (McAfee et al. 1989). Because a separable mechanism cannot use such an instrument and given that the ex post utility constraint restricts the use of advance payments, it is a priori not clear that a separable mechanism can be optimal. Indeed, as we discuss later, static mechanisms may be suboptimal with correlated values. Nevertheless, in our setting with independent values, we show that any mechanism can be converted to a separable one with the same profit. The main insight is that dynamic screening is a weakly more powerful instrument than bundling.

Figure 1. Static vs. Optimal Mechanisms

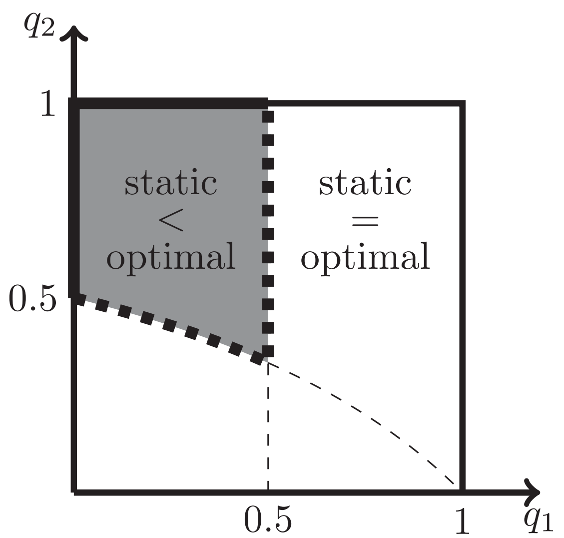

Notes. The dark-shaded region is the set of (q_1, q_2) for which static mechanisms are suboptimal. The right and bottom boundaries are not in the set.

The fact that separable mechanisms are optimal significantly simplifies the problem because optimal separable mechanisms can be characterized via standard recursive methods (Green 1987, Spear and Srivastava 1987, Thomas and Worrall 1990). In particular, an optimal mechanism maintains a state variable, the *promised utility*, which is the buyer's expected utility. The promised utility affects the allocation and is updated in each period. For any given period and any promised utility, the optimal allocation can be characterized via backward induction. In particular, in each period, the optimal allocation maximizes the seller's expected revenue given how much revenue the seller can extract in future periods for any updated promised utility.

We use our recursive characterization to identify conditions under which static mechanisms are (strictly) suboptimal. With two types and two values, we provide a complete characterization. To describe this characterization, suppose that the value for each product is either 1 or 2. Let q_1 be the probability that that the first product's value equals 2 and q_2 be the probability that the second product's value equals 2. The set of possible pairs (q_1, q_2) for which static mechanisms are suboptimal is specified in Figure 1. Roughly speaking, static mechanisms are suboptimal if and only if q_1 is low and q_2 is high: that is, the first product is ex ante less valuable than the second product.

We generalize this insight to any number of products and values. In particular, we show that static mechanisms are suboptimal if the first product is ex ante less valuable than the last product in the sense that it has lower monopoly prices.[1] To see the connection between the two results, consider again two products with values of 1 and 2. If $q_1 < 0.5 < q_2$, the optimal monopoly price for the first product is 1, and the optimal monopoly price for the second product is 2. In this case, as shown in Figure 1, static mechanisms are suboptimal. Thus, the result for any number of products and values partially generalizes the result for the case of two values and two products. Under this condition, the ability to bundle products is strictly less profitable than the ability to screen types dynamically.

We study the robustness of our results to the case of correlated values via numerical calculations with two products and two values. These calculations suggest that our main results extend if values are positively correlated but fail if they are negatively correlated. In particular, if values are positively correlated, separable mechanisms remain optimal, and static mechanisms are suboptimal if the first product has a lower monopoly price than the second product. Both of these two conclusions fail if values are negatively correlated. We leave a thorough analysis of correlated values to future work.

1.1. Related Work

Closest to our work are the papers that consider selling multiple products with ex post participation constraints.

These papers assume that the allocation of each product must be specified when the product arrives (because the product would perish otherwise). As a result, unlike ours, these papers are not concerned with the performance of static bundling mechanisms. Papadimitriou et al. (2016) show that when the buyer's values are correlated, finding optimal mechanisms is computationally hard. Mirrokni et al. (2016) characterize approximately optimal mechanisms with multiple buyers recursively. Mirrokni et al. (2020) consider the design of approximately optimal mechanisms when buyers have different expectations of future distributions. Balseiro et al. (2018) impose a martingale constraint on the buyer's utility and show that the seller's profit approaches first best (full surplus extraction) as the number of products grows.

The ex post utility constraint is related to limited liability constraints in dynamic principal agent models. Krishna et al. (2013), Krähmer and Strausz (2017), and Grillo and Ortner (2018) study contracts in which the agent's stage utility is nonnegative. In our setting, a mechanism with nonnegative stage utilities also has nonnegative ex post utilities. Nonetheless, the two constraints are equivalent when solving for optimal mechanisms because a separable mechanism satisfies the nonnegative stage utility constraint. Relatedly, Sappington (1983), Clementi and Hopenhayn (2006), and DeMarzo and Sannikov (2006) assume that the agent cannot make monetary transfers to the principal.

Ex post participation constraints have also been studied for selling a single product. Krähmer and Strausz (2015) consider a problem where the seller has a single item to sell and the buyer sequentially receives signals about her valuation. They show that assuming a monotone hazard rate condition, static mechanisms are optimal. Bergemann et al. (2017) consider the same setting and provide necessary and sufficient conditions for optimality of static mechanisms. Krähmer and Strausz (2016) consider a multiunit extension of the problem. When the buyer's utility is linear in quantity but the seller's costs are nonlinear, static mechanisms are suboptimal. The main question studied in these papers, namely optimality of static mechanisms, is similar to ours. Nevertheless, the settings and results are different.

More broadly, our work relates to two well-studied branches of literature on mechanism design, namely multiproduct bundling and dynamic mechanism design.

The literature on multiproduct bundling goes back to Stigler (1963) and Adams and Yellen (1976). This literature considers static mechanisms. That is, the buyer walks into the store knowing her values (alternatively, the buyer learns no new information about her values). McAfee et al. (1989) and Manelli and Vincent (2007) show that optimal screening mechanisms typically involve mixed bundling (i.e., offering a menu of bundles and prices). More generally, the literature shows that optimal mechanisms are complex. The optimal menu may include unboundedly many randomized bundles (Manelli and Vincent 2007). As such, characterizations of optimal mechanisms are rare. Exceptions exist, such as Rochet and Chone (1998) and Daskalakis et al. (2017). Rochet and Chone (1998) characterize optimal mechanisms via a sweeping procedure that generalizes ironing. Daskalakis et al. (2017) characterize optimal mechanisms via a dual measure that satisfies certain stochastic dominance conditions. To apply either characterization, one must be able to identify sweeping procedure or the dual measure for which no general construction is known. In contrast, optimal mechanisms can be characterized recursively in our dynamic setting.

The literature on dynamic mechanism design is similarly broad. The main thrusts in this literature study dynamic arrivals and departures of agents, such as in Pai and Vohra (2008) and Gershkov and Moldovanu (2009, 2010), and agents whose private information evolves, such as in Courty and Li (2000), Eső and Szentes (2007), Bergemann and Välimäki (2010), Boleslavsky and Said (2012), Kakade et al. (2013), and Pavan et al. (2014). Garrett (2016) combines these two branches by considering a setting with dynamic arrival and evolving values. A main difference with our paper is in the ex post nonnegative utility constraint. Prior literature, with the exceptions we discussed before, considers weaker notions of individual rationality requiring that in the beginning of each period, the expected utility from all future periods must be nonnegative.

2. The Model

A seller has k products to sell to a single buyer. The cost of production is normalized to zero. The buyer's value for product $i \in \{1, \ldots, k\}$ is $v_i \in V_i \subseteq \mathbb{R}^+$. Assume that V_i is finite. Each value v_i is drawn independently from all other values with probability $f_i(v_i) > 0$. The distributions f_1 to f_k are commonly known to the seller and the buyer. We refer to $v = (v_1, \ldots, v_k)$ as the ex post type of the buyer. The utility of an ex post type v for receiving a set of products $S \subseteq \{1, \ldots, k\}$ and transferring $t \in \mathbb{R}$ units of money to the seller is $(\sum_{i \in S} v_i) - t$. The buyer is risk neutral; that is, the utility of receiving each product i with probability a_i and a random monetary transfer with expectation t to the seller is $v \cdot a - t = (\sum_i v_i a_i) - t$.

