Procurement Mechanisms for Differentiated Products*

PRELIMINARY; NEW VERSION COMING SOON

Daniela Saban and Gabriel Y. Weintraub †
Graduate School of Business, Columbia University

November 2014

Abstract

We consider the problem faced by a procurement agency that runs an auction-type mechanism to construct a menu (assortment of products with posted prices), from a set of differentiated products offered by strategic suppliers. Heterogeneous consumers then buy their most preferred alternative from the menu as needed. Framework agreements (FAs), widely used in the public sector, take this form; the central government runs the initial auction and then the public organizations (hospitals, schools, etc.) buy from the selected assortment as required. This type of mechanism is also relevant in other contexts, such as the design of medical formularies and group buying. When evaluating the bids, the procurement agency must consider the optimal tradeoff between offering a richer menu of products for consumers versus offering less variety, hoping to engage the suppliers in a more aggressive price competition. We develop a mechanism design approach to study this problem, and provide a characterization of the optimal menus. The optimal mechanism balances the tradeoff between product variety and price competition, in terms of suppliers' costs, products' characteristics, and consumers' characteristics. We then use the optimal mechanism as a benchmark to evaluate the performance of the Chilean government procurement agency's current implementation of FAs, used to acquire US$2 billion worth of goods per year. We show how simple modifications to the current mechanism, which increase price competition among close substitutes, can considerably improve performance.

*We would like to thank the executives from the Chilean government procurement agency, Dirección ChileCompra, in particular, David Escobar and Guillermo Burr, for the insights and data shared with us. The opinions expressed here are those of the authors and do not necessarily reflect the positions of Dirección ChileCompra nor of its executives. We thank Omar Besbes, Marcelo Olivares, Sasa Pekec, Jay Sethuraman, and seminar participants at various conferences and institutions for useful comments. This work was supported by the Chazen Institute of International Business and the Deming Center at Columbia Business School, as well as by a Chile-Columbia Fund grant.

†dhs2131@columbia.edu, gweintraub@columbia.edu

1 Introduction

During the last decades, governments and firms alike have opted for procurement mechanisms in which the purchasing decisions are shared by a central authority and local divisions, who will ultimately consume the goods. Typically, the central authority selects an assortment of differentiated products through competitive bidding to satisfy the demand arising from the divisions that have heterogeneous preferences. The rationale behind adopting such a procurement mechanism is to exploit the purchasing power of a central buyer, while still providing individual consumers with some flexibility in product selection. These mechanisms are relevant for many real-world applications, such as medical formularies and group purchasing in the healthcare industry (see, for example, [26]). Also, framework agreements (FAs), widely used in the public sector, take this form.[1]

Roughly speaking, a FA works as follows. First, the central government specifies a broad category (e.g., computers), and a succinct description of products and/or services that are needed within the category (e.g., laptops of certain size and specifications). Suppliers are allowed to submit bids for *any* product fitting the description. Then, an auction-type mechanism is run to select an assortment of differentiated products with unit prices. Once the government decides on the winning bids, the public organizations (e.g., hospitals, schools, etc.) buy their most preferred product at the agreed price as needed, without undergoing any additional public tendering process.

This paper is one of the first in the literature to provide a formal economic analysis of this type of procurement mechanisms. Our contribution is three-fold: we first introduce a model for the problem faced by the procurement agency, we then characterize the optimal mechanism for this setting and, finally, we use these results to study the design of simpler mechanisms that are commonly used in practice. While our main motivation is to improve our understanding of FAs, these results also shed light on buying mechanisms that could be used in the other settings mentioned above. We describe our main contributions in more detail next.

Our first main contribution is introducing a model that incorporates the following fundamental trade-off faced by a procurement agency buying differentiated products. On one hand, consumers buying from the assortment usually have heterogeneous preferences; for example, while a public school may want to buy laptops with attractive graphics features, the department of treasury may need laptops with high processing power. Therefore, increasing product variety in the assortment may increase consumer satisfaction, as it becomes more likely they will find a better product for their needs. On the other hand, price competition in the auction stage may be depressed if too many products are included in the assortment. Our model extends the classic auction and mechanism design models to study this trade-off between product variety and price competition.

[1] For example, in 2010 the European Union awarded €80 billion using FAs, accounting for 15% of the total value of all public procurement [8].

In our model, there is a set of risk-neutral suppliers offering differentiated products, which are imperfect substitutes of each other. In the tradition of the auctions literature, we assume that suppliers have private information about their costs. The central procurement agency (designer) uses an auction-type mechanism to determine a *menu*, that is, an assortment of differentiated products together with the unit prices. Then, consumers with private heterogeneous preferences buy their most preferred alternative in the menu, which induces aggregate demands over products. In the tradition of the assortment literature, we assume that the designer can predict the aggregate demands for a given menu. Given the demand model, the designer chooses a mechanism with the objective of maximizing expected consumer surplus.

Our second main contribution is the characterization of the optimal direct-revelation posted-price mechanism for a broad class of affine demand models. This class includes the classic horizontal Hotelling demand model and a pure vertical demand model as particular cases, as well as more general specifications with both horizontal and vertical sources of product differentiation. Affine demand models are commonly used in competition models (see, e.g., [27]) and we think they provide a reasonable balance between tractability and generality in our setting. The optimal mechanism quantifies the optimal tradeoff between variety and price competition in terms of suppliers' costs, product characteristics, and substitution patterns. For example, the mechanism may optimally choose to restrict the entry of some products to the assortment; this decreases expected payments to suppliers at the expense of reducing variety for consumers.

Relative to a classic mechanism design problem, a distinctive feature of our formulation is that the auctioneer cannot directly decide how to allocate demand across the products. Instead, the auctioneer selects the menu and demands are then determined by the underlying preferences of the organizations, which introduces significant complexities in the analysis. In addition, for the most part, previous auction and mechanism design work assumes homogeneous products (with some notable exceptions discussed in Section 2). Our work advances the theory of auctions and mechanism design by accounting for an endogenous demand system for differentiated products.

Our third main contribution is to improve our understanding of the performance of simple mechanisms used in practice. The optimal mechanisms previously characterized are rarely implemented in applications due to their complexity. However, they serve as a powerful tool to study practical mechanisms: optimal mechanisms provide a benchmark on what is achievable and their structure provide insights on how to improve current practice. We are particularly interested in the type of FAs run by our collaborator in this project, the Chilean government procurement agency (Dirección ChileCompra), that bought US$2 billion worth of goods using FAs during 2013.[2]

An important first observation that arises by looking at the data from ChileCompra's FAs

[2]This represented a 21% of the value of all public procurement in Chile [10].

is that, because product definitions are narrow and auctions for different products are run independently, there is a single supplier bidding and winning for many products. Hence, while these suppliers may compete for demand once in the assortment, there is little to none competition *for the market*. Thus motivated, we study whether the current FA performance can be improved by creating thicker markets, making imperfect substitute products compete to be in the menu.

To this end, we provide an extensive theoretical analysis of the current implementation of ChileCompra's FAs in a simple model. Then, using the insights gained from the optimal mechanism, we explore possible changes to ChileCompra's first-price auction FA implementation with regards to the set of suppliers to include in the menu. We show how, in general, using rules that restrict the entry of close substitute products can significantly increase price competition across suppliers, which translates into a significant increase in expected consumer surplus. We provide a detailed analysis that illustrates when is it profitable to restrict the entry, as a function of the market primitives namely, suppliers' costs distributions and consumers' demand characteristics. Using numerical experiments, we validate the robustness of our results in more general settings. Overall, our results show that simple modifications to current practice can significantly increase performance.

The rest of the paper is organized as follows. Section 2 describes related literature. In Section 3 we formulate the mechanism design problem faced by the designer. In Section 4, we describe the general solution approach that we use to solve for the optimal mechanism. In Section 5, we characterize the optimal mechanism for affine demand models. In Section 6, we discuss the design of practical mechanisms using ChileCompra as a case study. We conclude and provide extensions in Section 7. All proofs are deferred to the Appendix.

2 Related literature

Our work is related to several streams of literature in economics and operations. First, as previously mentioned, our work extends classic work in mechanism design in the tradition of Myerson [22], in which a mechanism is specified by a payment and an allocation function. In our problem instead, the designer selects an assortment of products together with their unit prices, and then demands are allocated according to consumer preferences for the given menu. Hence, the designer does not have complete freedom to choose the allocations (these must respect the underlying demands from consumers), adding significant difficulties when solving for the optimal mechanism. Another difference with the classic framework is that, in our problem, the designer maximizes consumer surplus (as oppose to just minimizing payments to suppliers), and consumer surplus also depends on the underlying preferences of consumers.

Our work is also related to oligopoly pricing models that have studied the effects of entry and competition in consumer surplus (see, e.g., Tirole [25]). The main difference is that, in our setting, the decision to enter the market is not freely made by firms, but is decided by the designer based on the information elicited in the auction. Further, in our setting there is asymmetric information about firms' costs.

In that sense, our work is more related to previous papers in procurement and regulation economics. For example, Dana and Spier [9] that studies how to allocate production rights to firms that have private cost information. In their paper, however, the auction only determines the market structure and lump-sum fees, as opposed to our case in which it also determines unit prices. Similarly, Anton and Gertler [4] and McGuire and Riordan [21] study the optimal mechanism with an endogenous market structure in a Hotelling model of product differentiation. However, in these models demand is not endogenously determined like in our case. A general insight of this body of work is that the designer may single-source more frequently if firms have private cost information, to be able to exert more pressure on efficient suppliers to reveal their costs; this is similar to some of our insights.

Closer to our work, Wolinsky [29] studies a spatial duopoly model, where firms firms compete in both prices and quality. While this paper considers an endogenous demand, the analysis is restricted to "interior" solutions (i.e. when both firms have positive demands). Instead, we are particularly interested in "border" solutions, in which some firms may be left out of the assortment. Overall, to the best of our knowledge, our work is the first to characterize the optimal mechanism with an endogenous market structure, endogenous demand, and in which prices are determined in the auction.

Another stream of related work that considers endogenous market structures is that of split-award auctions or dual sourcing in economics and operations [7, 18, 11, 5]. Split-award auctions have been studied in a variety of settings. However, these papers do not assume an underlying set of heterogeneous consumers; instead, purchases are decided by the auctioneer to maximize his own goal.

Our work is related to the operations literature that studies assortment planning decisions [15]. In these settings, decisions are made by one retailer that carries all products. In our case instead, an assortment is built using an auction that elicits private cost information from many different suppliers. The assortment literature typically assumes some kind of parametric discrete choice model as a demand system like, for example, the multinomial logit model. In our setting, this model is not appropriate because of its inability to capture substitution patterns due to the IIA property. An alternative that resolves this issue is the multinomial logit model with random coefficients; however, this model is hard to solve even in the standard assortment problem, let alone in our auction setting. Another option that is typically more tractable is the nested logit model (see, e.g., [19]). We are currently exploring whether our framework can be applied to this demand system.

Finally, to the best of our knowledge, only two prior papers study framework agreements (FAs), which is one of the main objectives of our work. Albano and Sparro [2] consider a Hotelling model of horizontal differentiation, in which firms are located equidistantly and the subset of potential suppliers with lowest bids are selected in the assortment. In our case, we consider a richer set of rules in which the assortment can depend on product characteristic or location. Further, their analysis assumes complete information about firms' costs. Gur et al. [13] consider a model of FAs that studies the cost uncertainty faced by a supplier over the FA time horizon when selling a single-item, but does not consider differentiated products nor heterogenous consumers.

3 Model and Problem Formulation

In this section, we present our model and a formulation of the auctioneer's problem as a mechanism design problem.

3.1 Model

We introduce a model of procurement mechanisms for differentiated products demand systems. The agents of the model are (*i*) an auctioneer (or designer); (*ii*) suppliers; and (*iii*) consumers. The designer runs an auction-type mechanism to construct a menu (i.e., an assortment of products with posted prices) based on the suppliers' offers. Then, consumers purchase their most preferred product from this menu at the agreed price. We describe the main elements of the model next.

Suppliers. There is an exogenous set N of n potential *suppliers* indexed by i. Suppliers offer differentiated products that are imperfect substitutes to each other; the characteristics of these products are common-knowledge. To simplify the exposition, we initially assume that each supplier offers exactly one product. Hence, unless otherwise stated, firms and products share the same indexes. In Section 7 and Appendix C, we discuss the extension to the multiproduct setting; it is worth highlighting that our main results also hold under this extension. We assume suppliers are risk-neutral, so they seek to maximize expected profits.

Following the tradition in the auctions' literature (see, for example, Krishna [16]), we assume that suppliers have production costs drawn independently from common-knowledge distributions, whose realizations are the private information of each supplier. Formally, supplier i has a private cost $\theta_i \in \Theta_i$, a finite set of strictly positive real numbers. We index the elements of Θ_i, such that

$\theta_i^j < \theta_i^k$ whenever $j < k$, for all $\theta_i^j, \theta_i^k \in \Theta_i$. We say that supplier i is of type θ_i if his cost is θ_i. Let f_i be a probability mass function over Θ_i, where $f_i(\theta_i)$ represents the probability that supplier i is of type θ_i. Let $F_i(\theta_i^j) = \sum_{k \leq j} f_i(\theta_i^k)$ be the cumulative probability distribution. Let $\Theta = \Pi_i \Theta_i$ denote the type space. Because suppliers' types are independent, the joint probability of $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$ is equal to $f(\boldsymbol{\theta}) = \Pi_{i=1}^n f_i(\theta_i)$. We denote the probability that all suppliers other than i have type $\boldsymbol{\theta}_{-i}$ by $f_{-i}(\boldsymbol{\theta}_{-i})$. We use boldfaces to denote vectors and matrices throughout the paper.

Further, we assume that suppliers have constant marginal costs of production and do not face capacity constraints. Therefore, the products included in the assortment are always available and their production cost does not depend on the quantity demanded. In the reminder of the paper, we use agents and suppliers interchangeably.

Consumers. There is a (possible continuum) set J of consumer types, indexed by j. For each $j \in J$, we denote by $h(j)$ the density of consumers of type j. The total mass of consumers is normalized to 1. Consumers have quasi-linear utilities that depend on the product characteristics, the price, and the consumer type. Formally, the utility a consumer of type j obtains by consuming product i at price p_i is given by:

$$u_{ji}(p_i) = V + v_{ji} - p_i, \tag{1}$$

where V represents the value a consumer has for consuming any good in the set N and v_{ji} is the consumption benefit for a consumer of type j given by product i. Each consumer wants to buy exactly one unit of the product offered in the menu to maximize her own utility.[3] We assume that V is large enough so that the consumer market is covered.

Suppose that from the set of potential suppliers N, we fix a subset $Q \subseteq N$ of active suppliers. Let $\boldsymbol{p}_Q = \{p_i\}_{i \in Q}$, be the vector of their unit prices. We define for all $i \in Q$, the set $A_i(Q, \boldsymbol{p}_Q) = \{j : u_{ji}(p_i) \geq u_{jk}(p_k),\ \forall k \in Q\}$. That is, $A_i(Q, \boldsymbol{p})$ is the set of types of consumers that would buy from supplier $i \in Q$ given assortment Q at prices $\boldsymbol{p}_Q$.[4] We define $A_i(Q, \boldsymbol{p}_Q) = 0$ for all $i \notin Q$. Then, the expected demand for product $i \in Q$ is given by:

$$d_i(Q, \boldsymbol{p}_Q) = \int_{j \in A_i(Q, \boldsymbol{p}_Q)} h(j) dj, \tag{2}$$

and $d_i(Q, \boldsymbol{p}_Q) = 0$, for $i \notin Q$.

In the tradition of the assortment literature (see, for example Kök et al. [15]) and the work in oligopoly pricing (see, for example Tirole [25]), we assume that these demand functions are common

[3] Many of our results extend to the case in which each product is also offered by an outside supplier at a fixed cost. In Section 7 and Appendix C.2, we discuss an extension to the case of elastic demand.

[4] We assume that ties occurred are broken randomly.

knowledge. This implies that, even though the preferences of a specific consumer may be private information, the designer can predict the aggregate demand for every fixed set of products and prices. These aggregate demands can usually be summarized by what is known as a consumer-choice model or a demand system. The assumption is plausible, in particular in the contexts discussed in the introduction, because a demand system can typically be estimated using historical data or consumer surveys (Ackerberg et al. [1]).

