Pareto Improving Segmentation of Multi-Product Markets

Nima Haghpanah and Ron Siegel*

August 20, 2021

Abstract

We consider the welfare-improving potential of a platform that collects consumer data and supplies it to a multi-product seller. Such data may lead the seller to optimally offer consumers more suitable products but may also enable her to set higher prices. We ask whether some regulation of data collection and usage can benefit all consumers and the seller. In reduced form, any regulation gives rise to some market segmentation, which may allow the seller to combine second- and third-degree price discrimination. We show that markets for which profit-maximization implies allocative inefficiency can, generically, be segmented in a way that increases the surplus of all consumers weakly and of some consumers and the seller strictly. Our proof is constructive and based on deriving implications of binding incentive compatibility constraints when it is optimal for the seller to serve some consumers inefficiently. Our results imply that with proper regulation, data-collecting platforms can increase any strongly monotone function of consumer and seller welfare.

---

*Department of Economics, the Pennsylvania State University, University Park, PA 16802 (e-mail: nuh47@psu.edu and rus41@psu.edu). We thank Nageeb Ali, Dirk Bergemann, Wouter Dessein, Harry Di Pei, Piotr Dworzcak, Wioletta Dziuda, Teddy Kim, Stephen Morris, Alessandro Pavan, Eduardo Perez Richet, Andy Skrzypacz, Asher Wolinsky, and Jidong Zhou for very helpful comments and suggestions. We also thank the audiences in various seminars and conferences.

1 Introduction

Technological advances and the explosion of online activity over the last decade have enabled platforms such as Google, Facebook, and Amazon to collect unprecedented amounts of detailed consumer data. Access to such data enables sellers to segment markets in novel ways that may have both positive and negative effects on consumers. On the one hand, consumers may be offered more suitable products; on the other hand, they may be charged higher prices. Unmitigated access to detailed consumer data coupled with the ability to make individual, targeted offers could allow a seller to engage in first-degree price discrimination, which would eliminate all consumer surplus. Policies that prohibit data collection by platforms, bar the provision of such data to sellers, or forbid the use of such data to segment markets would prevent this, but also prevent different market segments from being offered more suitable products. Other policies, which restrict but do not prohibit the collection or usage of consumer data, resolve the apparent trade off between efficiency and consumer surplus in other ways. This raises the possibility that, at least in some markets, some policies could simultaneously lead to increases in efficiency, consumer welfare, and the seller's surplus (relative to the unsegmented market).

We are interested in identifying such markets for a broad range of welfare functions that aggregate consumer utilities, while abstracting from the details of the data collection process and how various regulations or policies are implemented. We consider a setting in which a multi-product monopolistic seller faces a mass of heterogeneous consumers with preferences over subsets of products.[1] A combination of a data collection process and regulations regarding which data the seller can access and use may enable the seller to segment the market into several segments. She can then offer a potentially different menu of products in each market segment, thereby combining second- and third-degree price discrimination.

To accommodate a variety of welfare functions that aggregate consumers' utilities, we consider *Pareto improving* segmentations. We say that a market segmentation is Pareto improving if the surplus that every consumer obtains when choosing from the menu that the seller offers in his segment is no lower than the surplus the consumer would obtain in the unsegmented market, and is strictly higher for some consumers.

[1] Consumer preferences are quasi-linear in money but are otherwise quite general and need not satisfy conditions like increasing differences.

A Pareto improving segmentation increases the value of any welfare function that strictly increases in all consumers' utilities. If a Pareto improving segmentation exists for a market, we say that the market is Pareto improvable. In reduced form, different data collection and usage policies lead to different market segmentations by the seller, ranging from a single segment (the original market) to targeting each consumer individually. The heart of our analysis identifies markets that are both Pareto improvable and for which a Pareto improving segmentation also increases the seller's profit.

Our main finding is that the set of such markets includes every market with a unique profit-maximizing menu, provided that this menu leads to some inefficiency. Moreover, this set of markets is generic (in a sense we formalize in Section 6). This shows that, properly regulated, data-collecting platforms can have unambiguously positive welfare effects in a wide range of market settings. Regulation may take various forms, including restricting some data collection, mandating other data collection, regulating the seller's use of data to segment markets, mandating consumer consent for data collection, and providing legal recourse for consumers hurt by the resulting segmentation and pricing. 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 also show that replacing Pareto improvability with weaker welfare criteria, such as increasing average consumer surplus, does not substantially strengthen our results; the set of Pareto improvable markets is generic, and the set of markets that cannot be segmented in a way that improves average consumer surplus is often non-empty (see the example in Section 3).

We begin by considering environments with two consumer types and any finite number of products. We show that in any such environment Pareto improving segmentations exist for all but a finite number of markets (each market is defined by the proportions of the two types). To prove this result we provide a partial characterization of the profit-maximizing menus by identifying the binding incentive constraint.

We then turn to environments with more than two consumer types. The main difficulty is that with more than two types it is not clear which incentive constraints optimally bind. We therefore develop a different approach. This novel approach is based on understanding the interaction between binding incentive constraints and

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

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 profit in this segment to the profit-maximizing menu in the unsegmented market. Our proof is constructive and shows that the Pareto dominating segment may have to include numerous types. We then show that for every market with a unique profit-maximizing menu, a small perturbation of the market does not change the profit-maximizing menu.

Combining these findings delivers our main result: for any finite number of types and products, every inefficient market with a unique profit-maximizing menu can be segmented into two segments in a way that is Pareto improving and also increases the seller's profit. We also show that the set of such markets is generic in an appropriate sense, which reduces to "all but a finite number of markets" when there are only two types. This provides another proof for our first result.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 provides an illustrative example. Section 4 describes the model. Section 5 provides the two-type analysis. Section 6 studies the case of more than two types and provides the main result. Section 7 concludes. The Online Appendix discusses a single-agent interpretation, environments with a continuum of products, and the case of any number of types with linear valuations.

2 Related literature

Our work connects second- and third-degree price discrimination. The literature that studies third-degree price discrimination and its effects on producer and consumer surplus is broad. Pigou (1920) provides examples in which 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 offer a menu of products. Second, instead of considering expected consumer surplus we use the Pareto criterion. Third, most of the literature assumes that the segmentation is exogenously fixed.

A recent literature on third-degree price discrimination studies surplus across all possible segmentations of a given market. Bergemann et al. (2015) identify the set of producer and consumer surplus pairs that result from all segmentations of a given market. It follows from their analysis 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 (2020) and Hidir and Vellodi (2020) consider maximum consumer surplus when a multi-product seller offers a single product in each market segment. Ichihashi (2020) 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 (2020) 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 papers the seller in our setting may offer multiple products in each market segment. 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). In addition, our focus is on determining whether, for any given market, a segmentation exists that has certain welfare properties. If it does, then regulation may be designed to effectuate a segmentation with these properties.

