

The Limits of Multi-Product Price Discrimination

Nima Haghpanah

Ron Siegel*

December 1, 2021

We consider a multi-product seller who has access to information about consumer preferences that he can use for second- and third-degree price discrimination. We characterize markets for which such information can lead to the efficient allocation with consumers obtaining the entire surplus gain relative to the profit-maximizing allocation without the additional information. This benchmark is achievable for all markets with a given set of consumer types if and only if it is optimal for the seller to offer only the best product in each market. Analogous results characterize when the “surplus triangle” of Bergemann, Brooks, and Morris (2015) is achievable.

Understanding the welfare effects of price discrimination in a monopoly setting is a basic economics question. These effects are easy to see with first-degree price discrimination, which leads to efficiency and eliminates all consumer surplus. But with coarser market segmentations, based on age, location, or other data, followed by profit-maximizing pricing in each segment, the effects are less clear. Such segmentations may not achieve efficiency and consumer surplus may be positive, but precisely how much total surplus is generated and how it is divided between consumers and the seller is not immediately obvious.

A breakthrough in understanding these effects was achieved by Bergemann, Brooks, and Morris (2015). They showed that with a *single* product, *any* amount of total surplus and *any* division of the total surplus between consumers and the seller can be achieved by some segmentation, provided that total surplus is not higher than the efficient surplus, consumer

*Department of Economics, Pennsylvania State University, haghpanah@psu.edu, ron.siegel@gmail.com. We thank five referees, Dirk Bergemann, Wouter Dessein, Harry Di Pei, Piotr Dworzczak, Wioletta Dziuda, Teddy Kim, Stephen Morris, Alessandro Pavan, Eduardo Perez Richet, Andy Skrzypacz, Asher Wolinsky, Jidong Zhou, and participants in various seminars for very helpful comments and suggestions.

surplus is non-negative, and the seller’s surplus is no lower than in the unsegmented market. These constraints form a “surplus triangle” of achievable consumer-producer surplus pairs. In particular, efficiency can be achieved with consumers obtaining all of the surplus gains relative to the unsegmented market. We refer to this outcome as “first-best consumer surplus.”

We consider a multi-product setting and investigate the achievability of first-best consumer surplus and the surplus triangle when the seller may offer more than one product in each segment.¹ Our environment thus combines second- and third-degree price discrimination.² Each consumer is characterized by his valuations for the various products (his type), and a market is a distribution over consumer types. To facilitate the comparison with the single-product case we assume that there is a “best product,” it is efficient to allocate this product to all consumers, and consumer types are ranked so that, for each product, the valuation of any consumer type is higher than the valuations of lower types. We also assume zero production costs. A leading example is digital goods, such as streaming services, where the best product corresponds to the “premium” or “full feature” version of the service.

Our first result is that first-best consumer surplus and the surplus triangle are not achievable for any market in which the seller finds it optimal to sell more than one product. Of course, such “screening” implies inefficiency because it is efficient for all consumers to obtain the best product; the result shows that it is impossible to achieve efficiency via segmentation without the seller appropriating some of the surplus gains. In contrast, if the inefficiency is only caused by the seller *excluding* some consumers, then efficiency *can* be achieved via segmentation without the seller appropriating any of the surplus gains, as in the single-product case. Thus, the source of inefficiency matters for the achievable divisions of the efficient surplus.

¹Bergemann, Brooks, and Morris (2015) provide a parametric example with two types and non-linear valuations in which the seller offers more than one product in a single market. Ichihashi (2020) and Hidir and Vellodi (2021) consider maximum consumer surplus when a multi-product seller offers only one product in each market (but possibly different products in different markets).

²In Haghanah and Siegel (2021) we also study market segmentation with multiple products, but focus on segmentations that benefit every consumer and the seller. Such Pareto improving segmentation do not achieve first-best consumer surplus because the seller’s profit increases and the outcome need not be efficient. Identifying Pareto improving segmentations raises different challenges and requires different techniques from the ones in this paper. Daskalakis, Papadimitriou, and Tzamos (2016) and Cai et al. (2020) also study information provision in a multi-product setting. But in their settings the consumer (not the seller) receives information about products so there is no third-degree price discrimination.

Our second result, which follows from the first, is that, given the set of consumer types, first-best consumer surplus and the surplus triangle are achievable for *all* markets if and only if screening is not optimal for any market. In contrast, our third result shows that achievability fails for *all* non-trivial markets with a given set of consumer types if and only if for any market for which the seller finds it optimal to sell only the best product, there is a single optimal price for the product. These results imply that with two consumer types either achievability holds for all markets or for no non-trivial market. With three or more types, achievability may hold for some markets and fail for others.

Taken together, our results show that screening interferes with the achievability of first-best consumer surplus and make progress toward understanding the welfare effects of price discrimination in multi-product settings.