Let $\phi : J \to N$ be a function such that $\phi(j)$ denotes the product consumed by type j. Then, for such function ϕ and prices $\boldsymbol{p}$, the consumer surplus is equal to:

$$\begin{aligned} CS(\phi, \boldsymbol{p}) &= \int_{j \in J} u_{j\phi(j)}(p_{\phi(j)}) h(j) dj \\ &= \int_{j \in J} \left(v_{j\phi(j)} - p_{\phi(j)} \right) h(j) dj \\ &= \sum_{i \in N} \left[\int_{j \in J} v_{j\phi(j)} \mathbb{I}[\phi(j) = i] h(j) dj - \int_{j \in J} p_{\phi(j)} \mathbb{I}[\phi(j) = i] h(j) dj \right] , \end{aligned} \tag{3}$$

where $\mathbb{I}[\cdot]$ denotes the indicator function. To simplify notation, we have omitted the terms associated to the reservation value V in the previous expressions and will do so throughout the rest of the paper.

Let x_i be the mass of consumers buying product i, and define $\boldsymbol{x} = (x_1, ..., x_n)$. We make the following assumption that we keep throughout the paper.

Assumption 3.1. *For all $i \in N$, and all functions ϕ, there exists a function $k_i(\boldsymbol{x})$ of the mass of consumers buying each product under ϕ, such that:*

$$k_i(\boldsymbol{x}) = \int_{j \in J} v_{j\phi(j)} \mathbb{I}[\phi(j) = i] h(j) dj \ .$$

The assumption states that we can write consumer gross surplus associated to each product i as a function of the demand vector $\boldsymbol{x}$. We highlight that Assumption 3.1 holds for all demand models that are considered in the paper. Under this assumption we can re-write consumer surplus as:

$$CS(\boldsymbol{x}, \boldsymbol{p}) = \sum_{i=1}^{n} \left[k_i(\boldsymbol{x}) - p_i x_i \right]. \tag{4}$$

Now, note that demands are induced by consumers' utility maximization decisions. Hence, by aggregating these decisions, it is simple to observe that:

$$(d_1(N,\boldsymbol{p}),\ldots,d_n(N,\boldsymbol{p})) \in \underset{\boldsymbol{x}}{\operatorname{argmax}} \; CS(\boldsymbol{x},\boldsymbol{p}) \;, \qquad \text{s.t.} \quad \sum_{i=1}^{n} x_i = 1, \quad x_i \geq 0 \quad \forall i \in N \;. \tag{5}$$

That is, the demands derived from Eq. (2) given prices $\boldsymbol{p}$ and when all products are part of the assortment maximize consumer surplus given those prices. Note that the solution of this maximization problem may set some of the demands equal to zero. Problem (5) implies that demands can be also interpreted as the consumption decisions made by a single 'representative consumer' that maximizes consumer surplus.[5] We come back to this interpretation in our general treatment of affine demand models presented in Section 5.3.

To illustrate the concepts defined so far, we present a canonical Hotelling demand model of horizontal differentiation with two suppliers and linear 'transportation costs'. A detailed analysis of our problem using the Hotelling model as the demand model can be found in Section 5.2.

Example 3.1 (Hotelling model with two suppliers)**.** *Consider the unit interval as the product space, with two potential suppliers located at the extremes of the interval. There is a continuum of consumers uniformly distributed on the product space. Each consumer demands one unit of good and incurs transportation costs which are linear in the distance between the consumer and the supplier. Therefore, utility functions are given by:*

$$u_{j1}(p_1) = -\left(\delta\ell_j + p_i\right) \qquad \textit{and} \qquad u_{j2}(p_2) = -\left(\delta(1-\ell_j) + p_2\right),$$

where supplier 1 *(resp.* 2*) is assumed to be located at* 0 *(resp.* 1*),* δ *is the transportation cost, and* ℓ_j *is the position of consumer* j *in the unit line. As consumers are uniformly distributed on the* $[0,1]$ *segment, consumer surplus is given by:*

$$CS(\boldsymbol{x},\boldsymbol{p}) = -\left(\frac{\delta}{2}\left(x_1^2 + x_2^2\right) + p_1x_1 + p_2x_2\right),$$

where the first terms represent the transportation costs and the latter terms the monetary costs. Note that in this example, $k_i(\boldsymbol{x}) = -\frac{\delta}{2}x_i^2$*, which is equivalent to the total transportation cost incurred by those consumers buying from* i*. Further, assuming both firms are active, the demand functions are given by* $d_i(N,\boldsymbol{p}) = \frac{p_j - p_i + \delta}{2\delta}$ *for* $i,j \in \{1,2\}$ *and* $j \neq i$*. It is simple to observe that these expressions maximize consumer surplus whenever* $|p_1 - p_2| \leq \delta$*.*

[5] See Chapter 3 in Anderson et al. [3] for a formal discussion on the representative consumer approach.

Auctioneer. The role of the auctioneer is to select or design an auction-type mechanism to construct the menu of products based on the suppliers' offers. As previously mentioned, the menu consists of a subset of suppliers and prices for their products. Once selected, the rules of the auction are common-knowledge. The auctioneer is risk-neutral and her objective is to maximize expected consumer surplus; this objective incorporates both variety considerations and payments to suppliers. Throughout the rest of the paper, we use auctioneer and designer interchangeably.

3.2 Mechanism Design Problem Formulation

We provide a mechanism design formulation of the auctioneer's problem. We consider mechanisms implemented in Bayes Nash equilibria. By invoking the revelation principle, we restrict attention to direct revelation mechanisms without loss of optimality. Hence, for given cost declarations, the designer selects a menu which consists of an assortment of products (or suppliers) and their unit prices. Formally, a direct revelation mechanism can be specified by (a) the *'assortment' functions* $q_i : \Theta \to \{0, 1\}$ that are equal to 1 if and only if supplier i is included in the assortment when cost declarations are $\boldsymbol{\theta}$; and (b) the *price functions* $p_i : \Theta \to \mathbb{R}$, where $p_i(\boldsymbol{\theta})$ is the unit price for the item offered by supplier i when cost declarations are $\boldsymbol{\theta}$. Note that this formulation allows for multiple suppliers to be in the menu. We define $\boldsymbol{q} = (q_1, ..., q_n)$ and $\boldsymbol{p} = (p_1, ..., p_n)$. For given cost declarations $\boldsymbol{\theta}$, the menu is given by $(\boldsymbol{q}(\boldsymbol{\theta}), \boldsymbol{p}(\boldsymbol{\theta}))$.

We also define the *allocation functions* $x_i : \Theta \to [0, 1]$, where $x_i(\theta_i, \boldsymbol{\theta}_{-i})$ is the fraction of demand allocated to bidder i if his cost declaration is θ_i and his competitors' declaration are $\boldsymbol{\theta}_{-i}$. Let $\boldsymbol{x} = (x_1, \ldots, x_n)$. For each realization of $\boldsymbol{\theta}$, given the menu $(\boldsymbol{q}(\boldsymbol{\theta}), \boldsymbol{p}(\boldsymbol{\theta}))$, consumer demand is determined by the underlying demand model. Hence, for given $(\boldsymbol{q}, \boldsymbol{p})$, the allocation function $\boldsymbol{x}$ is restricted by the demand constraints in Eq. (2). This is in sharp contrast with classic mechanism design theory, in which the designer specifies a payment (or transfer) function and an allocation function. In our case, the designer selects an assortment and unit prices and, given these, allocations are decided by consumers. As discussed below, these demand constraints on the allocations introduce significant additional complexities to the mechanism design problem.

In the optimal mechanism design problem, the designer maximizes its objective (in our case, expected consumer surplus) subject to the usual constraints in mechanism design theory: incentive compatibility (IC), individual rationality (IR), and feasibility of allocations (Feas). To write these constraints, we define the *interim expected utility* for a supplier of type θ_i and report θ_i' as:

$$U_i(\theta_i'|\theta_i) = \sum_{\boldsymbol{\theta}_{-i} \in \Theta_{-i}} f_{-i}(\boldsymbol{\theta}_{-i}) \left(\left(p_i(\theta_i', \boldsymbol{\theta}_{-i}) - \theta_i \right) x_i(\theta_i', \boldsymbol{\theta}_{-i}) \right) \tag{6}$$

In addition, the problem must also have constraints to ensure that the allocations are consistent with the underlying demand system (Demand 1 and 2). Using the above definitions, the optimal mechanism design problem can be formulated as follows:

$$
\begin{aligned}
[P_0] \qquad & \max_{\boldsymbol{q},\boldsymbol{p},\boldsymbol{x}} \quad \mathbb{E}_{\boldsymbol{\theta}}[CS(\boldsymbol{x}(\boldsymbol{\theta}),\boldsymbol{p}(\boldsymbol{\theta}))] && \\
& \text{s.t.} \quad U_i(\theta_i|\theta_i) \geq U_i(\theta_i'|\theta_i) \qquad \forall i \in N,\ \forall \theta_i, \theta_i' \in \Theta_i && \text{(IC)} \\
& \qquad U_i(\theta_i|\theta_i) \geq 0 \qquad \forall i \in N,\ \forall \theta_i \in \Theta_i && \text{(IR)} \\
& \qquad \sum_{i \in N} x_i(\boldsymbol{\theta}) = 1 \qquad \forall \boldsymbol{\theta} \in \Theta, \qquad x_i(\boldsymbol{\theta}) \geq 0 \qquad \forall i \in N, \forall \boldsymbol{\theta} \in \Theta && \text{(Feas)} \\
& \qquad x_i(\boldsymbol{\theta}) = d_i(\boldsymbol{q}(\boldsymbol{\theta}),\boldsymbol{p}(\boldsymbol{\theta})) \qquad \forall i \in N,\ \forall \boldsymbol{\theta} \in \Theta && \text{(Demand 1)} \\
& \qquad d_i(\boldsymbol{q}(\boldsymbol{\theta}),\boldsymbol{p}(\boldsymbol{\theta})) \text{ is given by equation (2)} \qquad \forall i \in N,\ \forall \boldsymbol{\theta} \in \Theta && \text{(Demand 2)}
\end{aligned}
$$

In the next section we discuss our approach to solve the optimal mechanism problem P_0.

4 General Solution Approach

Problem P_0 is a mixed integer mathematical program. More specifically, demand equations can be typically written as complementarity conditions, and therefore, even if one relaxes the integrality of the variables $\boldsymbol{q}$, the program is typically non-convex. Our solution approach relies on relaxing these demand constraints and solving the *relaxed problem*. The advantage of doing this is that the relaxed problem admits an analytical solution by extending standard mechanism design arguments based on the envelope theorem [22] adapted for the setting of discrete distributions [28]. Then, we provide conditions that guarantee the existence of unit prices $\boldsymbol{p}$ that are consistent with the optimal solution of the relaxed problem and satisfy the demand constraints. If such prices $\boldsymbol{p}$ exist, the optimal solution to the relaxed problem can be achieved by the original problem P_0. We formalize this argument next.

To do so, we introduce a new set of variables $t_i : \Theta \to \mathbb{R}$, where $t_i(\boldsymbol{\theta}) = p_i(\boldsymbol{\theta})x_i(\boldsymbol{\theta})$ represents the total *transfer* (or payment) to supplier i for a given cost declaration $\boldsymbol{\theta}$. Recall that $\boldsymbol{p}(\boldsymbol{\theta})$ is the vector of unit prices for reported costs $\boldsymbol{\theta}$. Relaxing the demand constraints from $[P_0]$ and noting that interim utilities (Eq. (6)) can be written in terms of total transfers $\boldsymbol{t}$, we obtain the relaxed problem:

$$
\begin{aligned}
[P_1] \qquad & \max_{\boldsymbol{x},\boldsymbol{t}} \quad \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n}\left[k_i(\boldsymbol{x}(\boldsymbol{\theta})) - t_i(\boldsymbol{\theta})\right]\right] \\
& \text{s.t.} \quad U_i(\theta_i|\theta_i) \geq U_i(\theta_i'|\theta_i) \qquad \forall i \in N,\ \forall \theta_i, \theta_i' \in \Theta_i \qquad \text{(IC)} \\
& \qquad\quad U_i(\theta_i|\theta_i) \geq 0 \qquad \forall i \in N,\ \forall \theta_i \in \Theta_i \qquad \text{(IR)} \\
& \qquad\quad \sum_{i \in N} x_i(\boldsymbol{\theta}) = 1 \qquad \forall \boldsymbol{\theta} \in \Theta, \qquad x_i(\boldsymbol{\theta}) \geq 0 \qquad \forall i \in N,\ \boldsymbol{\theta} \in \Theta. \qquad \text{(Feas)}
\end{aligned}
$$

Problem $[P_1]$ only differs from the classic mechanism design formulation in the objective function; while the traditional objective is to minimize transfers, we aim to maximize consumer surplus here. The objective, therefore, contains the k_i terms associated to gross consumer surplus. Similarly to the setting of continuous cost distributions, we introduce the following definition of the virtual cost function for cost distributions with discrete support.[6]

Definition 4.1 (Virtual costs)**.** *For $\theta_i \in \Theta_i$, let $\rho_i(\theta_i) = \max\{\theta' \in \Theta_i : \theta' < \theta_i\}$, that is, $\rho_i(\theta_i)$ is the predecessor of θ_i in Θ_i.*[7] *Let $v_i(\theta_i) = \theta_i + \frac{F_i(\rho_i(\theta_i))}{f_i(\theta_i)}(\theta_i - \rho_i(\theta_i))$ be the* virtual cost *of supplier i when he has type θ_i.*

We make the standard regularity assumption in mechanism design that we keep throughout the paper:

Assumption 4.1 (Increasing virtual costs)**.** *The function $v_i(\theta_i)$ is strictly increasing for all $i \in N$.*

Finally, we also define the *interim expected allocations* and *interim expected transfers* as follows:

$$
\begin{aligned}
X_i(\theta_i) &\equiv \sum_{\boldsymbol{\theta}_{-i} \in \Theta_{-i}} f_{-i}(\boldsymbol{\theta}_{-i}) x_i(\theta_i, \boldsymbol{\theta}_{-i}), \\
T_i(\theta_i) &\equiv \sum_{\boldsymbol{\theta}_{-i} \in \Theta_{-i}} f_{-i}(\boldsymbol{\theta}_{-i}) t_i(\theta_i, \boldsymbol{\theta}_{-i}).
\end{aligned}
$$

The advantage of solving the relaxed problem $[P_1]$ is that we can extend standard mechanism design arguments to characterize its optimal solution, as we formalize next.

Proposition 4.1. *Suppose that $(\boldsymbol{x}, \boldsymbol{t})$ satisfy the following conditions:*

[6] Note that, in the limit, this definition of the virtual cost agrees with the well-known definition of virtual costs for a continuous support, i.e., $v_i(\theta_i) = \theta_i + \frac{F_i(\theta_i)}{f_i(\theta_i)}$.

[7] If θ_i is the lowest in the support, we define $\rho_i(\theta_i) = \theta_i$.

1. The allocation function satisfies for all $\boldsymbol{\theta} \in \Theta$,

$$\begin{aligned} \boldsymbol{x}(\boldsymbol{\theta}) \quad &\in \quad \text{argmax} \sum_{i=1}^{n} (k_i(x(\boldsymbol{\theta})) - x_i(\boldsymbol{\theta}) v_i(\theta_i)) \\ s.t. \quad & \sum_{i=1}^{n} x_i(\boldsymbol{\theta}) = 1, \qquad x_i(\boldsymbol{\theta}) \geq 0 \quad \forall i \in N \ . \end{aligned} \tag{7}$$

2. Interim expected allocations are monotonically decreasing for all $i \in N$, that is, $X_i(\theta) \geq X_i(\theta')$ for all $\theta, \theta' \in \Theta_i$ such that $\theta \leq \theta'$.

3. Interim expected transfers satisfy for all $i \in N$ and $\theta_i^j \in \Theta_i$:

$$T_i(\theta_i^j) = \theta_i^j X_i(\theta_i^j) + \sum_{k=j+1}^{|\Theta_i|} (\theta_i^k - \theta_i^{k-1}) X_i(\theta_i^k) \tag{8}$$

Then, $(\boldsymbol{x}, \boldsymbol{t})$ is an optimal mechanism for problem P_1.

