

Pareto Improving Segmentation of Multi-product Markets

Nima Haghpanah and Ron Siegel*

November 22, 2019

Abstract

We investigate whether a market served by a multi-product monopolistic seller can be segmented in a way that benefits all consumers. In the unsegmented market, the seller makes the same menu of products and product bundles available to all consumers. In a segmented market, the seller can offer a potentially different menu in each market segment. We show that for generic markets there exists a segmentation in which the surplus of each consumer is weakly higher, and the surplus of some consumer is strictly higher, than in the unsegmented market.

1 Introduction

Recent technological advances enable sellers to segment markets based on detailed consumer data and make segment-specific offers. An extreme example is first-degree price discrimination: every consumer is offered his most preferred product at a price equal to his willingness to pay. This eliminates all consumer surplus. Coarser market segmentations make it optimal for the seller to offer consumers relevant products at prices lower than their willingness to pay. This may increase some consumers' surplus relative to what they

*Department of Economics, The Pennsylvania State University, University Park, IL 16802 (e-mail: nuh47@psu.edu and rus41@psu.edu). We thank Dirk Bergemann, Ben Brooks, Wouter Dessein, Harry Di Pei, Piotr Dworzak, Wioletta Dziuda, Stephen Morris, Eduardo Perez Richet, Andy Skrzypacz, Jidong Zhou for very helpful comments and suggestions. We thank seminar participants at Johns Hopkins University, Pennsylvania Economic Theory Conference, and University of Toronto.

would obtain in the unsegmented market, but other consumers may be hurt. This paper investigates which markets can be segmented in a way that benefits all consumers when a multi-product seller maximizes profits in each market segment.

This question is relevant to regulatory discussions regarding consumer privacy and sellers' use of consumer data. A regulator interested in increasing consumer welfare may be able to control the data that the sellers access or use to make targeted offers, or consumers may be able to jointly decide what data to provide to sellers.¹ As a 2012 report by the Federal Trade Commission puts it, "The Commission recognizes the need for flexibility to permit [...] uses of data that benefit consumers."²

We consider a setting in which a multi-product monopolistic seller faces a mass of heterogeneous consumers with preferences over subsets of products. The seller offers a menu of products and product bundles to maximize his profit. In particular, the seller may engage in second-degree price discrimination. If the market is segmented, the seller can offer a potentially different menu in each market segment, thereby combining second- and third-degree price discrimination. We say that a market segmentation is Pareto improving if the surplus that every consumer obtains when choosing from the menu offered in his segment is no lower than the surplus he would obtain when choosing from the menu that would be offered in the unsegmented market, and is strictly higher for some consumers. Our goal is to understand for which markets Pareto improving segmentations exist.

To illustrate, suppose the seller can produce only a single product at no cost and consumers have unit demand. A quarter of the consumers are willing to pay 1 for the product (type 1 consumers), and the rest are willing to pay 2 for the product (type 2 consumers). In this unsegmented market the seller optimally sells the product at a price of 2, only type 2 consumers buy the product, and the surplus of every consumer is 0. This market can be segmented into two market segments in a way that is Pareto improving: the first segment includes all type 1 consumers and a small mass of type 2 consumers, and the second segment includes the remaining type 2 consumers. The seller optimally sells the product at a price of 1 in the first segment and a price of 2 in the second segment. Type t_2

¹In the single-agent interpretation of our model, discussed in Section 3.2, the agent can commit to an information-disclosure policy prior to learning his type, as in Ichihashi (2018).

²"Protecting Consumer Privacy in an Era of Rapid Change, Recommendations for Businesses and Policymakers", FTC report, March 2012.

consumers in the first segment obtain a surplus of 1, and all the other consumers continue to obtain a surplus of 0, so the segmentation is Pareto improving. Similar segmentations are Pareto improving for any “inefficient market,” in which the seller optimally sells the product at a price of 2 so type 1 consumers are not served (these are markets in which type 2 consumers are a majority).³

Things are different when the seller can offer multiple products. Continuing with the example, suppose that the seller can also offer a low-quality version of the original product, and consumers still have unit demand. Type 1 consumers are willing to pay 0.75 for the low-quality product and type 2 consumers are willing to pay 1 for it. In the unsegmented market (in which a quarter of the consumers are of type 1), the seller optimally screens consumers: she offers the low-quality product at a price of 0.75 and the original product at a price of 1.75. Type 1 consumers buy the low-quality product and type 2 consumers buy the original product. Unlike the single-product setting, even though the market is inefficient (because type 1 consumers buy the low-quality product), the market cannot be segmented in a way that is Pareto improving. Indeed, any segmentation of the market into multiple segments that are not all identical to the original market must include a segment in which more than three quarters of the consumers are of type 2. In this segment the seller optimally sells only the original product at a price of 2, so the surplus of type 2 consumers is 0 whereas their surplus in the unsegmented market is 0.25. In fact, it can be shown that every segmentation also lowers the average consumer surplus.⁴ Perhaps surprisingly, however, this is the only market where this phenomenon arises: any inefficient market in which the proportion of type 2 consumers is not three quarters can be segmented in a way that is Pareto improving. This is true despite the fact that screening is profit-maximizing for all markets in which the proportion of each type is at least one quarter.

Our first result shows that this is not a coincidence: in environments with two consumer types and any finite number of products, Pareto improving segmentations exist for all but

³This follows from the analysis of Bergemann et al. (2015), who considered markets with a single product (see Section 3.1). Conversely, since setting segment-specific prices can only benefit the seller, no segmentation can increase average consumer surplus for any “efficient market.” In our example, efficient markets are ones in which type 2 consumers are a minority where the seller optimally sells the product at a price of 1.

⁴This can be seen by considering the concavification of the graph of average consumer surplus as a function of the proportion of high types in the market.

a finite number of markets (each market is defined by the ratio of types). We prove this result by characterizing properties of the seller’s optimal menu for each market. This characterization shows that the surplus of one of the types is 0 in every market, and the surplus of the other type is piecewise constant and weakly decreasing in that type’s proportion in the market. This implies that only markets at the endpoints of the intervals on which the second type’s surplus is constant cannot be segmented in a way that is Pareto improving.

With more than two consumer types things are more complicated. The main difficulty is that no characterization of profit-maximizing menus exists when the seller can offer multiple products. In particular, the seller may find it strictly optimal to offer randomized bundles of products. Since we cannot characterize the optimal menus in each market and the resulting consumer surplus for each type, we develop a different approach. This novel approach is based on understanding what drives market inefficiency: the only reason a seller serves some consumer types inefficiently is to reduce the information rents of other types. This simple observation has far-reaching implications. We show that for every inefficient market with any finite number of types and products there is a market segment that is Pareto dominating: every consumer in this segment weakly prefers, and some strictly prefer, the menu that maximizes the seller’s profits in this segment to the profit-maximizing menu in the unsegmented market.

Our proof is constructive and shows that the Pareto improving segment may have to include numerous types. We then show that for every market in a generic set of markets, a small perturbation of the market does not change the profit-maximizing menu. Combining these two results delivers our main result: for any finite number of types and products, every inefficient market in a generic set of markets can be segmented into two segments in a way that is Pareto improving. Our definition of a generic set of markets reduces to “all but a finite number of markets” when there are only two types. This provides another proof for our first result.

The assumption of a finite number of types and products is important for our main result. With a continuum of products, it may be that for every market a small perturbation changes the profit-maximizing menu. Consequently, in some environments with a continuum of products, there may be infinitely many markets for which no Pareto improving segmentation exists. However, we show that an appropriate generalization of our main

result holds, and even becomes stronger, if the assumption that the number of products finite is relaxed. In particular, we provide a characterization of markets for which a Pareto improving segmentation exists in a setting where valuations are linear (as in Mussa and Rosen, 1978) and where the number of products may be finite or infinite. We show that a Pareto improving segmentation exists for a market if and only if certain perturbations of the market leave the optimal mechanism unchanged. The relevant perturbations are ones in which the proportions of the consumers of all types above some type with inefficient allocation are increased. If the number of products is indeed finite, then for a generic set of markets, such a permutation leaves the optimal mechanism unchanged. Thus this characterization generalizes our main result in the special case where valuations are linear.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 defines the model. Section 4 provides the two-type analysis. Section 5 studies any number of types and provides the main result. We discuss environments with a continuum of products in Section 6. We conclude the paper in Section 7

2 Related literature

Our work connects the literature on second and third degree price discrimination. The literature that studies third degree price discrimination and its effects and producer and consumer surplus is broad. Pigou (1920) provides examples where a segmentation may decrease total and hence consumer surplus. Follow up work provides conditions for a segmentation to increase or decrease total surplus or consumer surplus (Robinson, 1969; Schmalensee, 1981; Varian, 1985; Aguirre et al., 2010; Cowan, 2016). Our work differs from this literature in three significant ways. First, with third degree price discrimination, the seller offers a single product to all consumers in a market, whereas the seller in our setting may screen consumers in each market by offering a menu of products and bundles. Second, instead of considering expected consumer surplus we use the Pareto criterion. Third, with the exceptions we now discuss, the literature assumes that the segmentation is exogenously fixed.

A growing part of the literature on third degree price discrimination studies surplus across all possible segmentations of a given market for a single product. Bergemann et al. (2015) identify the set of all producer and consumer surplus pairs that can result from

some segmentation of a given market. Their results imply that in environments with a single product any inefficient market can be segmented in a way that is Pareto improving. Glode et al. (2018) study optimal disclosure by an informed agent in a bilateral trade setting, and show that the optimal disclosure policy leads to socially efficient trade, even though information is revealed only partially. Ichihashi (2018) and Hidir and Vellodi (2018) consider maximum consumer surplus when a multi-product seller offers a single product to each market. Ichihashi (2018) considers a finite number of products and compares two regimes, one in which the seller may offer the same product at different prices to different segments, and another one in which the seller fixes the price in advance. Hidir and Vellodi (2018) characterize optimal segmentations with a continuum of products. Braghieri (2017) studies market segmentation with a continuum of firms each producing a single differentiated product. In contrast to these works, the seller in our setting may offer multiple products in a market in order to screen consumers.⁵ The only instance of this we are aware of is a parametric example with two types and non-linear valuations in Bergemann et al. (2015).

The literature on multi-product bundling goes back to Stigler (1963) and Adams and Yellen (1976), who study bundling as an instrument to engage in second degree price discrimination. Theoretical findings on welfare effects of bundling are inconclusive.⁶ The main hurdles are the difficulty with identifying optimal menus and their complexity. Thanassoulis (2004) and Daskalakis et al. (2017) show that optimal menus may have to include randomized bundles. Vincent and Manelli (2007) and Hart and Nisan (2013) show that optimal menus may have to include infinitely many bundles. Daskalakis et al. (2014) and Chen et al. (2015) show that the problem of finding optimal menus is computationally intractable. A more recent literature empirically estimates the welfare effects of bundling (Ho et al., 2012; Crawford and Yurukoglu, 2012).

Our analysis can also be cast in a Bayesian persuasion framework (Kamenica and Gentzkow, 2011). As we discuss in Section 3.2, the market is replaced by a single agent

⁵We use the term screening to mean that there are at least two bundles in the seller’s menu. A menu that offers a single product at a high price and therefore excludes certain consumers is not a screening mechanism.

⁶Adams and Yellen (1976) show that bundling may be inefficient as it leads to oversupply or undersupply of certain goods. Salinger (1995) argues that bundling may result in lower or higher prices and therefore may increase or decrease consumer surplus.

(the sender) who faces the seller (the receiver). A market segmentation corresponds to a distribution over posteriors. One important difference is that in the usual persuasion setting one considers the agent’s expected utility, whereas we consider the agent’s ex-post utility. Another reason that standard persuasion techniques do not help is that the agent’s utility for a given posterior depends on the seller’s optimal menu, for which no characterization exists when there are multiple products.

3 Setup

A monopolistic seller faces a continuum of consumers (Section 3.2 discusses the interpretation of a single consumer). The *environment* includes a finite set T of consumer types and a finite set A of alternatives, where alternative $0 \in A$ is consumers’ outside option of not purchasing from the seller. We will refer to $k = |A|$ as the number of alternatives. A consumer type specifies a valuation for every alternative: type t ’s valuation for alternative a is $v(t, a)$. Type t ’s valuation for a random alternative $x \in \Delta(A)$ is $v(t, x) = E_{a \sim x}[v(t, a)]$. Type t ’s surplus for random alternative x and payment p to the seller is $v(t, x) - p$. Assume that the valuation for the outside option is zero, $v(t, 0) = 0$. We assume that each type t has a unique efficient alternative $\bar{a}(t)$. The seller’s cost of producing each alternative is normalized to 0 without loss of generality.⁷ Notice that place no restrictions on consumers’ valuations: different consumer types may rank the alternatives differently, and consumers’ valuations need not be ordered by their types or satisfy a condition like increasing differences.

Each alternative $a \neq 0$ corresponds to a product or a set of products. This captures horizontal and vertical differentiation, allows for multi-unit demand, and accommodates bundling. To illustrate this, suppose that the seller can produce a product p and a second product that has a low-quality version q_L and a high-quality version q_H . Suppose that consumers may want to buy one or both products but not both versions of the second product. The alternatives correspond to the relevant subsets of $\{p, q_L, q_H\}$: $\emptyset, \{p\}, \{q_L\}, \{q_H\}, \{p, q_L\}, \{p, q_H\}$. Alternatively, we could specify an alternative for every

⁷A non-zero cost $c(a)$ for alternative a can be accommodated by redefining valuations as $\tilde{v}(t, a) = v(t, a) - c(a)$ without changing the analysis or results. Notice that $\tilde{v}(i, a)$ may be negative even if all valuations $v(i, a)$ are non-negative. Thus, throughout the paper we allow for negative valuations.

subset of $\{p, q_L, q_H\}$ and reflect in consumers' types the fact that consumers do not want to buy both versions of the second product. If some consumers demand multiple units of a single product, that would be captured by additional alternatives.

An *allocation rule* $x : T \rightarrow \Delta(A)$ is a mapping from types to random alternatives, where $x(t)$ is the allocation of type t . The allocation rule is efficient if the allocation of each type is efficient, that is, $x(t) = \bar{a}(t)$ with probability one. A (direct) *mechanism* $M = (x, p)$ consists of an allocation rule x and a payment rule $p : T \rightarrow \mathbb{R}$. The interpretation is that the consumer reports his type and receives the corresponding random alternative in return for the specified payment. The mechanism is incentive compatible (IC) if no type benefits from misreporting, that is,

$$v(t, x(t)) - p(t) \geq v(t, x(t')) - p(t')$$

for all types t, t' . The mechanism is individually rational (IR) if every type obtains at least 0 by reporting truthfully, that is,

$$v(t, x(t)) - p(t) \geq 0,$$

for all types t . Any mechanism we will refer to will be IC-IR unless otherwise stated. A mechanism is efficient if its allocation rule is efficient. Every mechanism can be represented by a menu of (random alternative, price) pairs, where each type chooses a pair that maximizes his surplus. If a type is indifferent between two or more pairs, he chooses the one with a higher price (or uses any rule if the prices are the same).