Our model can also be cast in a Bayesian persuasion framework (Kamenica and Gentzkow, 2011) with a single consumer, the sender, who faces the seller, the receiver (Online Appendix B.1 provides details). The persuasion literature has developed techniques to study such problems with two states (concavification) or a large number of states (Dworczak and Martini, 2019). These techniques are not applicable to our setting for two reasons. First, whereas the usual persuasion settings consider the agent's expected utility, we consider the agent's ex-post utility. Second, and more importantly, these techniques require a specification of the sender's utility for inducing any given posterior. In our setting the consumer's utility depends on the seller's optimal menu, for which no characterization exists when there are multiple products.

3 An Example

To illustrate our framework and results, 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. Such a segmentation could correspond, for example, to a situation in which all consumers in a certain income bracket are of type 2, most of the type 2 consumers are in this income bracket, and the seller is able to price discriminate based on consumers' income bracket. 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 2 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. The seller's profit also increases, so she would rather segment the market in this way than offer a uniform price of 2. 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

[3]The existence of Pareto improving segmentations in this example follows from the analysis of Bergemann et al. (2015), who considered markets with a single product (see Section 4.1). Conversely, since setting segment-specific prices can only benefit the seller, no segmentation can increase even average consumer surplus for any "efficient market." In our example, efficient markets are ones in which type 2 consumers are a minority so the seller optimally sells the product at a price of 1.

and type 2 consumers buy the original product. But 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.

4 Setup

A monopolistic seller faces a continuum of consumers (Online Appendix B.1 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. We will refer to $k = |A| - 1$ as the number of alternatives (excluding the outside option). 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 from a random alternative x and payment p to the seller is $v(t, x) - p$. The valuation for the outside option is 0, $v(t, 0) = 0$. The seller's cost of producing each alternative is normalized to 0 without loss of generality.[5] We assume that each type t has a unique efficient alternative $\bar{a}(t)$ that maximizes the type's valuation over all alternatives. Notice that different consumer types may rank the alternatives differently, and consumers' valuations need not be ordered by their types or satisfy a condition like increasing

[4] This 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.

[5] A non-zero cost $c(a)$ for alternative $a \neq 0$ can be accommodated by redefining valuations as $\tilde{v}(t, a) = v(t, a) - c(a)$ without changing the analysis or results. Notice that $\tilde{v}(t, a)$ may be negative even if all valuations $v(t, a)$ are non-negative. Thus, throughout the paper we allow for negative valuations.

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 two products, 1 and 2, where product 2 has a low-quality version L and a high-quality version H. Suppose that consumers may want to buy one or both products but not both versions of product 2. This setting can modeled by an environment with six alternatives, which correspond to the relevant subsets of $\{1, L, H\}$: $0, \{1\}, \{L\}, \{H\}, \{1, L\}, \{1, H\}$. Alternatively, we could specify an alternative for every subset of $\{1, L, H\}$ and reflect in consumers' types the fact that consumers do not want to buy both versions of product 2. If some consumers demand multiple units of a single product, that would be captured by additional alternatives.

An *allocation rule* $x : T \to \Delta(A)$ 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 t 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 $p : T \to \mathbb{R}$. A 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'. A 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 and price pairs such that 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* $f \in \Delta(T)$ is a distribution over types, where $f(t)$ is the fraction of consumers of type t. An IC-IR mechanism is optimal for a market f if it maximizes the seller's expected revenue among all IC-IR mechanisms. A market may have multiple optimal mechanisms, and a type's surplus may vary across these mechanisms. Thus, to compare consumer surplus across different markets we fix a rule that selects an optimal mechanism for each market. The selection rule should satisfy a mild consistency requirement but is otherwise arbitrary. The requirement is that if two markets have the same set of optimal mechanisms, then the selected mechanisms are the same. We henceforth fix such a selection rule, and refer to the selected optimal mechanism for a market as *the* optimal mechanism for that market. 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, and is inefficient otherwise.

Data about consumers in the market may be collected by a platform and provided to the seller, who may use it to segment the market. Whether and to what extent data can be collected and used may be restricted by various regulations. We model this in reduced form by realizing that any combination of data collection and regulation leads to some a segmentation of the market by the seller. Formally, a *segmentation* of market f is a distribution $\mu \in \Delta(\Delta(T))$ over a finite set of markets that averages to f, that is, $E_{f' \sim \mu}[f'] = f$.[6] 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. Given a segmentation, the seller offers in each market segment the optimal mechanism for that segment.

4.1 Pareto improvements

Our goal is to understand when properly regulated data collection and usage can increase aggregate measures of consumer welfare and the seller's profit. That is, given an environment, which markets can be segmented in such a welfare- and profit-enhancing way. 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 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. Such a

[6] The restriction to a finite set of markets is without loss of generality for all our results and examples since the number of types is finite.

segmentation increases the value of any monotone function of all consumers' surplus. Notice that any segmentation weakly increases the seller's revenue since the seller can use the optimal mechanism for the original market in all segments. The Pareto improving segmentations that we construct also *strictly* increase the seller's revenue, thus ensuring that the seller has a strict incentive to segment the market. This will be the case if the optimal mechanism for the original market is no longer optimal for some segment. The rest of paper focuses on identifying, for every environment, which markets are Pareto improvable, and constructing a Pareto improving segmentation for these markets that also increases the seller's revenue.

We begin by observing that if a market is efficient then it is not Pareto improvable. We already argued that segmenting the market cannot lower the seller's revenue. And the total surplus that any segmentation generates is at most the surplus generated by the efficient allocation. Thus, segmenting an efficient market weakly decreases average consumer surplus, so no segmentation is Pareto improving.

Observation 1 *Any Pareto improvable market is inefficient.*

Which inefficient markets are Pareto improvable? The example in Section 3 showed that in an environment with a single alternative and two types all inefficient markets are Pareto improvable. This is in fact true for all environments with a single alternative and any number of types.

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.[7] However, that proof relies heavily on there being a single alternative and does not generalize to multiple alternatives. As we have seen in Section 3, with more than one alternative not all inefficient markets are Pareto improvable. In addition,

[7]A 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 in the rest of the paper do, and they apply in particular to markets with a single alternative.

the main results of Bergemann et al. (2015) do not generalize to environments with multiple alternatives.[8] So, it is not obvious which markets are Pareto improvable. In the following sections of the paper we study environments with any number of alternatives, first for the case of two types and then for any number of types.

5 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 Section 3 shows, not every inefficient market is necessarily Pareto improvable. However, as we now show, all but at most a finite number of inefficient markets are Pareto improvable by a two-market segmentation. The Pareto improving segmentation also strictly increases the seller's revenue.