The proof can be found in Appendix A. Condition (1) in Proposition 4.1 states that, for each $\boldsymbol{\theta} \in \Theta$, the optimal vector of allocations $x(\boldsymbol{\theta})$ must be a maximizer of the consumer surplus when prices are set to be the virtual costs, subject to the feasibility constraints (see Eq. (4)). Further, by Eq. (5), the optimal solution is of the form $x_i(\boldsymbol{\theta}) = d_i(N, v(\boldsymbol{\theta}))$. Therefore, optimal allocations in $[P_1]$ have an intuitive form: they coincide with the demand functions given in Eq. (2) when the unit price of each supplier is exactly his virtual cost. This follows because, like in classic mechanism design, the equilibrium ex-ante expected payment that the auctioneer makes to a bidder is equal to the ex-ante expectation of the virtual cost times the allocation.

It is important to note that, while the optimal demands are completely characterized, the optimal transfers are not. The only constraint imposed on transfers by the optimal solution is over interim expected transfers. As transfers are equal to unit price times demand, this implies that the optimal prices in the relaxed problem are underspecified. This freedom in the definition of optimal prices becomes useful later on, when we characterize the optimal solution to the original problem.

To illustrate the result, consider Example 3.1 and suppose both suppliers have the same cost distribution. Let θ_1 and θ_2 be the cost realizations of supplier 1 and 2 respectively. In this case, the relaxed problem P_1 yields an optimal allocation characterized by: (1) if $\delta > |v(\theta_1) - v(\theta_2)|$, the demand is splitted between the two suppliers with $x_1 = (v(\theta_2) - v(\theta_1) + \delta)/(2\delta)$ and $x_2 = (v(\theta_1) - v(\theta_2) + \delta)/(2\delta)$; and (2) if $\delta < |v(\theta_2) - v(\theta_1)|$, all the demand is awarded to the supplier with the lowest cost realization. Note that the decision of whether to split or not the demand depends on the cost realizations. The optimal solution single-awards if and only if the transportation cost is small relative to the differences in virtual costs. In this case, it is worth paying the cost of having less variety in the assortment with the upside of decreasing the expected payments to bidders. By single-awarding in some scenarios, the auctioneer can reduce these expected payments while still providing incentives for truthful cost revelation.

Because problem P_1 is a relaxation of P_0, the optimal objective of the former is an upper bound on the optimal objective of the latter. The next corollary provides necessary and sufficient conditions under which P_0 indeed attains the optimal objective of P_1.

Corollary 4.1. *Let $(\boldsymbol{x}, \boldsymbol{t})$ be the unique optimal solution to the relaxed problem P_1.*[8] *Define*

$$q_i(\boldsymbol{\theta}) = 1 \text{ if and only if } x_i(\boldsymbol{\theta}) > 0,\ \forall i \in N,\ \boldsymbol{\theta} \in \Theta. \tag{9}$$

Suppose that for all $\boldsymbol{\theta} \in \Theta$, there exists prices $\boldsymbol{p}(\boldsymbol{\theta})$ such that

$$\begin{aligned} x_i(\boldsymbol{\theta}) &= d_i(\boldsymbol{q}(\boldsymbol{\theta}), \boldsymbol{p}(\boldsymbol{\theta})) \qquad \forall i \in N,\ \forall \boldsymbol{\theta} \in \Theta && (10) \\ & \quad \text{where } d_i(\boldsymbol{p}) \text{ is given by Eq. } (2), \text{ and} \\ \sum_{\boldsymbol{\theta}_{-i} \in \Theta_{-i}} p_i(\theta_i, \boldsymbol{\theta}_{-i}) x_i(\theta_i, \boldsymbol{\theta}_{-i}) f_{-i}(\boldsymbol{\theta}_{-i}) &= T_i(\theta_i), \qquad \forall i \in N,\ \forall \theta_i \in \Theta_i\ , && (11) \end{aligned}$$

where, for all $i \in N$, $T_i(\cdot)$ is the expected interim transfer function given $t_i(\cdot)$. Then, the optimal objective of P_0 is equal to the optimal objective of P_1. Moreover, an optimal solution of P_0 is given by $(\boldsymbol{q}, \boldsymbol{p})$ characterized by Eqs. (9), (10), *and* (11), *and the corresponding optimal allocation $\boldsymbol{x}$ of P_0. Furthermore, the optimal objective of P_0 is equal to the optimal objective of P_1 if and only if such solution $(\boldsymbol{q}, \boldsymbol{p})$ exists.*

The corollary suggests the following approach to solving the optimal mechanism design problem. First, solve the relaxed problem, the solution of which has an appealing structure. Then, find unit prices that support the optimal relaxed solution.

5 Affine Demand Models

The optimal mechanism design problem takes an underlying consumer demand model as an input. To obtain analytical solutions we will restrict attention to a general class of affine demand models, that is, models in which for every set Q, demands d as given by Eq. (2) are (piece-wise) affine functions of prices. The advantage of these models is that they admit a convex and closed-form expression for consumer surplus. Even under affine demand models, the demand constraints are

[8] Problem P_1 admits a unique optimal solution for all demand systems considered in the paper. If P_1 admits more than one solution, our arguments can easily be extended accordingly.

piece-wise linear, and problem P_0 remains non-convex.[9] However, the approach described above of relaxing these constraints will allow solving the problem. Further, affine demand models capture a vast array of substitution patterns including both horizontal and vertical dimensions of differentiation.

In the remainder of this section, we discuss the solution to the optimal mechanism problem when we assume affine demand models. We first explain how to apply the general solution approach introduced in Section 4 to affine demand models. Next, we characterize the optimal mechanisms for specific linear consumer-choice models. We start by analyzing a popular affine-demand model: the Hotelling model of horizontal differentiation. Then, we provide the analysis of a general affine demand model that includes the Hotelling model (and a pure vertical model) as particular cases.

5.1 Applying the Solution Approach to Affine Demand Models

We now discuss how to adapt the general solution approach described in Section 4 to affine demand models. Equations (10) require that unit prices $\boldsymbol{p}$ induce the optimal allocations $\boldsymbol{x}$ of P_1 through the demand system. If the demand function is affine prices, these equations yield linear constraints in prices. Note that the equations are linear because they require to find prices to generate a *given* vector of demands $\boldsymbol{x}$. Equations (11) require that unit prices $\boldsymbol{p}$ induce the expected interim transfers T_i in the optimal solution of P_1. Given an optimal mechanism for P_1, $(\boldsymbol{x}, \boldsymbol{t})$, these equations are also linear in prices. Therefore, when an affine demand system is assumed, showing that P_0 attains the optimal objective of P_1 amounts to proving that a system of linear equations is consistent and admits a solution. We formalize this next.

Let $(\boldsymbol{x}, \boldsymbol{t})$ be an optimal solution to the relaxed problem P_1. By Proposition 4.1 and the discussion that follows the proposition, we have $x_i(\boldsymbol{\theta}) = d_i(N, \boldsymbol{v}(\boldsymbol{\theta}))$ where $\boldsymbol{v}(\boldsymbol{\theta})$ is defined as the vector of virtual costs, i.e., $\boldsymbol{v}(\boldsymbol{\theta}) = (v_1(\theta_1), \ldots, v_n(\theta_n))$. We denote $Q(\boldsymbol{\theta})$ as the set of active suppliers with strictly positive demands in the optimal solution under cost realizations $\boldsymbol{\theta}$. To satisfy the conditions of Corollary 4.1 we need to find unit prices such that $d_i(N, v(\boldsymbol{\theta})) = d_i(\boldsymbol{q}(\boldsymbol{\theta}), \boldsymbol{p}(\boldsymbol{\theta}))$ for all $i \in Q(\boldsymbol{\theta})$ and $\boldsymbol{\theta} \in \Theta$. This imposes $|Q(\boldsymbol{\theta})|$ constraints over the prices $\boldsymbol{p}(\boldsymbol{\theta})$, corresponding to firms with strictly positive demands. However, as the allocations must add up to one, one of these constraints is redundant; the demands for $|Q(\boldsymbol{\theta})| - 1$ suppliers determines the demand for the remaining active supplier. Therefore, the equations in (10) impose $|Q(\boldsymbol{\theta})| - 1$ constraints over prices $\boldsymbol{p}(\boldsymbol{\theta})$. The redundancy of one constraint plays an important role because it induces degrees of freedom that can be used to satisfy the constraints on expected interim transfers.

[9] For instance, consider the simple Hotelling model described in Example 3.1. There, the demand constraints for agent $i \in \{1, 2\}$ should be expressed as $x_i(\boldsymbol{\theta}) = \max\{0, \min\{1, \frac{p_j(\boldsymbol{\theta}) - p_i(\boldsymbol{\theta}) + \delta}{2\delta}\}\}$ with $j \in \{1, 2\}, j \neq i$, which yield a non-convex problem.

Let $\boldsymbol{A}_{ij}(\boldsymbol{\theta})$ denote the coefficient of $v_j(\theta_j)$ in the equation $d_i(N, v(\boldsymbol{\theta}))$. In all demand models considered in the paper, $\boldsymbol{A}_{ij}(\boldsymbol{\theta}) = 0$ for every $i \in Q(\boldsymbol{\theta})$ and $j \notin Q(\boldsymbol{\theta})$. This property is natural: if a supplier has zero demand, then its price does not play a role in the demand equations of competitors. Let $\iota(Q(\boldsymbol{\theta})) = \max\{i \in N : i \in Q(\boldsymbol{\theta})\}$. For a given $\boldsymbol{\theta}$ and a given $i \in Q(\boldsymbol{\theta})$ with $i \neq \iota(Q(\boldsymbol{\theta})$, the constraints imposed by Eqs. (10) can be expressed as:

$$\sum_{j \in Q(\boldsymbol{\theta})} \boldsymbol{A}_{ij}(\boldsymbol{\theta}) p_j(\boldsymbol{\theta}) = \sum_{j \in Q(\boldsymbol{\theta})} \boldsymbol{A}_{ij}(\boldsymbol{\theta}) v_j(\boldsymbol{\theta}) \qquad (M_i(\boldsymbol{\theta}))$$

We refer to the constraint associated with costs $\boldsymbol{\theta}$ and supplier $i \in Q(\boldsymbol{\theta})$ ($i \neq \iota(Q(\boldsymbol{\theta}))$ as $M_i(\boldsymbol{\theta})$. Note that any set of prices $\boldsymbol{p}(\boldsymbol{\theta})$ (for all $\boldsymbol{\theta} \in \Theta$) that satisfy all constraints in the set $\{M_i(\boldsymbol{\theta}) : \boldsymbol{\theta} \in \Theta,\ i \in Q(\boldsymbol{\theta}),\ i \neq \iota(Q(\boldsymbol{\theta})\}$ implement the optimal allocations given by the solution of P_1.

In addition, by Corollary 4.1, we need to guarantee that the expected interim transfers coincide with the optimal ones from P_1. We abuse notation and refer to the equality constraint on the expected transfers corresponding to supplier i and cost $\theta_i^j \in \Theta_i$ by $T_i(\theta_i^j)$. This constraint can be expressed as:

$$\sum_{\boldsymbol{\theta}_{-i} \in \Theta_{-i}} f_{-i}(\boldsymbol{\theta}_{-i}) x_i(\theta_i^j, \boldsymbol{\theta}_{-i}) p_i(\theta_i^j, \boldsymbol{\theta}_{-i}) \ = \ T_i(\theta_i^j) \qquad \forall i \in N,\ \forall \theta_i^j \in \Theta_i, \qquad (T_i(\theta_i^j))$$

where the right hand side corresponds to the optimal expected interim transfer of P_1 (given by Eq. (8)). Observe that, if in the optimal solution we have $x_i(\theta_i^j, \boldsymbol{\theta}_{-i}) = 0$ for all $\boldsymbol{\theta}_{-i} \in \Theta_{-i}$, then it must be that $T_i(\theta_i^j) = 0$. This follows by conditions (2) and (3) in Proposition 4.1. Hence, the previous equation imposes $\sum_{i \in N} \sum_{\theta_i \in \Theta_i} \mathbb{I}[\exists\ \boldsymbol{\theta}_{-i} :\ i \in Q(\theta_i, \boldsymbol{\theta}_{-i})]$ constraints. Note that these number of constraints is less than or equal to $\sum_{i \in N} |\Theta_i|$.

By imposing the constraints in Eqs. $(M_i(\boldsymbol{\theta}))$, the allocations $x_i(\theta_i^j, \boldsymbol{\theta}_{-i})$ are fixed and equal to the optimal allocations of P_1; therefore, the equations described in $(T_i(\theta_i^j))$ are linear in unit prices. Hence, verifying whether $OPT(P_0) = OPT(P_1)$ is equivalent to establishing whether the linear system of equations defined by Eqs. $(M_i(\boldsymbol{\theta}))$ (for all $\boldsymbol{\theta} \in \Theta$ and all $i \in Q(\theta)$ with $i \neq \iota(Q(\boldsymbol{\theta})$) and Eqs. $(T_i(\theta_i^j))$ (for all $i \in N$ and $\theta_i^j \in \Theta_i$) admits a solution.

Let $\boldsymbol{M}$ and $\boldsymbol{m}$ be the coefficient matrix and the corresponding RHS respectively defined by the linear equations in $(M_i(\boldsymbol{\theta}))$ and $(T_i(\theta_i^j))$, where each column is associated with a price $p_i(\boldsymbol{\theta})$. We can safely discard the columns corresponding to prices $p_i(\boldsymbol{\theta})$ such that $i \notin Q(\boldsymbol{\theta})$, as all the coefficients of such columns are zero. The resulting matrix $\boldsymbol{M}$ will have $\sum_{\theta \in \Theta} |Q(\boldsymbol{\theta})|$ columns as we have one price variable per active supplier and per profile of costs. In addition, for each $\theta \in \Theta$, there will be $|Q(\boldsymbol{\theta})| - 1$ rows given by the constraints in Eqs. $(M_i(\boldsymbol{\theta}))$ and $\sum_{i \in N} \sum_{\theta_i \in \Theta_i} \mathbb{I}[\exists\ \boldsymbol{\theta}_{-i} :\ i \in Q(\theta_i, \boldsymbol{\theta}_{-i})] \leq |\Theta|$ rows given by Eqs. $(T_i(\theta_i^j))$. The preceding observations are summarized by the following remark:

Remark 5.1 (Dimension of the coefficient matrix)**.** *The coefficient matrix* $\boldsymbol{M}$ *has* $\sum_{\theta\in\Theta}|Q(\boldsymbol{\theta})|$ *columns and* $\sum_{\theta\in\Theta}|Q(\boldsymbol{\theta})| - \Theta + \sum_{i\in N}\sum_{\theta_i\in\Theta_i}\mathbb{I}[\exists\ \boldsymbol{\theta}_{-i}:\ i\in Q(\theta_i,\boldsymbol{\theta}_{-i})]$ *rows. Further, the number of columns is greater or equal than the number of rows.*

By the Rouché-Frobenius theorem, a system of linear equations $\boldsymbol{Mp}=\boldsymbol{m}$ is consistent (has a solution) if and only if the rank of its coefficient matrix $\boldsymbol{M}$ is equal to the rank of its augmented matrix $[\boldsymbol{M}|\boldsymbol{m}]$. Note that whenever the rows of $\boldsymbol{M}$ are linearly independent the system is trivially consistent.

In the remainder of this section we show that (under additional conditions) we can guarantee that the associated system of equations is consistent. Hence, we can characterize the optimal mechanism.

5.2 Optimal Mechanism for Hotelling Demand Model

Having described the general solution approach, we now discuss the structure of the optimal mechanism when the consumer demand is given by a Hotelling model. Recall that a simple version of the Hotelling model was introduced in Example 3.1. We now briefly discuss a general Hotelling demand model with an arbitrary number n of suppliers in the unitary segment. The n potential suppliers are located at $0\leq\ell_1<\ell_2<\ldots<\ell_n\leq 1$ respectively; the location represents the horizontal characteristic of the product offered relative to the product space. The closer two suppliers are in the product space, the closer substitutes the products they offer are. The locations of the suppliers are assumed to be common-knowledge. A continuum of consumers, all of whom must buy one unit of product, are distributed on the product space. To simplify the exposition, we assume that consumers are uniformly distributed. However, our results can be easily extended to arbitrary distributions.

The utility consumer j obtains from buying the product offered by i is given by:

$$u_{ji}(p_i) = -\left(\delta|\ell_i-\ell_j|+p_i\right), \tag{12}$$

where δ is the transportation cost and ℓ_j is the position of j in the unit line.