A *market* $f \in \Delta(I)$ is a distribution over types, where $f(t)$ is the fraction of consumers of type t . The *optimal mechanism* $M(f)$ for market f maximizes the seller's revenue among all IC-IR mechanisms.⁸ Type t 's surplus $CS(t, f)$ in market f is the type's surplus from the optimal mechanism. A market is efficient if the optimal mechanism is efficient; otherwise, the market is inefficient. A *segmentation* of market f is a distribution $\mu \in \Delta(\Delta(I))$ over markets that averages to f , that is, $E_{f' \sim \mu}[f'] = f$. We refer to a market in the support of a segmentation as a (market) segment. A segmentation is non-trivial if not all segments are identical to the original market.

⁸Fix an arbitrary selection rule if there are multiple optimal mechanisms.

3.1 Pareto improvements

Our goal is to understand for each environment which markets can be segmented in a way that benefits all consumers. To formalize this, we say that market f' *weakly Pareto dominates* market f if every type t in market f' prefers the optimal mechanism for market f' to the one for market f , that is, $CS(t, f') \geq CS(t, f)$ for all types t such that $f'(t) > 0$. If, in addition, the preference is strict for some type t with $f'(t) > 0$, then we say that f' *Pareto dominates* f . A segmentation μ of market f is *Pareto improving* if every segment weakly Pareto dominates f and some segment Pareto dominates f . A market is *Pareto improvable* if it has a Pareto improving segmentation.

We begin by observing that if a market is efficient then it is not Pareto improvable. Segmenting the market can never lower the seller's revenue, since she can offer the optimal mechanism for the original market in every segment. And the total surplus any segmentation generates is at most the surplus generated by the efficient allocation. Thus, segmenting an efficient market weakly decreases average consumer surplus.

Observation 1 *Any Pareto improvable market is inefficient.*

Which inefficient markets are Pareto improvable? The example in the introduction showed that in an environment with a single alternative and two types all inefficient markets are Pareto improvable. This is in fact true for any environment with a single alternative.

Proposition 1 *In any environment with a single alternative all inefficient markets are Pareto improvable.*

Proposition 1 follows from the proof of Theorem 1 in Bergemann et al. (2015). Their result implies that any inefficient market with a single alternative can be segmented in a way that achieves efficiency and provides the entire expected gains to consumers, but their proof in fact shows that a Pareto improving segmentation exists.⁹ However, that proof relies heavily on there being a single alternative and does not generalize to multiple

⁹A technical point is that their proof, and our Proposition 1, require selecting the efficient mechanism if it is optimal, whereas our definition of the optimal mechanism allows for any selection rule when there are multiple optimal mechanisms. Proposition 1 does not hold for any selection rule, but our results for any number of types and alternatives in the rest of the paper do, and they apply in particular to markets with a single alternative.

alternatives. As we have seen in the introduction, with more than one alternative not all inefficient markets are Pareto improvable. In the following sections of the paper we study environments with any number of products, first for the case of two types and then for any number of types. Before proceeding to this analysis, however, we comment on how our setup and question can be cast in a single-agent setting.

3.2 Single-agent interpretation

Consider an agent whose type is drawn from the set T according to a prior distribution f . Before learning his type, the agent commits to an information disclosure policy, which maps every type in T to a distribution over signals. The seller observes the policy and the realized signal and forms a posterior f' over the agent's type. The seller then offers a mechanism (or menu) to maximize her revenue, and the agent responds by reporting his type optimally.¹⁰ For which prior distributions f does there exist an information disclosure policy that increases the agent's ex-post utility relative to a policy that discloses no information? This model and question are equivalent to those described earlier. Following Aumann et al. (1995) and Kamenica and Gentzkow (2011), we can describe the process as the agent choosing a distribution μ over posteriors f' that averages to f , that is, $E_{f' \sim \mu}[f'] = f$.

The single-agent model corresponds to a Bayesian persuasion setting (Kamenica and Gentzkow, 2011) in which the agent is the sender and the seller is the receiver. The state is the sender's type, the receiver's set of actions is the set of IC-IR mechanisms, and the sender's state-dependent utility from the receiver's chosen mechanism (action) is the sender's utility from responding optimally to the mechanism. However, existing results and techniques in the Bayesian persuasion literature concern the sender's expected utility, whereas our focus is on ex-post utility.¹¹ In addition, no analytical description exists of the sender's state-dependent utility as a function of the receiver's action because there is no characterization of optimal mechanisms in our environment.

¹⁰One motivating example is an online purchase setting in which the seller may be better than the consumer at determining which products are most appropriate for the consumer based on personal data the consumer discloses (see Ichihashi (2018) for a discussion).

¹¹This ex-post criterion may be relevant when would like to find improvements that work for all possible social welfare functions, which assign possibly different weights to different types.

4 Environments with two types

Suppose that there are only two consumer types (but any number $k \geq 1$ of alternatives). As the two-alternative example in the introduction shows, not every inefficient market is necessarily Pareto improvable. However, as we now show, all but at most a finite number of markets are Pareto improvable by a two-market segmentation.

Proposition 2 *In any environment with two types and $k \geq 1$ alternatives, all but at most $k + 1$ inefficient markets are Pareto improvable by a two-market segmentation.*

The first step is to characterize optimal mechanisms, which is possible because there are only two types. We do this in the appendix. The characterization shows that for any environment with two types, there is a low type and a high type with the following properties. First, the surplus of the low type is optimally 0 in any market. Second, the optimal allocation of the high type is efficient in any market. These properties are not obvious since we do not assume any ranking over the valuations of types. The high type is the type t_H whose valuation for the efficient alternative of the other type t_L exceeds the other type's valuation for that alternative, $v(t_H, \bar{a}(t_L)) > v(t_L, \bar{a}(t_L))$. If there is no such type, then every market is efficient and both types' surplus is 0 in any market.

To identify the optimal mechanism in a market it suffices to identify the alternative allocated to the low type, since incentive compatibility then pins down the payment of the high type.¹² The less valuable the low type's alternative is to the high type, the more surplus can be extracted from the high type. Thus, the low type's alternative optimally balances the surplus extraction from the low type with the reduction in the surplus of the high type.

The higher the fraction q of the high type in the market, the more important is the reduction in their surplus. Therefore, the high type's surplus optimally decreases in q . For small enough q , the low type is allocated his efficient alternative, so the allocation is efficient; for large enough q the low type is allocated an alternative that reduces the high type's surplus to 0, so the overall consumer surplus is 0.¹³ For intermediate values of q the

¹²The high type must be indifferent between reporting truthfully and misreporting.

¹³If the high type has a non-positive valuation for the low type's efficient alternative, then in every market each type optimally gets his efficient alternatives and overall consumer surplus is 0.

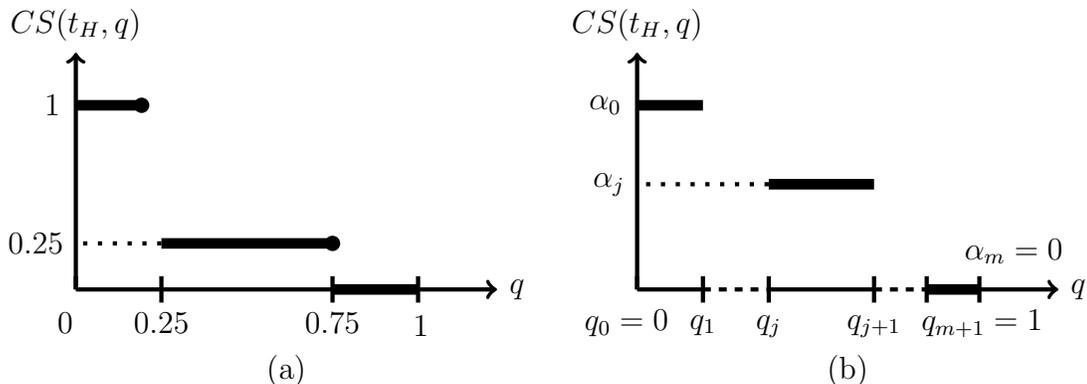


Figure 1: (a) The example from the introduction with two types t_L and t_H and two alternatives L and H , where $v(t_L, L) = 0.75, v(t_L, H) = 1, v(t_H, L) = 1, v(t_H, H) = 2$. (b) The surplus of the high type is constant within each interval, and decreases between intervals.

low type may be allocated an inefficient alternative that does not reduce the high type's surplus to 0.

With two alternatives, for example, the set of markets q can be divided into at most three intervals: an efficient low interval, a zero-surplus high interval, and an inefficient positive-surplus intermediate interval. The high type's surplus is constant on each interval, since it depends only on the optimal mechanism, and decreases on higher intervals. This is what happens in the two-alternative example from the introduction, which is depicted in Figure 1, (a).

More generally, with k alternatives there are up to $k + 1$ intervals, where the optimal mechanism, and therefore the high type's surplus, is constant on each interval, and the high type's surplus decreases on higher intervals. This is depicted in Figure 1, (b). For the following lemma, which summarizes the discussion, we denote by q the market with a fraction q of high types.

Lemma 1 *Consider an environment with two types. For one of the types, denoted t_L , the surplus is 0 in any market: $CS(t_L, q) = 0$ for any q in $[0, 1]$. For the other type, denoted t_H , there exists some $m \leq k$, thresholds $q_0(= 0) < \dots < q_{m+1}(= 1)$, and surpluses $\alpha_0 > \dots > \alpha_m(= 0)$ such that $CS(t_H, q) = \alpha_j$ for q in (q_j, q_{j+1}) .*

Lemma 1 implies Proposition 2. The idea is that all the efficient markets with $q < 1$ are in the first interval $[q_0, q_1]$, where the fraction of low types is sufficiently high to make serving them efficiently optimal for the seller. For any other interval $[q_j, q_{j+1}]$ ($1 \leq j \leq m$),

any market q in the interior of the interval can be segmented into two segments $q' < q''$ such that q'' is also in the interior of the interval $[q_j, q_{j+1}]$, so the surplus of the high type is unchanged, and q' is in the interior of the lower interval $[q_{j-1}, q_j]$, so the surplus of the high type is increased. This shows that the segmentation into q' and q'' is Pareto improving. The endpoints of the intervals, however, may not be Pareto improvable since for any segment $q'' > q$ it may be that the surplus of the high type is decreased. This is the case for market $q = 0.75$ in the two-alternative example in the introduction, which is depicted in Figure 1, (a).

5 Environments with any number of types

With more than two types and multiple alternatives there is no general characterization of optimal mechanisms for different markets. The reason is that it is not clear which IC constraints should optimally bind. Thus, we cannot hope to compute the surplus of each type in each market and identify Pareto improving segmentations directly, as we did for environments with two types. Setting aside the difficulty with characterizing optimal mechanisms, a second difficulty is that it is not clear whether for any Pareto improvable market there exists a two-market Pareto improving segmentation, as is the case when there are two types.¹⁴ Consequently, when looking for Pareto improving segmentations for a given market, we cannot restrict attention to two-market segmentations. Moreover, even the set of two-market segmentations for a given market is quite rich: with two types the set of markets is single dimensional, but with more than two types there is a continuum of directions along each of which a market can be segmented into two markets.

Despite these difficulties, we now show that, for any environment, a two-market Pareto improving segmentation exists for a generic set of inefficient markets. To make this precise, we define a notion of non-genericity below. For every non-zero vector b in $\mathbb{R}^{|T|}$, where $|T| \geq 2$ is the number of types, let $H(b) = \{f \in \Delta(T) : \sum_{t \in T} b_t f(t) = 0\}$ be the set of markets contained in the hyperplane through the origin that is perpendicular to b . By definition, $H(b)$ is contained in a hyperplane of dimension $n - 2$, since $H(b)$ is defined by

¹⁴This is true with two types because by Carathéodory's theorem any market in the convex hull of a set of markets is in the convex hull of two markets from the set. This property does not hold with more than two types.

two independent linear equations.

Definition 1 *Given an environment with n types, a set F of markets is non-generic if it is contained in a finite number of sets $H(b)$: There exists some $L \geq 0$ and non-zero vectors b_1, \dots, b_L in \mathbb{R}^n such that $F \subseteq H(b_1) \cup \dots \cup H(b_L)$.*

We can now state the main result of the paper. A two-market Pareto improving segmentation exists for a generic set of inefficient markets, in the sense that the set of inefficient markets for which no such segmentation exists is non-generic.

Theorem 1 *For any environment with $n \geq 2$ types and $k \geq 1$ alternatives, the set of inefficient markets that are not Pareto improvable by a two-market segmentation is non-generic.*

To get a sense for Theorem 1, let us apply it to an environment with two types. With two types, the set of markets $\Delta(T)$ is the one-dimensional simplex, that is, the straight line in \mathbb{R}^2 that connects the points $(0, 1)$ and $(1, 0)$. A hyperplane through the origin is a straight line through the point $(0, 0)$. Therefore, for any non-zero vector b in \mathbb{R}^2 the set $H(b)$ is either empty or is a singleton. Theorem 1 then shows the following result, which is a slightly weaker version of Proposition 2.

Corollary 1 *In any environment with two types, all but a finite number of inefficient markets are Pareto improvable by a two-market segmentation.*

Figure 2 illustrates this result for the two-alternative example in the introduction.

The proof of Theorem 1 relies on a new, two-step approach. The first step is to construct, for any inefficient market f , a Pareto dominating market f' whose support is a subset of the support of f . This is achieved by understanding what makes inefficient mechanisms optimal, and is the key to Theorem 1. The second step shows that except for a non-generic set of markets, perturbing a market slightly does not change its optimal mechanism. Combining the two steps leads to Theorem 1: segment market f by assigning probability ε to the Pareto dominating market f' and probability $1 - \varepsilon$ to the remaining market f'' , so that $f = \varepsilon f' + (1 - \varepsilon) f''$. If ε is small, then f'' is a small perturbation of f ; thus, as long as market f does not belong to the set of non-generic markets, market f'' weakly Pareto dominates market f . Therefore, the segmentation of market f into f' and f'' is Pareto improving. We now describe the approach in greater detail.

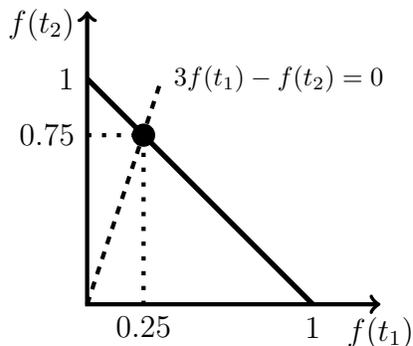


Figure 2: In the example from the introduction, the set of inefficient markets for which no Pareto improving segmentation exists, is the intersection of two sets, each satisfying a linear equality: the set of markets $f(t_1) + f(t_2) = 1$, and the hyperplane $3f(t_1) - f(t_2) = 0$.

5.1 Step 1 - constructing a Pareto dominating segment

The first step is formalized as follows.

Proposition 3 *For any environment and any market f , there exists an efficient Pareto dominating market whose support is a subset of the support of f if and only if f is inefficient.*

5.1.1 Simple environments

Before proving Proposition 3 for general environments, let us consider the relatively easy case of *simple environments*, in which (1) there exists a “lowest type” t_L such that $v(t, a) > v(t_L, a)$ for every other type $t \neq t_L$ and every non-trivial alternative $a \neq 0$, and (2) there is a “best alternative” $\bar{a} \neq 0$ such that $v(t, \bar{a}) > v(t, a)$ for every type t and every other alternative $a \neq \bar{a}$. Note that (2) implies that the efficient allocation assigns alternative \bar{a} to every type. We have the following result.

