Pareto Improving Segmentation of Multi-product Markets

Nima Haghpanah and Ron Siegel*

July 30, 2019

Abstract

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

1 Introduction

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

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

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

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

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

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

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

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

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

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

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

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

[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.

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

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

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

The assumption of a finite number of types and products is important for our main result. With a continuum of products, it may be that for every market a small perturbation changes the profit-maximizing menu. Consequently, in some environments with a continuum of products, there may be infinitely many markets for which no Pareto improving segmentation exists. However, using a two-type example, we show that an appropriate generalization of our results holds, and even becomes stronger, with a continuum of products. Namely, a Pareto improving segmentation exists for a market if and only if increasing the proportion of the consumers of the higher type leaves the optimal mechanism unchanged.

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

2 Related literature

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

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

Our work is also related to Roesler and Szentes (2017) and Condorelli and Szentes (2018) who consider a consumer who learns about his own preferences. Roesler and Szentes (2017) assume that the distribution of the consumer's valuations is fixed, but the consumer can choose to learn a noisy signal about her valuations. They show that the optimal learning structure for the buyer achieves efficient trade. Condorelli and Szentes (2018) study a bilateral trade setting in which the consumer can choose any distribution for her valuations at no cost prior to trade. They show that trade is ex post efficient.

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

Our analysis can also be cast in a Bayesian persuasion framework (Kamenica and Gentzkow, 2011). As we discuss in Section 3.2, the market is replaced by a single agent (the sender) who faces the seller (the receiver). A market segmentation corresponds to a

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

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

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

Dworczak and Martini (2019) study a Bayesian persuasion problem with a large number of states, where the geometric approach of concavification is of limited power. They use duality techniques to provide a price-theoretic interpretation of the problem. In their setting the value of the sender for inducing any given posterior distribution can be calculated straightforwardly. In our setting, on the other hand, calculating the value requires characterizing optimal screening mechanisms, for which the current mechanism design tools are insufficient. We develop tools to overcome this difficulty.

3 Setup

A monopolistic seller faces a continuum of consumers (Section 3.2 discusses the interpretation of a single consumer). The *environment* includes a finite set I of consumer types 1 to n and a finite set A of alternatives 0 to $k \geq 1$, where alternative 0 is consumers' outside option of not purchasing from the seller. We will refer to k as the number of alternatives. A consumer type specifies a valuation for every alternative: type i's valuation for alternative a is $v(i,a)$. Type i's valuation for a random alternative $x \in \Delta(A)$ is $v(i,x) = E_{a \sim x}[v(i,a)]$. Type i's surplus for random alternative x and payment p to the seller is $v(i,x) - t$. Assume that the value for the outside option is zero, $v(i,0) = 0$. We denote by $\bar{a}(i)$ the alternative with the highest valuation for type i, and refer to it as the efficient alternative for type i. The seller's cost of producing each alternative is normalized to 0 without loss of generality.[7] Notice that place no restrictions on consumers' valuations: different consumer types may rank the alternatives differently, and consumers' valuations need not be ordered by their types or satisfy a condition like increasing differences.

Each alternative $a \neq 0$ corresponds to a product or a set of products. This cap-

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

tures horizontal and vertical differentiation, allows for multi-unit demand, and accommodates bundling. To illustrate this, suppose that the seller can produce a product p and a second product that has a low-quality version q_L and a high-quality version q_H. Suppose that consumers may want to buy one or both products but not both versions of the second product. The alternatives correspond to the relevant subsets of $\{p, q_L, q_H\}$: $\phi, \{p\}, \{q_L\}, \{q_H\}, \{p, q_L\}, \{p, q_H\}$. Alternatively, we could specify an alternative for every subset of $\{p, q_L, q_H\}$ and reflect in consumers' types the fact that consumers do not want to buy both versions of the second product. If some consumers demand multiple units of a single product, that would be captured by additional alternatives.

An *allocation rule* $x : I \to \Delta(A)$ is a mapping from types to random alternatives, where $x(i)$ is the allocation of type i. The allocation rule is efficient if the allocation of each type is equal to its efficient alternative: $x(i) = \bar{a}(i)$ with probability one. A (direct) *mechanism* $M = (x, p)$ consists of an allocation rule x and a payment rule $p : I \to \mathbb{R}$. The interpretation is that the consumer reports his type and receives the corresponding random alternative in return for the specified payment. The mechanism is incentive compatible (IC) if no type benefits from misreporting. The mechanisms is individually rational (IR) if every type obtains at least 0 by reporting truthfully. Any mechanism we will refer to will be IC-IR unless otherwise stated. A mechanism is efficient if it's allocation is efficient. Every mechanism can be represented by a menu of (random alternative, price) pairs, where each type chooses a pair that maximizes its utility. If a type is indifferent between two or more pairs, it chooses the one with a higher price (or uses any rule if the prices are the same).

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

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

3.1 Pareto improvements

Our goal is to understand for each environment which markets can be segmented in a way that benefits all consumers. To formalize this, we say that market f' *weakly Pareto dominates* market f if every type i in market f' prefers the optimal mechanism for market f' to the one for market f', that is, $CS(i, f') \geq CS(i, f)$ for all types i such that $f'(i) > 0$. If, in addition, the preference is strict for some type i with $f'(i) > 0$, then we say that f' Pareto dominates f. A segmentation μ of market f is *Pareto improving* if every segment weakly Pareto dominates f and some segment Pareto dominates f. A market is *Pareto improvable* if it has a Pareto improving segmentation.

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

Observation 1 *Any Pareto improvable market is inefficient.*

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

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

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

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

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

3.2 Single-agent interpretation

Consider an agent whose type is drawn from the set I according to a prior distribution f. Before learning his type, the agent commits to an information disclosure policy, which maps every type in I to a distribution over signals. The seller observes the policy and the realized signal and forms a posterior f' over the agent's type. The seller then offers a mechanism (or menu) to maximize his revenue, and the agent responds optimally.[10] For which prior distributions f does there exist an information disclosure policy that increases the agent's ex-post utility relative to a policy that discloses no information? This model and question are equivalent to those described earlier. Following Kamenica and Gentzkow (2011), we can describe the process as the agent choosing a distribution μ over posteriors f' that averages to f, that is, $E_{f' \sim \mu}[f'] = f$.