1 Model

There is a monopolistic seller, a mass 1 of consumers, and a set $T = 1, \dots, n$ of consumer *types*. There is a set $A = 0, 1, \dots, k$ of *products*, where $k \geq 1$ and product 0 is the outside option. A product can correspond to a particular quantity or quality of a good or a service or to a bundle of goods or services. For example, if a streaming service offers a movie subscription, a series subscription, and full-access subscription that combines both, then there are four products (including the outside option). The cost of production is 0. Type i 's valuation for a product a is $v_a^i \geq 0$, with $v_0^i = 0$. We assume that some product \bar{a} is the “best product” that all consumers prefer, that is, $v_{\bar{a}}^i > v_a^i$ for all types i and products $a \neq \bar{a}$. In the streaming setting the best product would be the full-access subscription. We place no restrictions on how the other products are ranked by different types. We assume that types are ranked so that a higher type has a higher valuation for any product, that is, $v_a^1 < v_a^2 < \dots < v_a^n$ for any product $a \neq 0$. In the streaming setting, the higher the consumer's type, the more he likes to watch shows and movies, so the higher is his valuation for every kind of subscription. But some types of consumers may prefer a movie subscription to a series subscription while other types may have the opposite preference.

An *allocation* $x \in X = \Delta(A)$ is a distribution over products, where x_a denotes the probability of product a . An allocation x is empty if $x_0 = 1$, and is non-empty otherwise. For each type the efficient allocation x satisfies $x_{\bar{a}} = 1$. The (expected) utility of a type i consumer from an

allocation x and a payment p is $v^i \cdot x - p = (\sum_a v_a^i x_a) - p$.

A *mechanism* consists of an allocation rule $x : T \rightarrow X$ and a payment rule $p : T \rightarrow R$. Mechanism $M = (x, p)$ is incentive compatible (IC) if for all types i and j ,

$$v^i \cdot x(i) - p(i) \geq v^i \cdot x(j) - p(j).$$

Mechanism M is individually rational (IR) if for all types i ,

$$v^i \cdot x(i) - p(i) \geq 0.$$

Henceforth, by “mechanism” will refer to an IC and IR mechanism, unless otherwise stated.

We often represent a mechanism indirectly by a *menu* of allocation-price pairs, where each type chooses a pair that maximizes his utility. If a type is indifferent between two allocation-price pairs, he chooses the one with a higher price. If, further, the prices are identical, then the tie breaking can be arbitrary.

A mechanism is a *non-screening* mechanism if it can be represented by a menu with a single allocation-price pair, in addition to the outside option at price 0. Of particular interest is the set of non-screening mechanisms $\{N^i\}_{i \in T}$, where mechanism N^i offers the best product \bar{a} at price v_a^i . Types $j < i$ are *excluded* ($x_0(j) = 1$ and $p(j) = 0$) by N^i and types $j \geq i$ obtain the best product and pay v_a^i . A mechanism is a *screening* mechanism if it is not a non-screening mechanism, that is, every menu that represents it includes at least two positive allocation-price pairs.

A *market* $f \in \Delta(T)$ is a distribution over types, where f_i denotes the fraction of consumers with type i . The consumer surplus in market f with mechanism $M = (x, p)$ is $CS(f, M) = E_{i \sim f}[v^i \cdot x(i) - p(i)]$. A mechanism (x, p) is *optimal* for market f if it maximizes revenue

$$E_{i \sim f}[p(i)]$$

across all mechanisms. For a market f , let $ER(f)$ be the revenue in an optimal mechanism, $\mathcal{M}(f)$ be the set of optimal mechanisms, and $CS(f)$ be the highest consumer surplus across all optimal mechanisms,

$$CS(f) = \max_{M \in \mathcal{M}(f)} CS(f, M).$$

Market f is a *non-screening* market if for some i , mechanism N^i is optimal for f . Otherwise, f is a *screening* market, for which every optimal mechanism is a screening mechanism. Market

f is *efficient* if $N^{\underline{i}(f)}$ is an optimal mechanism for the market, where $\underline{i}(f)$ is the lowest type in the support of f . Otherwise, the market is *inefficient*

A *segmentation* $\mu \in \Delta(\Delta(T))$ of a market f is a distribution over markets that average to f , that is, $E_{f' \sim \mu}[f'] = f$. A market f' in the support of a segmentation μ is a market *segment* (or simply a segment). Abusing notation, let $CS(\mu)$ be the consumer surplus in segmentation μ ,

$$CS(\mu) = E_{f' \sim \mu}[CS(f')].$$

When discussing segmentations of some market f , we refer to f as the unsegmented market.

Another interpretation of the model is that there is a single consumer and the seller's prior over the consumer's type is f . The seller receives information about the consumer's type, and μ represents the distribution over the seller's posterior beliefs.