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

The first step is to partially characterize optimal mechanisms by identifying the binding IC constraint, which is possible because there are only two types. We do this in the appendix. The characterization shows that in any environment with two types, one of the types can be thought of as the "low type" and the other type as the "high type." In every market, the IC constraint of the low type is slack but the IC constraint of the high type binds. In addition, the surplus of the low type is optimally 0, and the optimal allocation of the high type is efficient. These properties follow from implications of binding IC constraints and are not obvious since we do not assume any ranking over the valuations of the types for different alternatives. Which type is high and which is low is determined as follows. The high type, t_H, is the type whose valuation for the efficient alternative of the other type, t_L, exceeds t_L's valuation for that alternative, that is, $v(t_H, \bar{a}(t_L)) > v(t_L, \bar{a}(t_L))$. There can be at most one such type. If there is no such type, then every market is efficient and both types' surplus is 0 in every market.

The low type's optimal alternative identifies the optimal mechanism, since incentive compatibility then pins down the payment of the high type.[9] The less valuable

[8] Specifically, in many markets with multiple alternatives there exist consumer-producer surplus pairs in the "surplus triangle" of all theoretically possible consumer-producer surplus pairs that are not achievable by any segmentation (see Haghpanah and Siegel, 2021)

[9] The high type must be indifferent between reporting truthfully and misreporting.

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 the high type's surplus. Therefore, the high type's surplus optimally decreases in q. For small enough q, the low type is assigned his efficient alternative, so the allocation is efficient. For high enough q the low type is assigned an alternative that reduces the high type's surplus to 0, so the overall consumer surplus is 0. For intermediate values of q the low type may be assigned an inefficient (and possibly random) 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 is lower on higher intervals. This is what happens in the two-alternative example in Section 3, 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 is lower on higher intervals.[10] This is depicted in Figure 1, (b). These intervals are independent of the selection rule of the optimal mechanism for markets with more than one optimal mechanism. 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 every market: $CS(t_L, q) = 0$ for every q in $[0, 1]$. For the other type, denoted t_H, there exists some $m \leq k$, thresholds $q_0(=0) < \cdots < q_{m+1}(=1)$, and surpluses $\alpha_0 > \cdots > \alpha_m(=0)$ such that $CS(t_H, q) = \alpha_j$ for q in (q_j, q_{j+1}).*

Lemma 1 implies Proposition 2. Indeed, 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

[10]That the number of intervals is $k+1$ is not obvious because in some markets it may be strictly optimal to assign a random alternative to the low type.

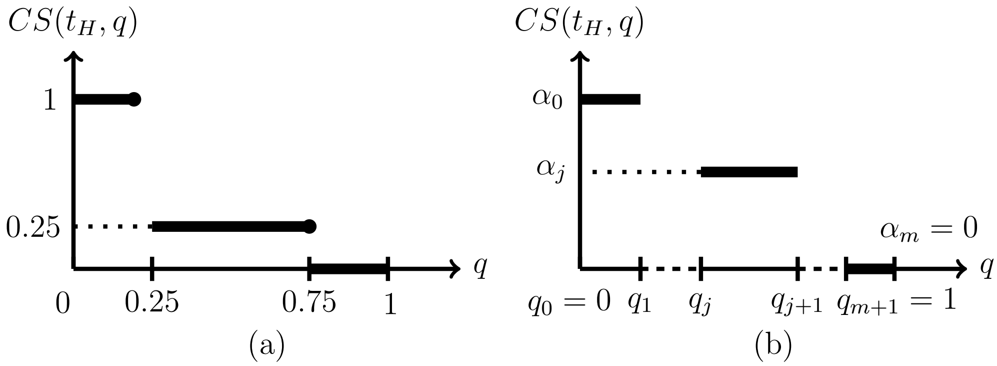

Figure 1: (a) The example from Section 3) with two types t_L and t_H and two alternatives L and H. Market $q = 0.25$ is efficient. Market $q = 0.75$ is the only inefficient market that is not Pareto improvable. (b) The surplus of the high type is constant within each interval, and decreases between intervals.

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 surplus of the seller also increases, since the higher surplus for the high type implies that the seller's revenue from the optimal mechanism on the interior of the lower interval must be different (and therefore higher) from the one on the interior of the higher interval. The endpoints of the intervals $q_1, \ldots, q_m$, however, may not be Pareto improvable since for any segment $q'' > q$ it may be that the surplus of the high type decreases (market q_{m+1} contains only the high type and is efficient). This is the case for market $q = 0.75$ in the two-alternative example in Section 3, which is depicted in Figure 1, (a).[11]

6 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 optimally bind. 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. Consequently, when looking for Pareto improving segmentations for a given market, we cannot a priori restrict attention to two-market segmentations. Moreover, even the set of two-market seg-

[11] Market $q = 0.25$ also has multiple optimal mechanisms that are inefficient. If such a mechanism is selected for market $q = 0.25$, then this market is not Pareto improvable.

mentations 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 segmented into two markets.

Despite these difficulties, we now show that, for any environment, a two-market Pareto improving segmentation exists for any inefficient market that has a unique optimal payment rule.[12] Moreover, we show the set of these markets is generic. Notice that this set of markets is independent of the selection rule of the optimal mechanism.

To make this precise, we define a notion of non-genericity. For every non-zero vector b in $\mathbb{R}^n$, where $n \geq 2$ is the number of types, let $H(b) = \{f \in \Delta(T) : \Sigma_{t \in T} b_t f(t) = 0\}$ 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 the two linear equations $\sum_{t \in T} f(t) = 1$ and $\sum_{t \in T} b_t f(t) = 0$.

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

We can now state the main result of the paper. 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, and has multiple optimal payment rules otherwise.

Theorem 1 *For any environment with $n \geq 2$ types and $k \geq 1$ alternatives, any inefficient market that is not Pareto improvable by a two-market segmentation has multiple optimal payment rules, and the set of such markets 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.*

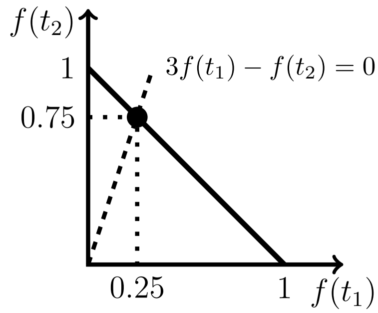

Figure 2: In the example from Section 3), the set of inefficient markets for which no Pareto improving segmentation exists, is the intersection of two sets, each satisfying a linear equality: the set of markets $f(t_1) + f(t_2) = 1$, and the hyperplane $3f(t_1) - f(t_2) = 0$.

Figure 2 illustrates this result for the two-alternative example in Section 3.

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 slightly perturbing a market with a unique optimal payment rule does not change the set of optimal mechanisms, and the set of markets with multiple optimal payment rules is non-generic. 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 with multiple optimal payment rules, market f'' has the same optimal mechanism as f and thus weakly Pareto dominates market f.[13] Therefore, the segmentation of market f into f' and f'' is Pareto improving. The segmentation also strictly increases the seller's revenue because the optimal mechanisms for f are not optimal for f'. We now describe the approach in greater detail.

6.1 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*

[12]Recall that a payment rule maps types to the payment they make.

[13]Recall that if two markets have the same set of optimal mechanisms the selection rule selects an arbitrary but identical optimal mechanism for both.