The buyer privately learns her values over time. In particular, in each period i from one to k, v_i is privately revealed to the buyer. Thus, in period i, the buyer knows values v_1 to v_i. Define $\Theta^i = \prod_{j=1}^{i} V_j$. For a given v, let $v^i = (v_1, \ldots, v_i) \in \Theta^i$ be the first i components of v. If the buyer's ex post type is v, her interim type in period i is v^i.

We focus on direct incentive compatible mechanisms. A (direct) mechanism (a, t) consists of an allocation rule $a_i : \Theta^1 \times \ldots \times \Theta^k \to [0,1]$ for each $i \in \{1, \ldots, k\}$ and a transfer rule $t : \Theta^1 \times \ldots \times \Theta^k \to \mathbb{R}$. The interpretation is that in each period i, upon realizing each value v_i, the buyer reports an interim type $\theta^i \in \Theta^i$ to the mechanism. At the end of the last period k, the buyer receives each product i with probability $a_i(\theta^1, \ldots, \theta^k)$ and transfers $t(\theta^1, \ldots, \theta^k)$ units of money to the mechanism. Notice that our mechanisms allow the buyer to "rereport" all the values she has observed so far. The reason is that we would like to define the class of all mechanisms generally so that it contains several interesting classes as special cases. For instance, as we see shortly, two special cases are static mechanisms and ones in which the agent only reports her value v_i in each period i.

A mechanism is periodic incentive compatible (PIC) if the agent maximizes her expected utility in each period by reporting her type truthfully, regardless of past reports. Formally, a mechanism (a, t) is PIC if for each period i, interim type $(v_1, \ldots, v_i)$, history of reports $\theta^1, \ldots, \theta^{i-1}$, and possible report θ^i in period i, we have

$$\begin{aligned} &\mathop{\mathbf{E}}_{v_{i+1}, \ldots, v_k}\Big[v \cdot a\Big(\theta^1, \ldots, \theta^{i-1}, v^i, v^{i+1}, \ldots, v^k\Big) \\ &\qquad - t\Big(\theta^1, \ldots, \theta^{i-1}, v^i, v^{i+1} \ldots, v^k\Big)\Big] \\ &\geq \mathop{\mathbf{E}}_{v_{i+1}, \ldots, v_k}\Big[v \cdot a\Big(\theta^1, \ldots, \theta^{i-1}, \theta^i, v^{i+1}, \ldots, v^k\Big) \\ &\qquad - t\Big(\theta^1, \ldots, \theta^{i-1}, \theta^i, v^{i+1}, \ldots, v^k\Big)\Big]. \end{aligned}$$

(Recall that for a given v, $v^j = (v_1, \ldots, v_j)$.) The left-hand side is the agent's expected utility from reporting her type truthfully in periods i to k, following the history of reports $\theta^1, \ldots, \theta^{i-1}$. The right-hand side is the expected utility from reporting θ^i in period i and reporting truthfully in periods $i+1$ to k, following the history of reports $\theta^1, \ldots, \theta^{i-1}$. Notice that backward induction implies that, regardless of what the agent reports in period i, reporting truthfully in periods $i + 1$ to k is indeed the optimal strategy in those periods. Therefore, PIC implies that the agent maximizes her expected utility by reporting her types truthfully over all possible strategies that may involve misreporting her types in the future periods.

A mechanism is ex post individually rational (ex post IR) if it guarantees nonnegative utility for the buyer. Let us abuse notation and denote by $a(v)$ and $t(v)$ the outcome of the mechanism if the buyer reports all of her interim types consistent with an ex post type v (that is, $a(v) = a(v^1, v^2, \ldots, v^k)$) and similarly for t. A mechanism (a, t) is ex post IR if at the end of period k, given the buyer's optimal strategy (reporting truthfully), the expected utility of the buyer is nonnegative,

$$v \cdot a(v) - t(v) \geq 0,$$

for all ex post types v. Note that a_i denotes the probability of allocation. Thus, the ex post individual rationality states that the utility of the buyer is nonnegative for all ex post types v but *in expectation* over the random choices of the mechanism. Even though the ex post IR constraint is written in expectation, it is possible to guarantee nonnegative utility for all random choices of the mechanism by appropriately correlating transfers with allocation. We defer the argument to Section A in the e-companion. Following that argument, we abuse terminology and refer to the constraint as the ex post IR constraint, even though it is written in expectation over the randomization of the mechanism. In addition, we refer to $a(v), t(v)$, and $v \cdot a(v) - t(v)$ as the buyer's ex post allocation, transfer, and utility, respectively.

The problem is to find a mechanism (a, t) that maximizes the (expected) revenue

$$\mathbf{E}_{v_1, \ldots, v_k}[t(v)],$$

subject to the PIC and ex post IR constraints.

A special class of mechanisms is the class of static mechanisms. A static mechanism is a mechanism where the outcome depends only on the report in the last period k. Formally, a mechanism (a, t) is static if $(a, t)(\theta^1, \ldots, \theta^k) = (a, t)(\hat{\theta}^1, \ldots, \hat{\theta}^k)$ whenever $\theta^k = \hat{\theta}^k$. We can, therefore, represent such a mechanism more succinctly by its allocation rule $a^{ST} : \Theta^k \to X$ and transfer rule $t^{ST} : \Theta^k \to \mathbb{R}$ in the last period. The interpretation is that in the last period k, having learned all her values, the buyer makes a report v to the mechanism. The buyer then receives each product i with probability $a_i^{ST}(v)$ and transfers $t^{ST}(v)$ to the mechanism. Because the reports in the periods before k are irrelevant, a static mechanism trivially satisfies all periodic incentive compatibility constraints before the last period k. Therefore, a static mechanism is PIC if it satisfies the last period incentive compatibility condition,

$$v \cdot a^{ST}(v) - t^{ST}(v) \geq v \cdot a^{ST}(\hat{v}) - t^{ST}(\hat{v}),$$

for all $v, \hat{v} \in \Theta^k$. Similarly, a static mechanism (a^{ST}, t^{ST}) is ex post IR if

$$v \cdot a^{ST}(v) - t^{ST}(v) \geq 0.$$

This formulation is used in the multiproduct mechanism design literature (e.g., in Manelli and Vincent 2007, Daskalakis et al. 2014). Thus, our model nests the optimal mechanism design problem for selling k products with static mechanisms as a special case.

Another special case is when the agent only reports v_i in period i. This is captured by requiring the allocation and the transfer to depend on the report θ^i in period i only through θ^i_i. That is, $(a,t)(\theta^1,\ldots,\theta^k) = (a,t)(\hat{\theta}^1,\ldots,\hat{\theta}^k)$ if $\theta^i_i = \hat{\theta}^i_i$ for all i. Even though this is a very natural class of mechanisms, it does not contain the class of all static mechanisms because the extensive form games they represent are different. By defining the class of mechanisms generally so that the agent reports her interim type in every period, we ensure that both static mechanisms and ones where the agent only reports her value are included as special cases.

The optimal revenue among all mechanisms is at least as high as the revenue from any static mechanism. This observation immediately follows from the fact that static mechanisms are a subclass of all mechanisms. A question we ask is whether the optimal revenue is strictly higher than that from static mechanisms. To this end, we first identify the optimal revenue and then, ask whether it can be achieved by a static mechanism.

3. Recursion, Separability, and Promised Utility

The periodic incentive compatibility constraints are complex. In each period, the buyer may misreport different dimensions of her interim type. Even for the special case of static mechanisms where all incentive constraints before the last period are trivially satisfied, the incentive constraints are complex. Nevertheless, we show that the optimization problem can be solved by making two observations. First, to maximize revenue, it is sufficient to focus on a simple class of *separable* mechanisms. Second, it is possible to optimize over separable mechanisms recursively.

A separable mechanism satisfies two properties. First, no rereporting is required. That is, in each period i, the buyer only reports her value v_i for product i (instead of her interim type). Second, the allocation of product i is based on the reports made up to (and including) period I but does not depend on reports made in periods $i+1$ to k. Formally, we have Definition 1.

Definition 1. A mechanism (a,t) is *separable* if for all $\theta^1,\ldots,\theta^k$ and $\hat{\theta}^1,\ldots,\hat{\theta}^k$,

1. $t(\theta^1,\ldots,\theta^k) = t(\hat{\theta}^1,\ldots,\hat{\theta}^k)$ if $\theta^i_i = \hat{\theta}^i_i$ for all i and
2. for all i, $a_i(\theta^1,\ldots,\theta^k) = a_i(\hat{\theta}^1,\ldots,\hat{\theta}^k)$ if $\theta^j_j = \hat{\theta}^j_j$ for all $j \le i$.