1.1 Upper Bound on the Maximum Consumer Surplus

Given a market f , the maximum consumer surplus across all segmentations of f is at most the expected surplus of an efficient allocation minus the seller's revenue in f . This is because, for any segmentation, the seller can offer a mechanism in $\mathcal{M}(f)$ in all market segments. The following lemma formalizes this observation.

Lemma 1 *For any segmentation μ of a market f , $CS(\mu) \leq E_{i \sim f}[v_a^i] - ER(f)$.*

We study the conditions under which this bound is achieved.

Definition 1 *A segmentation μ of market f achieves first-best consumer surplus if $CS(\mu) = E_{i \sim f}[v_a^i] - ER(f)$. If such a segmentation exists then first-best consumer surplus is achievable for market f .*

By definition, first-best consumer surplus is achievable for any efficient market.

1.2 The Surplus Triangle

Given a segmentation μ and a selection S of an optimal mechanism $S(f') \in \mathcal{M}(f')$ for each market segment f' , we denote the resulting consumer surplus by $CS(\mu, S)$. Given a market f , denote by $\Gamma(f)$ the set of consumer-producer surplus pairs resulting from all possible segmentations of f and selections of an optimal mechanism for each segment. Abusing notation, let

$ER(\mu) = E_{f \sim \mu}[ER(f)]$ be the producer surplus resulting from segmentation μ , and consider a consumer-producer surplus pair $(CS(\mu, S), ER(\mu))$. Since $CS(\mu, S) \geq 0$, $ER(\mu) \geq ER(f)$, and $CS(\mu, S) + ER(\mu) \leq E_{i \sim f}[v_a^i]$, the set $\Gamma(f)$ is a subset of the “surplus triangle”

$$\Delta(f) = \{(a, b) : a \geq 0, b \geq ER(f), a + b \leq E_{i \sim f}[v_a^i]\}.$$

If $\Gamma(f)$ coincides with the surplus triangle we say that the surplus triangle is *achievable* for f .

Bergemann, Brooks, and Morris (2015) showed that the surplus triangle is achievable for any market f with a single product. The surplus triangle is also obviously achievable for any “singleton market,” which consists only of consumers of some single type i . In this case, the surplus triangle is $\{(0, v_a^i)\}$.

To determine whether the surplus triangle is achievable for a non-singleton market, it is enough to determine whether each of its vertices is generated by some segmentation. The vertex $(0, E_{i \sim f}[v_a^i])$ is generated by first-degree price discrimination. The vertex $(E_{i \sim f}[v_a^i] - ER(f), ER(f))$ is generated by segmentations that achieve first-best consumer surplus. The vertex $(0, ER(f))$ generates the lowest possible total surplus of $ER(f)$.

Definition 2 *A segmentation μ of market f achieves the lowest possible total surplus if for some selection of an optimal mechanism for each segment the resulting consumer-producer surplus pair is $(0, ER(f))$. If such a segmentation exists then the lowest possible total surplus is achievable for market f .*

The discussion above shows the following.

Lemma 2 *The surplus triangle is achievable for a market if and only if first-best consumer surplus and the lowest possible total surplus are achievable for the market.*

1.3 Conditions for Achieving First-best Consumer Surplus

We specify two conditions which are together necessary and sufficient for a segmentation to achieve first-best consumer surplus. First, because the resulting allocation is efficient, every segment must be efficient. Second, the seller should not benefit from the segmentation, that is, every optimal mechanism for the unsegmented market must be optimal for every segment.³

³Otherwise, there is a segment f' such that the seller can benefit by offering in f' an optimal mechanism for f' and offering in all other segments an optimal mechanism for the unsegmented market.

Lemma 3 *For any segmentation μ of a market f , the following are equivalent:*

1. μ achieves first-best consumer surplus.
2. For some optimal mechanism M of f and every segment f' of μ , f' is efficient and M is optimal for f' .
3. For every optimal mechanism M of f and every segment f' of μ , f' is efficient and M is optimal for f' .

2 Two Types

We first consider markets with two types of consumers (and any number of products), and identify each market by its fraction $q \in [0, 1]$ of type 2 consumers. The following lemma divides the set of markets $[0, 1]$ into at most three regions. The first region includes markets in which the fraction of type 1 consumers is high, so mechanism N^1 is optimal for these markets and they are efficient. The second region includes markets in which the fraction of type 2 consumers is high, so mechanism N^2 is optimal for these markets and they are inefficient (except for market 1). The third region, which may be empty, includes the remaining, intermediate markets. These markets are screening markets, that is, the allocations of the two types are different and non-empty. Moreover, the optimal mechanisms may vary across markets in this region. To formalize this, denote by $\mathcal{F}(M)$ the (possibly empty) set of markets for which a particular mechanism M is optimal.