Lemma 2 *For any simple environment and any inefficient market, (1) the optimal mechanism gives type t_L surplus 0, (2) the optimal mechanism assigns alternative \bar{a} to some type $t_{\bar{a}}$, and (3) type $t_{\bar{a}}$ pays strictly more than $v(t_L, \bar{a})$.*

Proof. Suppose that the optimal mechanism gives type t_L a strictly positive surplus. Then every type has a strictly positive surplus, since every type t can report he is type t_L and obtain surplus $v(t, x(t_L)) - p(t_L) \geq v(t_L, x(t_L)) - p(t_L) > 0$. This contradicts optimality

since a mechanism in which all payments are uniformly and slightly increased is IC-IR, proving (1).

For (2), suppose that the optimal mechanism M does not assign alternative \bar{a} to any type. For every type, consider assigning \bar{a} for a price that gives that type the same surplus that he has in mechanism M . This price is strictly higher than what is specified by the mechanism, since the mechanism is IR and \bar{a} is the best alternative. Let type $t_{\bar{a}}$ be the type with the highest such price p . Modify mechanism M by assigning type $t_{\bar{a}}$ alternative \bar{a} for a price of p . This strictly increases the revenue from the mechanism, a contradiction.

For (3), suppose first that mechanism M assigns type $t_{\bar{a}}$ alternative \bar{a} for a price strictly lower than $v(t_L, \bar{a})$. Then every type has positive surplus in mechanism M , since every type t can report he is type $t_{\bar{a}}$ and obtain surplus $v(t, x(t_L)) - p(t_L) \geq v(t_L, x(t_L)) - p(t_L) > 0$, which contradicts optimality. Finally, suppose that mechanism M assigns type $t_{\bar{a}}$ alternative \bar{a} for a price of $v(t_L, \bar{a})$. Then the mechanism must be efficient: no type can be charged more than $v(t_L, \bar{a})$ (since every type t can report he is type $t_{\bar{a}}$ and obtain the best alternative for a price of $v(t_L, \bar{a})$), and the only IC-IR mechanism that charges $v(t_L, \bar{a})$ from every type is that one that assigns alternative \bar{a} to every type. ■

Lemma 2 implies Proposition 3 for simple environments. To see this, take an inefficient market and its corresponding types t_L and $t_{\bar{a}}$, and consider a two-type market that consists of a fraction ε of type t_L and a fraction $1 - \varepsilon$ of type $t_{\bar{a}}$. If ε is large enough, then the unique optimal mechanism is the efficient one, which assigns alternative \bar{a} to both types for a price of $v(t_L, \bar{a})$. This is because the reduction in the price charged from type t_L when assigning him an alternative $a \neq \bar{a}$ instead is at least $v(t_L, \bar{a}) - v(t_L, a) > 0$, whereas the most that can be charged from type $t_{\bar{a}}$ is $v(t_{\bar{a}}, \bar{a})$, and $\varepsilon(v(t_L, \bar{a}) - v(t_L, a)) > (1 - \varepsilon)v(t_{\bar{a}}, \bar{a})$ for ε sufficiently close to 1. Among the mechanisms that assign alternative \bar{a} to type t_L the efficient one is the one with the highest revenue. This two-type market Pareto dominates the original market: By Lemma 2 type t_L has surplus 0 in both markets, but type $t_{\bar{a}}$ pays more than $v(t_L, \bar{a})$ in the original market.

5.1.2 General environments, an example

A two-type Pareto dominating market does not always exist for inefficient markets in environments that are not simple. We show this in an environment with three types and two alternatives, illustrated Figure 3, (a). Alternatives are a_1 and a_2 , and types are $t_1, t_2,$

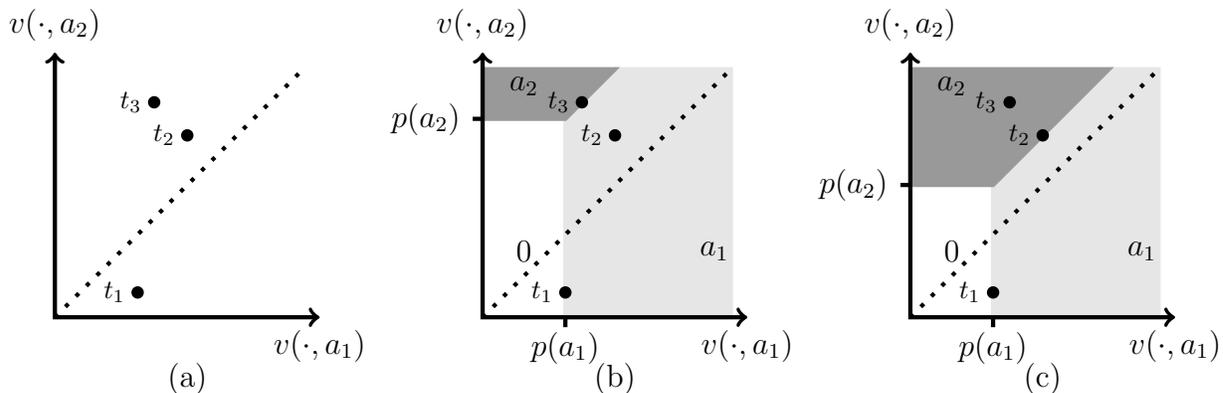


Figure 3: (a) An environment with three types and two alternatives. The environment is not simple since types have different efficient alternatives. Panel (b) depicts the optimal prices $p(a_1), p(a_2)$ for market f , and panel (c) depicts the optimal prices for Pareto dominating market f' . In each panel, types in the lightly shaded region choose alternative a_1 , in the darkly shaded region choose alternative a_2 , and in the unshaded region choose alternative 0.

and t_3 . Each type is described by the dot with the type's label to its left. The horizontal axis shows the valuation for alternative a_1 , and the vertical axis shows the valuation for alternative a_2 . This environment is not simple: there is a lowest type (type t_1) but not a best alternative (type t_1 prefers alternative a_1 but types 2 and 3 prefer alternative a_2).

Figure 3, (b) depicts a mechanism in which the two alternatives are offered at prices $p(a_1)$ and $p(a_2)$. The prices are $p(a_1) = v(t_1, a_1)$ and $p(a_2) = v(t_3, a_2) - v(t_3, a_1) + v(t_1, a_1)$. At these prices, the lightly shaded region contains the set of types that prefer alternative a_1 , the darkly shaded region contains the set of types that prefer alternative a_2 , and the unshaded region contains the set of types that prefer alternative 0. In particular, type t_1 is indifferent between alternative a_1 and the outside option (and chooses a_1), and type t_3 is indifferent between a_1 and a_2 (and chooses a_2). Type t_2 has a higher marginal valuation for a_1 compared to a_2 than does type t_3 , and therefore type t_2 chooses alternative a_1 . This mechanism gives surplus 0 to type t_1 and strictly positive surplus to the other types. The mechanism is inefficient, since type t_2 is assigned alternative a_1 . There exists a market f for which this mechanism is optimal. In such a market, the fraction of type t_1 is large enough that it is optimal to assign him his preferred alternative a_1 for a price that is equal to his valuation; among the remaining consumers the fraction of type t_3 is large enough that it is optimal to assign him his preferred alternative a_2 for the maximal price that maintain IC.

There is a market f' that contains all three types that Pareto dominates f . In this market the fraction of type t_1 is large enough that it is optimal to assign him his preferred alternative a_1 for a price that is equal to his valuation; among the remaining consumers the fraction of type t_2 is large enough that it is optimal to assign him his preferred alternative a_2 for the maximal price that maintains IC. Type t_3 is also assigned alternative a_2 for this price. This mechanism is illustrated in Figure 3, (b). Since the price of alternative a_2 is lower than in the optimal mechanism for market f , market f' Pareto dominates market f .

There is, however, no two-type market that Pareto dominates f : in any market without type t_1 the surplus of one of the other types is 0 (any optimal mechanism gives surplus 0 to some type); in any market without type t_2 either the allocation and surpluses of the other types is unchanged or the surplus of type t_3 is 0; in any market without type t_3 either the surplus of type t_1 is 0 and the surplus of type t_2 is strictly lowered (he is allocated alternative a_2 for the maximal price that maintains IC) or the surplus of type t_2 is 0.

The reason that all three types are needed to form a Pareto dominating market is that in order to increase the surplus of type t_3 (who is already assigned his efficient alternative a_2 in market f), type t_2 must be present in sufficient proportion to make it optimal to lower the price of alternative a_2 in order to extract more surplus from type t_2 . But type t_2 's surplus in market f is positive; in order to maintain this surplus in the Pareto dominating market, type t_1 must be present in sufficient proportion to make it optimal for the seller to assign alternative a_1 to type t_1 , thereby providing information rents to type t_2 .

5.1.3 Proof of Proposition 3

The proof of Proposition 3 generalizes the idea from the construction discussed above. In an inefficient market f , some type t is assigned an inefficient alternative. The reason for this inefficiency is to lower the surplus (information rents) of some other type t' . In a new market that includes only type t and t' and in which the proportion of type t is much higher than that of type t' it is optimal to assign type t his efficient alternative; this increases the marginal surplus that type t' obtains from being able to mimic type t . But the surplus of type t may decrease; to prevent this, we identify an “information rents path” in market f that begins with type t and ends with some type t'' that has surplus 0, and add to the new market all the types in the path in the correct proportions. This generates a market that Pareto dominates market f . We now describe this procedure in more detail.

Take an inefficient market f that (without loss of generality) has full support, and let t be some type that is assigned an inefficient alternative in the optimal mechanism. We will inductively construct a set of types S that contains i such that for every type t' in S there is a directed path of types in S from type t to type t' in which the IC constraint from each type t_j to the next type t_{j+1} in the path binds (type t_j is indifferent between reporting truthfully and misreporting that he is type t_{j+1}). The construction stops when a type that has surplus 0 is added to S . If type t has surplus 0 we are done. Otherwise, given the set S so far constructed, there is a type t' not in S such that the IC constraint from some type in S to type t' binds. Otherwise the revenue can be increased by increasing the payments of all types in S by the same small amount. This concludes the construction of S .

Consider the binding IC path in S that begins with type t and ends with the type that has surplus 0. Suppose without loss of generality that type t is the only type in the path that is assigned an inefficient alternative (otherwise denote by t the last type in the path that is assigned an inefficient alternative). Denote by T the set of types in the path, and notice that the payments of the types in T weakly decrease along the path (otherwise the revenue in market f can be increased by replacing a type's assigned alternative and payment with those of the next type in the path, without violating incentive compatibility).

Now, modify the optimal mechanism M for f by assigning type t its efficient alternative and increasing his payment to leave his surplus unchanged. The modified mechanism M^1 violates IC, otherwise mechanism M^1 would generate more revenue than mechanism M in market f . Therefore, when faced with the modified mechanism some type $t' \neq t$ strictly prefers to misreport that he is type t . Type t' is not in T , since every type in T other than type t is assigned his efficient alternative and is paying less than type t does in M (and therefore in M^1). Modify mechanism M^1 by replacing the assigned alternative and payment of type t' with those of t . This modified mechanism M^2 satisfies IC and IR for the set of types $T \cup \{t'\}$, and type t' has strictly higher surplus than in mechanism M . Finally, if t' is not assigned the efficient alternative in mechanism M^2 , modify M^2 by assigning type t' his efficient alternative and increasing his payment to leave his surplus unchanged. Denote the resulting mechanism by M^* . Notice that in mechanism M^* every type in T is assigned his efficient alternative and pays less than does t' , and thus would not benefit from misreporting that he is type t' . Consider the restricted environment with types $T \cup \{t'\}$. Mechanism M^* is efficient and IC-IR in this environment. Further, the surplus of every

type in $T \cup \{t'\}$ is weakly higher than in mechanism M , and the surplus of type t' is strictly higher.

It remains to show that M^* is the (unique) optimal mechanism for some full-support market in the restricted environment. We only provide the intuition here and defer the formal proof. A market for which mechanism M^* is optimal exists because M^* is efficient. Such a market can be constructed iteratively. Take the path that defined T and add type t' to its beginning (so type t follows type t'). Begin with a large enough fraction of the last type in the path so that it is strictly optimal for the seller to assign this type his efficient alternative for a price that is equal to his valuation. Add a large enough fraction of the second-to-last type in the path so that it is strictly optimal for the seller to assign this type his efficient alternative for the maximal price that maintains IC, etc. The resulting mechanism is M^* , which is the unique optimal mechanism for the resulting market.¹⁵ This proves Proposition 3.

Let us apply the procedure to the example in Section 5.1.2. Consider the optimal mechanism for market f shown in Figure 3, (b). Figure 4, (a) illustrates the binding IC and IR constraints in the optimal mechanism and the allocation of each type in the optimal mechanism. Beginning with type $t = t_2$, we have $S = \{t_1, t_2\}$. The binding IC path in S that begins with type t_2 and ends with type t_1 , whose surplus is 0, is depicted in Figure 4, (b). The appended path with type $t' = t_3$ at its beginning and the modifications to the optimal mechanism for the original market are illustrated Figure 4, (c). The allocation in the resulting mechanism is the one in Figure 3, (c). The payments in the resulting mechanism can be obtained via the allocation and the binding IC constraints.

5.2 Step 2 - perturbing a generic market leaves the optimal mechanism unchanged

In this section we show that for a generic set of markets, the optimal mechanism remains optimal if the market is slightly perturbed. The intuition is as follows. Consider the set P

¹⁵The construction uses the assumption that each type has a unique efficient alternative. If types have multiple efficient alternatives, it is still true that there exists a market for which it is optimal to assign each type an efficient alternative. However, that alternative may be different than the efficient alternative prescribed by mechanism M^* . Thus if there are multiple efficient alternatives, a more complicated construction may be required that allows us to argue about which efficient alternative is selected.

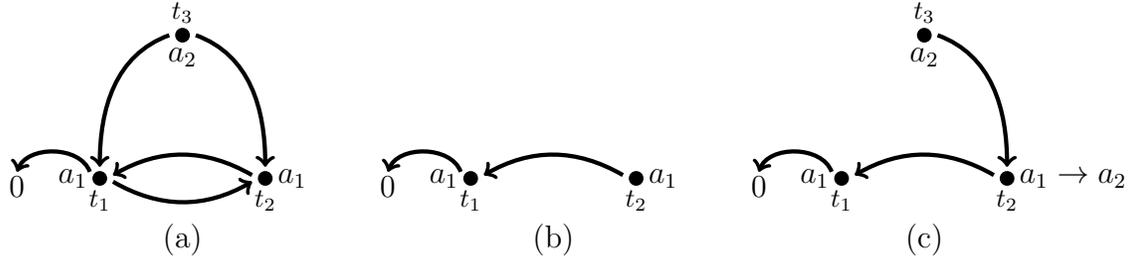


Figure 4: The execution of the procedure that constructs a Pareto dominating market. (a) Binding IC and IR constraints for market f . (b) The path of binding IC constraints that starts from a type with inefficient allocation, type t_2 , and ends with a type with binding IR constraint, type t_1 . (c) The appended path with type t_3 who strictly benefits in the Pareto dominating market.