The first property states that the payment rule depends on the report θ^i in each period i only through the value learned in that period θ^i_i. The second property states that the allocation of product i depends on the report θ^j in period $j \le i$ only through the value learned in that period θ^j_j and does not depend on the report $\theta^{j'}$ in period $j' > i$. We will henceforth represent a separable mechanism more succinctly with functions $a^{SP}_i : \Theta^i \to [0,1]$ and $t^{SP} : \Theta^k \to \mathbb{R}$ (as opposed to $a_i : \Theta^1 \times \ldots \times \Theta^k \to [0,1]$ and $t : \Theta^1 \times \ldots \times \Theta^k \to \mathbb{R}$ for a general mechanism). The interpretation is that the buyer reports v_i in each period i. Given reports $(v_1,\ldots,v_k)$, product i is allocated with probability $a^{SP}_i(v_1,\ldots,v_i)$, and the transfer is $t^{SP}(v_1,\ldots,v_k)$.

We now show that to maximize revenue, it is without loss of generality to restrict attention to separable mechanisms. Notice that a separable mechanism is handicapped. It does not have the ability to bundle the products together because it cannot tie the allocation of a product to the allocation of future products (and the buyer's reports about those values). In contrast, a static mechanism does have the ability to bundle (we later return to this comparison and provide examples). As a result, it is a priori not clear that the optimal separable mechanism obtains at least as much revenue as all static mechanisms, let alone all mechanisms (that include separable and static mechanisms as special cases).

Example 1. There are two products, and the value for each product is either 1 or 2. Consider a static mechanism that only offers the bundle of both products for a price of 3. The allocation probabilities and transfers are shown in Table 1 for all types.

Notice that the allocation of product 1 depends on the value for product 2, $a_1(1,2) \ne a_1(1,1)$, and vice versa for product 2. Thus, no separable mechanism can implement this allocation.

To argue that restricting to separable mechanisms is without loss of generality for maximizing revenue, we convert any mechanism to a separable mechanism with the same revenue (but with a different allocation rule). In particular, given a mechanism (a,t), define its *induced separable mechanism* (a^{ISP}, t^{ISP}) as follows. The allocation probability a^{ISP}_i is the expectation of the allocation probability a_i assuming truthful reporting in all future periods. That is, for any $v_1,\ldots,v_i$,

$$a^{ISP}_i(v_1,\ldots,v_i) := \mathbf{E}_{v_{i+1},\ldots,v_k}[a_i(v)]. \tag{1}$$

(Recall that $a_i(v)$ is the shorthand for the allocation when the buyer reports $(v_1,\ldots,v_j)$ in each period j.)

Table 1. The Static Mechanism in Example 1

v_1	v_2	a_1	a_2	t
1	1	0	0	0
1	2	1	1	3
2	1	1	1	3
2	2	1	1	3

For $v = (v_1, \ldots, v_k)$, define the transfer as follows:

$$t^{ISP}(v) := t(v) - v \cdot a(v) + v \cdot a^{ISP}(v). \tag{2}$$

(Recall similarly that $t(v)$ is the shorthand for the transfer when the buyer reports $(v_1, \ldots, v_j)$ in each period j.)

Let us verify properties of this construction. First, if a mechanism is ex post IR, then so is its induced separable mechanism. This is because the transfer rule of the induced separable mechanism is defined such that the two mechanisms have the same ex post utility. That is, rearranging Equation (2), we have

$$v \cdot a^{ISP}(v) - t^{ISP}(v) = v \cdot a(v) - t(v) \tag{3}$$

for all v. Second, the two mechanisms have the same revenue. The reason is that the two mechanisms have the same ex post utility and also, create the same surplus. More precisely, take the expectation of Equation (2),

$$\begin{aligned}\mathbf{E}_v\big[t^{ISP}(v)\big] &= \mathbf{E}_v[t(v)] + \mathbf{E}_v\big[v \cdot a^{ISP}(v) - v \cdot a(v)\big]\\ &= \mathbf{E}_v[t(v)] + \sum_i \mathbf{E}_v\big[v_i a_i^{ISP}(v_1, \ldots, v_i) - v_i a_i(v)\big]\\ &= \mathbf{E}_v[t(v)] + \sum_i \mathbf{E}_{v_1,\ldots,v_i}\big[v_i(a_i^{ISP}(v_1, \ldots, v_i)\\ &\quad - \mathbf{E}_{v_{i+1},\ldots,v_k}[a_i(v)])\big]\\ &= \mathbf{E}_v[t(v)],\end{aligned} \tag{4}$$

where the last equality follows from Equation (1). It only remains to verify that these adjustments do not violate the PIC constraints.

To see that the construction preserves incentive compatibility, let us first verify incentive compatibility on path: that is, following a history of truthful reports. More precisely, the PIC constraint for a separable mechanism requires that for each period i, interim type $(v_1, \ldots, v_i)$, history $(\hat{v}_1, \ldots \hat{v}_{i-1})$, and report $\hat{v}_i$ in period i, we have

$$\begin{aligned}&\mathop{\mathbf{E}}_{v_{i+1},\ldots,v_k}\big[v \cdot a^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, v_i, v_{i+1}, \ldots, v_k)\\ &\qquad - t^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, v_i, v_{i+1}, \ldots, v_k)\big]\\ &\geq \mathop{\mathbf{E}}_{v_{i+1},\ldots,v_k}\big[v \cdot a^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, \hat{v}_i, v_{i+1}, \ldots, v_k)\\ &\qquad - t^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, \hat{v}_i, v_{i+1}, \ldots, v_k)\big].\end{aligned}$$

A mechanism satisfies PIC on path if the above inequality holds for all i, $(v_1, \ldots, v_i)$, $(\hat{v}_1, \ldots \hat{v}_{i-1}) = (v_1, \ldots, v_{i-1})$, and $\hat{v}_i$. Consider the utility of a buyer with value v_i from reporting $\hat{v}_i$. By Equation (3), the expected utility of the buyer in the induced separable mechanism is equivalent to the utility she would get in mechanism (a, t) if she reports $\hat{v}_i$ instead of v_i in *every period* i to k (recall that the buyer reports her full interim type in each period and that $a(v)$ and $t(v)$ stand for the outcome if the buyer reports $v_1, \ldots, v_i$ in each period i). However, by incentive compatibility of (a, t), the buyer is better off if she reports v_i instead of $\hat{v}_i$ in every period. Thus, the incentive constraint is satisfied on path.

The equivalence discussed no longer holds off path (i.e., following a history of nontruthful reports $(\hat{v}_1, \ldots \hat{v}_{i-1}) \neq (v_1, \ldots, v_{i-1})$). To establish incentive compatibility off path, we notice that a separable mechanism is PIC if it is PIC on path. Indeed, in a separable mechanism, the report in period i does not affect the allocation of products 1 to $i-1$. Thus, because future values are independent of the interim type, the incentive constraint at period i for an interim type $(v_1, \ldots, v_i)$ following a history of reports $\hat{v}_1, \ldots, \hat{v}_{i-1}$ is identical, up to a constant, to the incentive constraint for an interim type $(\hat{v}_1, \ldots, \hat{v}_{i-1}, v_i)$ following a history of truthful reports $(\hat{v}_1, \ldots, \hat{v}_{i-1})$. Thus, if a separable mechanism is incentive compatible for all histories of truthful reports, then it is incentive compatible for all histories. Formally, we have the following lemma.