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

[10]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 (2018) for a discussion).

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

4 Environments with two types

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

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

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

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

The higher the fraction q of the high type in the market, the more important is the reduction in their surplus. Therefore, the high type's surplus optimally decreases in q. For small enough q, the low type is allocated his efficient alternative, so the allocation is efficient; for large enough q the low type is allocated an alternative that reduces the high type's surplus to 0, so the overall consumer surplus is 0.[13] For intermediate values of q the

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

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

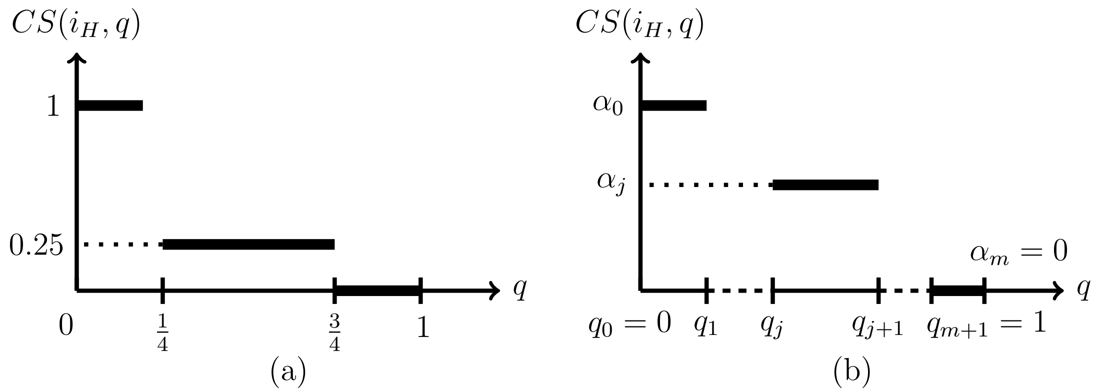

Figure 1: (a) The example from the introduction with two types i_L and i_H and two alternatives L and H, where $v(i_L, L) = 0.75, v(i_L, H) = 1, v(i_H, L) = 1, v(i_H, H) = 2$. (b) The surplus of the high type is constant within each interval, and decreases between intervals.

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

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

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

Lemma 1 *Consider an environment with two types. For one of the types, denoted i_L, the surplus is 0 in any market: $CS(i_L, q) = 0$ for any q in $[0, 1]$. For the other type, denoted i_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 $CS(i_H, q) = \alpha_j$ for q in (q_j, q_{j+1}).*

Lemma 1 implies Proposition 2. The idea is that all the efficient markets $q < 1$ are in the first interval $[q_0, q_1]$, where the fraction of low types is sufficiently high to make serving them efficiently optimal for the seller. For any other interval $[q_j, q_{j+1}]$ $(1 \leq j \leq m)$, any market q in the interior of the interval can be segmented into two segments $q' < q''$ such that q'' is also in the interior of the interval $[q_j, q_{j+1}]$, so the surplus of the high type is unchanged, and q' is in the interior of the lower interval $[q_{j-1}, q_j]$, so the surplus of the high type is increased. This shows that the segmentation into q' and q'' is Pareto improving. The endpoints of the intervals, however, may not be Pareto improvable since for any segment $q'' > q$ it may be that the surplus of the high type is decreased. This is the case for market $q = 0.75$ in the two-alternative example in the introduction, which is depicted in Figure 1, (a).

5 Environments with any number of types

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

Despite these difficulties, we now show that, for any environment, a two-market Pareto improving segmentation exists for "almost all" inefficient markets. To make this precise, we define a notion of non-genericity below. For every non-zero vector b in $\mathbb{R}^{|I|}$, where $|I| \geq 2$ is the number of types, let $H(b) = \{f \in \Delta(I) : \Sigma_{i \in I} b_i f(i) = 0\}$ be the set of

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

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 two independent linear equations.

Definition 1 *Given an environment with n types, a set F of markets is* non-generic *if it is contained in finite number 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. A two-market Pareto improving segmentation exists for almost all inefficient markets, in the sense that the set of inefficient markets for which no such segmentation exists is non-generic.

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

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

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

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

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

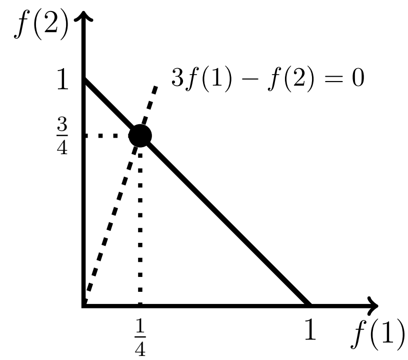

Figure 2: In the example from the introduction, the set of inefficient markets for which no Pareto improving segmentation exists, is the intersection of two sets, each satisfying a linear equality: the set of markets $f(1) + f(2) = 1$, and the hyperplane $3f(1) - f(2) = 0$.

f; thus, as long as market f does not belong to the set of non-generic markets, market f'' weakly Pareto dominates market f. Therefore, the segmentation of market f into f' and f'' is Pareto improving. We now describe the approach in greater detail.

5.1 Step 1 - constructing a Pareto dominating segment

The first step is formalized as follows.

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

5.1.1 Simple environments

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

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

Proof. Suppose that the optimal mechanism gives type i_L a strictly positive surplus. Then every type has a strictly positive surplus, since every type i can report he is type i_L and obtain surplus $v(i, x(i_L)) - p(i_L) \geq v(i_L, x(i_L)) - p(i_L) > 0$. This contradicts optimality since a mechanism in which all payments are uniformly and slightly increased is IC-IR, proving (1).