Suppose that suppliers have fixed unit prices $\boldsymbol{p}=\{p_i\}_{i\in N}$. Then, the set of active suppliers with strictly positive demand is given by $Q(\boldsymbol{p}) = \{i\in N:\ p_i\leq\min_{k\neq i}\{p_k+\delta|\ell_k-\ell_i|\}\}$. If this condition is satisfied for a supplier i, then the consumers located in a neighborhood of ℓ_i prefer to buy from him than from any other supplier; hence, supplier i will be active.

For unit prices $\boldsymbol{p}$ and supplier $i\in Q(\boldsymbol{p})$, let $\varrho_{\boldsymbol{p}}(i)$ (resp. $\vartheta_{\boldsymbol{p}}(i)$) denote the supplier preceding (resp. following) i in $Q(\boldsymbol{p})$, that is, $\varrho_{\boldsymbol{p}}(i) =\ \max\{j\in Q(\boldsymbol{p}):\ j<i\}$ and $\vartheta_{\boldsymbol{p}}(i) =\ \min\{j\in$

$Q(\boldsymbol{p}): \ j > i\}$. Also, let $\iota(Q(\boldsymbol{p}))$ (resp. $\eta(Q(\boldsymbol{p}))$) denote the rightmost (resp. leftmost) supplier in $Q(\boldsymbol{p})$. Then, the aggregate demand for product i is given by:

$$d_i(\boldsymbol{p}) = \begin{cases} 0 & \text{if } i \notin Q(\boldsymbol{p}) \\ \ell_i + \frac{1}{2\delta}\left(p_{\vartheta_{\boldsymbol{p}}(i)} - p_i + \delta(\ell_{\vartheta_{\boldsymbol{p}}(i)} - \ell_i)\right) & \text{if } i = \eta(Q(\boldsymbol{p})) \\ \frac{1}{2\delta}\left(p_{\varrho_{\boldsymbol{p}}(i)} - p_i + \delta(\ell_i - \ell_{\varrho_{\boldsymbol{p}}(i)})\right) + & \text{if } i \in Q(\boldsymbol{p}),\ i \neq \eta(Q(\boldsymbol{p})),\ \iota(Q(\boldsymbol{p})) \\ \qquad \frac{1}{2\delta}\left(p_{\vartheta_{\boldsymbol{p}}(i)} - p_i + \delta(\ell_{\vartheta_{\boldsymbol{p}}(i)} - \ell_i)\right) & \\ \frac{1}{2\delta}\left(p_{\varrho_{\boldsymbol{p}}(i)} - p_i + \delta(\ell_i - \ell_{\varrho_{\boldsymbol{p}}(i)})\right) + (1 - \ell_i) & \text{if } i = \iota(Q(\boldsymbol{p})) \end{cases} \tag{13}$$

Note that by Proposition 4.1 and the discussion that follows, the optimal allocations in the relaxed problem P_1 for a cost realization $\boldsymbol{\theta}$ are given by the demand characterization (13) with prices equal to the vector of virtual costs $\boldsymbol{v}(\boldsymbol{\theta})$. Similarly to the Hotelling example with 2 suppliers, the auctioneer may optimally restrict participation of bidders in the assortment to decrease expected payments.

We now study in which cases it is possible to achieve the same optimal objective in both the original problem and the relaxed problem, that is, in which cases $OPT(P_0) = OPT(P_1)$. Consider the optimal solution of the relaxed problem as described by Proposition 4.1. Let q be defined as Corollary 4.1 as follows:

$$q_i(\boldsymbol{\theta}) = \begin{cases} 1 & \text{if } i \in Q(\boldsymbol{\theta}) \\ 0 & \text{otherwise} \end{cases}$$

By comparing the Hotelling demands as described by Eq. (13) with the optimal allocations of P_1 as defined in Proposition 4.1, it should be clear that the constraints given by Eqs. $(M_i(\boldsymbol{\theta}))$ can be summarized as:

$$p_{\vartheta_{\boldsymbol{\theta}}(i)}(\boldsymbol{\theta}) - p_i(\boldsymbol{\theta}) = v_{\vartheta_{\boldsymbol{\theta}}(i)}(\theta_{\vartheta_{\boldsymbol{\theta}}(i)}) - v_i(\theta_i) \qquad \forall \boldsymbol{\theta} \in \Theta,\ i \in Q(\boldsymbol{\theta}),\ i \neq \iota(\boldsymbol{\theta}). \tag{14}$$

These constraints will implement the optimal allocations of P_1 using prices $\boldsymbol{p}(\boldsymbol{\theta})$. In words, the difference in prices between adjacent active suppliers must be equal to the difference in virtual costs. By Corollary 4.1, we must also guarantee that the expected transfers agree with the optimal ones, that is, unit prices should satisfy constraints $T_i(\theta_i^j)$, for all $\theta_i^j \in \Theta_i$ and all $i \in N$. Hence, if we can find a feasible pair $(\boldsymbol{q}, \boldsymbol{p})$ for P_0 such that the optimal allocations for P_1 can be supported and the constraints on the expected interim transfers are maintained, then we have found an optimal solution for the original problem. To show that the system of linear equations is consistent, we exploit the fact that Eq. (14) imposes a very particular structure on the coefficient matrix of the system.

We start by analyzing the setting in which suppliers have IID costs and are located at equidistant intervals. Even in this context, the problem is asymmetric whenever we have three or more suppliers, as the most central agent has an advantage to capture demand. We have the following result.

Theorem 5.1 (IID costs)**.** *Consider the setting in which for all $i \in N$ we have $\ell_i = \frac{i-1}{n-1}$ (agents are located at equidistant intervals), $\Theta_i = \tilde{\Theta}$ and $f_i = f$ for some support $\tilde{\Theta}$ and pdf f. Then, $OPT(P_0) = OPT(P_1)$.*

The proof of Theorem 5.1 can be found in the companion appendix. We show that there is no gap between the optima of the original and the relaxed problem by showing that the system of linear equations $\boldsymbol{Mp} = \boldsymbol{m}$ is consistent. Ideally, one would like to show that the rows of the coefficient matrix are linearly independent. However, this need not be the case. Indeed, the reader can verify that in the simple case of $n = 2$, $\tilde{\Theta} = \{\theta_L, \theta_H\}$ and $\delta \geq \frac{1}{f(\theta_H)}(\theta_H - \theta_L)$ the rows of the associated matrix of coefficients are linearly dependent.

We now turn our attention to the more general case in which the cost functions are not IID and locations are arbitrary. Unfortunately, as opposed to our result in Theorem 5.1, the optima of both problems might not agree in the general case. This situation is illustrated by the following example.

Example 5.1 ($OPT(P_0) > OPT(P_1)$)**.** *Consider an instance with only two players located at the extremes of the unit segment. Let $\delta = 1$ be the transportation cost. Let $\Theta_1 = \{1, 2.5\}$, $\Theta_2 = \{1, 2, 2.3\}$. The probability functions f_1, f_2, and v_1, v_2 are described in the following tables.*

Θ_1	1	2.5
f_1	1/2	1/2
v_1	1	4

Θ_2	1	2	2.3
f_2	1/2	1/3	1/6
v_2	1	3.5	3.8

To show that a gap exists between both problems, we show that it is not possible to find item prices satisfying the conditions in Corollary 4.1. To that end, note that the set of possible outcomes is $\Theta = \{(1, 1), (1, 2), (1, 2.3), (2.5, 1), (2.5, 2), (2.5, 2.3)\}$. Whenever $\theta_1 = 1$ or $\theta_2 = 1$ (but not both), only the agent with cost 1 will be active in the optimal solution. Therefore, whenever agent 2 has cost $\theta_2 = 2$ he is only active in one profile, that is, in profile $(2.5, 2)$. By Eq. (11), the price $p_2(2.5, 2)$ is completely determined. In addition, Eq. (10) now complete determines price $p_1(2.5, 2)$. Similarly, when agent 2 has cost $\theta_2 = 2.3$ he is also active only in profile $(2.5, 2.3)$. Using the same arguments as before, Eq. (11) pins-down $p_2(2.5, 2.3)$ and hence Eq. (10) fixes price $p_1(2.5, 2.3)$. However, once the values of $p_1(2.5, 2)$ and $p_1(2.5, 2.3)$ are fixed as explained above, the expected transfer constraint for $T_1(2.5)$ fails to hold and a gap between both problems must exist. In the case, the optimal objective value of the relaxed and orginal problems are 2.0638 and 2.0645 respectively.

Given that, in general, the optima of the relaxed problem and the original problem may not agree, we next focus on providing sufficient conditions under which $OPT(P_0) = OPT(P_1)$. In particular, we provide sufficient conditions for the associated system of linear equations to be consistent. This is summarized by the following theorem.

Theorem 5.2. *Consider the general setting in which agents have arbitrary locations and costs distributions. Let* $c^* = \min_{1 \le i \le n-1}(\ell_{i+1} - \ell_i)$*. Suppose that the following conditions are simultaneously satisfied:*

1. *There is at least one profile* $\boldsymbol{\theta} \in \Theta$ *such that* $|v_{i+1}(\theta_{i+1}) - v_i(\theta_i)| \le \delta(\ell_{i+1} - \ell_i)/2$ *for all* $i \in N$*; and*

2. $|\Theta_i| \ge 3$ *for all* $i \in N$*, and for every* $i \in N$ *and every* $\theta^j \in \Theta_i$*, we have* $v_i(\theta_i^{j+1}) - v_i(\theta_i^j) \le \frac{\delta c^*}{4}$.

Then, we have $OPT(P_0) = OPT(P_1)$*.*

The second condition essentially requires the difference in the virtual costs between adjacent points in the support to be bounded by a function of δ. The smaller the δ, the closer the virtual costs should be. If we think of the discrete distribution as an approximation of an underlying continuous distribution, then this is equivalent to require the discretization to be thin enough with respect to δ. The intuition behind the first condition is to require the existence of an 'interior solution'. First, note that every agent is active in profile $\boldsymbol{\theta}$. Further, using the second condition, it is simple to observe that all agents are active in a neighborhood of $\boldsymbol{\theta}$.

The complete proof of Theorem 5.2 can be found in Appendix E.[10] In the proof, we show that the rows of the associated coefficient matrix $\boldsymbol{M}$ are linearly independent and, therefore, there must exist prices that support the optimal allocation and satisfy the expected interim transfer constraints.

It is easy to verify that condition (2) in the theorem is violated in Example 5.1. In particular, $|\Theta_1| = 2$ and, furthermore, the difference between consecutive virtual costs in general exceeds $\frac{\delta c^*}{4} = \frac{1}{4}$. Intuitively, the support of the cost distributions in the example are coarse and, therefore, the dimensionality of the price vectors is low. As a result, there are not enough degrees of freedom to find prices that simultaneously satisfy the demand and the expected interim transfers constraints. The second condition of the theorem guarantees this is always the case. In particular, by requiring adjacent virtual costs to be "close", the optimal allocations do not vary much if we replace the cost of an agent by one of his adjacent costs. Then, for a pair $\theta_i^j, \theta_i^{j+1} \in \Theta_i$, there exists at least some profile $\boldsymbol{\theta}_{-i}$ for which we have $i \in Q(\theta_i^j, \boldsymbol{\theta}_{-i})$ and $i \in Q(\theta_i^{j+1}, \boldsymbol{\theta}_{-i})$. This is crucial, as it guarantees a structural relationship between the expected transfers constraints (Eq. $(T_i(\theta_i^j))$) of adjacent costs (e.g., $(T_i(\theta_i^j))$ and $(T_i(\theta_i^{j+1}))$). Further, by imposing conditions (1) and (2), we guarantee the existence of several cost profiles for which all agents are active, which translates into a structural

[10] In Appendix E we prove a more general theorem. Then, we explain how the general theorem implies Theorem 5.2.

relationship between the expected transfers constraints (Eq. ($T_i(\theta_i^j)$)) of all the agents. As the prices become more related with each other, there are more degrees of freedom to find prices that satisfy both the optimal demand constraints and the expected transfer constraints.

In Appendix B, we provide a related characterization and result for a classic model of pure vertical differentiation.

5.3 Optimal mechanisms for general Affine Demand models

So far we considered the classic models of demand for products that are horizontally (or vertically) differentiated. We now study more general affine demand models, that allow us to combine both vertical and horizontal sources of differentiation. An affine demand function is one where the relation $\boldsymbol{d}(\boldsymbol{p}) = \boldsymbol{\alpha} - \boldsymbol{\Gamma}\boldsymbol{p}$ holds for all $\boldsymbol{p} \in \{\boldsymbol{p} \in \mathbb{R} : \boldsymbol{\alpha} - \boldsymbol{\Gamma}\boldsymbol{p} \geq 0\}$. Here, $\alpha \geq 0$ represents a quality (or vertical) component; $\boldsymbol{\Gamma}_{ij}$ represents the variation in the demand of product i as a result of a unit change in the price of product j, when all other prices remain constant. We assume that the products are gross substitutes, hence, $\boldsymbol{\Gamma}_{ij} \leq 0$ for $i \neq j$. Note that the Hotelling model presented in the previous section and the vertical model studied in the appendix are both particular cases of affine demand models.

For our purposes, it is important to consider the extension of this specification to price vectors under which some products get zero demand as introduced by Shubik and Levitan [23] and further analyzed by Soon et al. [24]. We formalize this extension in our setting in which demands must add up to one using a representative consumer approach (Farahat and Perakis [12] also use this approach to study oligopolistic pricing models under affine demand functions).[11] We consider a representative consumer with a strictly concave gross utility function given by $u(\boldsymbol{x}) = \boldsymbol{c}'\boldsymbol{x} - \frac{1}{2}\boldsymbol{x}'\boldsymbol{D}\boldsymbol{x}$, where $\boldsymbol{D}$ is a positive definite matrix and $\boldsymbol{D}^{-1}$ is symmetric positive definite. The vector $\boldsymbol{c}'$ denotes the transpose of vector $\boldsymbol{c}$. Here, $\boldsymbol{D} = \boldsymbol{\Gamma}^{-1}$ and $\boldsymbol{c} = \boldsymbol{\Gamma}^{-1}\boldsymbol{\alpha}$ have been renamed to avoid burdensome notation. The demand function is defined as the solution of the representative consumer's maximization problem, whose utility also corresponds to consumer surplus. That is, for any $\boldsymbol{p} \in \mathbb{R}^n$, let $\boldsymbol{d}(\boldsymbol{p})$ be defined as the solution of the following maximization problem:

$$
\begin{aligned}
\max_{\boldsymbol{x}} \quad & \boldsymbol{c}'\boldsymbol{x} - \frac{1}{2}\boldsymbol{x}'\boldsymbol{D}\boldsymbol{x} - \boldsymbol{p}'\boldsymbol{x} \\
s.t \quad & \mathbf{1}'\boldsymbol{x} = 1 \\
& \boldsymbol{x} \geq 0
\end{aligned}
\tag{LD($\boldsymbol{p}$)}
$$

Clearly, Problem (LD($\boldsymbol{p}$)) has a unique solution for very $\boldsymbol{p} \in \mathbb{R}^n$, and thus the demand function

[11] Alternatively, a general affine demand model can be micro-founded considering consumers' individual utilities like in the Hotelling and vertical models [20]. However, we think the representative consumer approach provides a cleaner analysis.

$\boldsymbol{d}(\boldsymbol{p})$ is well defined. To illustrate, we consider the following example:

Example 5.2. *We consider a duopoly where* $\boldsymbol{c} = (\alpha_1, \alpha_2)$ *and* $\boldsymbol{D} = \left(\begin{smallmatrix} \beta_1 & \gamma \\ \gamma & \beta_2 \end{smallmatrix}\right)$, *with all the parameters positive except possibly* γ *and with* $\beta_1 + \beta_2 \geq 2\gamma$. *For any given* $\boldsymbol{p}$, *the demand function* $\boldsymbol{d}(\boldsymbol{p})$ *is defined as:*

$$d_1(\boldsymbol{p}) = \max\left\{0,\ \min\left\{\frac{\alpha_1 - \alpha_2 - \gamma + \beta_2 - p_1 + p_2}{\beta_1 + \beta_2 - 2\gamma},\ 1\right\}\right\}$$

and

$$d_2(\boldsymbol{p}) = \max\left\{0,\ \min\left\{\frac{\alpha_2 - \alpha_1 - \gamma + \beta_1 - p_2 + p_1}{\beta_1 + \beta_2 - 2\gamma},\ 1\right\}\right\}.$$

These demand functions exhibit natural properties; they are decreasing in a firm's own price and increasing in the competitor's price. Also, depending on the price vector, there could be one or two firms active. As before, we note that the optimal allocations in the relaxed problem P_1 for a cost realization $\boldsymbol{\theta}$ are given by the demand characterization above with prices equal to the vector of virtual costs $\boldsymbol{v}(\boldsymbol{\theta})$. Again, the auctioneer may optimally restrict participation of bidders in the assortment when the difference in virtual costs is large enough to decrease expected payments. We generalize the demand specification for a larger number of firms next.