Lemma 1. *A separable mechanism (a^{SP}, t^{SP}) is PIC if it is PIC on path.*

Proof. The PIC constraint in period i is that for an interim type $v_1, \ldots, v_{i-1}$ and following a history of reports $\hat{v}_1, \ldots, \hat{v}_{i-1}$, the expected utility of the buyer

$$\begin{aligned}&\mathbf{E}\big[v \cdot a^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, \hat{v}_i, v_{i+1}, \ldots, v_k)\\ &\quad - t^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, \hat{v}_i, v_{i+1}, \ldots, v_k)\big]\end{aligned}$$

is maximized over all reports $\hat{v}_i$ by setting $\hat{v}_i = v_i$. Separability implies that the utility of the buyer from the allocation of products 1 to $i-1$, $\sum_{j<i} v_j a_j^{SP}(\hat{v}_1, \ldots, \hat{v}_j)$, does not depend on the report in period i. Therefore, the report that maximizes the expected utility does not change if $\sum_{j<i} v_j a_j^{SP}(\hat{v}_1, \ldots, \hat{v}_j)$ is replaced by $\sum_{j<i} \hat{v}_j$ $a_j^{SP}(\hat{v}_1, \ldots, \hat{v}_j)$, which also does not depend on $\hat{v}_i$. As a result, the incentive constraint holds if

$$\begin{aligned}&\mathbf{E}\big[(\hat{v}_1, \ldots, \hat{v}_{i-1}, v_i, \ldots, v_k) \cdot a^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, \hat{v}_i, v_{i+1}, \ldots, v_k)\\ &\quad - t^{SP}(\hat{v}_1, \ldots, \hat{v}_{i-1}, \hat{v}_i, v_{i+1}, \ldots, v_k)\big]\end{aligned}$$

is maximized over all $\hat{v}_i$ by setting $\hat{v}_i = v_i$. This constraint is the PIC constraint of interim type $(\hat{v}_1, \ldots, \hat{v}_{i-1}, v_i)$ following a truthful history of reports $\hat{v}_1, \ldots, \hat{v}_{i-1}$. Notice that this proof uses the assumption that values are independent. Without it, the two expectations should be conditioned on the interim type $(v_1, \ldots, v_i)$, and so, the last expectation does not represent the PIC constraint of interim type $(\hat{v}_1, \ldots, \hat{v}_{i-1}, v_i)$, which needs to be conditioned on $(\hat{v}_1, \ldots, \hat{v}_{i-1}, v_i)$. □

The following proposition summarizes the arguments made so far.

Proposition 1. *The revenue of any mechanism is equal to the revenue of its induced separable mechanisms. If a mechanism is PIC and ex post IR, then so is its induced separable mechanism.*

The PIC constraint for a separable mechanism is simpler than the PIC constraint for a general mechanism. Consider the incentive compatibility constraint at a period i. By Lemma 1, we need to only consider the PIC constraints on path. In a separable mechanism, the allocation of products 1 to $i-1$ does not depend on the report at period i. Therefore, to choose her report in period i, the buyer only takes into account the allocations of products i to k and the transfer. For reports $(v_1, \ldots, v_i)$, define the *continuation utility* CU_i of the buyer to be the expected utility from the allocation of products $i+1$ to k and the transfer, assuming truthful reporting in future periods:

$$CU_i(v_1, \ldots, v_i) = \mathbf{E}_{v_{i+1},\ldots,v_k}\left[\left(\left(\sum_{j>i} v_j a_j(v^j)\right) - t(v)\right)\right].$$

Note also that the continuation utility does not depend on the buyer's interim type in period i and instead, is only a function of the reports that the buyer makes. The PIC constraint on path at every period i is that for all $v_1, \ldots, v_i$ and $\hat{v}_i$, the buyer maximizes the sum of her stage utility in period i plus her continuation utility from the future periods by reporting her value truthfully:

$$\begin{aligned} &v_i a_i(v_1, \ldots, v_{i-1}, v_i) + CU_i(v_1, \ldots, v_{i-1}, v_i) \\ &\geq v_i a_i(v_1, \ldots, v_{i-1}, \hat{v}_i) + CU_i(v_1, \ldots, v_{i-1}, \hat{v}_i). \end{aligned} \tag{5}$$

This simplification allows us to recursively optimize over separable mechanisms, as discussed next.

3.1. Recursive Optimization: Separability and Promised Utility

We now present a recursive characterization of optimal separable mechanisms. In particular, we define a class of promised utility mechanisms and show that they are optimal. These promised utility mechanisms are separable mechanisms that maintain a scalar state variable, the promised utility to the agent. This state variable affects the allocation in each period and is updated in each period based on the report of the agent. Promised utility mechanisms are defined given solutions to a certain *one-product* mechanism design problem. We start by defining these one-product mechanism design problems recursively.

Definition 2 (The Continuation Revenue Problem). Define the *seller's continuation revenue* functions $CR_{k+1}, \ldots, CR_1$ recursively as follows. Let $CR_{k+1}(\mathsf{EU}) = -\mathsf{EU}$ for $\mathsf{EU} \in \mathbb{R}^+$. For all $i \leq k$ and $\mathsf{EU} \in \mathbb{R}^+$,

$$\begin{aligned} &CR_i(\mathsf{EU}) \\ &:= \max_{a: V_i \to [0,1],\, t: V_i \to \mathbb{R}} \mathbf{E}_{v_i}[v_i a(v_i) + CR_{i+1}(v_i a(v_i) - t(v_i))], \end{aligned} \tag{6}$$

$$\text{s.t.} \quad v_i a(v_i) - t(v_i) \geq v_i a(\hat{v}_i) - t(\hat{v}_i); \qquad \forall v_i, \hat{v}_i \in V_i, \tag{7}$$

$$v_i a(v_i) - t(v_i) \geq 0; \qquad \forall v_i \in V_i, \tag{8}$$

$$\mathbf{E}_{v_i}[v_i a(v_i) - t(v_i)] = \mathsf{EU}. \tag{9}$$

Define $(\mathcal{A}_i^{\mathsf{EU}}, \mathcal{T}_i^{\mathsf{EU}})$ to be the set of optimal solutions to this problem.

The continuation revenue problem in each period i is the problem of optimizing over one-product mechanisms (a, t) that map the report in period i to an allocation for that product and a transfer. Constraints (7) and (8) are the standard incentive compatibility and individual rationality constraints for a one-product mechanism. However, this problem has two nonstandard features. First, there is an expected utility Constraint (9). It requires that the expected utility of the agent is equal to a given constant EU. Second, the objective is to maximize the expected surplus from allocation in period i plus the continuation revenue from period $i+1$ (instead of the standard objective of maximizing revenue). We next use the solutions to the continuation revenue problem to define promised utility mechanisms. Because there may be multiple optimal solutions to the continuation revenue problem, there may be multiple promised utility mechanisms.

Definition 3 (The Promised Utility Mechanism). A promised utility mechanism is parameterized by any profile of optimal solutions $(a_i^{\mathsf{EU}}, t_i^{\mathsf{EU}}) \in (\mathcal{A}_i^{\mathsf{EU}}, \mathcal{T}_i^{\mathsf{EU}})$ to the continuation revenue problem (Definition 2) for all i and $\mathsf{EU} \in \mathbb{R}^+$. Set the initial promised utility PU_1 equal to any maximizer of $CR_1(\mathsf{EU})$, and set the agent's transfer $\mathsf{T} = 0$. At each period i, given the current promised utility PU_i and report v_i,

- the buyer gets product i with probability $a_i^{\mathsf{PU}_i}(v_i)$,
- the transfer is updated by setting $\mathsf{T} := \mathsf{T} + v_i a_i^{\mathsf{PU}_i}(v_i)$, and
- the promised utility is updated by setting $\mathsf{PU}_{i+1} := v_i a_i^{\mathsf{PU}_i}(v_i) - t_i^{\mathsf{PU}_i}(v_i)$.

At the end of the last period, the buyer pays $\mathsf{T} - \mathsf{PU}_{k+1}$.

A promised utility mechanism maintains a scalar state variable PU_i, which is the agent's promised (expected) utility in period i. When the agent reports v_i in period i, she gets product i with probability $a_i^{\mathsf{PU}_i}(v_i)$, and her final transfer T increases by a certain amount. An important feature is that this extra transfer is the surplus of allocation $v_i a_i^{\mathsf{PU}_i}(v_i)$ and is not $t_i^{\mathsf{PU}_i}(v_i)$. This means that the agent's stage utility from being truthful is zero. To incentivize truthfulness, the agent's promised utility is adjusted to $v_i a_i^{\mathsf{PU}_i}(v_i) - t_i^{\mathsf{PU}_i}(v_i)$. This feature of the promised utility mechanism explains Objective (6) of the continuation revenue problem. The objective is the extra transfer from the current period, which is equal to the surplus of allocation in that period plus the continuation revenue given the promised utility in the future periods.

To see that promised utility mechanisms satisfy PIC, consider any report $\hat{v}_i$. The agent gets the product with probability $a_i^{\mathsf{PU}_i}(\hat{v}_i)$, her payment increases by $\hat{v}_i a_i^{\mathsf{PU}_i}(\hat{v}_i)$, and her promised utility becomes $\hat{v}_i a_i^{\mathsf{PU}_i}(\hat{v}_i) - t_i^{\mathsf{PU}_i}(\hat{v}_i)$. So, she maximizes

$$v_i a_i^{\mathsf{PU}_i}(\hat{v}_i) - \hat{v}_i a_i^{\mathsf{PU}_i}(\hat{v}_i) + \hat{v}_i a_i^{\mathsf{PU}_i}(\hat{v}_i) - t_i a_i^{\mathsf{PU}_i}(\hat{v}_i) = v_i a_i^{\mathsf{PU}_i}(\hat{v}_i) - t_i a_i^{\mathsf{PU}_i}(\hat{v}_i),$$

which is achieved by reporting truthfully, $\hat{v}_i = v_i$, by the incentive constraint of the continuation revenue Problem (7). After the last period, the mechanism gives the agent a discount of PU_{k+1} on her final transfer. That is, the mechanism fulfills the final promised utility to the agent by paying her cash.