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

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

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

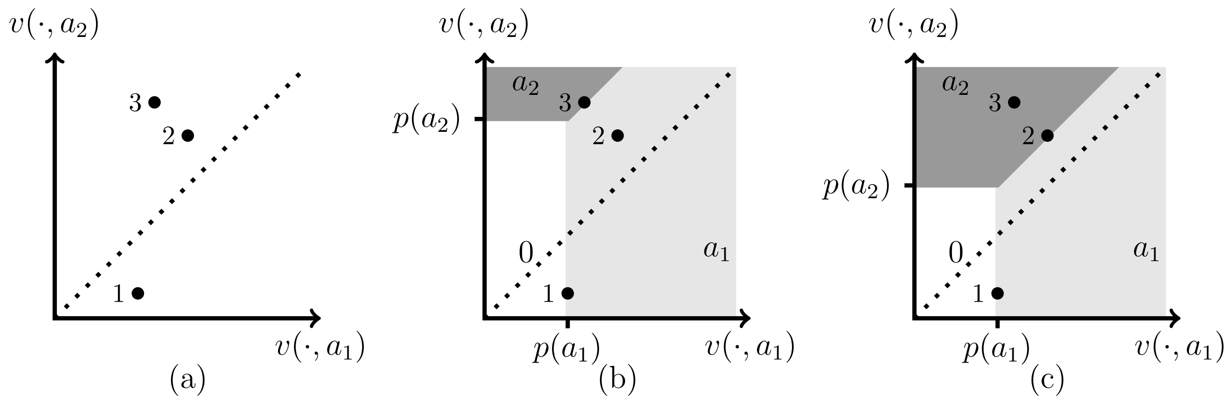

Figure 3: (a) An environment with three types and two alternatives. The environment is not simple since types have different efficient alternatives. Panel (b) depicts the optimal prices $p(a_1), p(a_2)$ for market f, and panel (c) depicts the optimal prices for Pareto dominating market f'. In each panel,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

5.1.2 General environments, an example

A two-type Pareto dominating market does not always exist for inefficient markets in environments that are not simple. We show this in an environment with three types and two alternatives, illustrated Figure 3, (a). Alternatives are a_1 and a_2, and types are $1, 2$, and 3. Each type is described by the dot with the type's number to its left (the horizontal axis shows the valuation for alternative a_1, and the vertical axis shows the valuation for alternative a_2). This environment is not simple: there is a lowest type (type 1) but not a best alternative (type 1 prefers alternative 1 but types 2 and 3 prefer alternative 2).

Figure 3, (b) describes a mechanism in which the two alternatives are offered at prices $p(a_1)$ and $p(a_2)$. The prices are $p(a_1) = v(1, a_1)$ and $p(a_2) = v(2, a(2)) - v(2, a(1)) + v(1, a(1))$. At these prices, the lightly shaded region contains the set of types that prefer alternative a_1, the darkly shaded region contains the set of types that prefer alternative a_2, and the unshaded region contains the set of types that prefer alternative 0. In particular, type 1 is indifferent between alternative a_1 and the outside option (and chooses a_1), and type 3 is indifferent between a_1 and a_2 (and chooses a_2). Type 2 has a higher marginal value for a_1 compared to a_2 than does type 3, and therefore type 2 chooses alternative a_1. This mechanism gives surplus 0 to type 1 and strictly positive surplus to the other types. The mechanism is inefficient, since type 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 1 is large enough that it is optimal to assign him his preferred alternative a_1 for a price that is equal to his valuation; among the remaining consumers the fraction of type 3 is large enough that is it optimal to assign him his preferred alternative a_2 for the maximal price that maintain IC.

There is a market f' that contains all three types that Pareto dominates f. In this market the fraction of type 1 is large enough that it is optimal to assign him his preferred alternative a_1 for a price that is equal to his valuation; among the remaining consumers the fraction of type 2 is large enough that is it optimal to assign him his preferred alternative a_2 for the maximal price that maintains IC. Type 3 is also assigned alternative a_2 for this price. This mechanism is illustrated in Figure 3, (b). Since the price of alternative 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 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 2 either the allocation and surpluses of the other types is unchanged or the surplus of type 3 is 0; in any market without type 3 either the surplus of type 1 is 0 and the surplus of type 2 is strictly lowered (he is allocated alternative a_2 for the maximal price that maintains IC) or the surplus of type 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 3 (who is already assigned his efficient alternative a_2 in market f), type 2 must be present in sufficient proportion to make it optimal to lower the price of alternative a_2 in order to extract more surplus from type 2. But type 2's surplus in market f is positive; in order to maintain this surplus in the Pareto dominating market, type 1 must be present in sufficient proportion to make it optimal for the seller to assign alternative a_1 to type 1, thereby providing information rents to type 2.

5.1.3 General environments, the general construction

The proof of Proposition 3 generalizes the idea from the construction discussed above. In an inefficient market f, some type i is assigned an inefficient alternative. The reason for this inefficiency is to lower the surplus (information rents) of some other type i'. In a new market that includes only type i and i' and in which the proportion of type i is much higher than that of type i' it is optimal to assign type i his efficient alternative; this increases the marginal surplus that type i' obtains from being able to mimic type i. But the surplus of type i may decrease; to prevent this, we identify an "information rents path" in market f that begins with type i and ends with some type i'' 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 i be some type that is assigned an inefficient alternative in the optimal mechanism. We will inductively construct a set of types S that contains i such that for every type i' in S there is a directed path of types in S from type i to type i' in which the IC constraint from each type i_j to the next type i_{j+1} in the path binds (type i_j is indifferent between reporting truthfully and misreporting that he is type i_{j+1}). The construction stops when a type that has surplus 0 is added to S. If type i has surplus 0 we are done. Otherwise, given the set S so far constructed, there is a type j not in S such that the IC constraint from some type in S to type j 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 i and ends with the type that has surplus 0. Suppose without loss of generality that type i is the only type in the path that is assigned an inefficient alternative (otherwise denote by i the last type in the path that is assigned an inefficient alternative). Denote by T the set of types in the path, and notice that the payments of the types in T weakly decrease along the path (otherwise the revenue in market f can be increased by replacing a type's assigned alternative and payment with those of the next type in the path).