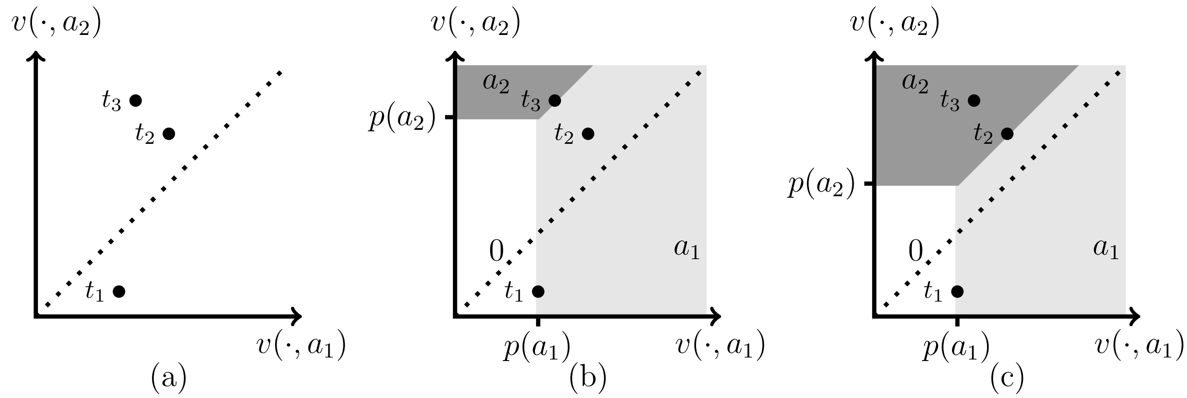

Figure 3: (a) An environment with three types and two alternatives. Panel (b) depicts the optimal prices $p(a_1)$ and $p(a_2)$ for market f, and panel (c) depicts the optimal prices for Pareto dominating market f'. In each panel, types in the lightly shaded region prefer alternative a_1, types in the darkly shaded region prefer alternative a_2, and types in the unshaded region prefer alternative 0.

Pareto dominating market whose support is a subset of the support of f if and only if f is inefficient.

6.1.1 An example

In environments with more than two types, a Pareto dominating market may necessarily include more than two types.[14] We show this in an example with three types and two alternatives, illustrated Figure 3, (a). Alternatives are a_1 and a_2, and types are t_1, t_2, 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.

Figure 3, (b) depicts a mechanism in which the two alternatives are offered at prices $p(a_1) = v(t_1, a_1)$ 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 (the outside option). 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 strictly prefers (and 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

[14]In Online Appendix B.1.1 we identify some environments with more than two types in which two-type dominating markets exist.

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 efficient 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 is it optimal to assign him his efficient 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 efficient 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 is it optimal to assign him his efficient 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, (c). 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 assigned 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 for the seller 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.

6.1.2 Proof of Proposition 3

The proof of Proposition 3 generalizes the idea in the previous example. 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 sufficiently high, it is optimal to assign type t his efficient alternative; this increases the 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 M. We inductively construct a set of types S that contains t such that for every type t' in S there is a directed path of types in S from type t to type t' such that the IC constraint from each type t_j to the next type t_{j+1} in the path binds (that is, type t_j is indifferent between reporting truthfully and misreporting that he is type t_{j+1}). The construction of S 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. Without loss of generality type t is the only type in S that is assigned an inefficient alternative (otherwise denote by t the last type in the path that is assigned an inefficient alternative, and remove from S all types that preceded t in the path). Notice that the payments of the types in S 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 IC and IR).

Now, modify the optimal mechanism M for market 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 mechanism M^1 some type $t' \neq t$ strictly prefers to misreport that he is type t. This type t' is not in S, since in M (and therefore in M^1), every type in S other than type t is assigned his efficient alternative and pays less than type t does in M (payments weakly decrease along the path). 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 $S \cup \{t'\}$, and type t' has strictly higher surplus than in mechanism M. Finally, if t' is not assigned his 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 S is assigned his efficient alternative and pays less than t' does, and thus would not benefit from misreporting that he is type t'. Consider the restricted environment with types $S \cup \{t'\}$. Mechanism M^* is efficient and IC-IR in this environment. Moreover, the surplus of every type in $S \cup \{t'\}$ 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 provide the intuition here and defer the formal proof to the appendix. Such a market can be constructed iteratively. Take the path that defined S and add type t' to its beginning (so type t follows type t'). Begin with a large enough fraction, smaller than 1, 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 6.1.1. Consider the optimal mechanism for market f shown in Figure 3, (b). Figure 4, (a) illustrates the binding IC and IR constraints 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

[15]This construction is the only part of the paper that 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 to argue about which efficient alternative is selected.

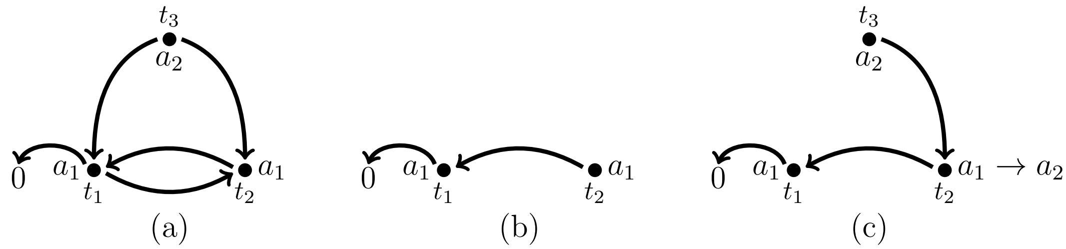

Figure 4: The execution of the procedure that constructs a Pareto dominating market. (a) Binding IC and IR constraints for market f. (b) The path of binding IC constraints that starts from a type with an inefficient allocation, type t_2, and ends with a type whose IR constraint binds, type t_1. (c) The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

in the resulting mechanism can be obtained via the allocation and the binding IC constraints.

6.2 Step 2 - perturbing the market

We now show that for a market with a unique optimal payment rule, perturbing the market leaves the set of optimal mechanisms unchanged. Moreover, the set of such markets is generic. The intuition is as follows. Consider the set P of all payment rules $p : T \to \mathbb{R}$ that are part of IC-IR mechanisms. A mechanism is optimal for a market f if and only 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 uniquely optimal for small enough perturbations of such a market. 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 payment rule, and therefore the optimal mechanism. The set of such markets is non-generic. Before formalizing this discussion, we point out that the intuition and argument are not completely self-evident; we allow for random allocations, of which there is a continuum, and the Online Appendix shows that in some environments with a continuum of allocations the perturbation argument fails. The finite number of alternatives leads to the polytope structure and facilitates the perturbation argument.