Giving the agent zero stage utility in all periods (except last) is useful because it means that a nonnegative promised utility is sufficient to ensure that the agent's ex post IR constraint is satisfied. This is because the agent's ex post utility is the cash she is offered PU_{k+1} at the end of the last period. So, ex post IR is satisfied if and only if the promised utility at the end of the last period is nonnegative.

The proposition shows that promised utility mechanisms are optimal.

Proposition 2. *A mechanism is optimal if and only if its induced separable mechanism is a promised utility mechanism.*

The main observation in the proof of Proposition 2 is that in an optimal separable mechanism, following any history, the "continuation mechanism" must be optimal over all possible continuation mechanisms with the same continuation utility. In particular, fix a history $(v_1, \ldots, v_i)$. Consider the continuation mechanism (that is, a mechanism that maps reports in periods after i, $(v_{i+1}, \ldots, v_k)$, to allocations for those products and a transfer). Notice that if we replace this continuation mechanism with another one that has the same continuation utility, then the PIC Constraint (5) in period i will remain satisfied. Thus, in an optimal mechanism, the continuation mechanism following the history must be optimal over all continuation mechanisms with the same continuation utility. Otherwise, the continuation mechanism can be replaced with one with higher revenue. We can thus maintain the "promised utility" $\mathsf{PU}_i = CU_i(v_1, \ldots, v_i)$ as a scalar state variable that summarizes the history. We can use this observation to recursively characterize the optimal continuation revenue for a given promised utility.

4. Optimality of Static Mechanisms

The fact that separable mechanisms are optimal does not necessarily mean that static mechanisms are suboptimal because there may be multiple optimal mechanisms. We now study whether static mechanisms can be optimal. To answer this question, we use Proposition 2 to identify optimal revenue and then, verify whether a static mechanism exists that achieves that optimal revenue. The interpretation of these results is that, under the specified conditions for suboptimality of static mechanisms, the ability to screen types dynamically is strictly more profitable for the seller than the ability to bundle products, even if we take away the seller's ability to charge advance payments (because of the ex post IR constraint). We start by providing necessary and sufficient conditions for optimality of static mechanisms with two products and two values. We then provide sufficient conditions for suboptimality of static mechanisms with any number of products and values.

4.1. Tight Conditions for Two Products and Two Values

Suppose that there are two products and two values, $V_1 = V_2 = \{\underline{v}, \overline{v}\}$, where $\underline{v} < \overline{v}$. The proposition specifies two conditions that are together necessary and sufficient for suboptimality of static mechanisms. For this result, let $q_1 = f_1(\overline{v})$ and $q_2 = f_2(\overline{v})$ denote the probability of high value in each period.

Proposition 3. *Assume that $k = 2$ and $V_1 = V_2 = \{\underline{v}, \overline{v}\}$, where $\underline{v} < \overline{v}$. Any static mechanism is suboptimal if and only if $q_1 < \underline{v}/\overline{v}$ and $q_2 > \underline{v}(1 - q_1)/(\overline{v} - q_1\underline{v})$.*

The set of parameters q_1 and q_2 for which static mechanisms are suboptimal is drawn in Figure 2. Notice that the condition $q_1 < \underline{v}/\overline{v}$ or equivalently, $\overline{v}q_1 < \underline{v}$ states that the unique optimal monopoly price for the first product is $\underline{v}$. That is, for selling only the first product, the seller obtains a strictly higher revenue by choosing a low price compared with a high price. The condition $q_2 > \underline{v}(1 - q_1)/(\overline{v} - q_1\underline{v})$ is more complex. Nevertheless, because $\underline{v}(1 - q_1)/(\overline{v} - q_1\underline{v})$ is decreasing in q_1, it is sufficient that $q_2 > \underline{v}/\overline{v}$ or equivalently, $\overline{v}q_2 > \underline{v}$. That is, the unique optimal monopoly price for the second product is $\overline{v}$. To summarize, static mechanisms are suboptimal if (but not only if) the optimal monopoly price for product 1 is strictly less than the optimal monopoly price for product 2. In the

Figure 2. Static vs. Optimal Mechanisms

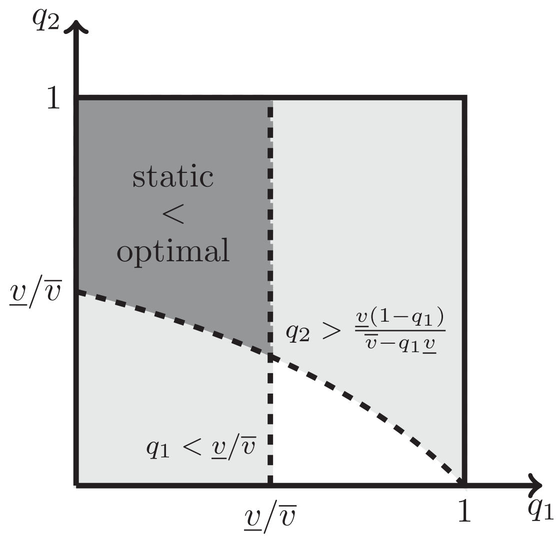

Note. The dark-shaded region is the set of (q_1, q_2) for which static mechanisms are suboptimal.

next section, we show that this statement generalizes to any number of products and values.

To prove Proposition 3, we provide a characterization of the optimal mechanisms stated. There are five cases. In four cases, a static mechanism is optimal. The four static mechanisms are simple. Three of them sell the products separately. That is, each product has a price, and the buyer can buy each product by paying its price. The fourth static mechanism is a bundling mechanism that only offers the two products as a bundle. The fifth mechanism is separable. This mechanism sells the second product via a take it or leave it price that depends on the reported value in the first period. The conditions of Proposition 3 for suboptimality of static mechanisms are precisely those under which this separable mechanism outperforms all four static mechanisms.

To state the proposition, recall that $q_1 = f_1(\overline{v})$ and $q_2 = f_2(\overline{v})$. Define the price $p^* = \underline{v} - (1-q_2)(\overline{v} - \underline{v})$. The price p^* is constructed such that the expected utility of the buyer from being offered a take it or leave it price p^* for the second product is equal to $\overline{v} - \underline{v}$. This price is low enough (below $\underline{v}$) to be accepted by both possible values.

Proposition 4. *Assume that $k = 2$ and $V_1 = V_2 = \{\underline{v}, \overline{v}\}$, where $\underline{v} < \overline{v}$. At least one of the following five mechanisms is optimal.*

1. *Sell each product separately at price $\underline{v}$.*
2. *Sell each product separately at price $\overline{v}$.*
3. *Sell each product separately at price $\overline{v}$ for the first product and $\underline{v}$ for the second product.*
4. *Sell only the grand bundle at price $\underline{v} + \overline{v}$.*
5. *In the first period, the buyer reports v_1 and receives the first product with probability one. In the second period, the buyer pays v_1, and in addition, she is offered the second product at price $\overline{v}$ if $v_1 = \underline{v}$ and at price p^* if $v_1 = \overline{v}$.*

If $q_1 < \underline{v}/\overline{v}$ and $q_2 > \underline{v}(1-q_1)/(\overline{v} - q_1\underline{v})$, then the fifth mechanism is the unique optimal separable mechanism. Otherwise, that is if $q_1 \geq \underline{v}/\overline{v}$ or $q_2 \leq \underline{v}(1-q_1)/(\overline{v} - q_1\underline{v})$, then at least one of the first four mechanisms is optimal.

Given Proposition 4, to identify an optimal mechanism, one needs only to compare the revenue of the five mechanisms. The revenue of each mechanism can be written in closed form. For the first four mechanisms, revenue is simply the prices times the probability of purchase. To calculate the revenue of the fifth mechanism, notice that the buyer pays her expected value of the first product, and she is in addition offered the second product at price $\overline{v}$ with probability $1 - q_1$ (if $v_1 = \underline{v}$) and at price p^* with probability q_1 (if $v_1 = \overline{v}$). Thus, revenue is

$$\mathbf{E}[v_1] + (1-q_1)q_2\overline{v} + q_1 p^*.$$

The comparisons between the revenues of these mechanisms are provided in the proof of Proposition 4 in the e-companion. The conditions of Proposition 3 are precisely those under which this dynamic mechanism outperforms all four static mechanisms identified in Proposition 4.