Now, modify the optimal mechanism M for f by assigning type i his efficient alternative and increasing his payment to leave his surplus unchanged. The modified mechanism M' violates IC, otherwise mechanism M' would generate more revenue than mechanism M in market f. Therefore, when faced with the modified mechanism some type $j \neq i$ strictly prefers to misreport that he is type i. Type j is not in T, since every type in T other than i is assigned his efficient alternative and is paying less than type i does in M (and therefore in M'). Modify mechanism M' by replacing type j's assigned alternative and payment with type i's. This modified mechanism M'' satisfies IC and IR for the set of types $T \cup \{j\}$, and type j has strictly higher surplus than in mechanism M. Finally, if j is not assigned the efficient alternative in mechanism M'', modify M'' by assigning type j his efficient alternative and increasing his payment to leave his surplus unchanged. Denote

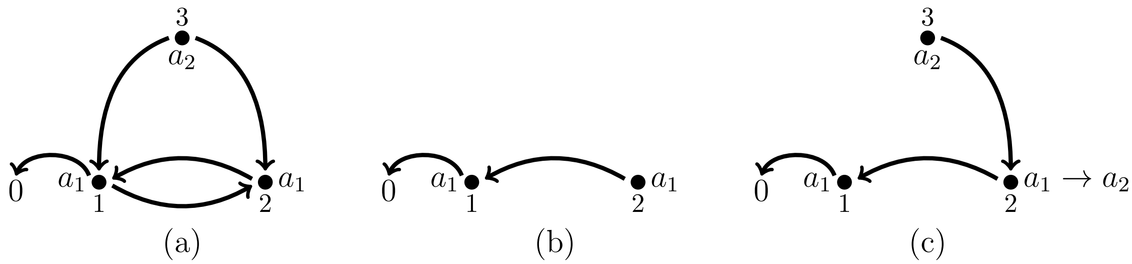

Figure 4: (a) The execution of the procedure that constructs a Pareto dominating market for market f in Section 5.1.2. (a) Binding IC and IR constraints for market f. (b) The path of binding IC constraints that starts from a type with inefficient allocation, type 2, and ends with a type with binding IR constraint, type 1. (c) The appended path with type 3 who strictly benefits in the Pareto dominating market.

the resulting mechanism by M^*. Consider the restricted environment with types $T \cup \{j\}$. Mechanism M^* is efficient and IC-IR in this environment, the surplus of every type in $T \cup \{j\}$ is weakly higher than in mechanism M, and the surplus of type j is strictly higher.

It remains to show that M^* is the (unique) optimal mechanism for some full-support market in the restricted environment. This is true because M^* is efficient. Such a market can be constructed iteratively. For this take the path that defined T and add type j to its beginning (so type i follows type j). Begin with a large enough fraction of the last type in the path so that it is strictly optimal for the seller to assign this type his efficient alternative for a price that is equal to his valuation. Add a large enough fraction of the second-to-last type in the path so that it is strictly optimal for the seller to assign this type his efficient alternative for the maximal price that maintains IC, etc. The resulting mechanism is M^*, which is the unique optimal mechanism for the resulting market. This proves Proposition 3.

Let us apply the procedure to the example in Section 5.1.2. Consider the optimal mechanism for market f shown in Figure 3, (b). Figure 4, (a) illustrates the binding IC and IR constraints in the optimal mechanism. Beginning with type $i = 2$, we have $S = \{1, 2\}$. The binding IC path in S that begins with type 2 and ends with type 1, whose surplus is 0, is depicted in Figure 4, (b). The appended path with type $j = 3$ at its beginning and the modifications to the optimal mechanism for the original market are illustrated Figure 4, (c). The resulting mechanism is the one in Figure 3, (c).

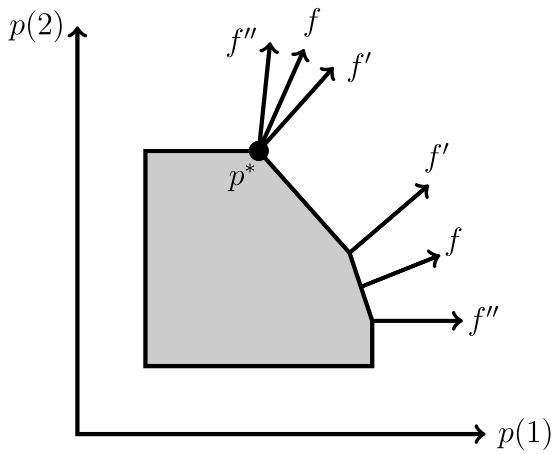

Figure 5: The set P of payment rules that are part of some IC-IR mechanism is a polytope. The optimal $p^* \in P$ for a market f maximizes $p(1)f(1)+p(2)f(2)$. If f is not orthogonal to a face of P, then any small enough perturbation of f leaves the optimal mechanism unchanged. If f is orthogonal to a fact of P, then different mechanisms are optimal for different perturbations.

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

In this section we show that for almost all inefficient markets, the optimal mechanism remains optimal if the market is slightly perturbed. The intuition is as follows. Consider the set P of all payment rules $p : X \to \mathbb{R}$ that are part of some IC-IR mechanism. A mechanism is optimal for a market f if its payment rule p is maximal in P in the direction specified by f, as depicted in Figure 5. We show that P is a polytope. Therefore, the payment rule that maximizes the expected revenue is unique and is a vertex of P for almost all markets. The same payment rule remains optimal for small enough perturbations of such markets. In contrast, if f is orthogonal to a face of P, then any small perturbation of the market may result in a change in the optimal mechanism. We formalize this discussion below.

The main result of this section states that, generically, perturbing a market leaves the optimal mechanism unchanged. Formally, markets f and f' are ε-close if $|f(i) - f'(i)| \leq \varepsilon$ for all types i. We say that perturbing market f leaves the optimal mechanism unchanged if there exists small enough ε such that the optimal mechanism for f is also optimal for all markets f' that are ε-close to f. Let $\mathcal{F}_P$ denote the set of markets f such that perturbing f leaves the optimal mechanism unchanged. The proposition below shows that $\mathcal{F} \backslash \mathcal{F}_P$ is non-generic (see Definition 1).