To formalize the argument, we say that 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 set of optimal

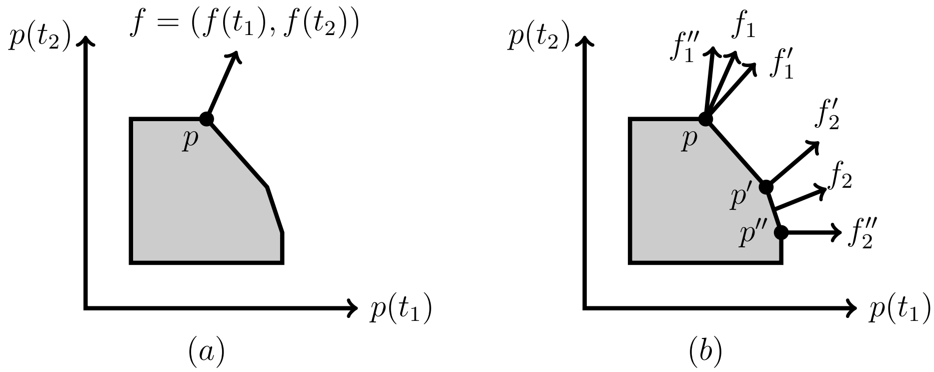

Figure 5: The set of all payment rules P in an environment with two types. (a) The optimal payment rule for market f is maximal in the direction specified by f. (b) f_1 is not orthogonal to a face of P. Any small enough perturbation of f_1 leaves the optimal payment rule unchanged. f_2 is orthogonal to a face of P. Different payment rules (and thus mechanisms) are optimal for different small perturbations of f_2.

mechanisms unchanged if there exists a small enough $\varepsilon > 0$ such that the set of optimal mechanisms for f is equal to the set of optimal mechanisms for any market f' that is ε-close to f. Notice that the consistency requirement from the selection of optimal mechanisms guarantees that the optimal mechanism selected for f is the same as the one selected for f'. Let F_P denote the set of markets f such that perturbing f leaves the set of optimal mechanisms unchanged. The proposition below shows that the complement of F_P is non-generic (see Definition 1).

Proposition 4 *The set of markets $\Delta(T)\backslash F_P$ 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 IC and IR 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 IC-IR mechanisms is a projection of the set of IC-IR mechanisms, and is therefore a polytope in $\mathbb{R}^n$. Consequently, set P has a finite set of vertices.

Lemma 2 *There exists a finite set $P_V \subseteq R^n$ 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 set of optimal mechanisms unchanged, that is, f is in F_P.

Lemma 3 *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. We first show that p is the unique payment rule for markets close to f. Indeed, since the set P_V is finite (by Lemma 2) and p is the unique optimal payment rule, there exists $\delta > 0$ such that $E_{t \sim f}[p(t)] > E_{t \sim f}[p'(t)] + \delta$ for all $p' \in P_V \backslash \{p\}$. By continuity of the expected revenue in f, there exists $\varepsilon > 0$ such that $E_{t \sim f'}[p(t)] > E_{t \sim f'}[p'(t)]$ for all $p' \in P_V \backslash \{p\}$ and all f' that are ε-close to f. Since all payment rules are convex combinations of the payment rules in P_V, we have $E_{t \sim f'}[p(t)] > E_{t \sim f'}[p'(t)]$ for all payment rules $p' \in P \backslash \{p\}$. That is, the payment rule p is also the unique optimal payment rule for all f' that are ε-close to f. Therefore f and any such f' have the same set of optimal mechanisms, since any two IC-IR mechanisms with the same payment rule generate the same revenue. ■

The second step in the proof of Proposition 4 is to show that, generically, a market has a unique optimal payment rule.

Lemma 4 *The set of markets with multiple optimal payment rules 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 exist two payment rules $p \neq p'$ in P_V that are both optimal for f. Thus f 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 with multiple optimal payment rules 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 with multiple optimal payment rules is non-generic. ■

To complete the proof of Proposition 4, note that by Lemma 3, the set F_P contains all markets f with a unique optimal payment rule. By Lemma 4, the set of markets with multiple optimal payment rules is non-generic. Therefore, $\Delta(T) \backslash F_P$ is non-generic.

7 Conclusions

This paper addresses the question of whether regulating data collected and used for market segmentation can increase consumer welfare and the seller's profit. In reduced form this translates to studying the existence of Pareto improving market segmentations, i.e., segmentations in which the surplus of each consumer is weakly higher and the surplus of some consumers is strictly higher than in the unsegmented market. We show that every inefficient market with a unique profit-maximizing menu can be segmented in a way that is Pareto improving and also increases the seller's profit, and that this set of markets is generic. This shows that with proper regulation, data-collecting platforms can lead to unambiguous welfare improvements simultaneously for all consumers and the seller. Our results may also contribute to discussions about regulating sellers, the use of information and ability to price discriminate, and consumers' privacy and control of their data. The results imply that, generically, consumers can provide (or allow to be collected) information to the seller in a way that benefits all consumers and the seller. In particular, the total gain from the increase in allocative efficiency is larger than the seller's gain from improved price discrimination.

To obtain these results we developed a novel methodology to investigate consumer surplus in the optimal menu without requiring a characterization of that menu. The methodology relies on binding incentive compatibility constraints and the implications of the seller optimally not serving some types efficiently. This methodology has the potential to be useful in addressing other mechanism design questions that so far have proven intractable.

Although we place few restrictions on consumers' preferences, an important assumption is that there are finitely many types and alternatives (while allowing for randomized alternatives). The interpretation is that there are substantially more consumers than tastes or products. In the Online Appendix C we discuss environments with possibly infinitely many alternatives, and show that our results do not always hold.[16] We also characterize the Pareto improvable markets in environments with linear valuations.

One important aspect of our analysis is that we do not impose any constraints on the set of segmentations we consider. In reality, the type of information that

[16]We do not explore the case of infinitely many types in general, because it raises additional technical issues. One issue is generalizing our notion of genericity. We leave infinite types for future work.

can be collected and the seller's ability to price discriminate based on the available information may be restricted because of technological and other limitations, and these limitations may vary across settings. Our results establish that in general there is scope for Pareto improvements via segmentation. One direction for future research is to investigate specific settings by identifying and incorporating the limitations they imply.

A Appendix

A.1 Proof of Lemma 1

Proof of Lemma 1. 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 0 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 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.

We first show that 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_2 is strictly positive. Thus neither the IR nor the IC constraint for type t_2 binds, implying that we can increase the payment of type t_2 without violating IC-IR, contradicting optimality.

Since the IC constraint of type t_2 binds, its allocation must be efficient as argued above. Now let (x, p) and (x', p') be optimal mechanisms for markets $q \leq q'$. Then we must have $p(t_1) \geq p(t_1)$ and $p(t_2) \leq p'(t_2)$. Since the allocation of type t_2 is efficient, the surplus of t_2 must be weakly lower in market q' than in market q.

We next 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 left hand side, 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 each type gets his efficient alternative, the IC constraint of type t_1 is slack.