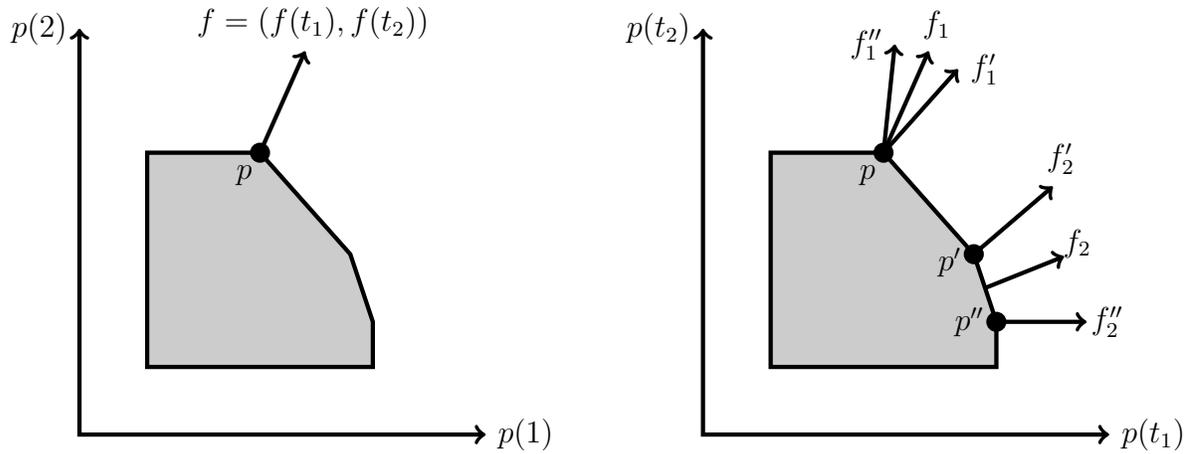


Figure 5: An environment with two types. (a) Each market f can be seen as a vector $(f(t_1), f(t_2))$ in \mathbb{R}^2 . The optimal payment rule for market f is maximal in the direction specified by f . (b) f_1 is not orthogonal to a face of P . Any small enough perturbation of f_1 leaves the optimal mechanism unchanged. f_2 is orthogonal to a face of P . Different mechanisms are optimal for different perturbations of f_2 .

of all payment rules $p : T \rightarrow \mathbb{R}$ that are part of some IC-IR mechanism. A mechanism is optimal for a market f if its payment rule p is maximal in P in the direction specified by f , as depicted in Figure 5, (a). We show that P is a polytope. If a market is not orthogonal to a face of P , then the payment rule that maximizes the expected revenue is unique and is a vertex of P , as shown in Figure 5, (b). The same payment rule remains optimal for small enough perturbations of such markets. In contrast, if f is orthogonal to a face of P , then any small perturbation of the market may result in a change in the optimal mechanism. We formalize this discussion below.

The main result of this section states that, generically, perturbing a market leaves the

optimal mechanism unchanged. Formally, markets f and f' are ε -close if $|f(t) - f'(t)| \leq \varepsilon$ for all types t . We say that perturbing market f leaves the optimal mechanism unchanged if there exists small enough ε such that the optimal mechanism for f is also optimal for all markets f' that are ε -close to f . Let F_P denote the set of markets f such that perturbing f leaves the optimal mechanism unchanged. The proposition below shows that the complement of F_P is non-generic (see Definition 1).

Proposition 4 *The set of markets $\Delta(T) \setminus F_P$ is non-generic.*

To prove the proposition, we first notice that the set of IC-IR mechanisms is a polytope in $\mathbb{R}^{(k+2)n}$, where k is the number of alternatives and n is the number of types. Indeed, a mechanism is a point in $\mathbb{R}^{(k+2)n}$ (for each of the n types it specifies the payment and the probability of being assigned each one of the k alternatives and the outside option), each of the finite number of incentive and individual rationality constraints corresponds to a half space, and the (linear) probability constraints together with the IR constraints guarantee that the set is bounded. The set P of payment rules that are part of some IC-IR mechanism is a projection of the set of IC-IR mechanisms, and is therefore a polytope in \mathbb{R}^n . Consequently, there set P has a finite set of vertices.

Lemma 3 *There exists a finite set $P_V \subseteq R^{|T|}$ such that P is the convex hull of P_V .*

We prove Proposition 4 in two steps. The first is to show that if a market f has a unique optimal payment rule, then perturbing f leaves the optimal mechanism unchanged, that is, f is in F_P . We say that a market f has a unique optimal payment rule if $p = p'$ for any two optimal mechanisms (x, p) and (x', p') of f .

Lemma 4 *It a market f has a unique optimal payment rule, then f is in F_P .*

Proof. Consider a market f with a unique optimal payment rule p . Since the set P_V is finite (by Lemma 3) and p is the unique optimal payment rule, there exists $\delta > 0$ such that $E_{t \sim f}[p(t)] \geq E_{t \sim f}[p'(t)] + \delta$ for all $p' \in P_V$. By continuity of the expected revenue in f , there exists $\varepsilon > 0$ such that $E_{t \sim f}[p(t)] \geq E_{t \sim f'}[p'(t)]$ for all $p' \in P_V$ and all f' that are ε -close to f . Since all payment rules are convex combinations of the payment rules in P_V (Lemma 3), we must have $E_{t \sim f}[p(t)] \geq E_{t \sim f'}[p'(t)]$ for all payment rules $p' \in P$, that is, the payment rule p is also optimal for all f' that are ε -close to f . ■

The second step in the proof of Proposition 4 is to show that almost all markets have unique optimal payment rules.

Lemma 5 *The set of markets without a unique optimal payment rule is non-generic.*

Proof. Consider a market f for which more than one payment rule in P maximizes revenue. Since P is the convex hull of P_V , there must exist two non-identical payment rules p and p' in P_V that are both optimal for f . Thus such a market is contained in a hyperplane $H_{p,p'}$ defined by the equation $\sum_t f(t)(p(t) - p'(t)) = 0$. Since P_V is finite, the set of markets without a unique optimal payment rule is contained in a finite union of hyperplanes, one for each pair of payment rules in P_V , i.e., $\cup_{p,p' \in P_V} H_{p,p'}$. Thus, by Definition 1, the set of markets without a unique optimal payment rule is non-generic. ■

To complete the proof of Proposition 4, note that by Lemma 4, the set F_P contains all markets f with a unique optimal payment rule. By Lemma 5, the set of markets without a unique optimal payment rule is non-generic. Therefore, $\Delta(T) \setminus F_P$ is generic.

6 Relaxing the assumption of finite alternatives

So far we have assumed that there are finitely many types and alternatives. The interpretation is that there are substantially more individuals than tastes or products. For instance, in a bundling application where each alternative represents a subset of indivisible products, there are finitely many alternatives. In this section we briefly discuss environments with possibly infinitely many alternatives.¹⁶

Theorem 1, which states that all inefficient markets in a generic set are Pareto improvable may fail with a continuum of alternatives. We can illustrate this in an environment with two types. Suppose that a product can be produced with a range of qualities $a \in [0, 1]$, where the cost of producing quality a is $C(a) = a^2/2$. Types t_1 and t_2 have valuations a and $2a$ for quality a , respectively. A market is identified by the proportion q of type t_2 . In any market, the surplus of type t_1 in the optimal mechanism is zero. The surplus

¹⁶We do not explore the case of infinitely many types. With infinitely many types, it is unclear what the appropriate generalization of our notion of genericity is. Recall that we defined genericity by viewing the set of all markets as a subset of a finite dimensional space, $\Delta(T) \subset \mathbb{R}^{|T|}$. We then defined a set of markets to be non-generic if it is contained in a lower dimensional space. This is no longer feasible if there are infinitely many types. We leave appropriate generalization of our result to infinite types to future work.

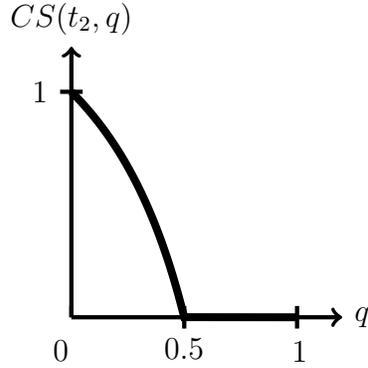


Figure 6: Surplus of type t_2 as a function of the fraction $q = f(t_2)$ of consumers of type t_2 .

of type t_2 in the optimal mechanism is illustrated in Figure 6. It strictly decreases in q if q is smaller than 0.5, and remains constant at 0 if q is larger than 0.5. Any non-trivial segmentation of a market $q < 0.5$ must include a segment $q' > q$ in which type t_2 is strictly worse off. Therefore, a market $q < 0.5$ is not Pareto improvable. The set of inefficient markets that are not Pareto improvable, $(0, 0.5)$, is not finite, unlike two-types environments with finitely many alternatives (Proposition 2). Going back to our two-type analysis in Section 4, however, this should not be surprising. With finitely many alternatives, as the number of alternatives grow, the intervals over which the surplus of the high type is constant shrink, and their number grows. It should thus be expected that “in the limit,” the number of intervals and thus the number of markets for which no Pareto improving segmentations exist, becomes infinite as well.

While Theorem 1 no longer holds, an appropriate reformulation of our analysis holds, and even becomes stronger, for the two-type example discussed above. In particular, recall that with finitely many alternatives, a Pareto improving segmentation exists for an inefficient market if perturbing the market leaves the optimal mechanism unchanged. In the two-type example, a Pareto improving segmentation exists for an inefficient market if, and only if, perturbing the market in a certain way leaves the optimal mechanism unchanged. Indeed, notice that for any market in $[0.5, 1)$, increasing the proportion of type t_2 leaves the optimal mechanism unchanged. Therefore, any such market q can be segmented in a Pareto improving way into two segments: a segment in which the fraction of type t_2 is slightly higher than q , and another segment in which the fraction of type t_2 is less than 0.5. Since we already argued that no Pareto improving segmentation exists for inefficient markets $(0, 0.5)$, we conclude that a Pareto improving segmentation exists

for an inefficient market q if and only if $q \in [0.5, 1)$, that is, increasing the proportion of type t_2 leaves the optimal mechanism unchanged. Notice that an arbitrary perturbation of market 0.5 may change the optimal mechanism. Nevertheless, increasing the proportion of type t_2 does not change the optimal mechanism, and a Pareto improving segmentation exists. We generalize this analysis in the following subsection to obtain a characterization of Pareto improvable markets.

Even though Theorem 1 may fail in some environments with infinitely many alternatives, it holds in certain other environments with infinitely many alternatives. In fact, notice that in our setting with finitely many alternatives, there are infinitely many randomized alternatives. Therefore, Theorem 1 holds with infinitely many random alternatives. What is the difference between environments with randomizations over finite alternatives and the two-type example discussed above? The distinction can be clarified by considering the set P of payment rules that are a part of some IC-IR mechanism. With randomization over finite alternatives, the set P is a polytope defined by finitely many halfspaces, as illustrated in Figure 5. Such a set P has a finite number of vertices, and therefore generically, perturbing a market f leaves the optimal mechanism unchanged. On the other hand, consider the set P of payment rules (net of the cost of production) that are part of some IC-IR mechanism for the two-type example discussed above, illustrated in Figure 7. Even though P is convex, it has infinitely many vertices. Perturbing any market in which the proportion of type t_2 is less than 0.5 changes the optimal mechanism. If, instead of assuming that the cost function is $C(a) = a^2/2$, we assume that the cost function is convex but piecewise linear, then Theorem 1 can be recovered. For such a cost function, the set P has finitely many vertices. Therefore, generically, perturbing a market leaves the optimal mechanism unchanged.

6.1 Characterizing Pareto improvable markets in environments with linear valuations

In this section we characterize Pareto improvable markets in a setting with linear valuations, as in Mussa and Rosen (1978). There is a finite set $T = \{1, \dots, n\}$ of types and a compact set $A \subseteq \mathbb{R}_{\geq 0}$ of alternatives. The set of alternatives may be finite or infinite. The valuation of type t for alternative a is $v(t, a) = v(t)a$, where the types are ranked such that $v(t) \in \mathbb{R}$ is increasing in t . The cost of producing an alternative a is $C(a)$. Assume that

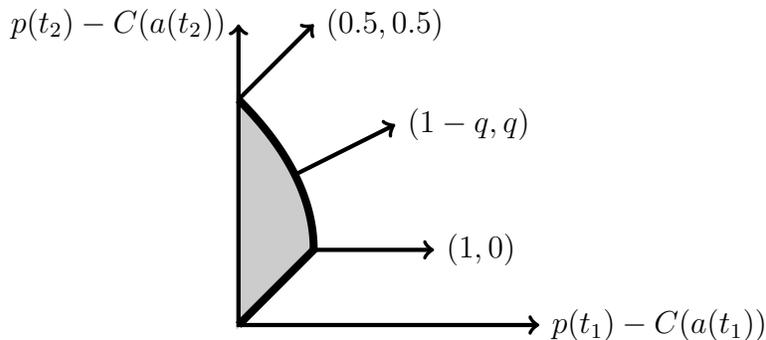


Figure 7: The profit obtained from each type in IC-IR mechanisms. The Pareto frontier is obtained from assigning alternative a to type t_1 and alternative 1 to type t_2 , which gives profit $a - a^2/2$ from type t_1 and $1.5 - a$ from type t_2 , respectively.

each type t has a unique efficient alternative $\bar{a}(t)$, i.e., one that maximizes $v(t)a - C(a)$ over all alternatives a .

We show that a market is Pareto improvable if and only if certain perturbations of the market leave the optimal mechanism unchanged. We start by defining the relevant perturbations. For a type t and a market f let the cumulative fraction $\tilde{F}(t) = \sum_{t' \geq t} f(t')$ be the fraction of consumers of types t or higher. The relevant perturbations for our characterization are ones in which the cumulative fractions of all types higher than some type t are changed proportionally, and the cumulative fractions of types t or below are unchanged.

Definition 2 A market f^1 with cumulative fractions \tilde{F}^1 is a (t, δ) -perturbation of a market f with cumulative fractions \tilde{F} if $\tilde{F}^1(t') = (1 + \delta)\tilde{F}(t')$ for all $t' > t$, and $\tilde{F}^1(t') = \tilde{F}(t')$ for all $t' \leq t$.

Our characterization states that a market f is Pareto improvable if and only if for some type t with inefficient allocation, any small enough (t, δ) -perturbation of f leaves the optimal mechanism unchanged.

Proposition 5 A market f is Pareto improvable if and only if there exists a type t with inefficient allocation and $\delta_0 > 0$ such that every (t, δ) -perturbation of f leaves the optimal mechanism unchanged for $\delta \leq \delta_0$.

The proof of Proposition 5 relies on the characterization of optimal mechanisms via virtual valuations, as in Myerson (1981). In a market f , each type t can be assigned a virtual

valuation $\phi(t)$ such that the allocation of type t in the optimal mechanism maximizes the virtual surplus $\phi(t)a - C(a)$. A key component of our proof is to notice how the perturbations defined in Definition 2 affect virtual valuations. Consider a (t, δ) -permutation with $\delta < 0$, i.e., a permutation in which the cumulative fraction of types above t is reduced proportionally. Assuming a regularity condition which we formalize later, such a permutation increases the virtual valuation of type t without changing the virtual valuation of any other type (if $\delta > 0$ the virtual valuation of type t decreases). Without regularity, such a permutation increases the virtual valuation of type t without decreasing the virtual valuation of any other type. Given this observation, we now outline the proof of Proposition 5.