Lemma 4 *There exist q_1 and q_2 , $0 \leq q_1 \leq q_2 \leq 1$, such that $\mathcal{F}(N^1) = [0, q_1]$, $\mathcal{F}(N^2) = [q_2, 1]$, and $\mathcal{F}(M) \subseteq [q_1, q_2]$ for any mechanism $M \neq N^1, N^2$.*

Proof. We first show that for any mechanism M , $\mathcal{F}(M)$ is a closed interval. Indeed, if M is optimal for two markets q, q' , then it is also optimal for any convex combination q'' of these markets, because for any mechanism, the revenue in q'' is the same convex combination of the revenues in q and in q' . And $\mathcal{F}(M)$ is closed because the revenue from any mechanism is continuous in the market q . We now argue that $q_1 \leq q_2$ and for any $M \neq N^1, N^2$, we have $\mathcal{F}(M) \subseteq [q_1, q_2]$. To see this, consider any two mechanisms M, M' with payment rules $p \neq p'$. Then there is at most a single market q where the two mechanisms have the same revenue,




Figure 1: (a) $q_1 = q_2$. For any market, either N^1 or N^2 is optimal. (b) $q_1 < q_2$. Neither N^1 nor N^2 is optimal for markets in the interval (q_1, q_2) . (c) $r_a^1 \leq r_a^2$. (d) $r_a^1 > r_a^2$.

$qp(1) + (1 - q)p(2) = qp'(1) + (1 - q)p'(2)$. Therefore, the intersection of $\mathcal{F}(M)$ and $\mathcal{F}(M')$ is at most a single market. The claim now follows from observing that for any mechanism $M \neq N^1, N^2$, the payment rules of M , N^1 , and N^2 are all different. ■

If $q_1 = q_2$, then all markets are non-screening markets, as shown in Figure 1 (a). Since the seller offers only the best product in each market, the setting is equivalent to one with a single product. Bergemann, Brooks, and Morris (2015)'s result then shows that the surplus triangle, and first-best consumer surplus in particular, are achievable for all markets. The proposition below shows that if $q_1 < q_2$, which is shown in Figure 1 (b), then first-best consumer surplus is unachievable for *any* inefficient market (the markets in $(q_1, 1)$).⁴ As mentioned in Section 1, first-best consumer surplus is achievable for any efficient market. We thus obtain a characterization of the achievability of first-best consumer surplus.

Proposition 1 *For any inefficient market q , first-best consumer surplus is achievable if and only if $q_1 = q_2$.*

⁴A market $q < 1$ is efficient if and only if N^1 is optimal for the market, and these are the markets $[0, q_1]$. Market 1 is clearly efficient.

Proof. Suppose that $q_1 = q_2$. For completeness, we replicate Bergemann, Brooks, and Morris (2015)'s result that first-best consumer surplus is achievable for all markets. This is true for markets in $[0, q_1] \cup \{1\}$ because they are efficient. Consider a market $q \in [q_2, 1]$, so mechanism N^2 is optimal for q , and a segmentation of q into $q' = 1$ and $q'' = q_1 = q_2$.⁵ Both q' and q'' are efficient and have N^2 as an optimal mechanism, so the segmentation achieves first-best consumer surplus by Lemma 3.

Now suppose that $q_1 < q_2$, and suppose that some segmentation μ of a market q achieves first-best consumer surplus. We show that q is efficient, that is, q is in $[0, q_1] \cup \{1\}$. By Lemma 3, every segment in μ is efficient, and any optimal mechanism for q is optimal for every segment. The only optimal mechanism for market 1 is mechanism N^2 . But since $q_1 < q_2$ and $\mathcal{F}(N^2) = [q_2, 1]$, N^2 is not optimal for any market in $[0, q_1]$. Therefore, either every segment of μ is equal to 1, in which case $q = 1$, or every segment is in $[0, q_1]$, in which case $q \in [0, q_1]$. Therefore, q is efficient. ■

We now turn to the achievability of the surplus triangle. As discussed in Section 1, the surplus triangle is a singleton and is achievable for any singleton market (markets 0 and 1). The following result provides a characterization for non-singleton markets.

Proposition 2 *For any non-singleton market q , the surplus triangle is achievable if and only if $q_1 = q_2$.*

Proof. Suppose that $q_1 = q_2$. As noted by Bergemann, Brooks, and Morris (2015), the same segmentation that achieves first-best consumer surplus also achieves the surplus triangle.

Now suppose that $q_1 < q_2$. By Proposition 1, first-best consumer surplus, and therefore the surplus triangle, is unachievable for any inefficient market $q > q_1$. Now consider an efficient non-singleton market $q \leq q_1$, so mechanism N^1 is optimal for q . In this case the lowest possible total surplus is unachievable. This is because if consumer surplus is 0 in some segment $q' \neq 0$, mechanism N^2 must be optimal for q' . Then mechanism N^1 is not optimal for $q' \neq 0$ and the segmentation increases producer surplus. ■

Haghpanah and Hartline (2021) characterize the two cases, $q_1 = q_2$ or $q_1 < q_2$, in terms of the valuations of the two types, which are a primitive of the model. The characterization shows that $q_1 = q_2$ if and only if for any product a , type 2 has a higher ratio of valuations of product

⁵The segmentation assigns probability α to q' and probability $1 - \alpha$ to q'' , where $\alpha = \frac{q - q_2}{1 - q_2}$.

a to \bar{a} , that is, $r_a^1 \leq r_a^2$, where $r_a^i = v_a^i/v_{\bar{a}}^i$. Figure 1 (c) and (d) illustrate this inequality and the reverse inequality for the case of two products.