Proposition 4 *The set of markets $\mathcal{F} \backslash \mathcal{F}_P$ is non-generic.*

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

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

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

Lemma 4 *It a market f has a unique optimal payment rule, then f is in $\mathcal{F}_P$.*

Proof. Consider a market f with a unique optimal payment rule p. Since the set $\mathcal{E}_P$ is finite (by Lemma 3) and p is the unique optimal payment rule, there exists $\delta > 0$ such that $E_{i\sim f}[p(i)] \geq E_{i\sim f}[p'(i)] + \delta$ for all $p' \in \mathcal{E}_P$. By continuity of the expected revenue in f, there exists $\varepsilon > 0$ such that $E_{i\sim f'}[p(i)] \geq E_{i\sim f'}[p'(i)]$ for all $p' \in \mathcal{E}_P$ and all f' that are ε-close to f. Since all payment rules are convex combinations of the payment rules in $\mathcal{E}_P$ (Lemma 3), we must have $E_{i\sim f'}[p(i)] \geq E_{i\sim f'}[p'(i)]$ for all payment rules $p' \in P$, that is, the payment rule p is also optimal for all f' that are ε-close to f. ■

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

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

Proof. Consider a market f for which more than one payment rule in P maximizes revenue. Since $\mathcal{E}_P$ is the set of extreme points of P, there must exist $p, p' \in \mathcal{E}_P, p \neq p'$ that are optimal for f. Thus such a market is contained in a hyperplane $H_{p,p'}$ defined by the equation $\sum_i f_i(p(i) - p'(i)) = 0$. Since $\mathcal{E}_P$ is finite, the set of markets f without a unique optimal payment rule is contained in a finite union of hyperplanes, one for each pair of payment rules in $\mathcal{E}_P$, i.e., $\cup_{p,p' \in \mathcal{E}_P} H_{p,p'}$. Thus, by Definition 1, the set of markets without a unique revenue maximizing payment rule is non-generic. ■

To complete the proof of Proposition 4, note that by Lemma 4, the set $\mathcal{F}_P$ contains all markets f with a unique optimal payment rule. By Lemma 5, the set of markets without a unique optimal payment rule is non-generic. Therefore, $\mathcal{F} \setminus \mathcal{F}_P$ is generic.

6 Environments with infinitely many alternatives

So far we have assumed that there are finitely many types and alternatives. The interpretation is that there are substantially more individuals than tastes or products. In this section we briefly discuss environments with infinitely many alternatives.[15]

Theorem 1, which states that Pareto improving segmentations exist for almost all markets, 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 1 and 2 and have values a and $2a$ for quality a, respectively. A market is identified by the proportion q of type 2. In any market, the surplus of type 1 in the optimal mechanism is zero. The surplus of type 2 in the optimal mechanism is illustrated in Figure 6. It strictly decreases in q if q is smaller than $\frac{1}{2}$, and remains constant at 0 if q is larger than $\frac{1}{2}$. Any non-trivial segmentation of a market $q < \frac{1}{2}$ must include a segment $q' > q$ in which type 2 is strictly worse off. Therefore, no Pareto improving segmentation exists for a market $q < \frac{1}{2}$. The set of inefficient markets for which no Pareto improving segmentations exists is not finite, unlike two-types environments with finitely many alternatives (Proposition 2).

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

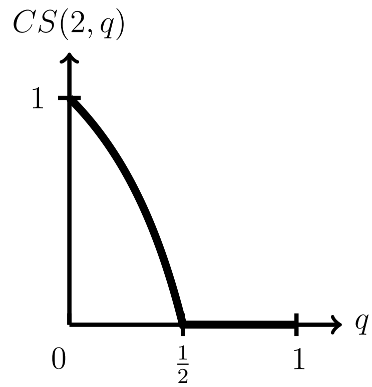

Figure 6: Surplus of type 2.

While Theorem 1 no longer holds, an appropriate reformulation of our analysis holds, and even becomes stronger, for the two-type example discussed above. In particular, recall that with finitely many alternatives, a Pareto improving segmentation exists for an inefficient market if perturbing the market leaves the optimal mechanism unchanged. In the two-type example, a Pareto improving segmentation exists for an inefficient market if, and only if, perturbing the market in a certain way leaves the optimal mechanism unchanged. Indeed, notice that for any market in $[\frac{1}{2}, 1)$, increasing the proportion of type 2 leaves the optimal mechanism unchanged. Therefore, any such market q can be segmented in a Pareto improving way into two segments: a segment in which the fraction of type 2 is slightly higher than q, and another segment in which the fraction of type 2 is less than $\frac{1}{2}$. Since we already argued that no Pareto improving segmentation exists for inefficient markets $(0, \frac{1}{2})$, we conclude that a Pareto improving segmentation exists for an inefficient market q if and only if $q \in [\frac{1}{2}, 1)$, that is, increasing the proportion of type 2 leaves the optimal mechanism unchanged. Notice that an arbitrary perturbation of market $\frac{1}{2}$ may change the optimal mechanism. Nevertheless, increasing the proportion of type 2 does not change the optimal mechanism, and a Pareto improving segmentation exists.

Even though Theorem 1 may fail in some environments with infinitely many alternatives, it holds in certain other environments with infinitely many alternatives. In fact, notice that in our setting with finitely many alternatives, there are infinitely many randomized alternatives. Therefore, Theorem 1 holds with infinitely many random alternatives. What is the difference between environments with randomizations over finite alternatives and the two-type example discussed above? The distinction can be clarified by considering

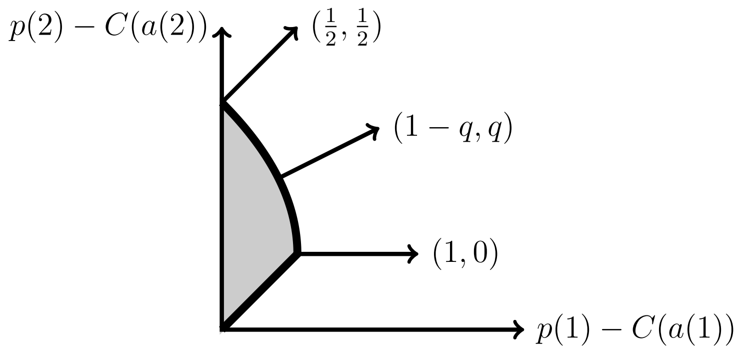

Figure 7: The profit obtained from each type in IC-IR mechanisms. The Pareto frontier is obtained from assigning alternative a to type 1 and alternative 1 to type 2, which gives profit $a - a^2/2$ from type 1 and $3/2 - a$ from type 2, respectively.