Proposition 4 is independently useful because it can be used to identify optimal static screening mechanisms. In particular, if $q_1 \geq \underline{v}/\overline{v}$ or $q_2 \leq \underline{v}(1-q_1)/(\overline{v} - q_1\underline{v})$, then one of the four static mechanisms identified in Proposition 4 is optimal among *all static mechanisms*. This is simply because under these conditions, one of these mechanisms is optimal in the larger class of all mechanisms (static or not). This suggests that our approach may be more generally useful for solving, either exactly or approximately, the notoriously difficult problem of selling multiple products using static mechanisms. In general even verifying the optimality of a given static mechanism is not straightforward because it requires the construction of appropriate dual certificates (Daskalakis et al. 2017, Carroll 2017, Cai et al. 2019). Additionally, even though the case of two products and two values can be solved in a static setting via case analysis, such analyses are typically tedious. For instance, Armstrong and Rochet (1999) solve a screening problem with four types. They consider all possible ways to relax subsets of the incentive constraints, and they identify conditions under which the solution to each relaxation satisfies all the constraints and therefore, is optimal. In comparison, our recursive formulation allows the incentive constraints in the two periods to be separated and solved using standard tools.

The proof of Proposition 4 relies on a characterization of the continuation revenue problem (in Definition 2) in the last period k. We provide this characterization generally with any number of values in the e-companion and use

Figure 3. The Continuation Revenue Function

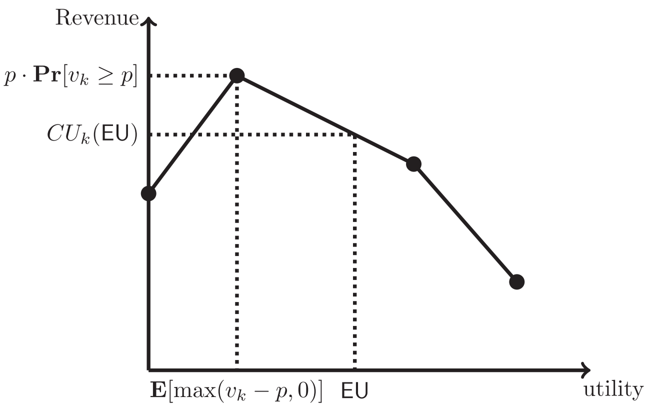

Note. The continuation revenue function in the last period CU_k is the concavification of a function that maps the expected utility $\mathbf{E}[\max(v_k - p, 0)]$ from Posting Any Price $p \in V_k$ to the Revenue of That Price $p \cdot \mathbf{Pr}[v_k \geq p]$

it also in the next section. The characterization shows that it is optimal to choose one of at most two prices at random and sell the product at that price as a take it or leave it offer. These prices are obtained from "concavifying" an appropriately constructed revenue function shown in Figure 3. The revenue function plots the expected utility to the buyer from posting any price $p \in V_k$, $\mathbf{E}[\max(v_k - p, 0)]$ against the revenue that the seller obtains from that price $p \cdot \mathbf{Pr}[v_k \geq p]$.

4.2. Sufficient Conditions for Any Number of Products and Values

In this section, we identify sufficient conditions for the suboptimality of static mechanisms with any number of products and values. We show that static mechanisms are suboptimal if the first product has lower monopoly prices than the second product, partially generalizing Proposition 3 to any number of values. To do so, we use the recursive characterization of optimal mechanisms in Proposition 2. To simplify exposition, we assume that $V_1 = \ldots = V_k$.

We start with defining the main condition of the result. For each i, let P_i be the set of optimal monopoly prices for selling product i. That is, $P_i = \arg\max_p p \cdot \mathbf{Pr}[v_k \geq p]$. Let $\overline{p}_i$ and $\underline{p}_i$ be the largest and smallest such prices. We say that product 1 has *lower monopoly prices* than product k if $\overline{p}_1 < \underline{p}_k$. If the optimal monopoly prices are unique, the condition simply means that the monopoly price for product 1 is lower than the monopoly price for product k. Notice that if there are two products and $V_1 = V_2 = \{\underline{v}, \overline{v}\}$, then product 1 has lower monopoly prices than product j if and only if $P_1 = \{\underline{v}\}$ and $P_2 = \{\overline{v}\}$. Thus, the condition is weaker than the conditions of Proposition 3 but allows for a generalization to any number of products and values.

Proposition 5. *Assume that $V_1 = \ldots = V_k$. Any static mechanism is suboptimal if product 1 has lower monopoly prices than product k.*

To outline the proof, suppose for simplicity that there are only two products and the optimal monopoly prices p_1 and p_2 are unique. Assume that $p_1 < p_2$. Assume for contradiction that a static mechanism (a, t) is optimal. By Proposition 1, its induced separable mechanism (a^{ISP}, t^{ISP}) must also be optimal. We show that this observation implies that $a(p_1, v_2) = (1,1)$ for all v_2. The intuition is that the optimal allocation in the continuation revenue problem is more efficient than in the standard problem of maximizing revenue for selling only product 1. This is because in the continuation revenue problem, the seller obtains some profit from giving information rents to the agent. So, any type with value $v_1 \geq p_1$ must receive product 1 with probability one, and we show that this implies that such a type must also receive product 2 with probability one. The fact that $a(p_1, v_2) = (1,1)$ for all v_2 means that the grand bundle is offered at a relatively low price. This implies that in the optimal separable mechanism, the promised utility to even the lowest interim type in the first period is relatively high. Then, we show that we can lower the promised utility to all types by the same amount and increase revenue, contradicting the optimality of the static mechanism.

5. Correlated Values

Extending our formal analysis to allow for correlated values requires a significantly different set of tools from those developed in this paper. Instead, we here use numerical calculations to study the robustness of our two main insights, namely the optimality of separable mechanisms and the conditions for suboptimality of static mechanisms. Our main finding is that these two insights extend with positively correlated values but fail with negatively correlated values. Throughout this section, we focus on two products and binary values where $V_1 = V_2 = \{1, 2\}$.

5.1. Our Parametrization

We start by discussing how we parameterize distributions. In order to facilitate comparison with the results in Section 4, we use $q_1 = \mathbf{Pr}[v_1 = 2]$ and $q_2 = \mathbf{Pr}[v_2 = 2]$ to denote the probabilities of high values. A third parameter $\lambda \in [-1, 1]$ pins down $\mathbf{Pr}[v_1 = 2, v_2 = 2]$ and thus, the whole distribution. This parameter λ measures the degree of correlation. For $\lambda \geq 0$, $\mathbf{Pr}[v_1 = 2, v_2 = 2]$ is a convex combination of this probability if values were independent, $q_1 q_2$, and the highest possible value it can take, $\min(q_1, q_2)$ (otherwise, either $\mathbf{Pr}[v_1 = 2, v_2 = 1]$ or $\mathbf{Pr}[v_1 = 1, v_2 = 2]$ becomes negative):

$$\mathbf{Pr}[v_1 = 2, v_2 = 2] = (1 - \lambda) q_1 q_2 + \lambda \min(q_1, q_2).$$

Figure 4. Positively Correlated Values

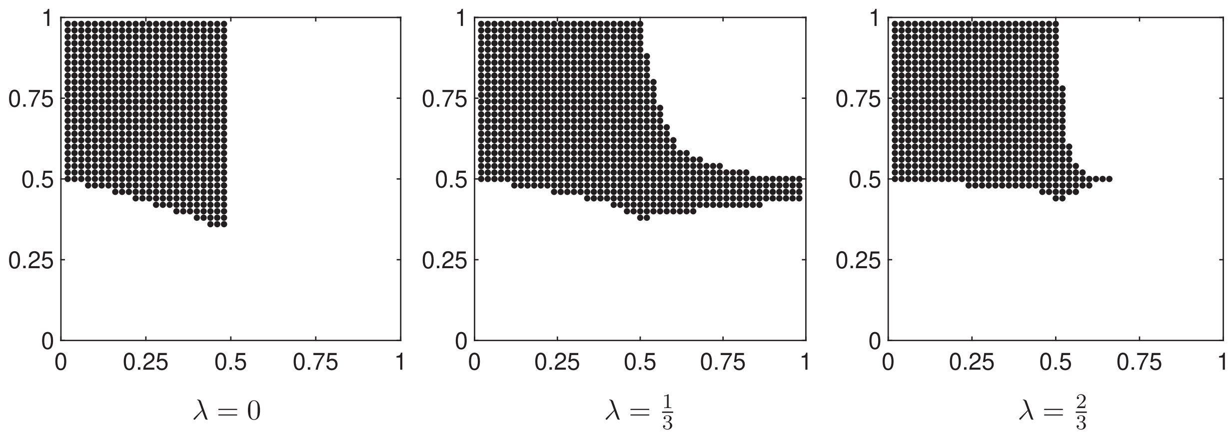

Notes. Values of q_1 are on the horizontal axis, and values of q_2 are on the vertical axis. Black circles indicate that separable mechanisms outperform static ones.