There must exist a threshold τ_2 (possibility equal to 1) such that the surplus of type t_2 is positive in markets $q < \tau_2$, and is zero in markets $q > \tau_2$. Thus in markets less than τ_2 the IR constraint of type t_2 does not bind. Since the IC constraint of type t_1 does not bind by the above argument, the optimal mechanism maximizes

$$(1-q)v(t_1, x(t_1)) + q\Big(v(t_2, \bar{a}(t_2)) - \big(v(t_2, x(t_1)) - v(t_1, x(t_1))\big)\Big), \tag{1}$$

over all distributions $x(t_1)$ over alternatives. Since the objective is linear, the set of markets $[0, \tau_2]$ can be divided into a finite number of intervals, such that the unique optimal mechanism assigns the maximizer of $v(t_1, a) - qv(t_2, a)$ to type t_1 in the interior of the interval, and an optimizer of $v(t_1, a) - qv(t_2, a)$ at the endpoints of the interval. Thus in the interior of each interval, the surplus of type t_2 is constant. ■

A.2 Completing the proof of Proposition 3

We complete the proof of Proposition 3 in two steps. First, we show that mechanism M^* is optimal in the restricted environment, completing the proof of the "if" direction. Second, we prove the "only if" direction.

Lemma 5 *Consider an environment with a set of types $\{t_1, \ldots, t_n\}$, 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 $\{t_1, \ldots, t_n\}$ for which mechanism M^* is the unique optimal mechanism.*

Proof. Consider any IC-IR mechanism (x, p) and any market f. Using the IC and IR constraints, we can write the expected revenue of the mechanism. 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 define $p(t_{n+1}) = 0$, $x(t_{n+1}) = 0$ (i.e., $x(t_{n+1})$ is a deterministic assignment of the outside option), and $F(t_0) = 0$ and write

$$\begin{aligned}
\sum_j p(t_j) f(t_j) &= \sum_j \big(p(t_j) - p(t_{j+1})\big) F(t_j) \\
&\leq \sum_j \Big(v(t_j, x(t_j)) - v(t_j, x(t_{j+1}))\Big) F(t_j) \\
&= \sum_j v(t_j, x(t_j)) F(t_j) - v(t_{j-1}, x(t_j)) F(t_{j-1}). \qquad (2)
\end{aligned}$$

Therefore, for any market f, the revenue of any IC-IR mechanism is at most the maximum of Expression (2) 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 $a \neq \bar{a}(j)$. As a result, for such a market, the unique maximizer of $v(t_j, x) F(t_j) - v(t_{j-1}, x) F(t_{j-1})$ 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 $\{t_1, \ldots, t_n\}$ 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 (2) 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 (2) over all allocation rules. Thus, mechanism M^* is the unique optimal mechanism for market f. ■

Proof of Proposition 3, the "only if" direction. 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

$$CS(t, M) - CS(t', M) \geq v(t, x(t')) - v(t', x(t')). \tag{3}$$

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' his 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'. This is because the IC constraint of all types not in the support of f' can be satisfied by allowing them to 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 $CS(t'', \hat{M}') = CS(t'', M') \geq CS(t'', M)$.

By the construction of the path S, $CS(t', \hat{M}') = CS(t', M') = CS(t', M)$. Thus

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

where the second inequality follows since M is IC-IR and efficient and using Inequality (3). Thus, the incentive constraint (3) 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. Thus 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 (3) 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.3 Proof of Theorem 1

Proof of Theorem 1. We show that if an inefficient market f is not Pareto improvable, then perturbing f does not maintain the optimal mechanism. That is, f is in $\Delta(T) \backslash F_P$. By Proposition 4, the latter set is non-generic. Since a subset of a non-generic set is non-generic, the theorem follows.

Thus 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. ∎

References

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*.

Cowan, S. (2016). Welfare-increasing third-degree price discrimination. *The RAND Journal of Economics*, 47(2):326–340.

Dworczak, P. and Martini, G. (2019). The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048.

Glode, V., Opp, C. C., and Zhang, X. (2018). Voluntary disclosure in bilateral transactions. *Journal of Economic Theory*, 175:652–688.

Haghpanah, N. and Siegel, R. (2021). The limits of multi-product price discrimination. *working paper*.

Hidir, S. and Vellodi, N. (2020). Privacy, personalization and price discrimination. *Journal of the European Economic Association*.

Ichihashi, S. (2020). Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–95.

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.

Schmalensee, R. (1981). Output and welfare implications of monopolistic third-degree price discrimination. *The American Economic Review*, 71(1):242–247.

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

B Online Appendix

B.1 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 selects a mechanism to maximize revenue, and the agent responds by reporting his type optimally.[17] For which prior distributions f does there exist an information disclosure policy that, for each signal, 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. Existing results and techniques in the Bayesian persuasion literature concern the sender's expected utility, whereas our focus is on ex-post utility.[18] 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.

B.1.1 Simple environments

We here prove Proposition 3 for the relatively easy case of *simple environments*, in which (1) there exists a "lowest type" t_L such that $v(t_L, a) < v(t, a)$ for every type $t \neq t_L$ and every alternative $a \neq 0$, and (2) there is a "best alternative" $\bar{a} \neq 0$ such

[17] One 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, 2020 for a discussion).

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

that $v(t, \bar{a}) > v(t, a)$ for every type t and every alternative $a \neq \bar{a}$.[19] Notice that (2) implies that the efficient allocation assigns alternative $\bar{a}$ to every type.[20] We have the following result.

Lemma 6 *For any simple environment and any inefficient market, (1) any optimal mechanism gives type t_L surplus 0, (2) any optimal mechanism assigns the best alternative $\bar{a}$ to some type, denoted $t_{\bar{a}}$, and (3) type $t_{\bar{a}}$ pays strictly more than the valuation $v(t_L, \bar{a})$ of the lowest type for the best alternative.*

Proof. For (1), 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.

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 $\bar{a}$ 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. The resulting mechanism is IC-IR and has strictly higher revenue, 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 the one that assigns alternative $\bar{a}$ to every type. ■

Lemma 6 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

[19] The environment in Section 6.1.1 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).

[20] Notice also that simple environments still do not impose a ranking on all types or all alternatives.

market that consists of a fraction $1-\varepsilon$ of type t_L and a fraction ε of type $t_{\bar{a}}$. If ε is small 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 $(1-\varepsilon)(v(t_L, \bar{a}) - v(t_L, a)) > \varepsilon v(t_{\bar{a}}, \bar{a})$ for ε sufficiently small. 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: type t_L has surplus 0 in both markets by Lemma 6, but type $t_{\bar{a}}$ pays more than $v(t_L, \bar{a})$ in the original market.

C Relaxing the assumption of finite alternatives

Our analysis in the main text assumed that there are finitely many types and alternatives (while allowing for randomized alternatives). In this section we briefly discuss environments with possibly infinitely many alternatives.