the set P of payment rules that are a part of some IC-IR mechanism. With randomization over finite alternatives, the set P is a polytope defined by finitely many halfspaces, as illustrated in Figure 5. Such a set P has a finite number of vertices, and therefore generically, perturbing a market f leaves the optimal mechanism unchanged. On the other hand, consider the set P of payment rules that are part of some IC-IR mechanism for the two type example discussed above, illustrated in Figure 7 (net of the cost of production). Even though P is convex, it has infinitely many vertices. Perturbing any market in which the proportion of type 2 is less than $\frac{1}{2}$ 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.

7 Discussion and conclusions

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

Our proof is constructive. We construct a segmentation that consists of two segments. The first segment Pareto dominates the unsegmented market. It is constructed by following a path of information rents from a type with an inefficient allocation, to a type with zero information rents. By decreasing the proportion of the types that gain information rents from being able to mimic a given type, the allocation of that type becomes more efficient. We show that, generically, perturbing a market leaves the optimal mechanism unchanged. As a result, if the Pareto dominating market is small enough, the optimal mechanism for the residual market will be identical to that of the unsegmented market.

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

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

References

Adams, W. J. and Yellen, J. L. (1976). Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics*, pages 475–498.

Aguirre, I., Cowan, S., and Vickers, J. (2010). Monopoly price discrimination and demand curvature. *American Economic Review*, 100(4):1601–15.

Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, 105(3):921–57.

Braghieri, L. (2017). Targeted advertising and price discrimination in intermediated online markets. *working paper*.

Chen, X., Diakonikolas, I., Orfanou, A., Paparas, D., Sun, X., and Yannakakis, M. (2015). On the complexity of optimal lottery pricing and randomized mechanisms. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1464–1479. IEEE.

Condorelli, D. and Szentes, B. (2018). Information design in the hold-up problem. *working paper*.

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

Crawford, G. S. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, 102(2):643–85.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. (2014). The complexity of optimal mechanism design. In *SODA*. To appear.

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

Dworczak, P. and Martini, G. (2019). The simple economics of optimal persuasion. *working paper*.

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

Hart, S. and Nisan, N. (2013). The menu-size complexity of auctions. *arXiv preprint arXiv:1304.6116*.

Hidir, S. and Vellodi, N. (2018). Personalization, discrimination and information revelation. *working paper*.

Ho, K., Ho, J., and Mortimer, J. H. (2012). The use of full-line forcing contracts in the video rental industry. *American Economic Review*, 102(2):686–719.

Ichihashi, S. (2018). Online privacy and information disclosure by consumers. *working paper*.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Pigou, A. (1920). *The economics of welfare*. London: Macmillan.

Robinson, J. (1969). *The economics of imperfect competition*. Springer.

Roesler, A.-K. and Szentes, B. (2017). Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–80.

Salinger, M. A. (1995). A graphical analysis of bundling. *Journal of Business*, pages 85–98.

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

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

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

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

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

8 Appendix

8.1 Optimal Mechanisms For Two Types

We characterize optimal mechanisms for two types 1 and 2 by considering three possible cases. (I) $v(1, \bar{a}(1)) \geq v(2, \bar{a}(1))$ and $v(1, \bar{a}(2)) \leq v(2, \bar{a}(2))$. In this case, extracting the full surplus is feasible since each type can be assigned its favorite alternative at price equal to its willingness to pay, without violating IC or IR. The remaining two cases consider $v(1, \bar{a}(1)) < v(2, \bar{a}(1))$ (the case where $v(1, \bar{a}(2)) > v(2, \bar{a}(2))$ is symmetric). (II) $v(1, \bar{a}(1)) < v(2, \bar{a}(1))$ and there exists a maximizer a of $v(1, a) - qv(2, a)$ such that $v(2, a) - v(1, a) \geq 0$. In this case, type 1 is assigned alternative a (possibly randomized if multiple such a exist) at price $v(1, a)$, and type 2 is assigned its efficient alternative $\bar{a}(2)$ and gets non-negative information rents $v(2, a) - v(1, a)$. (III) $v(1, \bar{a}(1)) < v(2, \bar{a}(1))$ and for any maximizer a of $v(1, a) - qv(2, a)$, we have $v(2, a) - v(1, a) < 0$. For some parameter λ and random assignment $x \in \Delta(\arg\max v(1, a)(1 - q + q\lambda) - v(2, a)q\lambda)$ that we identify below, type 1 is assigned x at price $v(1, x)$ and type 2 is assigned its efficient alternative $\bar{a}(2)$ at price equal to its willingness to pay $v(2, \bar{a}(2))$.

Lemma 6 *Assume that $v(1, \bar{a}(1)) < v(2, \bar{a}(1))$ and that for any maximizer a of $v(1, a) - qv(2, a)$, we have $v(2, a) - v(1, a) < 0$. There exists λ and $x \in \Delta(\arg\max v(1, a)(1 - q + q\lambda) - v(2, a)q\lambda)$ such that $v(2, x) - v(1, x) = 0$.*