Thus, $\lambda = 0$ represents the independent distribution, and $\lambda = 1$ represents the highest positive correlation (perfect positive correlation if $\lambda = 1$ and $q_1 = q_2$). For $\lambda \le 0$, $\mathbf{Pr}[v_1 = 2, v_2 = 2]$ is a convex combination of this probability if values were independent, $q_1 q_2$, and the lowest possible value it can take, $\max(0, q_1 + q_2 - 1)$ (otherwise, either $\mathbf{Pr}[v_1 = 2, v_2 = 2]$ or $\mathbf{Pr}[v_1 = 1, v_2 = 1]$ becomes negative):

$$\mathbf{Pr}[v_1 = 2, v_2 = 2] = (1 + \lambda) q_1 q_2 - \lambda \max(0, q_1 + q_2 - 1).$$

Thus, $\lambda = 0$ represents the independent distribution, and $\lambda = -1$ represents the highest negative correlation (perfect negative correlation if $\lambda = -1$ and $q_1 + q_2 = 1$).

To interpret λ, notice that $\lambda \ge 0$ means that values are positively correlated, $\mathbf{E}[v_1 v_2] \ge \mathbf{E}[v_1]\mathbf{E}[v_2]$ (equivalently, the Pearson correlation coefficient is nonnegative), and $\lambda \le 0$ means that values are negatively correlated, $\mathbf{E}[v_1 v_2] \le \mathbf{E}[v_1]\mathbf{E}[v_2]$. Further, the probability $\mathbf{Pr}[v_1 = v_2]$ that the two values are equal is increasing in λ. We thus interpret λ as a measure of correlation and use $(q_1, q_2, \lambda) \in [0,1] \times [0,1] \times [-1,1]$ to parameterize distributions. We use λ to measure correlation because it is orthogonal to q_1, q_2. That is, for any given $\lambda \in [-1,1]$, any $q_1, q_2 \in [0,1] \times [0,1]$ specifies a distribution. For other measures of correlation we are aware of, including the Pearson correlation coefficient, the possible values of q_1, q_2 depend on the value of the correlation measure. We now separately discuss positive and negative correlations.

5.2. Positive Correlation $\lambda \ge 0$

Our first observation is that, perhaps surprisingly, separable mechanisms seem to remain optimal with positive correlation. The intuition is that the PIC constraints for separable mechanisms become easier to satisfy with positive correlation. Indeed, an interim type $v_1 = 2$ assigns a higher conditional probability to $v_2 = 2$ than does $v_1 = 1$, and so, the interim type $v_1 = 2$ assigns a higher benefit to being offered a discounted price in the second period. Our second observation is that static mechanisms seem to remain suboptimal if product 1 has lower monopoly prices than product 2, $q_1 \le 0.5$ and $q_2 \ge 0.5$. This is shown in Figure 4 for three nonnegative values of λ. A black circle corresponds to a distribution where separable mechanisms outperform static ones. The intuition is that with positive correlation, the value to the seller of being able to bundle the products decreases, hence reducing the value of the bundling instrument relative to the dynamic screening instrument. As we see next, this second observation no longer holds with negative correlation.

5.3. Negative Correlation $\lambda \le 0$

With negative correlation, both of our main results fail, surprisingly quickly in the degree of correlation. In particular, first, separable mechanisms may be suboptimal. Second, static mechanisms may outperform separable mechanisms even if $q_1 \le 0.5$ and $q_2 \ge 0.5$. These findings are shown in Figure 5 for three nonpositive values of λ. A black circle corresponds to a distribution where separable mechanisms outperform static ones, and a white circle corresponds to a distribution where static mechanisms outperform separable ones. These findings suggest that the value of the bundling instrument increases relative to the dynamic screening instrument with negatively correlated values. The following example explores this intuition further.

Example 2. There are two products, and the value for a product is either 1 or 2. The probabilities of profiles (1, 2) and (2, 1) are 0.5 each, and the probabilities of profiles (1, 1) and (2, 2) are 0 each. Similar examples

Figure 5. Negatively Correlated Values

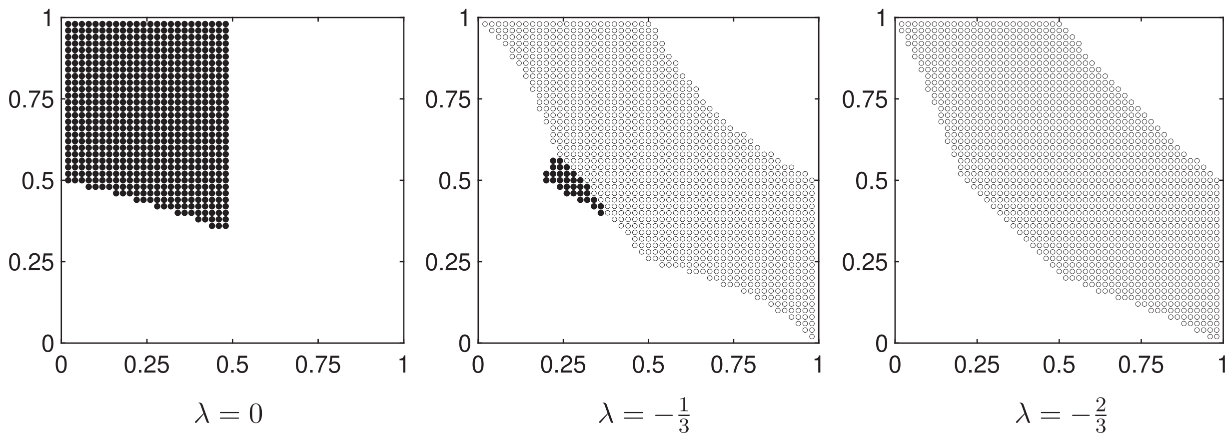

Notes. Values of q_1 are on the horizontal axis, and values of q_2 are on the vertical axis. Black circles indicate that separable mechanisms outperform static ones. White circles indicate that static mechanisms outperform separable ones.

can be constructed wherein the probabilities of profiles (1, 1) and (2, 2) are nonzero but small so that all four type profiles are in the support of the distribution.

The optimal mechanism is static. It extracts the full surplus by offering the bundle for a price of 3, thus obtaining a revenue of 3. However, we show that no separable mechanism has revenue of 3 as we show.

First consider a naive generalization of the construction in Section 3, where the expected allocation probabilities and the ex post utilities of the separable mechanism are equal to those of the static mechanism. Conditioned on $v_1 = 1$, v_2 is equal to 2 with probability one. Thus, $a_1(1) = 1$. Similarly, we have $a_1(2) = 1$. By definition, the allocation probabilities of product 2 are equal to those of the static mechanism. The separable mechanism is shown in Table 2.

Notice that the revenue of the separable mechanism is indeed 3. However, the separable mechanism is not PIC. Indeed, the expected utility of an interim type $v_1 = 2$ from truthfulness is zero because conditioned on $v_1 = 2, v_2 = 1$ with probability one. On the other hand, the expected utility from reporting $\hat{v}_1 = 1$ is 1 because by doing so, the buyer receives product 1 and pays 1.

We now argue that indeed no separable mechanism can obtain a revenue of 3. By ex post IR, for the revenue to be 3, both types (1, 2) and (2, 1) must receive both products and pay 3. Thus, in a separable mechanism, $a_1(v_1) = 1$ for all v_1. Further, the incentive compatibility constraint in the second period requires that the probability of allocation of product 2 for type (2, 2) must be no lower than that for type (2, 1). Thus, $a_2(2,2) \geq a_2(2,1) = 1$, and so, $a_2(2,2) = 1$. Because in the second period, the allocations of the types (2, 2) and (2, 1) are the same, their payments must be the same by incentive compatibility, and so, $t(2,2) = t(2,1) = 3$. We summarize our discussion in Table 3, in which the only free parameters are $a_2(1,1)$ and $t(1,1)$.