Through the rest of the section, given a matrix $\boldsymbol{A}$, we denote the i^{th} row of $\boldsymbol{A}$ by $\boldsymbol{A}_{i,*}$. Similarly, the j^{th} column is denoted by $\boldsymbol{A}_{*,j}$. For a subset of indices $Q \subset N$, $\boldsymbol{A}_Q$ denotes the principal submatrix of $\boldsymbol{A}$ obtained by selecting only the rows and columns in Q. Similarly, $\boldsymbol{c}_Q$ denotes the vector obtained by selecting only the components in Q and $\mathbf{1}_Q$ denotes the vector of ones of dimension $|Q|$. We have the following result that characterizes an affine demand function for the set of active suppliers.

Lemma 5.1. *Given a price vector* $\boldsymbol{p}$ *and the associated demand* $\boldsymbol{d}(\boldsymbol{p})$, *we denote by* $Q = Q(\boldsymbol{p}) = \{i \in N : d_i(\boldsymbol{p}) > 0\}$. *Then, demand* $\boldsymbol{d}(\boldsymbol{p})$ *can be expressed as:*

$$\boldsymbol{d}_Q(\boldsymbol{p}_Q) = (\boldsymbol{D}_Q)^{-1}\left(\boldsymbol{c}_Q - \boldsymbol{p}_Q + \left(\frac{1 - \mathbf{1}_Q'(D_Q)^{-1}\left(\boldsymbol{c}_Q - \boldsymbol{p}_Q\right)}{\mathbf{1}_Q'(\boldsymbol{D}_Q)^{-1}\mathbf{1}_Q}\right)\mathbf{1}_Q\right). \tag{15}$$

The proof is presented in Appendix D. The above demand specification exhibits a natural regularity property: if there is no demand for a particular product, the price of that product does not affect the demand for other products. In addition, it is simple to observe that any increase in price of a product with zero demand will not have an impact on the demand function either.

From Eq. (15), it should be clear that whenever two vector of prices $\boldsymbol{p}_Q$ and $\hat{\boldsymbol{p}}_Q$ satisfy

$$(\boldsymbol{D}_Q)^{-1}\left(\boldsymbol{p}_Q - \frac{\mathbf{1}_Q'(D_Q)^{-1}\boldsymbol{p}_Q}{\mathbf{1}_Q'(\boldsymbol{D}_Q)^{-1}\mathbf{1}_Q}\mathbf{1}_Q\right) = (\boldsymbol{D}_Q)^{-1}\left(\hat{\boldsymbol{p}}_Q - \frac{\mathbf{1}_Q'(D_Q)^{-1}\hat{\boldsymbol{p}}_Q}{\mathbf{1}_Q'(\boldsymbol{D}_Q)^{-1}\mathbf{1}_Q}\mathbf{1}_Q\right), \tag{16}$$

we must have that $\boldsymbol{d}_Q(\boldsymbol{p}_Q) = \boldsymbol{d}_Q(\hat{\boldsymbol{p}}_Q)$. This observation is useful because it states that demands only depend on price differences. This freedom in setting unit prices is essential to our proof technique, as we will find unit prices that satisfy the same differences induced by the virtual costs and that simultaneously satisfy the expected interim transfer constraints.

Hence, the coefficient matrix $\boldsymbol{M}$ as described in Section 4 will consist, for $\boldsymbol{\theta} \in \Theta$ and each $i \in Q(\boldsymbol{\theta})$ of at most $Q(\boldsymbol{\theta})$ non-zero rows: $Q(\boldsymbol{\theta}) - 1$ correspond to the demand equations[12] and the remaining one corresponding to the expected transfer constraint. Note that for given $\boldsymbol{\theta} \in \Theta$, the demand equations are given by Eq. (16) where we replace Q by $Q(\boldsymbol{\theta})$ and $\boldsymbol{p}_Q$ by $\boldsymbol{p}_{Q(\boldsymbol{\theta})}(\boldsymbol{\theta})$ in the left hand side. In the right hand side we replace prices $\hat{\boldsymbol{p}}_Q$ by virtual costs $\boldsymbol{v}_{Q(\boldsymbol{\theta})}(\boldsymbol{\theta})$.

Similarly to what we have previously done for the vertical and the horizontal differentiation models, we show that, under sufficient conditions, we can guarantee $OPT(P_0) = OPT(P_1)$, by showing that the rows of the associated matrix of coefficients of the system of linear equations $\boldsymbol{M}$ are linearly independent.

Theorem 5.3. *Consider the general setting in which agents have arbitrary costs distributions. Suppose that the following conditions are simultaneously satisfied:*

1. *There exists a profile $\boldsymbol{\theta} \in \Theta$ such that $Q(\boldsymbol{\theta}) = N$, and there exists a $d^* \in \mathbb{R}$ such that, for all $\boldsymbol{\theta}' \in \Theta$ with $|\boldsymbol{\theta} - \boldsymbol{\theta}'|_\infty \leq d^*$ we have $Q(\boldsymbol{\theta}') = N$.*

2. *$|\Theta_i| \geq 3$ for all $i \in N$, and for every $i \in N$ and every $\theta^j \in \Theta_i$, we have $v_i(\theta_i^{j+1}) - v_i(\theta_i^j) \leq d^*/3$.*

we have $OPT(P_0) = OPT(P_1)$.

We highlight that d^* depends on the primitives of the problem. However, the intuition agrees with that of the Hotelling and vertical models: we must guarantee the existence of an 'interior solution' and impose a 'thin enough' cost discretization. To provide more intuition, consider a duopoly where $\boldsymbol{c} = (\alpha, \alpha)$ and $\boldsymbol{D} = \left(\begin{smallmatrix} \beta & \gamma \\ \gamma & \beta \end{smallmatrix}\right)$. Note that this is a particular case of Example 5.2. In this case, the result will follow for any market satisfying the conditions with $d^* = \frac{\beta - \gamma}{2}$.[13]

6 Case Study: ChileCompra-Style Framework Agreements

In the previous section, we characterized the optimal directed-revelation posted-price mechanism. Procurement mechanisms used in the real-world, however, typically take simpler forms. In particular, FAs are usually implemented as first-price auctions with rules to decide the assortment

[12] Note that if we can find prices p_Q satisfying the constraints imposed by $x_1, \ldots, x_{|Q|-1}$, then the last constraint will also be satisfied as $x_Q = 1 - \sum_{j=1}^{|Q|-1} x_j$.

[13] Note that the choice of d^* imposes conditions on the cost distributions.

based on both the suppliers' bids and characteristics. Simple mechanisms are generally preferred in practice, because they are easier to explain to potential suppliers and require simpler management from the procurement agency.

In this section, we evaluate the type of framework agreements run by ChileCompra and provide concrete recommendations for their improvement. The optimal mechanism is crucial for this purpose, because it serves as a benchmark of what is achievable, and its structure also provides insights on how to modify the current practice to enhance performance.

In Section 6.1 we describe the FAs run by ChileCompra. Then, in Section 6.2 we provide analytical results in a simple model of horizontal differentiation evaluating the performance of the type of FAs run by ChileCompra and improvements thereof. Finally, in Section 6.3 we provide a large set of numerical experiments showing the robustness of the conclusions drawn from the analytical results in the simple model. Overall, our analysis shows that ChileCompra FAs creates thin markets and that, by emulating the optimal mechanism to make close-substitute products compete, consumer surplus can be significantly increased.

6.1 ChileCompra's Framework Agreements

Framework agreements have been playing an increasingly important role in the procurement strategy of the Chilean government since their introduction in 2004. To illustrate, in 2013 ChileCompra spent slightly more than US$ 2 billion in FAs, which corresponded to 21% of the total public expenditure in procurement and was twice the amount spent in 2010. Currently, more than 95 thousand products and services ranging from food to office supplies and computers, dialysis services and medicines can be acquired through FAs.

To award the FAs in a given category (e.g., food), ChileCompra runs a first-price-auction-type mechanism. Each supplier submits a bid for each *item*, which stands for a completely specified product. To illustrate, a box of Kellogg's Corn Flakes containing 15oz. and a box of Kellogg's Corn Flakes containing 17oz. are two different items. Suppliers can offer any subset of items they want, as long as the type of these products are among those required by the government. As an example, if "pasta" is among the types of products required, a bid for any type of pasta is allowed, regardless the brand, size, and so on.

Bids are evaluated using a scoring rule; all products whose scores are above a threshold are offered in the menu at the price specified by the supplier in his bid. Each item gets a score in the $0-100$ scale. *All items for which the score is at least* 75 *points will be in the menu.* Usually, around 70 points correspond to price and this is typically the main variable that influences the allocation rule.[14] Prices are compared only across *identical* items. The price-score for an item-supplier pair

[14]The remaining points correspond to supplier characteristics, such as certifications or transportation facilities.

is assigned by comparing his price to the minimum price of an identical item. If there is a unique supplier offering the item, he automatically obtains the maximum score *regardless* of the price.

The current FA implementation creates thin markets and fails to generate price competition to be included in the menu. Indeed, because auctions are run independently for each product and the product definitions are narrow, there is a single supplier bidding for many products. To illustrate we consider the FA for food products.[15] There, a total of 8091 products where offered by 116 suppliers. Out of those items, 4549 were offered by a *unique* supplier who got the maximum score for price in these item. The price score accounted for 72 points out of the minimum of 75 needed. As a result, *all* items with a single supplier were added to the menu. Hence, in these case bidders have hardly any incentives to aggressively compete in prices to be inside the menu. This issue notwithstanding, note that prices will affect competition inside the menu. In other words, there is competition "in the market", but not "for the market".

These observations motivated the following question: can the performance of the current FAs be improved if thicker markets are created by making imperfect substitute products compete to be in the menu? In other words, can competition for the market improve performance?

We highlight that other FAs, such as office supplies, prosthesis supplies, cleaning products, and personal care, among others, are similar to the FA for food products in that they create thin markets. In these agreements, it is typically important for the government to account for the heterogeneous preferences of agencies, by optimizing consumer surplus when deciding the menu. For example, some patients might find a prosthesis of a certain brand to be more comfortable than that of a competing brand, while for other patients it might work the other way around. Different organizations buying from the food FA might also have different needs, such as dietary constraints (e.g., hospitals and environments with kids). In addition, requirements for computers from schools may be very different to those from, say, the central bank. In all these cases, the government has a direct interest on providing variety to its agencies and our results provide insights into how to achieve it in a cost efficient way.[16]

6.2 Analytical Evaluation of ChileCompra-Style FAs in Simple Model

Supported both by the description of ChileCompra's mechanism and the analysis of their data, we propose the following first order approximation to their current FAs: we consider a procurement

[15] This FA corresponds to the public auction number $2239 - 20 - LP09$, titled "ALIMENTOS PERECIBLES Y NO PERECIBLES", which was valid 2010 through 2014.

[16] In other contexts, it may not be the government's responsibility to provide variety. For example, suppose agencies have idiosyncratic preferences for different types of soft drinks. In this case, our results provide a way of evaluating the cost of providing variety considering agencies idiosyncratic preferences, when perhaps it is not in the government's best interest to do it. Either way, in the rest of the section we assume the government's objective is to maximize consumer surplus.

mechanism in which there is no competition to be in the menu, but suppliers must compete for demand inside the menu. Every supplier whose price does not exceed the reserve price is added to the menu, and the bids of those suppliers are taken as posted prices. After, the demand is split amongst the agents in the menu according to the demand model.

Following the auction theory tradition, we assume that for a given mechanism bidders play a *pure strategy Bayesian Nash equilibrium* (BNE). Hence, to evaluate the performance of the FA we need to derive such equilibrium bidding strategies. Unfortunately, deriving such strategies analytically under general model primitives is challenging, because demands, and therefore profits, are a function of all bids. Therefore, to compute expected profits a bidder needs to integrate out over all possible demand realizations given competitors' bid functions.

Hence, to be able to derive analytical results we restrict our attention to a simple pure horizontal differentiation Hotelling model. We consider a problem with two IID potential sellers located at 0 and 1 respectively in the unit line and with two cost realizations. Let $\Theta_i = \{\theta_L, \theta_H\}$ for $i = 1, 2$ and let f_L and f_H denote $f(\theta_L)$ and $f(\theta_H)$, respectively. This simple model will provide essential insights. Then, we test the robustness of these insights with numerical experiments. All proofs in this section can be found in the companion appendix.

6.2.1 Analysis of ChileCompra-Style FAs

In this section we provide a theoretical analysis of the equilibrium bid functions and the performance of ChileCompra FA's in the two-by-two Hotelling model just described. In this setting, we say that the outcome of the mechanism is *single-award* if, whenever agents have different types, the low-cost agent obtains all the demand. Otherwise, we say that the outcome of the mechanism is *split-award*.[17]

We now analyze the performance of ChileCompra's mechanism in this setting. A full description of the mechanism is provided in Table 1. We can analytically calculate the optimal bidding strategies for the agents under the ChileCompra mechanism with reserve price θ_H. Using standard arguments, it is straightforward to verify that the equilibrium bid for a high-type agent is θ_H. The following proposition characterizes the bid for the low type.

Proposition 6.1. *The equilibrium bidding strategy for an agent of type θ_L in the unique BNE is as specified in Table 2 for different values of δ.*

[17] We highlight that the terms single-award and split-award have been used in the literature with a different meaning. In auctions where production is awarded, the outcome can be a sole-source (single) award, in which a single producer provides all of the required production, or in a split award, in which production is divided between two or more firms [5]. In the FA we are analyzing instead, the auctioneer cannot award demand but we can still think about the outcome of the mechanism as a single or split award, depending on how the final demand is distributed.

	ChileCompra's mechanism	Restricted-entry (RE) mechanism
Parameters	R: reserve price b_1, b_2: agents' bids	R: reserve price C: split b_1, b_2: agents' bids
Entry rule	If $b_i \leq R$, add i to the menu.	Only consider suppliers with bids at most R. If $\lvert b_1 - b_2 \rvert < C$, add both to the menu. Otherwise, just add the one with lowest bid.
Demand allocation	Split the demand among the suppliers in the menu according to the Hotelling demand with transportation cost δ	Split the demand among the suppliers in the menu according to the Hotelling demand with transportation cost δ.

Table 1: Description of the mechanisms considered: ChileCompra's mechanism and the restricted-enty mechanism

<table>
<tr><th rowspan="2">Value of δ</th><th colspan="2">Optimal</th><th colspan="2">ChileCompra</th></tr>
<tr><th>award</th><th>avg. low price</th><th>award</th><th>eq. strat. low</th></tr>
<tr><td>$\left[\frac{1}{f_H}(\theta_H - \theta_L), \infty\right)$</td><td>split</td><td>$\frac{(f_H/2+f_L(1-x))\theta_H+(x-1/2)\theta_L}{f_L/2+f_H x}$
where $x = \frac{1/f_H(\theta_H-\theta_L)+\delta}{2\delta}$</td><td rowspan="3">split</td><td rowspan="2">θ_H</td></tr>
<tr><td>$\left[(\theta_H - \theta_L), \frac{1}{f_H}(\theta_H - \theta_L)\right]$</td><td rowspan="5">single</td><td rowspan="5">$\frac{\theta_L+f_H\theta_H}{1+f_H}$</td></tr>
<tr><td>$\left[\frac{(\theta_H-\theta_L)}{2+f_H}, (\theta_H - \theta_L)\right]$</td><td>$\frac{\theta_L+f_H\theta_H+\delta}{1+f_H}$</td></tr>
<tr><td>$\left[\frac{f_L}{2}(\theta_H - \theta_L), \frac{(\theta_H-\theta_L)}{2+f_H}\right)$</td><td rowspan="2">single</td><td>$\theta_H - \delta$</td></tr>
<tr><td>$\left[\frac{f_H f_L(\theta_H-\theta_L)}{\frac{1}{2}(1+f_H)^2+f_H f_L}, \frac{f_L}{2}(\theta_H - \theta_L)\right]$</td><td>$\theta_L + \delta\frac{1+f_H}{f_L}$</td></tr>
<tr><td>$\left[0, \frac{f_H f_L(\theta_H-\theta_L)}{\frac{1}{2}(1+f_H)^2+f_H f_L}\right]$</td><td>no BNE</td><td>-</td></tr>
</table>

Table 2: Comparison Optimal mechanism and ChileCompra mechanism with reserve price θ_H. In all cases, the expected price for an item of cost θ_H is θ_H.