First suppose that there exists a type t with inefficient allocation such that every small enough (t, δ) -perturbation of f leaves the optimal mechanism unchanged. Similar to Theorem 1, we show that f is Pareto improvable by segmenting it into two markets f' and f'' , where f' Pareto dominates f and f'' weakly Pareto dominates f . Both f' and f'' are perturbations of f of the form defined in Definition 2. In particular, f' is a (t, δ') -perturbation of f with $\delta' < 0$ and f'' is a (t, δ'') -perturbation of f with $\delta'' > 0$. Market f' Pareto dominates market f since as argued above, the virtual valuations of all types is weakly higher (and virtual valuation of i is strictly higher), and therefore the allocation of all types is more efficient, in f' than in f . By setting the probability of f'' in the segmentation to be large enough, we can ensure that δ'' is small enough, and therefore the optimal mechanism for f'' is the same as the optimal mechanism for f by assumption. Therefore, f'' weakly Pareto dominates f , and the segmentation is Pareto improving.

Now suppose that market f is such that for any type t with inefficient allocation, any (t, δ) -perturbation of f changes the optimal mechanism. Assume for contradiction that a Pareto improving segmentation of f exists. Then there must exist a type t whose allocation is weakly more efficient in every segment, and strictly more efficient in some segment. Since any (t, δ) -perturbation of f changes the optimal mechanism, the virtual valuation of type t must be weakly higher in every segment, and strictly higher in some segment, than in f . A simple accounting analysis shows that it is impossible to increase the virtual valuation of some type weakly in all segments, and strictly in some segment. The analysis uses the fact that the average of all segments must be equal to the original market.

To prove Proposition 5 formally, we utilize the standard characterization of optimal mechanisms. The characterization is based on appropriately defined virtual valuations as

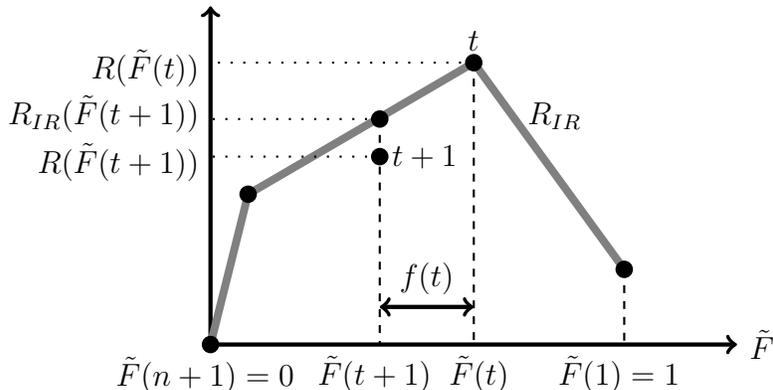


Figure 8: $\phi(t)$ is the slope of the concave hull of the revenue function to the left of t .

slopes of appropriately defined revenue functions. Define $\tilde{F}(n+1) = 0$. Define the revenue function $R : \{\tilde{F}(1), \dots, \tilde{F}(n+1)\} \rightarrow \mathbb{R}$ as follows. For each type t , $R(\tilde{F}(t)) = v(t)\tilde{F}(t)$ is the expected revenue from selling alternative $a = 1$ at price $v(t)$, and $R(\tilde{F}(n+1)) = 0$, as depicted in Figure 8. Notice that the domain of R is a subset of $[0, 1]$. Notice also that since a larger type t has a smaller cumulative fraction $\tilde{F}(t)$, type t is to the left of all smaller types in the graph of R . Let the ironed revenue function $R_{IR} : [0, 1] \rightarrow \mathbb{R}$ be the concavification of the revenue function, i.e., the lowest concave function with domain $[0, 1]$ that is weakly higher than R . The virtual valuation $\phi(t)$ of type t is

$$\phi(t) = \frac{R_{IR}(\tilde{F}(t)) - R_{IR}(\tilde{F}(t+1))}{f(t)}. \quad (1)$$

That is, the virtual valuation is the slope of the ironed revenue function to the left of t .¹⁷

The following proposition characterizes optimal mechanisms. A mechanism is optimal if it satisfies two properties. First, the allocation rule maximizes virtual surplus. That is, $a(t)$ maximizes $\phi(t)a - C(a)$ over all a . Second, the surplus $CS(t, f)$ of each type t is the area under the allocation rule of types lower than t , and the payment rule of the mechanism is defined to give each type the appropriate surplus. In particular, the second property implies that increasing the allocation of a type increases the surplus of all higher types.

¹⁷Equation (1) is well defined only if $f(t) > 0$. We allow for segments in which the proportion of consumers of some type is zero. Thus it is important to our analysis to extend the definition of virtual valuations to allow for zero fractions. We do so in our formal proof of Proposition 5 in the appendix.

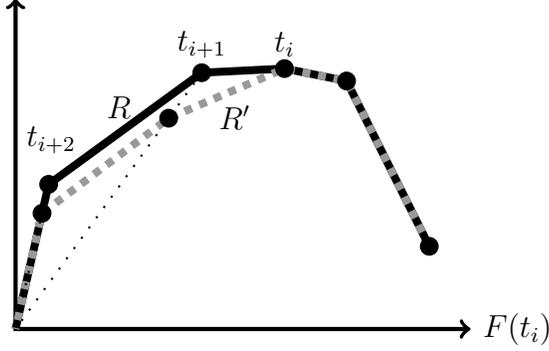


Figure 9: By lowering the fractions of types above t_i proportionally, the virtual valuation of all types weakly increase, while the virtual valuation of type t strictly decreases. Since the cumulative fractions of types t_{i+1} and t_{i+2} are reduced proportionally, the slope of the revenue function between them remains the same.

Proposition 6 *A mechanism (a, p) is optimal for market f if and only if the following holds for all t ,*

1. $a(t)$ maximizes $\phi(t)a - C(a)$ over all a , and
2. $p(t) = v(t)a(t) - CS(t, f)$ where $CS(t, f) = \sum_{1 < t' \leq t} (v(t') - v(t' - 1))a(t' - 1)$.

We now describe the proof of Proposition 5 in more detail. Suppose that there exists a type t with inefficient allocation such that a small enough (t, δ) -perturbation of f leaves the optimal mechanism unchanged. As outlined above, to argue that f is Pareto improvable, we show that there exists a (t, δ) -perturbation f' of market f that Pareto dominates f . The lemma below establishes this claim. The idea is as follows. Lowering δ decreases the cumulative fractions of types above t . Thus the revenue decreases for any type above t , and remains constant for type t and below. If f is regular, that is if the revenue function is concave, the virtual valuation of each type other than t remains unchanged, and the virtual valuation of t increases, as depicted in Figure 9. If f is not regular, the virtual valuation of types other than t weakly increases, and the virtual valuation of t strictly increases. Therefore, the optimal allocation in market f' is more efficient than the optimal allocation in market f . Property 2 of Proposition 6 then implies that all types above t have strictly higher surplus in f' than in f .

Lemma 6 *Consider a market f in which the allocation of some type t is inefficient. There exists δ such that a (t, δ) -permutation of market f Pareto dominates f .*

We now turn to proving the necessity of the condition of Proposition 5 for existence of Pareto improving segmentations. We here only provide the proof assuming that market f and all segments f' are regular. The complete proof is based on a careful generalization of the argument here and is deferred to the appendix.

Consider a market f such that for all types t with inefficient allocation and for all $\delta_0 > 0$, there exists $\delta \leq \delta_0$ such that a (t, δ) -permutation of the market changes the optimal mechanism. Assume for contradiction that there exists a Pareto improving segmentation μ of f . The allocation of some type in some segment must be more efficient than the allocation of that type in market f . Let t be the lowest such type. We show that the virtual valuation of type t must be weakly higher in every segment than in the original market, and strictly higher in some segment. We then show this contradicts the requirement that the expectation of all segment equals the original market.

The virtual valuation of type t must be weakly higher in every segment f' than in f . Otherwise, since the allocation of all type below t is the same in f and f' , the surplus of type $t + 1$ is lower in f' than in f by property 2 of Proposition 6. By the assumption that a (t, δ) -perturbation of f changes the optimal mechanism, the virtual valuation of type t must be weakly higher in every segment f' than in f . Since the allocation of type t is strictly more efficient in some segment, its virtual valuation is higher in that segment than in f . We conclude that the virtual valuation $\phi'(t)$ of type t in every segment f' satisfies $\phi'(t) \geq \phi(t)$, with strict inequality for some segment. Multiplying both sides by $f'(t)$ and taking the expectation over all segments, we have $E_\mu[\phi'(t)f'(t)] > E_\mu[\phi(t)f'(t)]$. We show that this is impossible.

For a regular market f , the revenue function R is equal to its concavification R_{IR} . Thus the virtual valuation of type t is the slope of the revenue function to the left of t ,

$$\phi(t) = \frac{R(\tilde{F}(t)) - R(\tilde{F}(t+1))}{f(t)}.$$

Rearranging and substituting the definition of the revenue function, we have

$$\phi(t)f(t) = v(t)\tilde{F}(t) - v(t+1)\tilde{F}(t+1).$$

In a segmentation μ , we have $\tilde{F}(t) = E_\mu[\tilde{F}'(t)]$ for all types t . Therefore,

$$\begin{aligned}\phi(t)f(t) &= E_\mu[v(t)\tilde{F}'(t) - v(t+1)\tilde{F}'(t+1)] \\ &= E_\mu[\phi'(t)f'(t)] \\ &> E_\mu[\phi(t)f'(t)] = \phi(t)E_\mu[f'(t)] = \phi(t)f(t).\end{aligned}$$

We have shown that $\phi(t)f(t) > \phi(t)f(t)$, which is a contradiction.

Proposition 5 does not directly identify conditions under which a (t, δ) -permutation of a market leaves the optimal mechanism unchanged. The conditions can indeed be specified via virtual valuations. In particular, recall that a (t, δ) -permutation of a market f decreases the virtual valuation of type t for $\delta > 0$. For such a permutation to leave the optimal mechanism unchanged, it must be that the allocation that maximizes $\phi(t)a - C(a)$ also maximizes $(\phi(t) - \epsilon)a - C(a)$ for every small ϵ . To formalize, consider $\phi \in \mathbb{R}$ and allocation a that maximizes $\phi a - C(a)$. Let $\Phi(\phi) = \{\phi' \in \mathbb{R} \mid a \in \arg \max_{a'} \phi' a' - C(a')\}$ denote the set of virtual valuations that have the same virtual surplus maximizer as ϕ . The set $\Phi(\phi)$ is a closed interval in \mathbb{R} (possibly a singleton). If decreasing $\phi(t)$ does not change the optimal allocation of type t , then $\phi(t)$ must be strictly higher than the smallest value in the interval $\Phi(\phi(t))$. We have the following corollary of Proposition 5.

Corollary 2 *A market f is Pareto improvable if and only if there exists a type t with inefficient allocation such that $\phi(t) > \min(\Phi(\phi(t)))$.*

We have the following two corollaries of Corollary 2. First, if $\Phi(\phi) = \{\phi\}$, then $\phi \not> \min(\Phi(\phi))$. Therefore, if the virtual valuation of a type with inefficient allocation is equal to ϕ , then a (t, δ) -permutation cannot be used to construct a Pareto improving segmentation. For instance, if C is strictly convex and differentiable, then any $a > 0$ is a maximizer of virtual surplus for a unique virtual valuations. Thus, if the optimal allocation of some type t is $a(t) > 0$, then $\phi(t) \not> \min(\Phi(\phi(t)))$. On the other hand, if $a(t) = 0$, then lowering the virtual valuation of type t does not change its optimal allocation. We thus have the following corollary.

Corollary 3 *If the cost function C is strictly convex and differentiable, then a market is Pareto improvable if and only if some type is excluded.*

In contrast to Corollary 3, there are environments in which Pareto improving segmentations exist for generic sets of inefficient markets. Notice that if the set A of alternatives

is finite, then the set of all virtual valuations can be partitioned into finitely many intervals such that the optimal allocation for all virtual valuations within an interval is the same. In this case, lowering the virtual valuation of a type leaves its optimal allocation unchanged unless the virtual valuation of the type belongs to the finite set of endpoints of intervals. Thus Corollary 2 generalizes Theorem 1 for the case of linear valuations.

Pareto improving segmentations may exist for generic markets even if the set of alternatives is infinite. In particular, if the cost function is piecewise linear, then similar to the case of finite alternatives, the set of all virtual valuations can be partitioned into finitely many intervals such that the optimal allocation for all virtual valuations within an interval is the same. We thus have the following corollary.

Corollary 4 *If the cost function C is convex and piecewise linear, then the set of inefficient markets that are not Pareto improvable is non-generic.*

7 Discussion and conclusions

This paper studies the existence of Pareto improving segmentations, i.e., segmentations in which each individual has a weakly higher payoff, and some individuals have a strictly higher payoff, than in the unsegmented market. In environments with a finite number of types and products, we show that every inefficient market in a generic set of markets can be segmented in a way that is Pareto improving. In environments with linear valuations, we characterize Pareto improvable markets.

Our work brings together second- and third-degree price discrimination. The literature on third-degree price discrimination assumes that the seller adjusts his selling strategy in different segments by only changing the price of the single product that he is selling. In our setting, the seller may offer different products and product bundles in different segments. Additionally, within each segment, the seller may not only exclude some consumers, but also distort the allocation of some consumers, in order to extract more information rents from other consumers. A main technical difficulty in our setting is that no general characterization is known for the seller's optimal menu with multiple products. We develop a novel methodology to investigate consumer surplus given the optimal menu without requiring a characterization of these menus.

Our results also contribute to discussions about regulating seller's use of information

and ability to price discriminate, and consumer privacy and their control of data. We show that consumers can provide information to the seller in a way that benefits all consumers. In particular, the total gain from the increase in allocative efficiency is larger than the seller's gain from improved price discrimination. We do not impose any constraints on the set of feasible segmentations. In reality, due to technological restrictions, the type of information that can be provided and the seller's ability to price discriminate based on the available information may be limited. Our result establishes that in general there is scope for improving consumer surplus via segmentation. A more specialized study of the limitations in relevant applications is left for future work.

References

- Adams, W. J. and Yellen, J. L. (1976). Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics*, pages 475–498.
- Aguirre, I., Cowan, S., and Vickers, J. (2010). Monopoly price discrimination and demand curvature. *American Economic Review*, 100(4):1601–15.
- Aumann, R. J., Maschler, M., and Stearns, R. E. (1995). *Repeated games with incomplete information*. MIT press.
- Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, 105(3):921–57.
- Braghieri, L. (2017). Targeted advertising and price discrimination in intermediated online markets. *working paper*.
- Chen, X., Diakonikolas, I., Orfanou, A., Paparas, D., Sun, X., and Yannakakis, M. (2015). On the complexity of optimal lottery pricing and randomized mechanisms. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1464–1479. IEEE.
- Cowan, S. (2016). Welfare-increasing third-degree price discrimination. *The RAND Journal of Economics*, 47(2):326–340.
- Crawford, G. S. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, 102(2):643–85.