3 More than Two Types

We now consider markets with any number of types and products. The logic of Bergemann, Brooks, and Morris (2015) shows that if for a given set of types all markets are non-screening markets, then the surplus triangle, and thus first-best consumer surplus, are achievable for every market with this set of types. We will show that this condition is in fact necessary by proving that first-best consumer surplus, and thus the surplus triangle, are not achievable for any screening market. Of course, a screening mechanism is inefficient; the result will show that if a market is inefficient because it is a screening market (and not only because some types are excluded), then it is impossible to achieve efficiency via segmentation without the seller appropriating some of the surplus gains.

Theorem 1 *First-best consumer surplus is not achievable for any screening market.*

Similarly to the proof of Proposition 1, the proof of Theorem 1 shows that if a screening market with a particular optimal mechanism is segmented into efficient markets, then this mechanism is not optimal for all segments. Combining this observation with Lemma 3 proves Theorem 1. This is straightforward when there are only two types because, as we have seen, with two types the only efficient market for which a screening mechanism may be optimal is the non-screening market q_1 . But with more than two types, the convex hull of the set of efficient markets for which a screening mechanism is also optimal may include screening markets. In Figure 2, this set is depicted in green and its convex hull is the shaded region, which includes screening markets. Such screening markets could thus conceivably be segmented in a way that achieves first-best consumer surplus.

Theorem 1 rules this out. Applied to Figure 2, Theorem 1 shows that any segmentation of a market in the shaded region that involves segment f^2 and some segment f^1 in the green curve increases the seller's revenue.

To prove Theorem 1, suppose that some screening market f with an optimal screening mechanism M is segmented into efficient markets. By definition, no non-screening mechanism is optimal for f and, in particular, mechanism N^j is not optimal for f , where j is the lowest




Figure 2: The set of markets with three types and the screening and non-screening regions. The convex hull (shaded gray) of the set of efficient markets for which a screening mechanism is also optimal (in green) includes some screening markets.

type not excluded in M . That N^j is not optimal for f implies that N^j is not optimal for some segment f' because the set of markets for which a given mechanism is optimal is convex.⁶ The following lemma, which completes the proof of Theorem 1, shows that mechanism M is not optimal for f' .

Lemma 5 *Consider an efficient market f' and a mechanism M in which the lowest type that is not excluded is j . If N^j is not optimal for f' , then M is not optimal for f' .*

To see how Lemma 5 works in Figure 2, notice that any segmentation of a screening market in the shaded region into efficient markets must include f^2 and one or more markets f^1 in the green curve as segments. And since each of N^1 , N^2 , and N^3 is not optimal for either f^2 or f^1 (or both), Lemma 5 shows that the optimal screening mechanism for the screening market is not optimal for at least one segment.

The proof of Lemma 5 is provided below. For some intuition for why Lemma 5 holds, suppose first that M does not exclude any type (so $j = 1$). If the support of f' does not include type 1, then M is not optimal for f' because every type in the support of f' obtains positive utility in M (since these types can mimic type 1). If the support of f' includes type 1, then $N^j = N^1$ is optimal for f' (because f' is efficient), so the lemma holds trivially. Now suppose that M

⁶Take a mechanism, a set of markets for which the mechanism is optimal, and a convex combination of these markets. Any mechanism that generates more revenue for the convex combination must increase the revenue for at least one of the markets.




Figure 3: Construction of mechanism M' from mechanism M in the proof of Lemma 5.

excludes at least one type (so $j > 1$). If the lowest type in the support of f' is at least j , then a modified version of the argument for $j = 1$ applies. If the lowest type in the support of f' is lower than j , a more elaborate argument is needed. We use the fact that N^j is not optimal for f' to modify M and obtain a mechanism M' that improves upon M in f' . Mechanism M' “screens more” than M by having a non-empty, inefficient allocation for some types that are excluded in M .

Proof. Denote by i the lowest type in the support of f' . If $i > j$, then mechanism M is not optimal for f' because every type in the support of f' obtains positive utility in M (because these types can mimic type j). If $i = j$, then efficiency means that N^j is optimal for f' so the lemma holds trivially.