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 0. The surplus 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 is 0 for q 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, markets $q < 0.5$ are not Pareto improvable. The set of inefficient markets that are not Pareto improvable, $(0, 0.5)$, is an interval, unlike in two-types environments with finitely many alternatives (Proposition 2). Going back to our two-type analysis in Section 5, however, this should not be surprising. With finitely many alternatives, as the number of alternatives grows, the intervals over which the surplus of the high type is constant shrink, and their number grows. In the limit, the number of intervals, and thus the number of markets for which no Pareto improving segmentations exist, becomes infinite as well.

It is instructive to see why Theorem 1 fails in this example. This is not simply

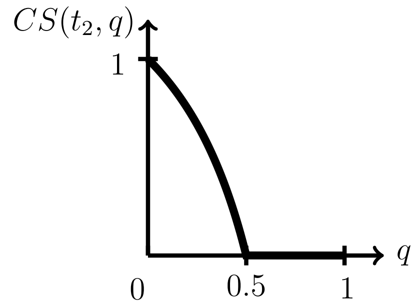

Figure 6: The surplus of type t_2 as a function of the fraction $q = f(t_2)$ of consumers of type t_2.

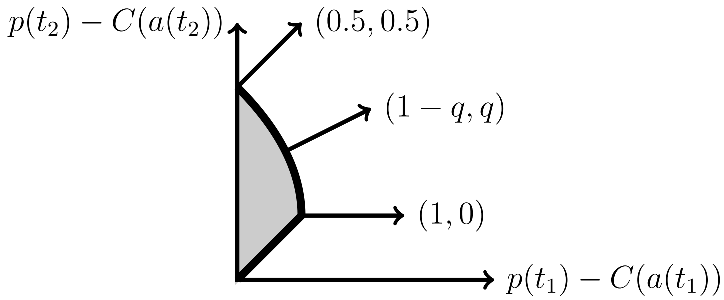

Figure 7: The profit obtained from each type in IC-IR mechanisms. The Pareto frontier is obtained from assigning alternative a to type t_1 and alternative 1 to type t_2, which gives profit $a - a^2/2$ from type t_1 and $1.5 - a$ from type t_2, respectively.

because there are infinitely many alternatives - even in our setting with finitely many alternatives, there are infinitely many randomized alternatives. The reason has to do with properties of the set P of payment rules that are a part of IC-IR mechanisms. With randomization over a finite set of alternatives, the set P is a polytope defined by finitely many halfspaces, as was illustrated in Figure 5. Such a set P has a finite number of vertices, and therefore generically, perturbing a market f leaves the optimal payment rule unchanged. On the other hand, consider the set P of payment rules (net of the cost of production) that are part of IC-IR mechanisms 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.

While Theorem 1 does not apply when there are infinitely many alternatives, key parts of our analysis can be used to provide even stronger results in certain environments with any number (finite or infinite) of alternatives. We demonstrate this in two settings. We first consider linear valuations and provide a tight characterization of Pareto-improvable markets. We then consider environments with two types and characterize all the Pareto-improving segmentations for each market.

C.1 Pareto improvable markets with linear valuations

In this section we characterize Pareto improvable markets in environments with linear valuations, as in Mussa and Rosen (1978). There is a finite set $T = \{1, \ldots, n\}$ 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 each type t has a unique efficient alternative $\bar{a}(t)$, i.e., $\bar{a}(t)$ 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 strictly 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 an inefficient allocation, any small enough (t, δ)-perturbation of f with $\delta > 0$ leaves the optimal mechanism unchanged.

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

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 that 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 an 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'', such that 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. Segment 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 t 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 an 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 this is impossible. The analysis uses the fact that the average of all segments must be equal to the original market. The formal proof is presented below.

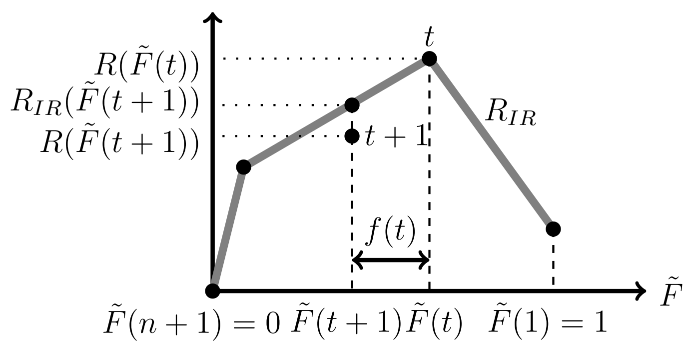

Figure 8: $\phi(t)$ is the slope of the concave hull of the revenue function to the left of t.

To prove Proposition 5 formally, we utilize the standard characterization of optimal mechanisms. The characterization is based on virtual valuations defined as slopes of appropriately defined revenue functions. Define $\tilde{F}(n+1) = 0$. Define the revenue function $R : \{\tilde{F}(1), \ldots, \tilde{F}(n+1)\} \to \mathbb{R}$ 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 $R_{IR} : [0, 1] \to \mathbb{R}$ 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 a type t where $f(t) > 0$ is

$$\phi(t) = \frac{R_{IR}(\tilde{F}(t)) - R_{IR}(\tilde{F}(t+1))}{f(t)}. \tag{4}$$

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

Equation (4) is well defined only if $f(t) > 0$. We allow for segments in which the proportion some type is 0. Thus it is important to extend the definition of virtual valuations to allow for zero proportions. For a type t, let $t_L \leq t$ be the highest type lower than t 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 (4) for types with positive proportion.

The proof of Proposition 5 uses a property of virtual valuations formalized below.

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 $R_{IR}(\tilde{F}(t)) = R_{IR}(\tilde{F}(t+1))$. Thus for such a type,

$$\phi(t)f(t) = R_{IR}(\tilde{F}(t)) - R_{IR}(\tilde{F}(t+1)), \tag{5}$$

since both sides of the above equality are 0 for such a type. Equation (5) also holds for any type t with $f(t) > 0$. This follows directly from Equation (4). We thus use Equation (5) in the proof of Proposition 5 as a property that the virtual valuation of any type t satisfies, whether $f(t) > 0$ or $f(t) = 0$.

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.

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. We show that there exists a (t, δ)-perturbation f' of market f that Pareto dominates f. The lemma below establishes this claim. 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, (a). 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

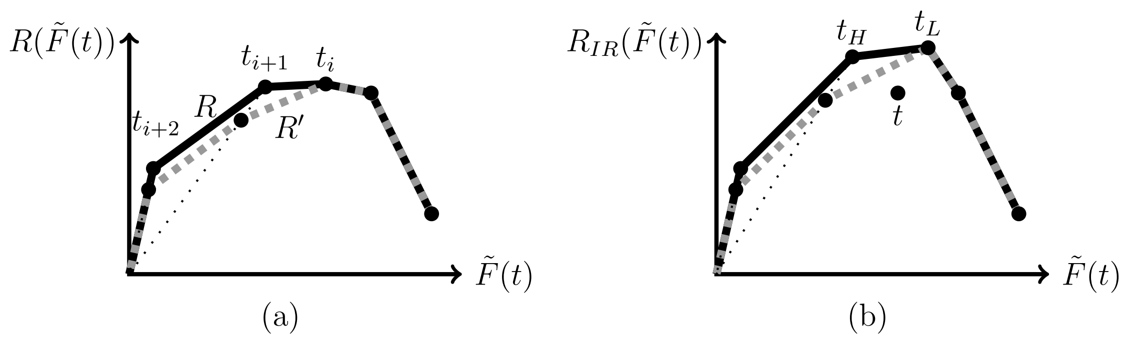

Figure 9: (a) A regular distribution. (b) Type t belongs to the ironed interval $\{t_L, \ldots, t_H\}$. A (t, δ)-perturbation for $\delta < 0$ increases the virtual valuations of types $t_L, \ldots, t_H - 1$, and keeps the virtual valuations of all other types unchanged.

market f. Property 2 of Proposition 6 then implies that all types above t have strictly higher surplus in f' than in f.