Proof. Notice that $v(1, a)(1 - q + q\lambda) - v(2, a)q\lambda$ is linear in λ. Therefore, the problem $\max v(1, a)(1-q+q\lambda)-v(2, a)q\lambda$ is to maximize over k linear functions. Thus $[0, 1]$ can be divided into intervals, each assigned a unique alternative, where alternative a maximizes $v(1, a)(1-q+q\lambda)-v(2, a)q\lambda$ for all λ within the interval assigned to alternative a. Further, $v(1, a) - v(2, a)$ is non-decreasing in λ. If $\lambda = 0$, then $\bar{a}(1)$ maximizes $v(1, a)(1 - q + q\lambda) - v(2, a)q\lambda$, for which we have $v(2, \bar{a}(1)) - v(1, \bar{a}(1)) > 0$ by assumption. Similarly, if $\lambda = 1$, then by assumption, for any maximizer a of $v(1, a) - qv(2, a)$, we have $v(2, a) - v(1, a) < 0$. Therefore, there must exist a threshold λ such that two alternative a, a' (possibly identical) maximize $v(1, a)(1 - q + q\lambda) - v(2, a)q\lambda$, such that $v(2, a) - v(1, a) < 0$, and $v(2, a') - v(1, a') > 0$. Thus there exists a distribution x over a, a' such that $v(2, a') - v(1, a') = 0$. ■

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

$$v(1,x(1)) - p(1) \geq 0, \tag{IR1}$$

$$v(2,x(2)) - p(2) \geq 0, \tag{IR2}$$

$$v(1,x(1)) - p(1) \geq v(1,x(2)) - p(2), \tag{IC1}$$

$$v(2,x(2)) - p(2) \geq v(2,x(1)) - p(1). \tag{IC2}$$

Lemma 7 *Assume that there are two types. Necessary and sufficient conditions for the optimality of a mechanism (x,p) are provided in each of the three cases below.*

1. $v(1,\bar{a}(1)) \geq v(2,\bar{a}(1))$ *and* $v(2,\bar{a}(2)) \geq v(1,\bar{a}(2))$. *Then* $x(1) = \bar{a}(1)$, $x(2) = \bar{a}(2)$, *and IR1 and IR2 are tight.*

2. $v(1,\bar{a}(1)) < v(2,\bar{a}(1))$ *and there exists* $a \in \arg\max v(1,a) - qv(2,a)$ *such that* $v(2,a) - v(1,a) \geq 0$. *Then* $x(1) \in \Delta(\arg\max v(1,a) - qv(2,a))$ *such that* $v(2,x(1)) - v(1,x(1)) \geq 0$, $x(2) = \bar{a}(2)$, *and IR1 and IC2 are tight.*

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

Proof. In what proceeds, it is more convenient to write valuations in dot product form. That is, let $v_a(i)$ denote the valuation of type i for alternative a. Using this notation we can write the valuation for a random alternative x as $v(i,x) = v(i) \cdot x = \sum_a v_a(i)x_a$.

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

$$\begin{aligned}(1-q)p(1) + qp(2) &= (1-q+q\lambda)p(1) + q\lambda(p(2) - p(1)) + q(1-\lambda)p(2) \\ &\leq (1-q+q\lambda)v(1,x(1)) + q\lambda(v(2,x(2)) - v(2,x(1))) + q(1-\lambda)v(2,x(2)) \\ &= (1-q+q\lambda)v(1) \cdot x(1) + q\lambda v(2) \cdot (x(2) - x(1))] + q(1-\lambda)v(2) \cdot x(2) \\ &= x(1) \cdot (v(1)(1-q+q\lambda) - v(2)q\lambda) + x(2) \cdot v(2)q \end{aligned} \tag{1}$$

Thus, an IC and IR mechanism is optimal if it satisfies two conditions. First, it maximizes (1) across all $x(1)$ and $x(2)$. Second, it satisfies complementary slackness conditions. That is, (IR1) is tight, if $\lambda < 1$ then (IR2) is tight, and if $\lambda > 0$ then (IC2) is tight. Conversely, if an IC mechanism that satisfies the aforementioned two properties exists, then a mechanism is optimal only if it satisfies those properties.

For future reference, we verify that (IC1) is satisfied if (IR1) and (IC2) hold with equality and $x(1) \in \Delta(\arg\max v(1,a)(1-q+q\lambda) - v(2,a)q\lambda)$. The utility of type 1 for mimicking type 2 is

$$\begin{aligned} v(1)\cdot x(2) - p(2) &= v(1)\cdot x(2) - p(1) - v(2)\cdot(x(2)-x(1)) \\ &= v(1)\cdot x(2) - v(1)\cdot x(1) - v(2)\cdot(x(2)-x(1)) \\ &= (v(2)-v(1))\cdot x(1) - (v(2)-v(1))\cdot x(2) \\ &= \sum_a x_a(1)\big((v(2,a)-v(1,a)) - (v(2)-v(1))\cdot x(2)\big) \end{aligned} \tag{2}$$

where the last equality followed since $\sum_a x_a(1) = 1$. Now consider any alternative a such that $x_a(1) > 0$. By definition, it must be that $a \in \arg\max v(1,a)(1-q+q\lambda) - v(2,a)q\lambda$. Thus, in particular,

$$v(1,a)(1-q+q\lambda) - v(2,a)q\lambda \geq v(1,\bar{a}(2))(1-q+q\lambda) - v(2,\bar{a}(2))q\lambda,$$

or equivalently,

$$(v(2,\bar{a}(2)) - v(2,a))q\lambda \geq (v(1,\bar{a}(2)) - v(1,a))(1-q+q\lambda).$$

Since $q\lambda \leq 1-q+q\lambda$, and $v(2,\bar{a}(2)) - v(2,a) \geq 0$, we have $v(2,\bar{a}(2)) - v(2,a) \geq v(1,\bar{a}(2)) - v(1,a)$. Therefore, for any a where $x_a(1) > 0$,

$$(v(2,a) - v(1,a)) - (v(2)-v(1))\cdot x(2) \leq 0.$$

Now note that for any a, either $x_a(1) = 0$, or $x_a(1) > 0$ in which case the above inequality holds. In either case, we have

$$x_a(1)\big((v(2,a)-v(1,a)) - (v(2)-v(1))\cdot x(2)\big) \leq 0.$$

Summing over all a, and given (2), we have

$$v(1)\cdot x(2) - p(2) = \sum_a x_a(1)\big((v(2,a)-v(1,a)) - (v(2)-v(1))\cdot x(2)\big) \leq 0.$$

Thus, the incentive constraint for type 1 is satisfied.

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

1. $v(1, \bar{a}(1)) \geq v(2, \bar{a}(1))$ and $v(2, \bar{a}(2)) \geq v(1, \bar{a}(2))$. Then full surplus extraction is optimal.