Suppose that $i < j$. Since f' is efficient, N^i is optimal for f' . We construct a mechanism M' that generates a higher revenue than M in f' . In M' , types below i get an empty allocation. Types i to $j - 1$ get product \bar{a} with a small probability $\epsilon > 0$ and the outside option with probability $1 - \epsilon$ and pay ϵv_a^i . Types j to n get the same allocation as in M , but their payment is decreased by $\epsilon(v_a^j - v_a^i)$ relative to their payment in M . This is depicted in Figure 3.

Mechanism M' generates a higher revenue than mechanism M in f' . Compared to M , M' gains ϵv_a^i from every type $i' \geq i$ and loses ϵv_a^j from every type $i' \geq j$. The difference in revenue is $\epsilon v_a^i \Pr[i' \geq i] - \epsilon v_a^j \Pr[i' \geq j]$, which is ϵ times the difference between the revenue of mechanism N^i and the revenue of mechanism N^j . This difference is strictly positive because N^i is optimal but N^j is not. It remains to show that M' is IR and IC for small enough $\epsilon > 0$.

IR holds for types $1, \dots, i - 1$ because they are excluded in M' . A type $i' = i, \dots, j - 1$ has utility $\epsilon v_a^{i'} - \epsilon v_a^i \geq 0$, and a type $i' \geq j$ has a higher utility in M' than in M . Thus, IR holds for any $\epsilon > 0$.

For IC, observe that M' coincides with M in the limit as ϵ goes to 0. Thus, if an IC constraint holds strictly in M , then it is satisfied in M' for small enough $\epsilon > 0$. In mechanism M a type i'

strictly prefers not to mimic another type i'' in two cases: (1) if $i' > j$ and $i'' < j$; (2) if $i' < j$ and $i'' \geq j$. In case (1), type i' has a strictly positive utility in M because he can mimic type j . Thus i' strictly prefers not to mimic type i'' (and get utility 0) in M . In case (2), type i' gets a strictly negative utility from mimicking i'' because $v^{i'} \cdot x(i'') - p(i'') < v^j \cdot x(i'') - p(i'') \leq 0$, where the last inequality follows since the utility of type j is 0 and incentive compatibility of mechanism M implies that the utility of type j from mimicking type i'' cannot be positive.

We next verify the remaining IC constraints in mechanism M' . Consider a type $i' < j$. As discussed in case (2) above, such a type i' does not benefit from mimicking types $i'' \geq j$. Type i' prefers the allocation of types $1, \dots, i-1$ (the outside option) to the allocation of types $i, \dots, j-1$ if and only if $\epsilon(v_a^{i'} - v_a^i) \leq 0$, that is, $i' \leq i$. Thus, truth-telling maximizes the utility of every type $i' < j$. For a type $i' \geq j$, note that mimicking a type j, \dots, n is not beneficial since M is IC and all such types get the same additional payment in M' . From case (1) above, a type $i' > j$ does not benefit from mimicking types $1, \dots, j-1$. Finally, the utility of type j in M' is at least $\epsilon(v_a^j - v_a^i) > 0$, which is the utility it would get by mimicking types $i, \dots, j-1$, and is no lower than the utility of 0 it would get by mimicking types $1, \dots, i-1$. ■

3.1 Achievability of First-Best and the Surplus Triangle

Theorem 1 and the logic of Bergemann, Brooks, and Morris (2015) imply that first-best consumer surplus and the surplus triangle are achievable for all markets with a given set of types T if and only if all markets with that set of types are non-screening markets, that is, $\cup_i \mathcal{F}(N^i) = \Delta(T)$. Whether non-screening is optimal for a *given* market is in general difficult to ascertain. But Haghpannah and Hartline (2021) show that non-screening is optimal for *all* markets if and only if the ratio $r_a^i = v_a^i / v_a^i$ is non-decreasing in i for every product a . This characterization is illustrated in Figure 4 (a) and (b). This characterization and Theorem 1 lead to the following result.

Theorem 2 *For any set of types T , the following are equivalent:*

1. *First-best consumer surplus is achievable for every market.*
2. *The surplus triangle is achievable for every market.*
3. *Every market is a non-screening market.*
4. *The ratio r_a^i is non-decreasing in i for all a .*




Figure 4: (a) Every market is a non-screening market (Statement (3) of Theorem 2). (b) The ratio of valuations increases in the valuation for the best product (Statement (4) of Theorem 2). (c) Regions $\mathcal{F}(N^1)$, $\mathcal{F}(N^2)$, and $\mathcal{F}(N^3)$ do not intersect (Statement (3) of Theorem 3). (d) The ratio of valuations decreases in the valuation for the best product (Statement (4) of Theorem 3).

Proof. (2) \rightarrow (1): By definition.

(3) \rightarrow (1) and (2): If offering only \bar{a} is optimal for all markets, the setting is equivalent to one with a single product \bar{a} . The results of Bergemann, Brooks, and Morris (2015) then imply (1) and (2).

(1) \rightarrow (3): Immediate from Theorem 1.