The ex post IR constraint for type (1, 1) is

$$1 + a_2(1,1) - t(1,1) \geq 0.$$

Now consider the PIC constraint in period 2 for ex post type (1, 2) following a history of truthful report $v_1 = 1$. The constraint is

$$0 \geq 1 + 2a_2(1,1) - t(1,1).$$

Given the two constraints, we must have $a_2(1,1) = 0$ and $t(1,1) = 1$. Thus, the mechanism is equal to the induced separable mechanism of the static mechanism that sells the bundle at price 3. As we argued, the separable mechanism is not incentive compatible.

6. Concluding Remarks

We study the problem of designing optimal ex post IR mechanisms for selling multiple products to a single buyer who learns her values sequentially. The ex post IR constraint takes away the seller's ability to charge advance payments and thus, allows us to compare static

Table 2. The Static Mechanism in Example 1

v_1	v_2	a_1	a_2	t
1	1	1	0	1
1	2	1	1	3
2	1	1	1	3
2	2	1	1	3

Table 3. The Static Mechanism in Example 1

v_1	v_2	a_1	a_2	t
1	1	1	$a_2(1,1)$	$t(1, 1)$
1	2	1	1	3
2	1	1	1	3
2	2	1	1	3

and dynamic mechanisms on an equal footing. We find that separable mechanisms are optimal and characterize optimal mechanisms via a recursive formulation. We find conditions under which static mechanisms are suboptimal. Interestingly, with two products and two values, static mechanisms are optimal for a relatively large set of distributions, even though the seller may use dynamic mechanisms. Obtaining sufficient conditions for optimality of static mechanisms beyond the case of two products and two values may help rationalize their widespread use.

Our analysis takes the arrival of information as given. In particular, the buyer learns her values in a fixed order. In many settings, sellers may be able to affect how information arrives to the buyer. For example, the seller may be able to choose the order in which the buyer learns the values of products. Even though this is not the focus of our paper, our results partially speak to this problem. In particular, consider two products A and B with possible values $V_A = V_B = \{1, 2\}$ such that $\mathbf{Pr}[v_A = 2] < 0.5$ and $\mathbf{Pr}[v_B = 2] > 0.5$. If the seller could choose the order with which the buyer learns the values, what should he do? Proposition 3 implies that if the buyer learns the value of A first, then static mechanisms are suboptimal, but if the buyer learns the value of B first, then static mechanisms are optimal. Because the set of static mechanisms is the same in either case, we conclude that the seller strictly prefers to reveal the value of product A, the one that is ex ante less valuable, to the buyer first.

Our analysis mostly assumes that the values are independent. This assumption is made for tractability and is in line with much of the literature on multiproduct mechanisms. Extending our analysis to allow for correlated values requires significantly different tools from those developed in this paper, and it is left for future work. Our numerical analysis suggests that our main results may hold with positively correlated values but fail with negatively correlated values.

Acknowledgments

A one-page abstract based on a preliminary version of this paper appeared in the *Association for Computing Machinery Conference on Economics and Computation 2016*.

Endnote

[1] An optimal monopoly price for a product is an optimal take it or leave it price for selling that product.

References

Adams WJ, Yellen JL (1976) Commodity bundling and the burden of monopoly. *Quart. J. Econom.* 90(3):475–498.

Armstrong M, Rochet J-C (1999) Multi-dimensional screening: A user's guide. *Eur. Econom. Rev.* 43(4-6):959–979.

Balseiro SR, Mirrokni VS, Leme RP (2018) Dynamic mechanisms with martingale utilities. *Management Sci.* 64(11):5062–5082.

Bergemann D, Välimäki J (2010) The dynamic pivot mechanism. *Econometrica* 78(2):771–789.

Bergemann D, Castro F, Weintraub GY (2020) The scope of sequential screening with ex post participation constraints. *J. Econom. Theory* 188:105055.

Boleslavsky R, Said M (2012) Progressive screening: Long-term contracting with a privately known stochastic process. *Rev. Econom. Stud.* 80(1):1–34.

Cai Y, Devanur NR, Weinberg SM (2019) A duality-based unified approach to Bayesian mechanism design. *SIAM J. Comput.* 50(3):160–200.

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

Clementi GL, Hopenhayn HA (2006) A theory of financing constraints and firm dynamics. *Quart. J. Econom.* 121(1):229–265.

Courty P, Li H (2000) Sequential screening. *Rev. Econom. Stud.* 67(4): 697–717.

Daskalakis C, Deckelbaum A, Tzamos C (2014) The complexity of optimal mechanism design. Preprint, submitted March 30, https://arxiv.org/abs/1211.1703.

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

DeMarzo PM, Sannikov Y (2006) Optimal security design and dynamic capital structure in a continuous-time agency model. *J. Finance* 61(6):2681–2724.

Eső P, Szentes B (2007) Optimal information disclosure in auctions and the handicap auction. *Rev. Econom. Stud.* 74(3):705–731.

Garrett DF (2016) Intertemporal price discrimination: Dynamic arrivals and changing values. *Amer. Econom. Rev.* 106(11): 3275–3299.

Gershkov A, Moldovanu B (2009) Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *Amer. Econom. J. Microeconomics* 1(2):168–198.

Gershkov A, Moldovanu B (2010) Efficient sequential assignment with incomplete information. *Games Econom. Behav.* 68(1): 144–154.

Green EJ (1987) Lending and the smoothing of uninsurable income. Prescott EC, Wallace N, eds. *Contractual Arrangements for Intertemporal Trade* (University of Minnesota Press, Minneapolis), 3–25.

Grillo E, Ortner J (2018) Dynamic contracting with limited liability constraints. Unpublished working paper, Collegio Carlo Alberto, Turin, Italy.

Kakade SM, Lobel I, Nazerzadeh H (2013) Optimal dynamic mechanism design and the virtual-pivot mechanism. *Oper. Res.* 61(4): 837–854.

Krähmer D, Strausz R (2015) Optimal sales contracts with withdrawal rights. *Rev. Econom. Stud.* 82(2):762–790.

Krähmer D, Strausz R (2016) Optimality of sequential screening with multiple units and ex post participation constraints. *Econom. Lett.* 142(May 2016):64–68.

Krähmer D, Strausz R (2017) Sequential vs. static screening: An equivalence result. *Games Econom. Behav.* 106(November 2017): 317–328.

Krishna RV, Lopomo G, Taylor CR (2013) Stairway to heaven or highway to hell: Liquidity, sweat equity, and the uncertain path to ownership. *RAND J. Econom.* 44(1):104–127.

Manelli AM, Vincent DR (2007) Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *J. Econom. Theory* 137(1):153–185.

McAfee RP, McMillan J, Whinston MD (1989) Multiproduct monopoly, commodity bundling, and correlation of values. *Quart. J. Econom.* 104(2):371–383.

Mirrokni VS, Paes Leme R, Tang P, Zuo S (2016) Optimal dynamic mechanisms with ex-post IR via bank accounts. Preprint, submitted July 6, https://arxiv.org/abs/1605.08840.

Mirrokni V, Paes Leme R, Tang P, Zuo S (2020) Non-clairvoyant dynamic mechanism design. *Econometrica* 88(5):1939–1963.

Pai M, Vohra RV (2008) Optimal dynamic auctions. Technical report, Center for Mathematical Studies in Economics and Management Science, Kellogg School of Management, Northwestern University, Evanston, IL.

Papadimitriou C, Pierrakos G, Psomas C, Rubinstein A (2022) On the complexity of dynamic mechanism design. *Games Econom. Behav.* 134:399–427.

Pavan A, Segal I, Toikka J (2014) Dynamic mechanism design: A Myersonian approach. *Econometrica* 82(2):601–653.

Rochet J-C, Chone P (1998) Ironing, sweeping, and multidimensional screening. *Econometrica* 66(4):783–826.

Sappington D (1983) Limited liability contracts between principal and agent. *J. Econom. Theory* 29(1):1–21.

Spear SE, Srivastava S (1987) On repeated moral hazard with discounting. *Rev. Econom. Stud.* 54(4):599–617.

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

Thomas J, Worrall T (1990) Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *J. Econom. Theory* 51(2):367–390.

Itai Ashlagi is an associate professor of management science and engineering, Stanford University. His research interest is in market design, matching, mechanism design, and game theory.

Constantinos Daskalakis is a professor at MIT's Electrical Engineering and Computer Science Department, a member of CSAIL, and affiliated with LIDS and ORC. His research interest is in theory of computation and its interface with economics, game theory, machine learning, statistics, and probability theory.

Nima Haghpanah is an assistant professor of economics, Pennsylvania State University. He studies the role of information in markets. His current interests are in mechanism design, information design, and price discrimination.