In Table 2, we compare the equilibrium bidding strategy for the low-type agent in ChileCompra with reserve price θ_H to the average price per unit payed to a supplier of type θ_L in the optimal mechanism.[18] Note that for low-values of δ a BNE does not exist for the same reasons a BNE does not typically exist in first-price auctions with discrete types [16]. The bidding strategy of a bidder of type θ_L in the ChileCompra FA is continuous as a function of δ. Intuitively, one would expect the bidding strategy to be increasing in δ; as the differentiation increases, demands are less affected by prices and therefore incentives to bid aggressively decrease. This intuition is correct for most values of δ except for the interval $\left[\frac{f_L}{2}(\theta_H - \theta_L), \frac{(\theta_H-\theta_L)}{2+f_H}\right)$, in which equilibrium bids are decreasing in δ. There, it is optimal to bid $\theta_H - \delta$; the low-type has incentives to submit this bid to gather demand against a high-type. In addition, for values of $\delta \geq \theta_H - \theta_L$ the equilibrium bid reaches the maximum allowed bid of θ_H.

Using the equilibrium bids, we can compute the expected consumer surplus (corresponding to the negative of expected supplier payments plus transportation cost) of the ChileCompra mecha-

[18]The prices given by the optimal mechanism are not unique. Therefore, we calculate the average price per unit payed to a supplier of type low as $T(\theta_L)/X(\theta_L)$.

nism and compare it to that of the optimal mechanism for different parameter values. To compare performance in this section, for a given a mechanism M, we define the *optimality gap* between the optimal mechanism and mechanism M as $(M/OPT-1)*100$, where we abuse notation and denote by M and OPT the total expected consumer surplus in mechanism M and the optimal mechanism, respectively. Optimality gaps are shown in Table 3 as a function of both f_L and δ.

A key difference between ChileCompra mechanism and the optimal mechanism is that the split-award outcome occurs more frequently in the former one, especially for higher values of f_L. In particular, whenever $\delta \geq \frac{(\theta_H-\theta_L)}{2+f_H}$, ChileCompra will split-award between both suppliers. In contrast, the optimal mechanism only split-awards once $\delta \geq \frac{1}{f_H}(\theta_H-\theta_L)$. As previously observed, the optimal mechanism single-awards to reduce expected payments to bidders. It is also interesting to observe that, in this setting, the full-information solution splits-award whenever $\delta \geq \theta_H - \theta_L$. Relative to this solution, ChileCompra mechanism split-awards for a wider range of differentiation costs, while the optimal mechanism does it for a smaller range.[19]

δ	$f_L = 0.1$		$f_L = 0.25$		$f_L = 0.5$		$f_L = 0.75$		$f_L = 0.9$	
	Chile	BRE	Chile	BRE	Chile	BRE	Chile	BRE	Chile	BRE
0.5	0.88	0.88	2.41	2.41	5.59	5.58	4.02	3.94	4.30	4.30
1	0.74	0.27	1.98	0.93	4.34	2.73	7.02	5.62	8.40	7.91
1.5	0.86	0.11	2.36	0.55	5.45	2.23	9.41	5.36	12.18	8.55
2	0.89	0.10	2.55	0.55	6.36	2.38	11.62	6.04	15.70	9.97
2.5	0.71	0.12	2.15	0.64	5.74	2.72	11.12	6.93	15.40	15.40
3	0.58	0.18	1.77	0.74	5.15	2.99	10.55	7.55	14.99	14.99
3.5	0.50	0.28	1.50	0.88	4.57	3.12	10.01	7.93	14.60	14.60
4	0.43	0.40	1.30	1.09	4.00	3.20	9.47	8.05	14.24	14.24
4.5	0.38	0.38	1.14	1.14	3.50	3.27	8.95	8.17	13.86	13.86
5	0.34	0.34	1.02	1.02	3.11	3.11	8.44	8.06	13.49	13.49
5.5	0.30	0.30	0.91	0.91	2.79	2.79	7.94	7.94	13.13	13.13
6	0.27	0.27	0.83	0.83	2.53	2.53	7.46	7.46	12.78	12.78

Table 3: Optimality gaps as a function of both the differentiation cost δ and f_L. The parameters are $\theta_L = 10, \theta_H = 12$. The horizontal lines indicate the point up to which restricting the entry outperforms ChileCompra's policy.

In the next section we explore how a simple change of the rules of ChileCompra mechanism that restricts entry can improve performance. This change emulates the optimal mechanism, in that it makes the single-award outcome more likely by inducing competition to be in the market.

6.2.2 Analysis of Mechanisms that Restrict Entry

Following ChileCompra's original design, we focus on first-price auction-type of mechanisms. We consider two possible changes in the auctions' rules. In the first case, entry is restricted ex-ante, that is, before observing the bids. In the second case, entry is restricted ex-post as a function of

[19] [9] also observe that by making a monopoly outcome more likely, the auctioneer is able to exert more pressure on low-cost suppliers to reveal their private information and therefore reduce the total expected procurement cost.

the observed bids.

Ex-Ante Restricted-Entry Mechanism. We start by analyzing what happens if we restrict entry before bids are placed. In particular, suppose that we decide how many agents will be in the menu before observing the bids and then run a first-price auction (FPA) type mechanism to decide the prices.

In our simple model, this amounts to deciding when does choosing a single winner using a FPA outperforms ChileCompra's mechanism. Recall that, in general, the FPA does not have an equilibrium in pure strategies when types are discrete. However, by allowing equilibria in mixed strategies, expected payments in the FPA are given by $\theta_H - f_L^2(\theta_H - \theta_L)$.[20] By adding the transportation cost, the total expected cost faced by a designer who chooses to run a FPA is $\theta_H - f_L^2(\theta_H - \theta_L) + \frac{\delta}{2}$.

Using these analytical expressions, we can characterize the set of parameters for which the FPA outperforms ChileCompra. To illustrate, for fixed $\theta_L = 10$ and $\theta_H = 12$, the relative performance of ChileCompra and FPA as a function of parameters (f_L, δ) can be seen in Figure 2. The black area is omitted from the analysis, as no equilibrium in pure strategies exists in ChileCompra's mechanism. As it can be observed, ChileCompra outperforms the FPA mechanisms when both f_L and the differentiation cost δ are relatively small (the white area).

As the differentiation cost increases beyond $\theta_H - \theta_L$ but f_L remains small, the FPA is still worse than ChileCompra. In that region (light gray area), the equilibrium strategy for the low-type in ChileCompra mechanism is to bid θ_H, which agrees with the bid a low-type agent will place if there was no competition. However, the designer cannot improve by switching to a FPA; in the light gray area, the reduction in purchasing costs that results from the price competition cannot compensate for the large transportation cost, even when bids in the ChileCompra mechanism are as high as possible.

On the other hand, as f_L increases, it is profitable to restrict the entry using a FPA even if that implies a higher transportation cost (gray area).[21] Overall, even a simple FPA where the number of winners is decided before observing the bids can sometimes improve over the current mechanism. However, there is still a large set of parameters for which this is not the case. We discuss the performance of more sophisticated mechanisms next.

Ex-Post Restricted-Entry Mechanism. The main issue with restricting entry ex-ante is that such mechanisms do not split-award when suppliers share the same cost. This causes an increase

[20]This follows from standard arguments. For completeness, the proof is provided in the companion appendix.

[21]We note that the non-convexity of the areas FPA and ChileCompra is due to the fact that, in ChileCompra, the equilibrium bidding strategy as a function of δ is decreasing in the interval $\left[\frac{f_L}{2}(\theta_H - \theta_L), \frac{1}{2+f_H}(\theta_H - \theta_L)\right]$.

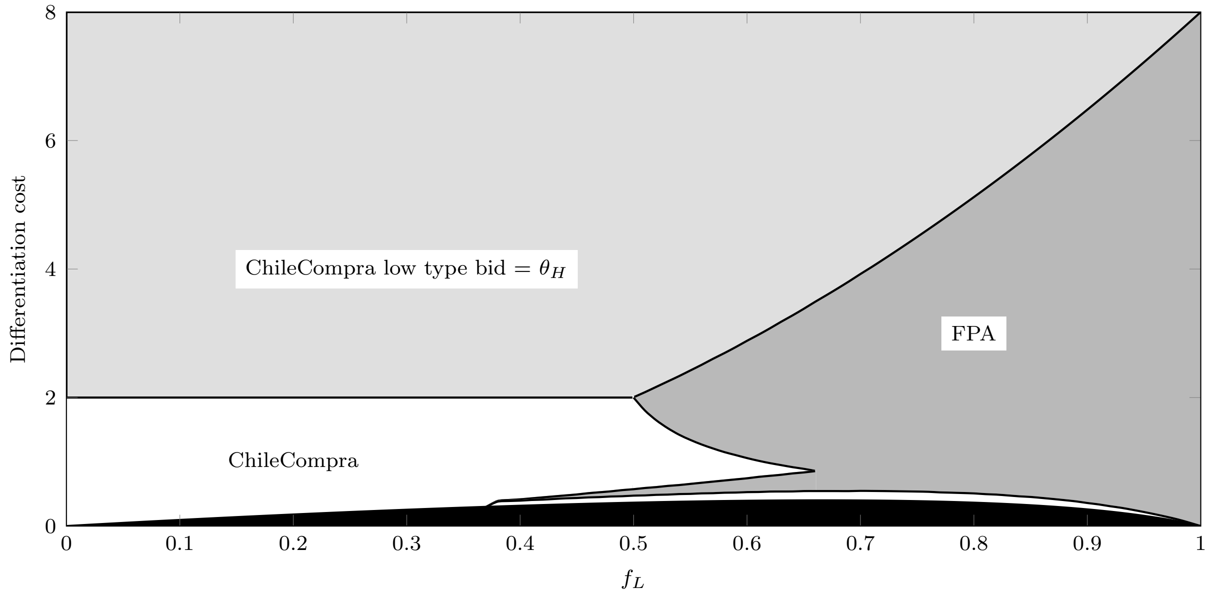

Figure 2: For $\theta_L = 10$, $\theta_H = 12$, we show when it is profitable to restrict the entry using a FPA as a function of f_L and δ. The black area is omitted from the analysis, as no equilibrium in pure strategies exists in the ChileCompra mechanism. ChileCompra outperfoms the FPA mechanisms only in the white area. The single-winner FPA is better in dark gray area. In the light-gray area, ChileCompra has the highest possible low-type bid, but it is still better than a single-winner FPA.

in the transportation cost. Therefore, we now study a class of mechanisms for which the decision on whom will be in the menu is contingent on the bids received by the auctioneer.

Using the intuition from the optimal mechanism, we propose the restricted-entry (RE) mechanism described in Table 1. This simple mechanism has two parameters: the reserve price R and the split parameter C. The only difference with ChileCompra's mechanism is that we restrict the entry to the menu by requiring the difference between bids to be at most C. The split parameter allow us to quantify how restrictive the entry to the market should be; whenever $C = \delta$, our mechanism coincides with ChileCompra's.[22] We assume $R = \theta_H$.

For the set of parameters in which ChileCompra single-awards, it can be shown that the performance of the mechanism cannot be improved by restricting entry.[23] Therefore, our focus is in the settings in which ChileCompra split-awards. For these cases, we find values of C for which the equilibrium bid of the low-type induces single-award. A candidate for such equilibrium bid for the low-type is $\theta_H - C$, because it is the highest bid that results in single-award. We have the following result.

[22]Whenever $C = 0$, our mechanism agrees with a FPA. However, in this section we are only going to consider split parameters C for which a BNE exists.

[23]The proof can be found in the companion appendix.

Proposition 6.2. *For every set of parameters f_L, θ_H, θ_L and δ, there exists a (possibly empty) interval $\mathcal{I}$ such that, for all $C \in \mathcal{I}$, we have that $\theta_H - C$ is the unique equilibrium bid for the low type in the RE mechanism with reserve price θ_H and split parameter C.*

In the companion appendix, we characterize the intervals referred to in the previous proposition as a function of f_L, θ_H, θ_L, and δ. Intuitively, if C is too small, the mechanism is similar to a FPA in which bidders have incentives to undercut each other and a BNE may not exist. On the other hand, if C is too big, an agent of type θ_L might prefer to place a bid greater that $\theta_H - C$ even if that implies splitting the demand with a high-type agent.

For given model primitives, the designer is interested in maximizing consumer surplus. If restricting entry is a helpful device to achieve this objective, then the auctioneer will choose the largest C for which a single-award equilibrium exists, because that induces the lowest bid for the low type. Hence, we define the "*best low-type bid*" to be $\theta_H - C^*$, where C^* is the highest C for which $\theta_H - C$ is an equilibrium. The characterization of the best low-type bids can be found in the companion appendix, but we briefly discuss the intuition. Whenever $f_L \leq \frac{2}{3}$, for every δ there exists (at least) one C for which $\theta_H - C$ is an equilibrium. As f_L increases beyond that point, the length of the interval for which we can guarantee this equilibrium decreases. Intuitively, the advantage of bidding at $\theta_H - C$ is to capture the whole demand when the other agent has a high cost. As f_L becomes close to one, this advantage vanishes. Furthermore, the best low-type bid is increasing in f_L and δ.

As previously pointed out, the RE mechanism may improve upon the ChileCompra mechanism in cases in which the latter split-awards and the former single-awards. However, the performance of a RE mechanism that single-awards can be worse than that of ChileCompra because single-award increases the transportation cost. We define the *best restricted-entry mechanism* (BRE) as the mechanism obtained by choosing the value of the split parameter C that maximizes consumer surplus. We obtain the following straightforward result.

Proposition 6.3. *For a given set of parameters, the BRE has one of two possible forms: (1) coincides with the ChileCompra mechanism; or (2) uses the value C^* associated to the best low-type bid.*

For a given set of parameters, if BRE improves over ChileCompra it must be by restricting entry. Then, case (2) above is optimal as the best low-type bids maximize the consumer surplus for the single-award case. If restricting entry does not improve over ChileCompra because of the increase in transportation cost, then (1) above is optimal. Note that the latter is equivalent to setting $C = \delta$.