(3) \rightarrow (4) and (4) \rightarrow (3): From Proposition 1 of Haghanian and Hartline (2021). ■

Using the notation from Section 2, Theorem 2 states that with two types first-best consumer surplus and the surplus triangle are achievable for every market if and only if $q_1 = q_2$. Thus, Theorem 2 generalizes parts of Proposition 1 and Proposition 2. However, Proposition 1 and Proposition 2 show that if $q_1 < q_2$, then first-best consumer surplus is unachievable for *all* inefficient markets and the surplus triangle is unachievable for *all* non-singleton markets. With more than two types, it may be that some markets are screening markets and yet first-best consumer surplus and the surplus triangle are achievable for some inefficient and non-singleton non-screening markets with the same set of types.

3.2 Unachievability of First-Best and the Surplus Triangle

We identify a condition for unachievability of first-best consumer surplus and the surplus triangle for *all* inefficient and non-singleton markets. For this, let us interpret the condition $q_1 < q_2$ in Proposition 1 and Proposition 2 as stating that the set of screening markets, (q_1, q_2) , *separates* the sets $[0, q_1]$ and $[q_2, 1]$ of non-screening markets. Our second main result shows that this is the correct condition for any number of types. The result is illustrated in Figure 4 (c) and (d).

Theorem 3 *For any set of types T , the following are equivalent:*

1. *First-best consumer surplus is unachievable for every inefficient market.*
2. *The surplus triangle is unachievable for every non-singleton market.*
3. *For every market, N^i is optimal for at most one i .*
4. *For every pair of types $i < j$, there exists some product a such that $r_a^i > r_a^j$.*

To see why statement (3) implies statement (1) in Theorem 3, suppose that that first-best consumer surplus is achievable for some inefficient market f . By Theorem 1, market f is a

non-screening market, so for some $i > \underline{i}(f)$ mechanism N^i is optimal for f , where $\underline{i}(f)$ is the lowest type in the support of f . At least one segment in any segmentation of f into efficient markets must include consumers of type $\underline{i}(f)$. By Lemma 3, both $N^{\underline{i}(f)}$ and N^i are optimal for that segment, so statement (3) does not hold.

To show that statement (4) implies statement (3) in Theorem 3, we cannot apply the results of Hagpanah and Hartline (2021) as we did in the proof of Theorem 2. Instead, we develop a new result that relates properties of type ratios to the set of non-screening mechanisms that may be optimal for any market. This is the content of the following lemma.

Lemma 6 *Consider a pair of types $i < j$ such that $r_a^i > r_a^j$ for some a . Then, for any market f , mechanisms N^i and N^j are not both optimal.*

The proof of Lemma 6 shows that, given a pair of types $i < j$ such that $r_a^i > r_a^j$ for some product a , if both N^i and N^j are assumed optimal, then there exists a mechanism that outperforms N^j .

Proof of Lemma 6. Assume for contradiction that $r_a^i > r_a^j$ for some $i < j$ and a , and N^i and N^j are both optimal for market f . Denote by q_i the fraction of types i and higher in market f , and by q_j the fraction of types j and higher in market f . For both N^i and N^j to be optimal, we must have $v_a^i q_i = v_a^j q_j$, that is, $q_i = \frac{v_a^j q_j}{v_a^i}$. Thus we can write

$$(1) \quad v_a^i q_i = v_a^i \left(\frac{v_a^j q_j}{v_a^i} \right) = \left(\frac{v_a^j v_a^i}{v_a^i} \right) q_j > v_a^j q_j,$$

where the inequality followed from the assumption that $r_a^i > r_a^j$ (that is, $v_a^i/v_a^i > v_a^j/v_a^j$).

We construct a mechanism M that improves upon N^j as follows. Types $i, \dots, j-1$ get product a with probability ϵ and pay ϵv_a^i . Types j, \dots, n get product \bar{a} and pay $v_a^j - \epsilon(v_a^j - v_a^i)$.

Let us compare the revenue of M with the revenue of N^j . Types $i, \dots, j-1$ pay ϵv_a^i more in M than in N^j . Types j and higher pay $\epsilon(v_a^j - v_a^i)$ less in M than in N^j . The difference in expected revenue is

$$\epsilon v_a^i (q_i - q_j) - \epsilon (v_a^j - v_a^i) q_j = \epsilon (v_a^i q_i - v_a^j q_j) > 0,$$

where the inequality followed from (1). To complete the proof, we show that M is IC and IR, which contradicts the assumption that N^j is optimal.

We begin by showing that mechanism M is IR. Types lower than i are excluded. Any type i' from i to $j-1$ has utility $\epsilon(v_a^{i'} - v_a^i) \geq 0$. Types j and higher have a higher utility in M than in N^j .