2. $v(1, \bar{a}(1)) < v(2, \bar{a}(1))$ and there exists $a \in \arg\max v(1, a) - qv(2, a)$ such that $x(1) \in \Delta(\arg\max v(1, a) - qv(2, a))$ such that $v(2) \cdot x(1) - v(1) \cdot x(1) \geq 0$, $x(2) = \bar{a}(2)$, and (IR1) and (IC2) are tight. (IC1) is satisfied by the analysis above for $\lambda = 1$. (IR1) is satisfied by the assumption that $v(2) \cdot x(1) - v(1) \cdot x(1) \geq 0$. The allocation x maximizes (1) for $\lambda = 1$. The complementary slackness conditions are satisfied since $\lambda = 1$ and (IR1) and (IC2) are tight. Therefore, the mechanism is optimal, and any optimal mechanism must satisfy the specified properties.

3. $v(1, \bar{a}(1)) < v(2, \bar{a}(1))$ and for all $a \in \arg\max v(1, a) - qv(2, a)$, we have $v(2, a) - v(1, a) < 0$. Consider mechanism (x, p) such that $v(1, \bar{a}(1)) < v(2, \bar{a}(1))$ and for all $a \in \arg\max v(1, a) - qv(2, a)$, we have $v(2, a) - v(1, a) < 0$. Then $x(1) \in \Delta(\arg\max v(1, a)(1 - q + q\lambda) - v(2, a)q\lambda)$ such that $v(2) \cdot x(1) - v(1) \cdot x(1) = 0$, $x(2) = \bar{a}(2)$, and (IR1), (IR2), and (IC2) are tight. The parameter λ and the assignment x exist by Lemma 6. The allocation x maximizes (1) for λ. The complementary slackness conditions are satisfied since (IR1), (IR2), and (IC2). (IC1) is satisfied by the analysis above. Therefore, the mechanism is optimal, and any optimal mechanism must satisfy the specified properties.

■

The following lemma reveals the structural properties of $\arg\max_a v(1, a) - qv(2, a)$, which identifies the allocation of type 1 by Lemma 7. The problem is one of maximizing over functions $v(1, a) - qv(2, a)$ that are linear in q. Thus, the set of markets $[0, 1]$ can be divided into intervals, each assigned a unique alternative, such that the alternative is the unique maximizer of $v(1, a) - qv(2, a)$ in the interior of the interval, and is an optimizer of $v(1, a) - qv(2, a)$ at the endpoints of the interval. Additionally, $\max_a v(1, a) - qv(2, a)$ is convex in q. See Figure 8. As q increases, as long as the maximizer a of $\arg\max_a v(1, a) - qv(2, a)$ satisfies $v(2, a) - v(1, a) \geq 0$, then the utility of type 2, $v(2, a) - v(1, a)$ decreases in q. If for some q, $v(2, a) - v(1, a) < 0$, then the utility of type 2 will be zero and will remain zero for all $q' \geq q$. Formally, we have the following lemma.

Lemma 8 *For some $m \leq k$, there exist thresholds $\tau_0 < \ldots < \tau_{m+1}$, $\tau_0 = 0$, $\tau_{m+1} = 1$ and an injective function $g : \{0, \ldots, m\} \to A$, $g(0) = \bar{a}(1)$, $g(m) = 0$, such that for*

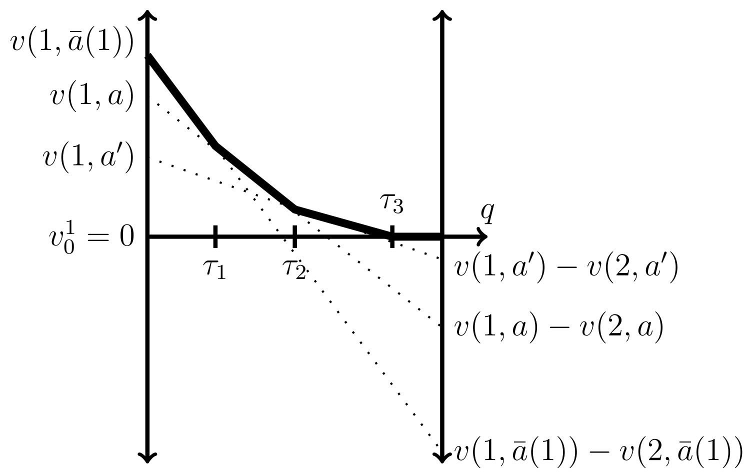

Figure 8: An example with $k = 4$, $m = 3$, $\tau_1 = a$, and $\tau_2 = a'$.

all $j \leq m$, $g(j) \in \arg\max_a v(1, a) - qv(2, a)$ *if and only if* $q = [\tau_j, \tau_{j+1}]$. *Additionally,* $v(2, g(j)) - v(1, g(j))$ *is strictly decreasing in* j.

Proof. Consider the problem of maximizing $v(1, a) - qv(2, a)$. For each q, the problem is to maximize over k linear functions. The graph is drawn in Figure 8. Note the properties of the solution. First, the set of markets $[0, 1]$ can be divided into intervals with thresholds $\tau_0 < \ldots < \tau_{m+1}$, with each τ_j assigned a unique alternative $g(j)$, such that for all $j \leq m$, $g(j) \in \arg\max_a v(1, a) - qv(2, a)$ if and only if $q = [\tau_j, \tau_{j+1}]$. Since $v(1, \bar{a}(1)) > v(1, a)$ for all $a \neq \bar{a}$, $g(0) = \bar{a}(1)$. Similarly, since $v(1, a) - v(2, a) < 0$ for all $a \neq 0$, $g(m) = 0$. Additionally, consider j such that $g_{i-1} = a$ and $g_j = a'$. We must have $v(1, a) - qv(2, a) > v(1, a') - qv(2, a')$ for all $q < \tau_j$ and $v(1, a) - qv(2, a) < v(1, a') - qv(2, a')$ for all $q > \tau_j$. Thus, $v(2, a) - v(1, a) > v(2, a') - v(1, a')$. ■