- Daskalakis, C., Deckelbaum, A., and Tzamos, C. (2014). The complexity of optimal mechanism design. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1302–1318. SIAM.
- Daskalakis, C., Deckelbaum, A., and Tzamos, C. (2017). Strong duality for a multiple-good monopolist. *Econometrica*, 85(3):735–767.
- Glode, V., Opp, C. C., and Zhang, X. (2018). Voluntary disclosure in bilateral transactions. *Journal of Economic Theory*, 175:652–688.
- Hart, S. and Nisan, N. (2013). The menu-size complexity of auctions. *arXiv preprint arXiv:1304.6116*.
- Hidir, S. and Vellodi, N. (2018). Personalization, discrimination and information revelation. *working paper*.
- Ho, K., Ho, J., and Mortimer, J. H. (2012). The use of full-line forcing contracts in the video rental industry. *American Economic Review*, 102(2):686–719.
- Ichihashi, S. (2018). Online privacy and information disclosure by consumers. *working paper*.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.
- Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):pp. 58–73.
- Pigou, A. (1920). *The economics of welfare*. London: Macmillan.
- Robinson, J. (1969). *The economics of imperfect competition*. Springer.
- Salinger, M. A. (1995). A graphical analysis of bundling. *Journal of Business*, pages 85–98.
- Schmalensee, R. (1981). Output and welfare implications of monopolistic third-degree price discrimination. *The American Economic Review*, 71(1):242–247.

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

Thanassoulis, J. (2004). Hagglng over substitutes. *J. Economic Theory*, 117(2):217–245.

Varian, H. R. (1985). Price discrimination and social welfare. *The American Economic Review*, 75(4):870–875.

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

A Appendix

A.1 Two-type analysis and the proof of Lemma 1

We first characterize optimal mechanisms in environments with two types, and then use the analysis to prove Lemma 1.

A.1.1 Characterization of optimal mechanisms for two types

We characterize optimal mechanisms for two types t_1 and t_2 . We consider three cases. (1) $v(t_1, \bar{a}(t_1)) \geq v(t_2, \bar{a}(t_1))$ and $v(t_1, \bar{a}(t_2)) \leq v(t_2, \bar{a}(t_2))$. In this case, extracting the full surplus is feasible since each type can be assigned his efficient alternative at price equal to his willingness to pay, without violating IC or IR.

The remaining two cases concern $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$ (the case where $v(t_1, \bar{a}(t_2)) > v(t_2, \bar{a}(t_2))$ is symmetric). In these cases, it is impossible to extract the full surplus since assigning type t_1 his efficient alternative at price equal to its willingness to pay gives type t_2 strictly positive surplus. In these two cases, type t_1 represents the low type in the sense that its surplus is zero, and type t_2 represents the high type in the sense that its allocation is efficient and it may have positive surplus (the surplus of type t_2 may be zero if type t_1 gets an inefficient alternative). The IR constraint for type t_1 and the IC constraint for type t_2 bind. Given the binding constraints and the allocation of type t_2 , the optimal mechanism is identified once the allocation of type t_1 is specified.

If $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$, we further consider two cases. (2) There exists a maximizer a of $v(t_1, a) - qv(t_2, a)$ such that $v(t_2, a) - v(t_1, a) \geq 0$. In this case, type t_1 is assigned

alternative a and the surplus of type t_2 is $v(t_2, a) - v(t_1, a)$. (3) For any maximizer a of $v(t_1, a) - qv(t_2, a)$, we have $v(t_2, a) - v(t_1, a) < 0$. In this case, for some parameter λ and random assignment $x \in \Delta(\arg \max v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda)$ which we identify below, type t_1 is assigned x . The parameters are chosen such that the surplus of type t_2 is zero, $v(t_2, x) - v(t_1, x) = 0$. The lemma below show that such parameters exist.

Lemma 7 *Assume that $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$ and that for any maximizer a of $v(t_1, a) - qv(t_2, a)$, we have $v(t_2, a) - v(t_1, a) < 0$. There exists λ and $x \in \Delta(\arg \max v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda)$ such that $v(t_2, x) - v(t_1, x) = 0$.*

Proof. Notice that the expression

$$v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda \tag{2}$$

is linear in λ for each alternative a . Therefore, the problem of maximizing Expression (2) over all a is to maximize over all linear functions corresponding to $a \in A$. Thus, the maximum of Expression (2) over all a is the convex envelope of all such functions. The slope of the line corresponding to a is $v(t_1, a) - v(t_2, a)$. Thus, $v(t_1, a) - v(t_2, a)$ is non-decreasing in λ for the maximizer a of Expression (2).

Further, $v(t_1, a) - v(t_2, a)$ switches sign between maximizers for $\lambda = 0$ and $\lambda = 1$. To see this, note that if $\lambda = 0$, then $\bar{a}(t_1)$ maximizes

$$v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda = v(t_1, a)(1 - q)$$

for which we have $v(t_2, \bar{a}(t_1)) - v(t_1, \bar{a}(t_1)) > 0$ by assumption. Similarly, if $\lambda = 1$, then by assumption, for any maximizer a of

$$v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda = v(t_1, a) - qv(t_2, a)$$

we have $v(t_2, a) - v(t_1, a) < 0$. Therefore, there must exist a threshold λ such that two alternative a, a' (possibly identical) maximize Expression (2), such that $v(t_2, a) - v(t_1, a) \leq 0$, and $v(t_2, a') - v(t_1, a') \geq 0$. Thus there exists a distribution x over a, a' such that $v(t_2, x) - v(t_1, x) = 0$. ■

Before stating the characterization of optimal mechanisms formally, we state another lemma that will be used later in the proof of the characterization. For future reference, we

label the four incentive constraints as follows

$$v(t_1, x(t_1)) - p(t_1) \geq 0, \quad (\text{IR1})$$

$$v(t_2, x(t_2)) - p(t_2) \geq 0, \quad (\text{IR2})$$

$$v(t_1, x(t_1)) - p(t_1) \geq v(t_1, x(t_2)) - p(t_2), \quad (\text{IC1})$$

$$v(t_2, x(t_2)) - p(t_2) \geq v(t_2, x(t_1)) - p(t_1). \quad (\text{IC2})$$

Lemma 8 *Constraint (IC1) is satisfied if constraints (IR1) and (IC2) hold with equality, $x(t_1) \in \Delta(\arg \max v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda)$ for some λ , and $x(t_2) = \bar{a}(t_2)$.*

Proof. Since $x(t_2) = \bar{a}(t_2)$ and constraints (IR1) and (IC2) hold with equality, the surplus of type t_1 for mimicking type t_2 is

$$\begin{aligned} v(t_1, \bar{a}(t_2)) - p(t_2) &= v(t_1, \bar{a}(t_2)) - v(t_2, \bar{a}(t_2)) + v(t_2, x(t_1)) - p(t_1) \\ &= v(t_1, \bar{a}(t_2)) - v(t_2, \bar{a}(t_2)) + v(t_2, x(t_1)) - v(t_1, x(t_1)) \\ &= E_{a \sim x(t_1)}[v(t_1, \bar{a}(t_2)) - v(t_2, \bar{a}(t_2)) + v(t_2, a) - v(t_1, a)]. \end{aligned} \quad (3)$$

Now consider any alternative a in the support of $x(t_1)$. By assumption, it must be that $a \in \arg \max v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda$. Thus, in particular,

$$v(t_1, \bar{a}(t_2))(1 - q + q\lambda) - v(t_2, \bar{a}(t_2))q\lambda \leq v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda.$$

Rearranging, we have

$$\left(v(t_1, \bar{a}(t_2)) - v(t_1, a) \right) (1 - q + q\lambda) \leq \left(v(t_2, \bar{a}(t_2)) - v(t_2, a) \right) q\lambda.$$

Notice that the right hand side is non-negative since $v(t_2, \bar{a}(t_2)) - v(t_2, a) \geq 0$. Moreover, $1 - q + q\lambda \geq q\lambda$. Thus we must have

$$v(t_1, \bar{a}(t_2)) - v(t_1, a) \leq v(t_2, \bar{a}(t_2)) - v(t_2, a).$$

Rearranging again, we conclude that for all a in the support of $x(t_1)$,

$$v(t_1, \bar{a}(t_2)) - v(t_2, \bar{a}(t_2)) + v(t_2, a) - v(t_1, a) \leq 0.$$

Taking the expectation all such a ,

$$E_{a \sim x(t_1)}[v(t_1, \bar{a}(t_2)) - v(t_2, \bar{a}(t_2)) + v(t_2, a) - v(t_1, a)] \leq 0.$$

Thus, by (3), the surplus of type t_1 for mimicking type t_2 is non-positive, and constraint (IC1) is satisfied. ■

The lemma below formally identifies optimal mechanisms in each of the three cases discussed above. Instead of giving the explicit formula for payments, the lemma identifies the allocation for each type and specifies which incentive constraints bind. The payments can then be recovered from the binding constraints.

Lemma 9 *Assume that there are two types. Necessary and sufficient conditions for the optimality of a mechanism (x, p) are provided in each of the three cases below.*

1. $v(t_1, \bar{a}(t_1)) \geq v(t_2, \bar{a}(t_1))$ and $v(t_2, \bar{a}(t_2)) \geq v(t_1, \bar{a}(t_2))$. Then $x(t_1)$ and $x(t_2)$ are distributions over such alternatives, and (IR1) and (IR2) are tight.
2. $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$ and there exists $a \in \arg \max v(t_1, a) - qv(t_2, a)$ such that $v(t_2, a) - v(t_1, a) \geq 0$. Then $x(t_1) \in \Delta(\arg \max v(t_1, a) - qv(t_2, a))$ such that $v(t_2, x(t_1)) - v(t_1, x(t_1)) \geq 0$, $x(t_2) = \bar{a}(t_2)$, and (IR1) and (IC2) are tight.
3. $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$ and for all $a \in \arg \max v(t_1, a) - qv(t_2, a)$, we have $v(t_2, a) - v(t_1, a) < 0$. Then there exists λ and $x(t_1)$ such that $x(t_1) \in \Delta(\arg \max v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda)$, $v(t_2, x(t_1)) - v(t_1, x(t_1)) = 0$, $x(t_2) = \bar{a}(t_2)$, and (IR1), (IR2), and (IC2) are tight.

Proof. Consider any IC and IR mechanism (x, p) . Substituting (IR1), (IR2), and (IC2), for any $\lambda, 0 \leq \lambda \leq 1$, consider the Lagrangian relaxation of the expected revenue

$$\begin{aligned}
p(t_1)(1 - q) + p(t_2)q &= p(t_1)(1 - q + q\lambda) + (p(t_2) - p(t_1))q\lambda + p(t_2)q(1 - \lambda) \\
&\leq v(t_1, x(t_1))(1 - q + q\lambda) + (v(t_2, x(t_2)) - v(t_2, x(t_1)))q\lambda + v(t_2, x(t_2))q(1 - \lambda) \\
&= (v(t_1, x(t_1))(1 - q + q\lambda) - v(t_2, x(t_1))q\lambda) + v(t_2, x(t_2))q. \tag{4}
\end{aligned}$$

Thus, an IC and IR mechanism is optimal if it satisfies two conditions for some $\lambda, 0 \leq \lambda \leq 1$. First, it maximizes (4) across all $x(t_1)$ and $x(t_2)$. Second, it satisfies complementary slackness conditions. That is, (IR1) is tight, if $\lambda < 1$ then (IR2) is tight, and if $\lambda > 0$ then (IC2) is tight. Conversely, if an IC mechanism that satisfies the aforementioned two properties exists, then a mechanism is optimal only if it satisfies those properties.

We now consider each of the three cases specified by the lemma.

1. Full surplus extraction is feasible and thus optimal.
2. Consider any distribution x' over maximizers a of $v(t_1, a) - qv(t_2, a)$ for which $v(t_2, x') - v(t_1, x') \geq 0$. Such a distribution exists since by assumption, there exists $a \in \arg \max v(t_1, a) - qv(t_2, a)$ such that $v(t_2, a) - v(t_1, a) \geq 0$. Consider mechanism (x, p) where $x(t_1) = x'$, $x(t_2) = \bar{a}(t_2)$, and (IR1) and (IC2) are tight. The constraint (IC1) is satisfied by Lemma 8 for $\lambda = 1$. The constraint (IR2) is satisfied by the assumption that $v(t_2, x') - v(t_1, x') \geq 0$. Thus the mechanism is IC-IR. We next argue that the mechanism is optimal. The allocation x maximizes (4) for $\lambda = 1$. The complementary slackness conditions are satisfied since $\lambda = 1$ and (IR1) and (IC2) are tight. Therefore, the mechanism is optimal, and any optimal mechanism must satisfy the specified properties.
3. By Lemma 7 there exists λ and $x' \in \Delta(\arg \max v(t_1, a)(1 - q + q\lambda) - v(t_2, a)q\lambda)$ such that $v(t_2, x') - v(t_1, x') = 0$. Consider mechanism (x, p) where $x(t_1) = x'$, $x(t_2) = \bar{a}(t_2)$, and (IR1) and (IC2) are tight. Notice that constraint (IR2) is satisfied with equality since $v(t_2, x') - v(t_1, x') = 0$. Constraint (IC1) is satisfied by Lemma 8. Thus the mechanism is IC-IR. We next argue that the mechanism is optimal. The allocation x maximizes (4) for λ . The complementary slackness conditions are satisfied since (IR1), (IR2), and (IC2) are tight. Therefore, the mechanism is optimal, and any optimal mechanism must satisfy the specified properties.

■

A.1.2 Proof of Lemma 1

We now use the characterization of optimal mechanisms in Lemma 9 to prove Lemma 1, restated below.

Lemma 1 *Consider an environment with two types. For one of the types, denoted t_L , the surplus is 0 in any market: $CS(t_L, q) = 0$ for any q in $[0, 1]$. For the other type, denoted t_H , there exists some $m \leq k$, thresholds $q_0(= 0) < \dots < q_{m+1}(= 1)$, and surpluses $\alpha_0 > \dots > \alpha_m(= 0)$ such that $CS(t_H, q) = \alpha_j$ for q in (q_j, q_{j+1}) .*

Proof. In case (1) of Lemma 9, the surplus of both types is zero and the lemma immediately follows. Thus for the rest of the proof assume that we are in cases (2) and (3) of

Lemma 9, that is, $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$. Type t_2 corresponds to the high type t_H in the statement of the lemma, and type t_1 corresponds to the low type t_L . Since the IR constraint of type t_1 binds, its surplus is zero. We next prove the lemma regarding the surplus of type t_2 .

The key to proof is to consider the maximization problem $\max v(t_1, a) - qv(t_2, a)$. If $v(t_2, a) - v(t_1, a) \geq 0$ for some maximizer a of $v(t_1, a) - qv(t_2, a)$, then by case (2), the surplus of type t_2 is equal to $v(t_2, a) - v(t_1, a)$. If $v(t_2, a) - v(t_1, a) < 0$ for all maximizers a of $v(t_1, a) - qv(t_2, a)$, then by case (3), the surplus of type t_2 is equal to zero.