For IC, observe similarly to the proof of Lemma 5 that if an incentive constraint holds strictly in N^j , then it is satisfied in M for small enough $\epsilon > 0$. In particular, (1) a type $i' > j$ does not benefit from mimicking a type $i'' < j$, and (2) a type $i' < j$ does not benefit from mimicking a type $i'' \geq j$.

We now verify the remaining incentive constraints. A type $i' < j$ prefers the allocation of types $i, \dots, j-1$ to the outside option if and only if $\epsilon(v_a^{i'} - v_a^i) \geq 0$, that is, $i' \geq i$. Thus the incentive constraints are satisfied for types $i' < j$. For types $i' \geq j$, note that mimicking any type j, \dots, n is not beneficial since all such types have the same allocation and payment. Finally, the utility of type j in M is $\epsilon(v_a^j - v_a^i)$, which is the utility it would receive by mimicking types $i, \dots, j-1$, and is strictly more than the utility he would receive by mimicking types $1, \dots, i-1$. ■

We now use Lemma 6 to prove Theorem 3.

Proof of Theorem 3. (3) \rightarrow (1): argued above.

(3) \rightarrow (2): That (3) implies (1) also shows that that if (3) holds, then the surplus triangle is not achievable for any inefficient market. It remains to show that the surplus triangle is not achievable for any non-singleton efficient market. Consider a non-singleton efficient market f and suppose that a segmentation μ achieves the lowest possible total surplus. Consider a segment f' whose support includes type $\bar{i}(f)$, the highest type in the support of f . Because μ achieves the lowest possible total surplus, consumer surplus in f' is 0, so $N^{\bar{i}(f)}$ is optimal for f' . And since f is efficient, $N^{\underline{i}(f)}$ is optimal for f , where $\underline{i}(f)$ is the lowest type in the support of f . By Lemma 3, $N^{\underline{i}(f)}$ is also optimal for f' .

(4) \rightarrow (3): Directly from Lemma 6.

(1) \rightarrow (4) and (2) \rightarrow (4): Suppose for contradiction that for some $i < j$, $r_a^i \leq r_a^j$ for all a . Haghpasand and Hartline (2021) show that either N^i or N^j is optimal for any market with support in $\{i, j\}$. By Proposition 1 and Proposition 2, first-best consumer surplus and the surplus triangle are achievable for every market with support in $\{i, j\}$. ■

4 Concluding remarks

Our analysis relies on two main assumptions. The first assumption is that consumer types are ranked, so for each product a higher type's valuation is higher than that of any lower type. This assumption implies that if a type is allocated a product, then all higher types obtain information

rents. This property is crucial for our key Lemma 5, and therefore for our key finding that first-best consumer surplus is not achievable for screening markets. The second assumption is that it is efficient to allocate a particular “best product” to all types, and production is costless. This assumption is convenient because, given the first assumption, it makes the definition of screening simple. One policy implication of these assumptions is that by prohibiting screening, a regulator can make first-best consumer surplus achievable for all markets (even screening markets) since the seller would then find it optimal to offer only the best product in each market.

In the Online Appendix we consider a relaxation of our second assumption by studying a setting with linear utility and positive production costs as in Mussa and Rosen (1978). In this setting, it is efficient to allocate different quantities to different types. Our results for two types extend to this setting, as does our key finding (for any number of types) that first-best consumer surplus is achievable only for non-screening markets. Our results extend to this setting because if a type is allocated a product, then all higher types obtain information rents.

We leave further investigations of the welfare effects of price discrimination in environments not covered by our results for future work. One direction is to investigate the maximum consumer surplus across all segmentations when first-best consumer surplus is not achievable. This will likely require a new approach; we consider a two-type, two-product example in the Online Appendix.

References

- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. “The limits of price discrimination.” *American Economic Review* 105 (3):921–57.
- Cai, Yang, Federico Echenique, Hu Fu, Katrina Ligett, Adam Wierman, and Juba Ziani. 2020. “Third-party data providers ruin simple mechanisms.” *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4 (1):1–31.
- Daskalakis, Constantinos, Christos Papadimitriou, and Christos Tzamos. 2016. “Does Information Revelation Improve Revenue?” In *Proceedings of the 2016 ACM Conference on Economics and Computation*. 233–250.
- Haghpanah, Nima and Jason Hartline. 2021. “When is pure bundling optimal?” *The Review of Economic Studies* 88 (3):1127–1156.
- Haghpanah, Nima and Ron Siegel. 2021. “Pareto Improving Segmentation of Multi-product markets.” *working paper* .
- Hidir, Sinem and Nikhil Vellodi. 2021. “Privacy, Personalization, and Price Discrimination.” *Journal of the European Economic Association* 19 (2):1342–1363.
- Ichihashi, Shota. 2020. “Online privacy and information disclosure by consumers.” *American Economic Review* 110 (2):569–95.
- Mussa, Michael and Sherwin Rosen. 1978. “Monopoly and product quality.” *Journal of Economic theory* 18 (2):301–317.