To illustrate, in Figure 3 we plot the outcome of the optimal, ChileCompra and BRE mechanisms

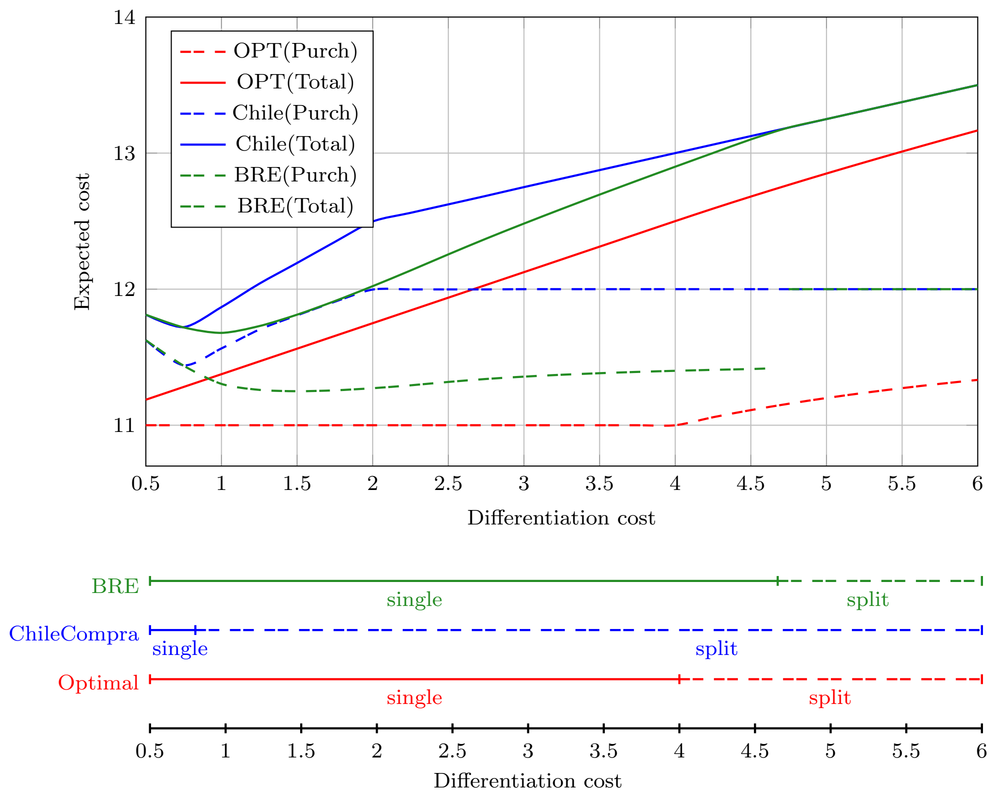

Figure 3: (*Top*) Expected purchasing and total (purchasing plus transportation) costs for optimal, ChileCompra and best restricted-entry (BRE) mechanisms as a function of the differentiation (transportation) cost δ. The parameters are $\theta_L = 10, \theta_H = 12$, $f_L = f_H = 1/2$. (*Bottom*) Single-award vs. split-award in optimal, ChileCompra's, and our best mechanism.

as a function of the differentiation (transportation) cost δ for parameters $\theta_L = 10$, $\theta_H = 12$, $f_L = f_H = 1/2$. As it can be observed, the BRE mechanism restricts the entry whenever $\delta \leq 4.675$. By doing so, an expected purchasing cost which is much closer to the optimal one can be obtained. However, when δ exceeds 4.675, the savings obtained in the purchases cannot compensate for the increase in transportation cost and, therefore, BRE and ChileCompra coincide beyond that point.

More generally, we study when RE outperforms ChileCompra as a function of the parameters. Recall that restricting entry achieves a substantial decrease in the low-type equilibrium bid. However, restricting entry also increases the transportation cost. Using the analytical expressions for ChileCompra's equilibrium bid and the best low-type bid, we characterize the boundary between when does restricting entry improves overall consumer surplus. We find that, for relatively small values of δ and regardless of the value of other parameters, restricting entry improves over ChileCompra's mechanism. Intuitively, the decrease in the low-type equilibrium bid results in a considerable decrease in the expected purchasing cost without a major increase in the expected

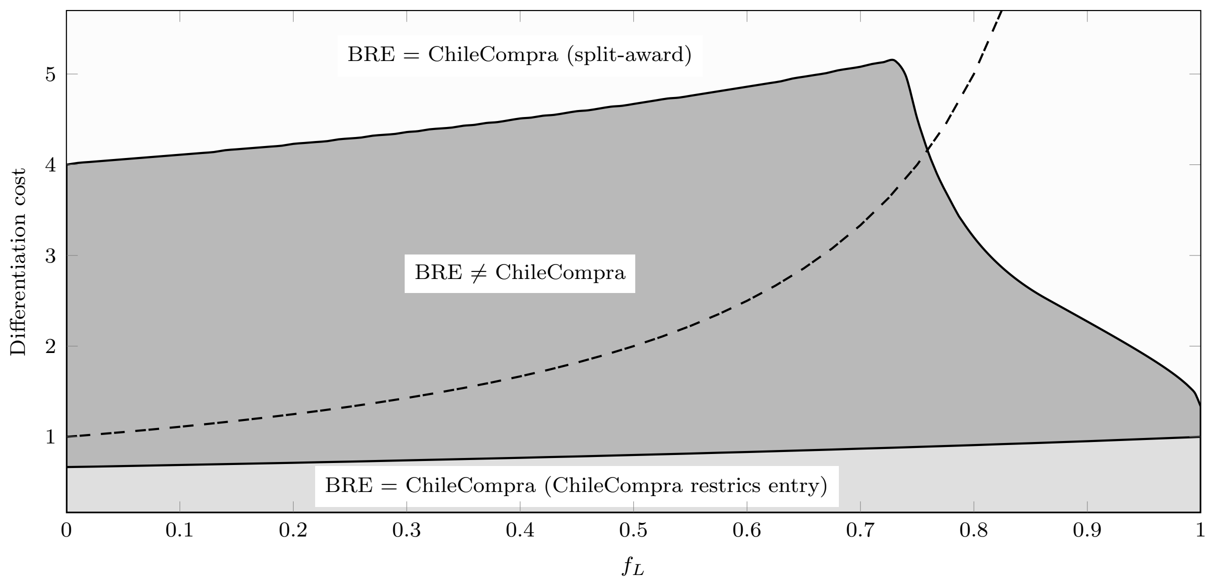

Figure 4: For $\theta_L = 10$, $\theta_H = 12$, we show when it is profitable to restrict the entry as a function of the differentiation cost δ and f_L. The dashed line represents the cutoff between single and split award in the optimal mechanism (i.e., $\delta = \frac{1}{f_H}(\theta_H - \theta_L)$).

transportation cost. In addition, as it can be observed in Table 3, restricting entry performs better for the middle-values of f_L. If f_L is too low, the savings are less likely to occur and therefore the potential impact is smaller. On the other hand, if f_L is too high, the best-low-type-bid tends to increase and the single-award becomes less profitable.

This is illustrated by Figure 4, where we fix $\theta_L = 10$, $\theta_H = 12$, and show when it is profitable to restrict entry as a function of δ and f_L. The graph is divided in three regions. In the white region, ChileCompra's mechanism cannot be improved upon by only restricting entry and, therefore, BRE and ChileCompra performances coincide. Similarly, whenever $\delta \leq \frac{1}{2+f_H}(\theta_H - \theta_L)$, the entry is also restricted in ChileCompra's mechanism; hence BRE and ChileCompra coincide again. The most relevant case corresponds to those combinations of (f_L, δ) which lie in the dark-gray colored area, for which the consumer surplus generated by ChileCompra can be increased by restricting entry. Finally, recall that the optimal mechanism will split award only if $\delta \geq \frac{1}{f_H}(\theta_H - \theta_L)$. This is represented by the dashed-line.[24]

We conclude this section with a note on the practical implementability of the restricted entry mechanisms. The BRE mechanism uses the best split-parameter C that depends on the problem primitives, and therefore, it may be hard to estimate in practice. However, we argue that even

[24]We note that these results are by just considering split parameters for which BNE exists. If we are also allowed to consider split parameters for which only equilibria in mixed strategies exist, then restricting the entry can improve upon ChileCompra for a wider range of parameters.

implementing the BRE mechanism with a rough estimate of the best C (but not the exact one) typically improves performance. In particular, for values of $f_L \leq 2/3$, the best C as a function of δ asymptotically converges to $\frac{f_H(\theta_H-\theta_L)}{1+f_H}$. Furthermore, any smaller C will induce the equilibrium bid $\theta_H - C$ as long as $C \geq \frac{f_L}{2+f_L}(\theta_H - \theta_L)$.[25] Hence, by choosing a conservative C the auctioneer should be able to increase consumer surplus for values of δ that are not too close to the boundary between the areas in which restricting entry improves performance (the boundary between gray and white regions in Figure 4). As the benefits of restricting entry are smaller close to that boundary, one should make sure that the C used is large enough as to guarantee a low-type bid that can compensate for the extra transportation cost. Therefore, in such areas, it is better to overestimate C (recall that any C larger than the best C yields the same outcome as the current mechanism, so it will not damage the worst-case performance).

6.3 Robustness Results: Numerical Experiments

To test the robustness of our intuition, in this section we numerically solve for the equilibrium strategies for ChileCompra and the restricted-entry mechanism and compare the total procurement cost of these mechanism with that of the optimal. We replicate this simulation exercise for a range of environments by varying the cost support, the cost distribution, the number of bidders and some of the other parameters of the model. The results are summarized next.

More General Cost Distributions. We first consider adding more points to the support of the cost distributions. To that end, we consider an initial interval and discretize it evenly into k costs, for $k = 2, 3, 5, 7$. We consider 4 types of distributions: uniform, left-skewed, right-skewed, and symmetric-unimodular (normal-like). We highlight that, even though now we have multiple costs in the support, the auctioneer still must pick a unique split-parameter that remains fixed throughout the mechanism.

The results of our simulation show that the intuition for the multiple-costs case coincides with that of the two-by-two simple model. Notably, restricting entry improves the performance of the current mechanism. In general, the optimality gap decreases by at least 40% for relatively small values of δ, and the differences in gap shrinks as δ increases. Similarly to the two-by-two case, the benefits are greater when the distribution is left-skewed or normal-like. There, restricting entry achieves greater reduction in the bids of the low-type. In addition, as the number of values in the support increases, restricting entry improves performance for even higher values of cost-differentiation, because the auctioneer can use a more refined splitting rule.

[25] Again, this is formally shown in the companion appendix.

Larger Number of Bidders. We now consider models with more than two agents. To that end, we consider n agents at equidistant locations with agent i located at $\ell_i = (i-1)/(n-1)$. We test our results for $n \in \{2, 3, 4, 5\}$. The costs are still assumed to be IID across agents; however, agents are not ex-ante symmetric due to their locations.

Regarding ChileCompra's performance, two conclusions can be drawn. First, the distribution of costs has the same impact in the performance as in the two-agent case; ChileCompra performs close to optimal for small values of f_L but, as f_L increases, the performance of ChileCompra quickly deteriorates. Second, the optimality gap increases with the number of agents. The intuition seems to be the same as in the two-agent case; ChileCompra fails to obtain competitive bids for the low-type (relative to the optimum), and this lack of competition has a higher impact as the number of suppliers increases. These results also generalize to more general supports.[26]

Whenever there are more than two agents, the auctioneer can choose whether to restrict entry as a function of bids or as a joint function of both bids and product characteristics. We discuss these two options next.

We first consider restricting the entry as a function of both the bids and product characteristics. The mechanism we consider is the restricted-entry mechanism with the following modification: for split parameter C and bids $b_1, \ldots, b_n$, supplier i will be in the menu only if $b_i - b_j < C * |\ell_i - \ell_j|$ for every $j \in N$ with $j \neq i$. This rule is intuitive: it induces more price competition for agents that are close-by in the product space. In accordance to what is observed in the two-agent case, restricting the entry is more efficient for the sets of parameters in which ChileCompra split-awards. However, the advantage of restricting the entry can only be observed for higher values of δ as the number of agents increases; this is due to the fact that, for a fixed δ, the expected demand for a high-type supplier deceases (and non-linearly) as the number of agents increases and therefore the role of the high-type suppliers becomes less significant. For the values of δ in which ChileCompra split-awards, restricting the entry performs better (with respect to the optimum) than in the two agent case. This intuition also generalizes for more general supports; the benefits are greater when the distribution is left-skewed or normal-like. In general, for fixed number of agents and cost distributions, the optimality gap decreases by at least 40%.

Next, we consider restricting the entry solely as a function of bids. Similarly to the case of two agents, for split parameter C and bids $b_1, \ldots, b_n$, supplier i will be in the menu only if $b_i - b_j < C$ for every $j \in N$ with $j \neq i$. As this rule is less sophisticated than the previous one, a poorer performance is to be expected. However, when suppliers can only be of one of two types, the performance is almost the same as that of the previous mechanism. However, as the number of

[26] In particular, ChileCompra's performance is closer to that of the optimal mechanism for right-skewed distributions in comparison to left-skewed distributions, and the performance in normal-like distributions is worse than in the uniform case but better than in the left-skewed case.

costs in the support increases, the advantage of taking into account the product characteristics also increases.

7 Conclusions and Extensions

In this paper we study procurement mechanisms for differentiated products demanded by heterogeneous consumers. First, we characterize the optimal mechanism for important classes of demand models. Second, we use these results to shed light on the FAs run by the Chilean government. Our results are useful to improve our understanding of FAs and, more generally, of buying mechanisms in similar contexts.

Our basic model can be extended in several interesting directions. First, to simplify the exposition, we assumed that each supplier offers one product. In Appendix C.1, we provide an extension to our model in which we allow for multi-product suppliers. We show that our main results extend to this setting, so we are able to characterize the optimal mechanism for the multi-product case.

In our basic model we assume an inelastic total demand, which may be reasonable for some products, like medicines, but perhaps less so for others. In Appendix C.2, we consider a model with an elastic total demand. We show that, in general, our main result fails to hold and a gap between the optima of the original and the relaxed problem exists. However, preliminary computational results show that this gap is typically small and that the assortments are usually similar in both the relaxed and the original problems.

We are currently further exploring the impact of demand elasticity in our results, as well as of assuming other demand systems, such as a nested logit model. Also, we plan to use econometric techniques to estimate important parameters of the model, such as those related to the underlying preferences of the agencies in the Chilean procurement setting, with the objective of sharpening our design recommendations. Overall, we hope that these extensions and our current results will help improve how FAs and other related mechanisms work in practice.

References

[1] D. Ackerberg, C. Benkard, S. Berry, and A. Pakes. Econometric tools for analyzing market outcomes. *Handbook of Econometrics*, 2006.

[2] G. L. Albano and M. Sparro. A simple model of framework agreements: Competition and efficiency. *Journal of Public Procurement*, 8(3):356–378, 2008.

[3] S. P. Anderson, A. De Palma, and J. F. Thisse. *Discrete choice theory of product differentiation.* MIT press, 1992.

[4] J. J. Anton and P. J. Gertler. Regulation, local monopolies and spatial competition. *Journal of Regulatory Economics*, 25(2):115–141, 2004.

[5] J. J. Anton and D. A. Yao. Split awards, procurement, and innovation. *The RAND Journal of Economics*, pages 538–552, 1989.

[6] T. F. Bresnahan. Competition and collusion in the american automobile industry: The 1955 price war. *The Journal of Industrial Economics*, pages 457–482, 1987.

[7] A. Chaturvedi, D. R. Beil, and V. Martínez-de Albéniz. Split-award auctions for supplier retention. *Management Science*, 2014.

[8] E. COMMISSION. Annual public procurement implementation review, 2012.

[9] J. J. Dana and K. E. Spier. Designing a private industry : Government auctions with endogenous market structure. *Journal of Public Economics*, 53(1):127–147, January 1994.

[10] A. de Estudios e Inteligencia de Negocios. Informe de gestión, diciembre 2013. *Dirección ChileCompra*, 2014.

[11] W. Elmaghraby. Supply contract competition and sourcing policies. *Manufacturing & Service Operations Management*, 2(4):350–371, 2000.

[12] A. Farahat and G. Perakis. A nonnegative extension of the affine demand function and equilibrium analysis for multiproduct price competition. *Operations Research Letters*, 38(4):280 – 286, 2010.

[13] Y. Gur, L. Lu, and G. Weintraub. Framework agreements in procurement: An auction model and design recommendation. Working paper, 2013.

[14] R. G. Hansen. Auctions with endogenous quantity. *The RAND Journal of Economics*, pages 44–58, 1988.

[15] A. G. Kök, M. L. Fisher, and R. Vaidyanathan. Assortment planning: Review of literature and industry practice. In *Retail supply chain management*, pages 99–153. Springer, 2009.

[16] V. Krishna. *Auction theory*. Academic press, 2009.

[17] J. Levin. An optimal auction for complements. *Games and Economic Behavior*, 18(2):176–192, 1997.

[18] C. Li and L. Debo. Second sourcing vs. sole sourcing with capacity investment and asymmetric information. *Manufacturing & Service Operations Management*, 11(3):448–470, 2009.

[19] G. Li and P. Rusmevichientong. A greedy algorithm for the two-level nested logit model. *Operations Research Letters*, 2014.

[20] S. Martin. Microfoundations for the linear demand product differentiation model. *Working Paper, Purdue*, 2009.

[21] T. G. McGuire and M. H. Riordan. Incomplete information and optimal market structure public purchases from private providers. *Journal of Public Economics*, 56(1):125–141, January 1995.

[22] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[23] M. Shubik and R. Levitan. *Market Structure and Behavior*. Harvard University Press, Cambridge, MA, 1980.

[24] W. Soon, G. Z. W., and J. Zhang. Complementarity demand functions and pricing models for multi-product markets. *Eur. J. Appl. Math*, 20(5):399–430, 2009.

[25] J. Tirole. *The Theory of Industrial Organization*. MIT press, 1988.

[26] V.-A. Truong. Optimal selection of medical formularies. *Journal of Revenue & Pricing Management*, 13(2):113–132, 2014.