Thus consider the problem $\max_a v(t_1, a) - qv(t_2, a)$. The problem is one of maximizing over functions $v(t_1, a) - qv(t_2, a)$, one for each alternative a , that are linear in q . Thus, the set of markets $[0, 1]$ can be divided into intervals, each assigned a unique alternative, such that the alternative is the unique maximizer of $v(t_1, a) - qv(t_2, a)$ in the interior of the interval, and is an optimizer of $v(t_1, a) - qv(t_2, a)$ at the endpoints of the interval. Additionally, the slope of the line that corresponds to the maximizer a increases in q . See Figure 10. The slope of the line is $v(t_1, a) - v(t_2, a)$. If the slope is negative, we are in case (2) of Lemma 9 and thus the slope is equal to the negative of the surplus of type t_2 . If the slope is positive, we are in case (3) of Lemma 9 and thus the surplus of type t_2 is equal to zero. As q increases, as long as the slope of the line is negative, the surplus of type t_2 remains constant and decreases at the thresholds where the maximizer changes. If for some q , the slope of the line is positive, then the surplus of type t_2 is zero and will remain zero for all $q' \geq q$. ■

A.2 Alternative proof for two types

Suppose first that $v(t_1, \bar{a}(t_1)) \geq v(t_2, \bar{a}(t_1))$ and $v(t_1, \bar{a}(t_2)) \leq v(t_2, \bar{a}(t_2))$. In this case, extracting the full surplus is feasible since each type can be assigned his efficient alternative at price equal to his willingness to pay, without violating IC or IR. Thus such a mechanism is optimal and both types have zero surplus in any market.

Now suppose $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$ (the case where $v(t_1, \bar{a}(t_1)) > v(t_2, \bar{a}(t_2))$ is symmetric). In what follows, we use the observation that if the IC constraint of type t to t' binds in an optimal mechanism, then the allocation of type t must be efficient. First notice that the payment of type t must be weakly higher than the payment of type t' . Otherwise, type t can be given the assignment of type t' and thus increase revenue. Now

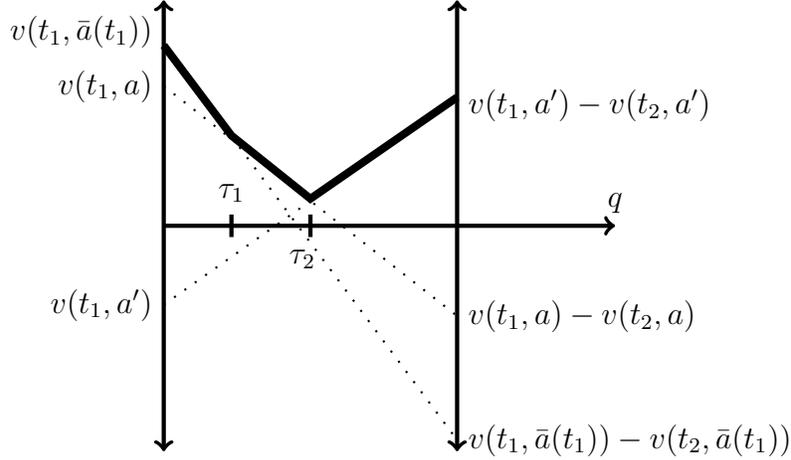


Figure 10: The slope $v(t_1, a) - v(t_2, a)$ of the line that corresponds to the alternative a that maximizes $v(t_1, a) - qv(t_2, a)$ increases in q .

consider modifying the optimal mechanism by giving type t his efficient alternative, and increasing his payment such that his surplus remains the same. If this change violates the IC constraint of type t' , give the same assignment to t' . The new mechanism is IC-IR and has strictly higher revenue. Thus the allocation of t must be efficient in any optimal mechanism. Given this observation, we prove properties of any optimal mechanism below.

The IR constraint for type t_1 binds Suppose for contradiction that the surplus of type t_1 is strictly positive. Since the IR constraint of type t_1 is slack, its IC constraint must bind. The argument above shows that the allocation of t_1 must be efficient. The assumption that $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$ implies that the surplus of type t_2 is strictly positive. Thus the surplus of both types is positive, contradicting optimality.

The IC constraint for type t_2 binds Suppose for contradiction that the IC constraint for type t_2 is slack. Suppose also that the allocation of type t_1 is inefficient. Consider making the allocation of type t_1 slightly more efficient (possible at random) and increasing the payment of type t_1 so that the surplus of type t_1 remains the same and the IC constraint of type t_2 remains satisfied. The mechanism has higher revenue, which contradicts optimality. Thus the allocation of type t_1 must be efficient. The assumption that $v(t_1, \bar{a}(t_1)) < v(t_2, \bar{a}(t_1))$ implies that the surplus of type t_1 is strictly positive, which means that the IC constraint of type t_1 must bind.

The allocation of type t_2 is efficient Since the IC constraint of type t_2 binds, its allocation must be efficient as argued above.

The surplus of type t_2 is non increasing in q In an optimal mechanism, the payments of the two types are pinned down by the fact that the IR constraint for type t_1 and the IC constraint for type t_2 bind. The allocation of type t_2 is efficient, which means that the only degree of freedom is the allocation of type t_1 . The optimal mechanism maximizes

$$(1 - q)v(t_1, x(t_1)) + q\left(v(t_2, \bar{a}(t_2)) - (v(t_2, x(t_1)) - v(t_1, x(t_1)))\right),$$

over all distributions $x(t_1)$ over alternatives subject to the remaining two constraints, namely the IR constraint of type t_2 and the IC constraint of type t_1 ,

$$0 \leq v(t_2, x(t_1)) - v(t_1, x(t_1)) \leq v(t_2, \bar{a}(t_2)) - v(t_1, \bar{a}(t_2)).$$

Notice that the objective is to maximize

$$(1 - q)v(t_1, x(t_1)) - q(v(t_2, x(t_1)) - v(t_1, x(t_1))) \tag{5}$$

plus a constant that does not depend on $x(t_1)$. Thus the problem is to maximize the weighted sum of the value $v(t_1, x(t_1))$ given to type t_1 minus the surplus $v(t_2, x(t_1)) - v(t_1, x(t_1))$ given to type t_2 . Therefore, as q increases, the surplus of type t_2 weakly decreases.

If $|A|$ is finite, then $CS(t_2, f)$ consists of at most $|A|$ constant pieces We first show that the IC constraint of type t_1 ,

$$v(t_2, x(t_1)) - v(t_1, x(t_1)) \leq v(t_2, \bar{a}(t_2)) - v(t_1, \bar{a}(t_2)).$$

does not bind. Since the surplus of type t_2 is decreasing in q , we only need to verify this claim for $q = 0$. In this case, type t_1 is assigned his efficient alternative. Since the payment of type t_2 is weakly higher than the payment of type t_1 and type t_1 is assigned his efficient alternative, the IC constraint of type t_1 is satisfied. Thus, the optimal mechanism maximizes (5) only subject to the IR constraint of type t_2 .

Consider the relaxed problem of maximizing the objective (5) over all $x(t_1)$. The solution is any distribution over alternatives that maximize the objective, that is, $v(t_1, a) -$

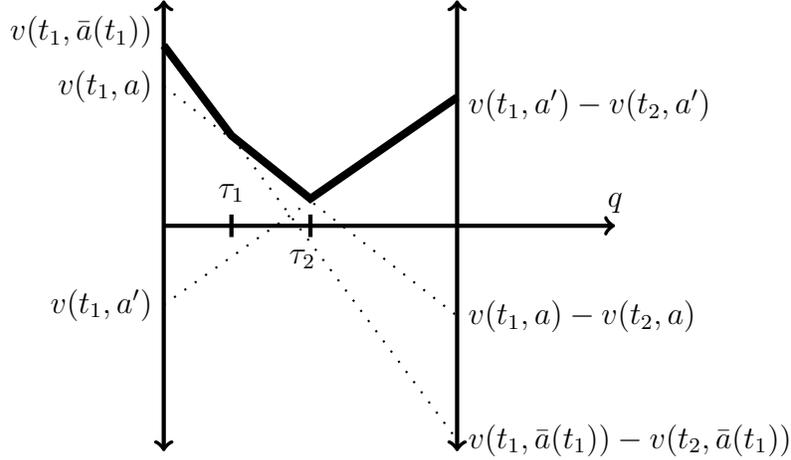


Figure 11: The slope $v(t_1, a) - v(t_2, a)$ of the line that corresponds to the alternative a that maximizes $v(t_1, a) - qv(t_2, a)$ increases in q .

$qv(t_2, a)$. The problem is one of maximizing over functions $v(t_1, a) - qv(t_2, a)$, one for each alternative a , that are linear in q . Thus, the set of markets $[0, 1]$ can be divided into intervals, each assigned a unique alternative, such that the alternative is the unique maximizer of $v(t_1, a) - qv(t_2, a)$ in the interior of the interval, and is an optimizer of $v(t_1, a) - qv(t_2, a)$ at the endpoints of the interval. Additionally, the slope of the line that corresponds to the maximizer a increases in q . See Figure 11. The slope of the line is $v(t_1, a) - v(t_2, a)$. If the slope is negative, then the IR constraint of type t_2 is satisfied by assigning alternative a to type t_1 . If the slope is positive, then any optimal solution must give zero surplus to type t_2 .

A.3 Completing the proof of Proposition 3

In order to complete the proof of Proposition 3, we show that there exists a market with full support in the restricted environment with types $T \cup \{j\}$ for which mechanism M^* is optimal.

Lemma 10 *Consider an environment with a set of types $\{i_1, \dots, i_n\}$, and an efficient mechanism M^* such that the IC constraint from each type t_j to the next type t_{j+1} and the IR constraint for type t_n bind. There exists a market with full support over $\{i_1, \dots, i_n\}$ for which mechanism M^* is the unique optimal mechanism.*

Proof. Consider any IC-IR mechanism (x, p) and any market f . The IC constraint corresponding to a deviation of type t_j to type t_{j+1} is

$$p(t_j) - p(t_{j+1}) \leq v(t_j, x(t_j)) - v(t_j, x(t_{j+1})),$$

and the IR constraint of type t_n is

$$p(t_n) \leq v(t_n, x(t_n)).$$

Using the above inequalities, we can write the expected revenue of mechanism (x, p) . For each type t_j , let $F(t_j) = \sum_{j' \leq j} f(t_{j'})$ be the cumulative fraction of types t_1 to t_j . Now write

$$\begin{aligned} \sum_j p(t_j) f(t_j) &= \sum_j (p(t_j) - p(t_{j+1})) F(t_j) \\ &\leq \sum_j \left(v(t_j, x(t_j)) - v(t_j, x(t_{j+1})) \right) F(t_j) \\ &= \sum_j v(t_j, x(t_j)) F(t_j) - v(t_{j-1}, x(t_j)) F(t_{j-1}). \end{aligned} \quad (6)$$

Therefore, for any market f , the revenue of any IC-IR mechanism is at most the maximum of Expression (6) over all allocation rules x .

By definition, the efficient alternative $\bar{a}(t_j)$ of type t_j satisfies $v(t_j, \bar{a}(t_j)) > v(t_j, a)$ for all alternatives $a \neq \bar{a}(t_j)$. Thus, if $F(t_{j-1})$ is small enough relative to $F(t_j)$, that is, $F(t_{j-1}) \leq \delta_j F(t_j)$ for some $\delta_j > 0$, then

$$v(t_j, \bar{a}(t_j)) F(t_j) - v(t_{j-1}, \bar{a}(t_j)) F(t_{j-1}) \geq v(t_j, a) F(t_j) - v(t_{j-1}, a) F(t_{j-1})$$

for all $a \neq \bar{a}(t_j)$. As a result, for such a market, the unique maximizer of $v(t_j, x) \tilde{F}(t_j) - v(t_{j-1}, x) \tilde{F}(t_{j-1})$ over all distributions x over alternatives is a distribution that assigns probability one to alternative $\bar{a}(t_j)$.

Now consider any market f with full support over the set of types $\{i_1, \dots, i_n\}$ such that $F(t_{j-1}) \leq \delta_j F(t_j)$ for all j . By the above discussion, the allocation rule of the mechanism M^* is the unique maximizer of Expression (6) over all allocation rules. In addition, since the IC constraint from each type t_j to the next type t_{j+1} and the IR constraint for type t_n bind, then the revenue of the mechanism M^* is equal to the maximum of Expression (6) over all allocation rules. Thus, mechanism M^* is the unique optimal mechanism for market f . ■

A.4 The “only if” direction of Proposition 3

For any mechanism $M = (x, p)$ and any type t , let $CS(t, M)$ denote the surplus of type t in mechanism M . Note for future reference that the IC constraint that states that type t should not benefit from reporting t' can be written as

$$CS(t, M) - CS(t', M) \geq v(t, x(t')) - v(t', x(t')). \quad (7)$$

Consider an efficient market f with optimal mechanism $M = (x, p)$. Assume without loss of generality that f has full support. Assume for contradiction that there exists a market f' with optimal mechanism M' that Pareto dominates f . Thus there exists a type t in the support of f' such that $CS(t, f) < CS(t, f')$. In market f' , either the IR constraint of type t binds, or its IC constraint to some other type t' binds. Thus, there exists a binding IC path S that starts with type t and ends with a type that has surplus 0. Suppose without loss of generality that for all types t' other than t on the path, $CS(t', f) = CS(t', f')$ (otherwise denote by t the last type in the path for which $CS(t, f) < CS(t, f')$).

We show that every type $t' \neq t$ on the path must be assigned his efficient alternative $\bar{a}(t')$ in market f' . Assume for contradiction that t' is assigned some alternative other than $\bar{a}(t')$. Modify the optimal mechanism M' for f' by assigning type t' its efficient alternative and increasing his payment to leave his surplus unchanged. Since type t' has the same surplus in the modified mechanism $\hat{M}' = (\hat{x}', \hat{p}')$ as in mechanism M' , it does not benefit from misreporting. So we need to only verify that another type t'' does not benefit from misreporting to type t' . It is without loss of generality to assume that type t'' is in the support of f' . The IC constraint of all types not in the support of f' can be satisfied by allowing them choose among all the allocation and price pairs that is available to the types in the support of f' , without affecting revenue. Since type t'' is in the support of f' and market f' Pareto dominates market f by assumption, we have $CS(t'', \hat{M}') = CS(t'', M') \geq CS(t'', M)$. Also by the construction of the path S , $CS(t', \hat{M}') = CS(t', M') \geq CS(t', M)$. Thus we have

$$\begin{aligned} CS(t'', \hat{M}') - CS(t', \hat{M}') &\geq CS(t'', M) - CS(t', M) \\ &\geq v(t'', \bar{a}(t')) - v(t', \bar{a}(t')) \\ &= v(t'', \hat{x}'(t')) - v(t', \hat{x}'(t')), \end{aligned}$$

where the second inequality follows since M is IC-IR and efficient and using Inequality (7). Thus, the incentive constraint (7) is satisfied and mechanism \hat{M}' is IC-IR. Since \hat{M}' has

higher expected revenue than mechanism M' , the optimality of M' is contradicted. We conclude that every type $t' \neq t$ on the path S must be assigned his efficient alternative $\bar{a}(t')$ in market f' .

We now show that the surplus of type t is no smaller in M than in M' , contradicting $CS(t, f) < CS(t, f')$. Indeed, since the IC constraint for all types t_j on path S and the IR constraint for the last type on the path bind in mechanism M' , we can write the surplus of type t in market f' as

$$\begin{aligned} CS(t, f') &= \sum_j CS(t_j, f') - CS(t_{j+1}, f') \\ &= \sum_j v(t_j, \bar{a}(t_{j+1})) - v(t_{j+1}, \bar{a}(t_{j+1})). \end{aligned}$$

Similarly, by applying the incentive constraint (7) for all types before the last type on path S and the IR constraint of the last type on the path in market f we have

$$\begin{aligned} CS(t, f) &\geq \sum_j CS(t_j, f) - CS(t_{j+1}, f) \\ &\geq \sum_j v(t_j, \bar{a}(t_{j+1})) - v(t_{j+1}, \bar{a}(t_{j+1})). \end{aligned}$$

We thus have $CS(t, f) \geq CS(t, f')$, which is a contradiction.