Lemma 7 *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 argue that $\phi(t') \leq \phi'(t')$ for all types t'. As a result, the optimal allocations a and a' of market f and f' satisfy $a(t') \leq a'(t')$ for all t'. 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 $v(t)$. Thus by choosing $-1 < \delta < 0$ small enough, the virtual valuation of type t becomes large enough so that $a(t) < a'(t)$. Property 2 of Proposition 6 then implies that all types $t' \leq t$ have weakly higher surplus in f' compared to f, and all types $t' > t$ have strictly higher surplus in f' compared to f. Since $\delta > -1$, all types above t have positive fractions in market f'. Therefore, f' Pareto dominates f.

To complete the proof, we thus argue that $\phi(t') \leq \phi'(t')$ for all types t' and $\phi(t) < \phi'(t)$, for any $\delta < 0$. This can be shown simply with a picture. Let $\{t_L, \ldots, t_H\} \ni t$ be the interval of types over which the revenue function of market f is ironed and to which type t belongs. Formally, $t_L \leq t$ is the highest type t' weakly smaller than t for which $R(t') = R_{IR}(t')$, and $t_H > t$ is smallest type t' larger than t for which $R(t') = R_{IR}(t')$, as depicted in Figure 9, (b). Lowering δ increases the virtual valuations of types $t_L, \ldots, t_H - 1$, and keeps the virtual valuations of all other types unchanged, as claimed. $\blacksquare$

Proof of Proposition 5. 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 7, 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.

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 t be the lowest such type. Let $\{t_L, \ldots, t_H\} \ni t$ be the interval of types over which the revenue function of market f is ironed and to which type t belongs. Formally, $t_L \leq t$ is the highest type t' weakly smaller than t for which $R(t') = R_{IR}(t')$, and $t_H > t$ is smallest type t' larger than t for which $R(t') = R_{IR}(t')$.

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 $a'(t) < a(t)$. We show that some type with positive fraction in f' must be worse off in f' compared to 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 $a'(t) < a(t)$. Now consider the lowest type t' that is higher than t and has positive fraction in market f', $f'(t') > 0$. All types in the set $\{t+1, \ldots, t'-1\}$ have 0 proportion in market f'. Therefore, the virtual valuation of all such types is equal to $\phi'(t)$. Thus we have $a'(t'') = a'(t) < a(t) \leq a(t'')$ for all $t'' \in \{t+1, \ldots, t'-1\}$ by monotonicity of a. Since $a'(t') = a(t')$ for all $t' < t$, 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 $a'(t) \geq a(t)$ for all segments f'. Since type t has a different allocation in some segment compared to f, we have $a'(t) > a(t)$ in some segment.

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

$$E_\mu[\sum_{t':t_L \leq t' < t_H} \phi(t')f'(t')] < E_\mu[\sum_{t':t_L \leq t' < t_H} \phi'(t')f'(t')]. \tag{6}$$

We next show that this is impossible.

Recall from Equation (5) that the virtual valuation of type t satisfies

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

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

$$\begin{aligned}\sum_{t':t_L \leq t' < t_H} \phi(t')f(t') &= \sum_{t':t_L \leq t' < t_H} R_{IR}(\tilde{F}(t')) - R_{IR}(\tilde{F}(t'+1)) \\ &= R_{IR}(\tilde{F}(t_L)) - R_{IR}(\tilde{F}(t_H)) = R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)).\end{aligned} \tag{7}$$

Similar to above, for every segment f' we have

$$\sum_{t':t_L \le t' < t_H} \phi'(t')f'(t') = \sum_{t':t_L \le t' < t_H} R'_{IR}(\tilde{F}'(t')) - R'_{IR}(\tilde{F}'(t'+1))$$
$$= R'_{IR}(\tilde{F}'(t_L)) - R'_{IR}(\tilde{F}'(t_H)). \tag{8}$$

Since R'_{IR} is the concavification of R', it is weakly higher than R', and thus in particular $R'_{IR}(\tilde{F}'(t_H)) \ge R'(\tilde{F}'(t_H))$. 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 t is chosen such that all types lower than t 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 (8), we conclude that

$$\sum_{t':t_L \le t' < t_H} \phi'(t')f'(t') = R'_{IR}(\tilde{F}'(t_L)) - R'_{IR}(\tilde{F}'(t_H)) \le R'(\tilde{F}'(t_L)) - R'(\tilde{F}'(t_H)). \tag{9}$$

To finish the proof, we show that inequalities (6), (7), and (9) cannot simultaneously hold. Indeed, by (7) we have

$$R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)) = \sum_{t':t_L \le t' < t_H} \phi(t')f(t').$$

Since μ is a segmentation of f, $f(t') = E_\mu[f'(t')]$ for all j and thus we have

$$R(\tilde{F}(t_L)) - R(\tilde{F}(t_H)) = E_\mu[\sum_{t':t_L \le t' < t_H} \phi(t')f'(t')]$$
$$< E_\mu[\sum_{t':t_L \le t' < t_H} \phi'(t')f'(t')]$$
$$\le E_\mu[R'(\tilde{F}'(t_L)) - R'(\tilde{F}'(t_H))],$$

where the inequalities followed from (7) and (9). 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. ■

Proposition 5 does not directly identify conditions under which a (t, δ)-permutation of a market leaves the optimal mechanism unchanged. The conditions can in fact 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 this, 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 thus 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 an 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 an 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 and the set of alternatives is $A = [0, \infty)$, then any $a > 0$ is a maximizer of virtual surplus for a unique virtual valuation. 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 and $A = [0, \infty)$, then a market is Pareto improvable if and only if in the optimal mechanism* *some type is excluded.*

Corollary 3, which applies to the example at the beginning of Appendix C, identifies some environments in which the set of inefficient markets that are not Pareto improvable is not generic. In other environments Pareto improving segmentations exist for generic sets of inefficient markets. For example 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, since Theorem 1 only provides the "if" direction.

Pareto improving segmentations also exist for generic markets when the cost function is piecewise linear. Then, similarly 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.*