A.5 Proof of Theorem 1

We here prove Theorem 1, restated below.

Theorem 1 *For any environment with $n \geq 2$ types and $k \geq 1$ alternatives, the set of inefficient markets that are not Pareto improvable by a two-market segmentation is non-generic.*

Proof. We show that if an inefficient market f is not Pareto improvable, then perturbing f does not maintain the optimal mechanism. Thus the set of inefficient markets that are not Pareto improvable markets contains is contained in the set of markets such that perturbing the market does not maintain the optimal mechanism. By Proposition 4, the latter set is non-generic. Since a subset of a non-generic set is non-generic, the theorem follows.

To complete the proof, we show that if perturbing an inefficient market f maintains the optimal mechanism, then f is Pareto improvable. By Proposition 3, there exists a

Pareto dominating market f' , whose support is a subset of the support of f . Define $f'' = (f - \varepsilon f') / (1 - \varepsilon)$, and consider a segmentation μ of f into f' and f'' , $f = \varepsilon f' + (1 - \varepsilon) f''$. Since perturbing f maintains the optimal mechanism, for small enough $\varepsilon > 0$, the optimal mechanism for f is also optimal for f'' . Therefore, f'' weakly Pareto dominates f and the segmentation μ is Pareto improving. ■

A.6 Proof of Lemma 6

We here prove Lemma 6, restated below.

Lemma 6 *Consider a market f in which the allocation of some type t is inefficient. There exists δ such that a (t, δ) -permutation of market f Pareto dominates f .*

Proof. Consider a market f and its (t, δ) -permutation f' for $\delta < 0$. Let ϕ and ϕ' be the virtual valuations in markets f and f' , respectively. We first argue that $\phi(j) \leq \phi'(j)$ for all types j . As a result, the optimal allocations a and a' of market f and f' satisfy $a(j) \leq a'(j)$ for all j . Further, we show that the virtual valuation of type t is strictly larger in market f' compared to f for any $\delta < 0$. As δ approaches -1 (from above), the virtual valuation of type t approaches his valuation $v(i)$. Thus by choosing $-1 < \delta < 0$ small enough, the virtual valuation of type t becomes large enough so that $a(i) < a'(i)$. Property 2 of Proposition 6 then implies that all types $j \leq i$ have weakly higher surplus in f' compared to f , and all types $j > i$ have strictly higher surplus in f' compared to f . Since $\delta > -1$, all types above i have positive fractions in market f' . Therefore, f' Pareto dominates f .

To complete the proof, we thus argue that $\phi(j) \leq \phi'(j)$ for all types j and $\phi(t) < \phi'(i)$, for any $\delta < 0$. This can be shown simply with a picture. Let $\{t_L, \dots, t_H\} \ni i$ be the interval of types over which the revenue function of market f is ironed and to which type t belongs. Formally, $t_L \leq i$ is the highest type weakly smaller than i for which $R(j) = R_{IR}(j)$, and $t_H > i$ is smallest type larger than i for which $R(j) = R_{IR}(j)$, as depicted in Figure 12. Lowering δ by a small amount increases the virtual valuations of types $t_L, \dots, t_H - 1$, and keeps the virtual valuations of all other types unchanged, as claimed. ■

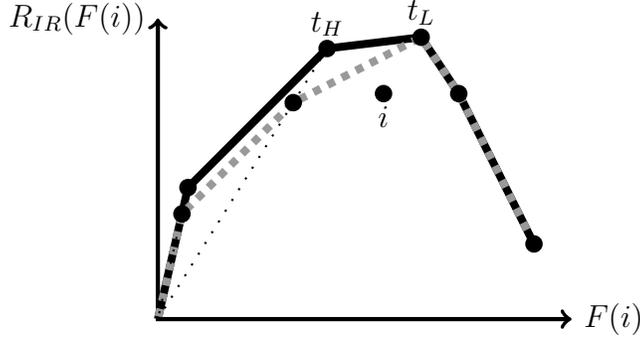


Figure 12: Type t belongs to the ironed interval $\{t_L, \dots, t_H\}$. A (t, δ) -perturbation for $\delta < 0$ increases the virtual valuations of types $t_L, \dots, t_H - 1$, and keeps the virtual valuations of all other types unchanged.

A.7 Proof of Proposition 5

Before proving Proposition 5, we first define virtual valuations more formally than Equation (1) such that they are well defined even for types with zero proportions.

A.7.1 Virtual valuations of types with zero proportions

For a type t , let $t_L \leq i$ be the highest type lower than i such that $R(\tilde{F}(t_L)) = R_{IR}(\tilde{F}(t_L))$. The virtual valuation of type t is defined to be the slope of the ironed revenue curve to the left of type t_L . This is well defined regardless of whether $f(t) > 0$ or $f(t) = 0$. The definition coincides with Equation (1) for types with positive proportion since the slope of the ironed revenue curve at i is equal to the slope of the ironed revenue curve at t_L .

The proof of Proposition 5 uses a property of virtual valuations shared by all types regardless of their proportions. Notice that if $f(t) = 0$ for some type t , then type t is vertically below type $t + 1$ in the graph of the revenue curve, and thus $R_{IR}(\tilde{F}(i)) = R_{IR}(\tilde{F}(i + 1))$. Thus for such a type we have

$$\phi(t)f(t) = R_{IR}(\tilde{F}(i)) - R_{IR}(\tilde{F}(i + 1)), \quad (8)$$

since both sides of the above equality are zero for such a type. Equation (8) also holds for any type t with $f(t) > 0$. This follows directly from Equation (1). We thus use Equation (8) in the proof of Proposition 5 as a property that the virtual valuation of any type t satisfies, whether $f(t) > 0$ or $f(t) = 0$.

A.7.2 The proof of Proposition 5

We here prove Proposition 5, restated below.

Proposition 5 *A market f is Pareto improvable if and only if there exists a type t with inefficient allocation and $\delta_0 > 0$ such that every (t, δ) -perturbation of f leaves the optimal mechanism unchanged for $\delta \leq \delta_0$.*

Proof. The “if” direction. Suppose that there exists a type t with inefficient allocation and $\delta_0 > 0$ such that a (t, δ) -perturbation of f leaves the optimal mechanism unchanged for $\delta \leq \delta_0$. By Lemma 6, there exists a market f' that is a (t, δ') -permutation of f and that Pareto dominates f . For any ϵ , consider a segmentation $f = \epsilon f' + (1 - \epsilon)f''$ of f into f' and f'' . Since f' is a (t, δ') -permutation of f , f'' is a (t, δ'') -permutation of f for some δ'' . Further, by choosing ϵ small enough, we have $\delta'' < \delta_0$. Thus, the optimal mechanism for f'' is identical to that of f , and f'' weakly Pareto dominates f . As a result, the segmentation is Pareto improving and f is Pareto improvable.

The “only if” direction. Suppose that for any type j with inefficient allocation, any (j, δ) -perturbation of f changes the optimal mechanism. Thus lowering the virtual valuation of type j changes its optimal allocation. We show that f is not Pareto improvable. Assume for contradiction that there exists a Pareto improving segmentation μ of f . The allocation of some type in some segment must be more efficient than the allocation of that type in market f . Let i be the lowest such type. Let $\{t_L, \dots, t_H\} \ni i$ be the interval of types over which the revenue function of market f is ironed and to which type t belongs. Formally, $t_L \leq i$ is the highest type weakly smaller than i for which $R(j) = R_{IR}(j)$, and $t_H > i$ is smallest type larger than i for which $R(j) = R_{IR}(j)$, as depicted in Figure 13.

We first argue that the allocation of type t must be weakly higher in every segment f' , and strictly higher in some segment, than in f . Assume for contradiction that the optimal allocation a' of some market f' satisfies $a'(i) < a(i)$. We show that some type with positive fraction in f' must be worse off in f' than in f , contradicting the requirement that all segments Pareto dominate f . Notice that there must exist a type higher than type t that has positive fraction in market f' . Otherwise, the allocation of type t must be efficient in market f' , which contradicts $a'(i) < a(i)$. Now consider the lowest type t' that is higher than i and has positive fraction in market f' , $f'(t') > 0$. All types in the set $\{i+1, \dots, t'-1\}$ have zero fraction in market f' . Therefore, the virtual valuation of all such

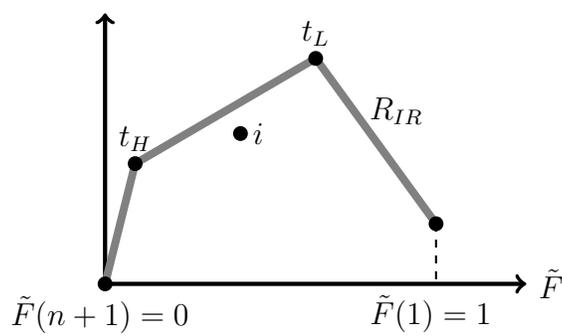


Figure 13: Type t belongs to the ironed interval $\{t_L, \dots, t_H\}$.

types is equal to $\phi'(i)$. Thus we have $a'(j) = a'(i) < a(i) \leq a(j)$ for all $j \in \{i+1, \dots, i'-1\}$ by monotonicity of a . Since $a'(j) = a(j)$ for all $j < i$, property 2 of Proposition 6 implies that type t' has strictly lower surplus in f' than in f , which contradicts the requirement that f' Pareto dominates f . We thus have $a'(i) \geq a(i)$ for all segments f' . Since type t has a different allocation in some segment compared to f , we have $a'(i) > a(i)$ in some segment.

We next argue that the virtual valuation of any type in $\{t_L, \dots, t_H - 1\}$ is weakly higher in every segment f' than in f , and is strictly higher for some such type and segment. Recall that lowering the virtual valuation of type t changes its optimal allocation, and notice that every type $j \in \{t_L, \dots, t_H - 1\}$ has the same virtual valuation in market f as does type t . Thus, if $\phi(j) > \phi'(j)$ for some segment f' , then type j has lower allocation than it does in market f . Type j cannot be lower than type t since by assumption, all types lower than i have the same allocation in any segment f' and in f . Type j can also not be weakly higher than i , since then have $a'(j) < a'(i)$ which violates monotonicity of a' . Thus $\phi(j) \leq \phi'(j)$ for all segments f' and types $j \in \{t_L, \dots, t_H - 1\}$. Further, since type t has higher allocation in some segment f' than in f , it must have a higher virtual valuation in f' compared to f . We conclude that the virtual valuations ϕ' in every segment f' satisfy $\sum_{j:t_L \leq j < t_H} \phi(j)f'(j) \leq \sum_{j:t_L \leq j < t_H} \phi'(j)f'(j)$, with strict inequality for some segment. Taking the expectation over all segments f' in the segmentation μ , we have

$$E_\mu \left[\sum_{j:t_L \leq j < t_H} \phi(j)f'(j) \right] < E_\mu \left[\sum_{j:t_L \leq j < t_H} \phi'(j)f'(j) \right]. \quad (9)$$

We next show that this is impossible.

Recall from Equation (8) that the virtual valuation of type t satisfies

$$\phi(i)f(t) = R_{IR}(\tilde{F}(i)) - R_{IR}(\tilde{F}(i+1)).$$

Summing over all types in $\{t_L, \dots, t_H - 1\}$, we have

$$\begin{aligned} \sum_{j:t_L \leq j < t_H} \phi(j)f(j) &= \sum_{j:t_L \leq j < t_H} R_{IR}(\tilde{F}(j)) - R_{IR}(\tilde{F}(j+1)) \\ &= R_{IR}(\tilde{F}(t_L)) - R_{IR}(\tilde{F}(t_H)) \\ &= R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)). \end{aligned} \tag{10}$$

Similar to above, for every segment f' we have

$$\begin{aligned} \sum_{j:t_L \leq j < t_H} \phi'(j)f'(j) &= \sum_{j:t_L \leq j < t_H} R'_{IR}(\tilde{F}'(j)) - R'_{IR}(\tilde{F}'(j+1)) \\ &= R'_{IR}(\tilde{F}'(t_L)) - R'_{IR}(\tilde{F}'(t_H)). \end{aligned} \tag{11}$$

Since R'_{IR} is the concavification of R' , it is weakly higher than R' , and thus in particular $R'_{IR}(\tilde{F}'(t_H)) \geq R'(\tilde{F}'(t_H))$. We argue that the inequality must hold with equality for type t_L , that is $R'_{IR}(\tilde{F}'(t_L)) = R'(\tilde{F}'(t_L))$. Otherwise, if $R'_{IR}(\tilde{F}'(t_L)) > R'(\tilde{F}'(t_L))$, then the virtual valuation of type t_L is equal to the virtual valuation of type $t_L - 1$. Thus the allocation of types $a'(t_L) = a'(t_L - 1)$. Moreover, $a(t_L) > a(t_L - 1)$ by the assumption that lowering the virtual valuation of type t changes its optimal allocation, and the observation that $\phi(t) = \phi(t_L) > \phi(t_L - 1)$. Recall that i is chosen such that all types lower than i has identical allocation in any segment f' and in f , and therefore $a(t_L) = a'(t_L)$ and $a(t_L - 1) = a'(t_L - 1)$. Therefore we must have $a'(t_L) = a'(t_L - 1)$ which contradicts $a'(t_L) > a'(t_L - 1)$. We conclude that $R'_{IR}(\tilde{F}'(t_L)) = R'(\tilde{F}'(t_L))$. Together with (11), we conclude that

$$\begin{aligned} \sum_{j:t_L \leq j < t_H} \phi'(j)f'(j) &= R'_{IR}(\tilde{F}'(t_L)) - R'_{IR}(\tilde{F}'(t_H)) \\ &\leq R'(\tilde{F}'(t_L)) - R'(\tilde{F}'(t_H)). \end{aligned} \tag{12}$$

To finish the proof, we show that inequalities (9), (10), and (12) cannot simultaneously hold. Indeed, by (10) we have

$$R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)) = \sum_{j:t_L \leq j < t_H} \phi(j)f(j).$$

Since μ is a segmentation of f , $f(j) = E_\mu[f'(j)]$ for all j and thus we have

$$\begin{aligned}
R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)) &= E_\mu\left[\sum_{j:t_L \leq j < t_H} \phi(j)f'(j)\right] \\
&< E_\mu\left[\sum_{j:t_L \leq j < t_H} \phi'(j)f'(j)\right] \\
&\leq E_\mu[R'(\tilde{F}'(t_L)) - R'(\tilde{F}'(t_H))],
\end{aligned}$$

where the inequalities followed from (10) and (12). From the definition of the revenue functions R and R' , and again employing the fact that μ is a segmentation of f , we have

$$\begin{aligned}
R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)) &< E_\mu[v(t_L)\tilde{F}'(t_L) - v(t_H)\tilde{F}'(t_H)] \\
&= v(t_L)\tilde{F}(t_L) - v(t_H)\tilde{F}(t_H) \\
&= R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)).
\end{aligned}$$

That is, $R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)) < R(\tilde{F}(t_L)) - R(\tilde{F}(t_H))$, which is a contradiction. ■