[27] X. Vives. *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press, 2001.

[28] R. Vohra. *Mechanism Design: A Linear Programming Approach*. Econometric Society Monographs. Cambridge University Press, 2011. ISBN 9781139499170.

[29] A. Wolinsky. Regulation of duopoly: managed competition vs regulated monopolies. *Journal of Economics & Management Strategy*, 6(4):821–847, 1997.

A Proof of Proposition 4.1

Proof of Proposition 4.1. This proof uses the standard arguments from mechanism design theory introduced in Myerson's seminal paper [22]. Since the supports of our cost distributions are discrete, we follow the version of these arguments presented by Vohra [28]. Throughout this proof, we define m_i to be the number of costs in the support of agent i, that is, $m_i = |\Theta_i|$.

We start by re-stating the IC and IR constraints in P_1 in terms of the expected allocations and transfers:

$$\begin{aligned} \max_{\boldsymbol{x},\boldsymbol{t}} \quad & \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n}\left[k_i(x(\boldsymbol{\theta})) - t_i(\boldsymbol{\theta})\right]\right] \\ \text{s.t.} \quad & T_i(\theta_i) - X_i(\theta_i)\theta_i \geq T_i(\theta_i') - X_i(\theta_i')\theta_i \quad \forall i,\ \forall \theta_i, \theta_i' \in \Theta_i \\ & T_i(\theta_i) - X_i(\theta_i)\theta_i \geq 0 \quad \forall i,\ \forall \theta_i \in \Theta_i \\ & \sum_{i\in N} x_i(\boldsymbol{\theta}) = 1 \qquad \forall \boldsymbol{\theta} \in \Theta, \qquad x_i(\boldsymbol{\theta}) \geq 0 \qquad \forall i \in N,\ \boldsymbol{\theta} \in \Theta, \end{aligned}$$

Recall that $\Theta_i = \{\theta_i^1, ..., \theta_i^{m_i}\}$. If we add a dummy type per agent $\theta_i^{m_i+1}$ such that $X_i(\theta_i^{m_i+1}) = 0$ and $T_i(\theta_i^{m_i+1}) = 0$, then we can fold the IR constraints into the IC constraints:

$$T_i(\theta_i^j) - X_i(\theta_i^j)\theta_i^j \geq T_i(\theta_i^k) - X_i(\theta_i^k)\theta_i^j \quad \forall j \in \{1, ..., m_i\},\ \forall k \in \{1, ..., m_{i+1}\}\ .$$

Applying Theorem 6.2.1 in [28] for our procurement setting we obtain that an allocation $\boldsymbol{x}$ is implementable in Bayes Nash equilibrium if and only if $X_i(\cdot)$ is monotonically decreasing for all $i = 1, ..., n$. [27] Further, by Theorem 6.2.2 in [28], all IC constraints are implied by the following local IC constraints:

$$\begin{cases} T_i(\theta_i^j) - X_i(\theta_i^j)\theta_i^j \geq T_i(\theta_i^{j+1}) - X_i(\theta_i^{j+1})\theta_i^j & (BNIC_{i,\theta}^d) \\ T_i(\theta_i^j) - X_i(\theta_i^j)\theta_i^j \geq T_i(\theta_i^{j-1}) - X_i(\theta_i^{j-1})\theta_i^j & (BNIC_{i,\theta}^u) \end{cases}$$

[27]Note that the results cited in Vohra are for IID bidders, but the extension to bidders with different distributions is straightforward.

Therefore, we can re-write the problem as:

$$
\begin{aligned}
\max_{\boldsymbol{x},\boldsymbol{t}} \quad & \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n} k_i(x(\boldsymbol{\theta}))\right] - \sum_{i=1}^{n}\sum_{j=1}^{m_i} f_i(\theta_i^j)T_i(\theta_i^j) && (obj)\\
\text{s.t.} \quad & T_i(\theta_i^j) - X_i(\theta_i^j)\theta_i^j \geq T_i(\theta_i^{j+1}) - X_i(\theta_i^{j+1})\theta_i^j \quad \forall i \in N,\ \forall j \in \{1,...,m_i\} && (BNIC_{i,j}^d)\\
& T_i(\theta_i^j) - X_i(\theta_i^j)\theta_i^j \geq T_i(\theta_i^{j-1}) - X_i(\theta_i^{j-1})\theta_i^j \quad \forall i \in N,\ \forall j \in \{2,...,m_i\} && (BNIC_{i,j}^u)\\
& 0 \leq X_i(\theta^{m_i}) \leq \ldots \leq X_i(\theta^1), \quad \forall i \in N && \text{(M)}\\
& \sum_{i=1}^{n} x_i(\boldsymbol{\theta}) = 1 \qquad \forall \boldsymbol{\theta} \in \Theta, \qquad x_i(\boldsymbol{\theta}) \geq 0 \qquad \forall i \in N,\ \boldsymbol{\theta} \in \Theta.
\end{aligned}
$$

In addition, using standard arguments, we can show that all downward constraints $(BNIC_{i,j}^d)$ bind in the optimal solution.[28] Hence,

$$
T_i(\theta_i^j) - X_i(\theta_i^j)\theta_i^j = T_i(\theta_i^{j+1}) - X_i(\theta_i^{j+1})\theta_i^j \quad \forall i \in N,\ \forall j \in \{1,...,m_i\}.
$$

Further, it is simple to show that in this case, the upward constraints $(BNIC_{i,j}^u)$ are satisfied. Applying the previous equation recursively we obtain:

$$
T_i(\theta_i^j) = \theta_i^j X_i(\theta_i^j) + \sum_{k=j+1}^{m_i} (\theta^k - \theta^{k-1})X_i(\theta_i^k)\ . \tag{17}
$$

[28] A formal proof can be obtained by trivially adapting the Lemma 6.2.4 in Vohra to the procurement case.

Replacing in the objective:

$$
\begin{aligned}
obj &= \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n} k_i(\boldsymbol{x}(\boldsymbol{\theta}))\right] - \sum_{i=1}^{n}\sum_{j=1}^{m_i} f_i(\theta_i^j)T_i(\theta_i^j) \\
&= \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n} k_i(\boldsymbol{x}(\boldsymbol{\theta}))\right] - \sum_{i=1}^{n}\sum_{j=1}^{m_i} f_i(\theta_i^j)\left(\theta_i^j X_i(\theta_i^j) + \sum_{k=j+1}^{m_i}(\theta_i^k - \theta_i^{k-1})X_i(\theta_i^k)\right) \\
&= \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n} k_i(\boldsymbol{x}(\boldsymbol{\theta}))\right] - \sum_{i=1}^{n}\sum_{j=1}^{m_i} f_i(\theta_i^j)\left(\theta^j X_i(\theta_i^j)\right) - \sum_{i=1}^{n}\sum_{j=1}^{m_i}\sum_{k=0}^{m_i-1} f_i(\theta_i^j)\left(\mathbb{I}\{k \geq j\}(\theta_i^{k+1} - \theta_i^k)X_i(\theta_i^{k+1})\right) \\
&= \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n} k_i(\boldsymbol{x}(\boldsymbol{\theta}))\right] - \sum_{i=1}^{n}\sum_{j=1}^{m_i} f_i(\theta_i^j)\left(\theta^j X_i(\theta_i^j)\right) - \sum_{i=1}^{n}\sum_{k=1}^{m_i} F_i(\theta_i^{k-1})(\theta_i^k - \theta_i^{k-1})X_i(\theta_i^k) \\
&= \sum_{\boldsymbol{\theta}\in\Theta} f(\boldsymbol{\theta})\left(\sum_{i=1}^{n} k_i(\boldsymbol{x}(\boldsymbol{\theta}))\right) - \sum_{i=1}^{n}\sum_{j=1}^{m_i} f_i(\theta_i^j)\left(\left(\theta^j + \frac{F_i(\theta_i^{j-1})}{f_i(\theta_i^j)}(\theta_i^j - \theta_i^{j-1})\right) X_i(\theta_i^j)\right) \\
&= \sum_{\boldsymbol{\theta}\in\Theta} f(\boldsymbol{\theta})\left(\sum_{i=1}^{n} k_i(\boldsymbol{x}(\boldsymbol{\theta}))\right) - \sum_{i=1}^{n}\sum_{\theta_i\in\Theta_i} f_i(\theta_i)v_i(\theta_i)X_i(\theta_i) \\
&= \sum_{\boldsymbol{\theta}\in\Theta} f(\boldsymbol{\theta})\left(\sum_{i=1}^{n} k_i(\boldsymbol{x}(\boldsymbol{\theta})) - v_i(\theta_i)x_i(\boldsymbol{\theta})\right)
\end{aligned}
$$

The equations follow by simple algebra. In particular, the fourth equation follows by changing the order of summations.

Therefore, if we find an allocation such that for all $\boldsymbol{\theta} \in \Theta$ and $i \in N$,

$$
\begin{aligned}
x(\boldsymbol{\theta}) \quad &\in \quad \operatorname{argmax} \sum_{i=1}^{n}\left(k_i(\boldsymbol{x}(\boldsymbol{\theta})) - v_i(\theta_i)x_i(\boldsymbol{\theta})\right) \\
&\text{s.t.} \qquad \sum_{i=1}^{n} x_i(\boldsymbol{\theta}) = 1, \qquad x_i(\boldsymbol{\theta}) \geq 0 \quad \forall i \in N\ ;
\end{aligned}
$$

and such that the interim expected allocations are monotonic for all $i \in N$, that is, $X_i(\theta) \geq X_i(\theta')$ for all $\theta \leq \theta' \in \Theta_i$; and that the interim expected transfers satisfy Eqs. (17), for all $i \in N$ and $\theta \in \Theta_i$, then we have found an optimal solution. □

B Optimal mechanisms for Vertical Demand Model

We now consider a classic model of pure vertical differentiation (see, e.g., [6]). There are n potential suppliers, supplier i offering a product of quality α_i. We assume, w.l.o.g., that $\alpha_1 < \ldots < \alpha_n$. The qualities of the products are common-knowledge. There is a continuum of consumers, all wishing to buy one unit of the good (so the market is covered), uniformly distributed on the consumer-type space $Z = [0, 1]$. The type of a consumer indicates her value for quality. In particular, the utility a consumer of type $j \in Z$ obtains from consuming the product offered by supplier i at price p_i is given by:

$$u_{ji}(p_i) = j\alpha_i - p_i, \tag{18}$$

Given a set of potential suppliers with fixed unit prices $\boldsymbol{p} = \{p_i\}_{i \in N}$, the set of active suppliers with strictly positive demand is given by:

$$Q(\boldsymbol{p}) = \left\{ i \in N : \max_{j \in Z} \min_{k \neq i} \{ j\,(\alpha_i - \alpha_k) - (p_i - p_k) \} > 0 \right\}.$$

Namely, a supplier $i \in N$ will be active only if there exists a $j \in Z$ for which $u_{ji}(p_i) > u_{jk}(p_k)$ for all $k \in N$ with $k \neq i$.

As in the previous section, for unit prices $\boldsymbol{p}$ and agent $i \in Q(\boldsymbol{p})$, let $\varrho_{\boldsymbol{p}}(i)$ (resp. $\vartheta_{\boldsymbol{p}}(i)$) denote the agent preceding (resp. following) i in $Q(\boldsymbol{p})$, that is, $\varrho_{\boldsymbol{p}}(i) = \max\{j \in Q(\boldsymbol{p}) : j < i\}$ and $\vartheta_{\boldsymbol{p}}(i) = \min\{j \in Q(\boldsymbol{p}) : j > i\}$. Also, let $\iota(Q(\boldsymbol{p}))$ (resp. $\eta(Q(\boldsymbol{p}))$) denote the rightmost (resp. leftmost) agent in $Q(\boldsymbol{p})$. Then, the expected demand for product i is given by:

$$d_i(\boldsymbol{p}) = \begin{cases} 0 & \text{if } i \notin Q(\boldsymbol{p}) \\ 1 & \text{if } Q(\boldsymbol{p}) = \{i\} \\ \frac{p_{\vartheta_{\boldsymbol{p}}(i)} - p_i}{\alpha_{\vartheta_{\boldsymbol{p}}(i)} - \alpha_i} & \text{if } i = \eta(Q(\boldsymbol{p})) \\ \frac{p_{\vartheta_{\boldsymbol{p}}(i)} - p_i}{\alpha_{\vartheta_{\boldsymbol{p}}(i)} - \alpha_i} - \frac{p_i - p_{\varrho_{\boldsymbol{p}}(i)}}{\alpha_i - \alpha_{\varrho_{\boldsymbol{p}}(i)}} & \text{if } i \in Q(\boldsymbol{p}),\ i \neq \eta(Q(\boldsymbol{p})),\ \iota(Q(\boldsymbol{p})) \\ 1 - \frac{p_i - p_{\varrho_{\boldsymbol{p}}(i)}}{\alpha_i - \alpha_{\varrho_{\boldsymbol{p}}(i)}} & \text{if } i = \iota(Q(\boldsymbol{p})) \end{cases} \tag{19}$$

The linear constraints imposed by Eq. (19) that the prices must satisfy so as to have $OPT(P_0) = OPT(P_1)$ agree with those of Hotelling demand case. That is, the prices must satisfy:

$$p_{\vartheta_{\boldsymbol{\theta}}(i)}(\boldsymbol{\theta}) - p_i(\boldsymbol{\theta}) = v_{\vartheta_{\boldsymbol{\theta}}(i)}(\theta_{\vartheta_{\boldsymbol{\theta}}(i)}) - v_i(\theta_i) \qquad \forall \boldsymbol{\theta} \in \Theta,\ i \in Q(\boldsymbol{\theta}),\ i \neq \iota(\boldsymbol{\theta}), \tag{20}$$

together with the constraints $T_i(\theta_i^j)$, $\forall i \in N$, $\forall \theta_i^j \in \Theta_i$. With this in mind, it is simple to derive a result analogous to that of Theorem 5.2.

Theorem B.1. *Consider the general setting in which agents have arbitrary qualities and costs distributions. Let $b^* = \min_{1 \leq i \leq n-1}(\alpha_{i+1} - \alpha_i)$. Suppose that the following two conditions are simultaneously satisfied:*

1. *There exists $\boldsymbol{\theta} \in \Theta$ and $c^* \in \mathbb{R}$ such that $\frac{v_{i+2}(\theta_{i+2}) - v_{i+1}(\theta_{i+1})}{\alpha_{i+2} - \alpha_{i+1}} > c^* + \frac{v_{i+1}(\theta_{i+1}) - v_i(\theta_i)}{\alpha_{i+1} - \alpha_i}$ for all $1 \leq i \leq n-2$, $\frac{v_2(\theta_2) - v_1(\theta_1)}{\alpha_2 - \alpha_1} > c^*$, and, $1 - c^* > \frac{v_n(\theta_n) - v_{n-1}(\theta_{n-1})}{\alpha_n - \alpha_{n-1}}$;*
2. *$|\Theta_i| \geq 3$ for all $i \in N$, and for every $i \in N$ and every $\theta^j \in \Theta_i$, we have $v_i(\theta_i^{j+1}) - v_i(\theta_i^j) \leq \frac{c^* b^*}{4}$.*

Then, we have $OPT(P_0) = OPT(P_1)$.

The intuition behind these two requirements is the same as that of Theorem 5.2. As usual, let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$. From the definition of vertical demands (Eq. (19)), it is easy to see that, by condition (1), for $n \geq 2$ we must have $Q(\boldsymbol{\theta}) = N$. Hence, the first condition guarantees the existence of an 'interior solution'. The second condition imposes a 'thin enough' cost discretization.

C Extensions to our model

We now discuss two important extensions to our model. The first one is related to the assumption that each supplier offers one product. In Section C.1, we provide a reasonable extension to our model under which suppliers can offer multiple products. We show that our main result extends accordingly, so we are able to characterize (under additional conditions) the optimal mechanisms for the multiproduct case.

The second extension is related to the constraint that demand is inelastic. In particular, we study what happens if we allow the total demand to be elastic in prices instead of requiring it to be constant. We show that, in general, our main result fails to hold and a gap between the optima of the original and the relaxed problem exists. However, preliminary computational results show that the market structures (i.e., which suppliers are in the menu) are usually similar in both the relax and the original problems.

C.1 Extension to multiple products per agents

We now show how to extend our model to the case where suppliers can offer more than one product. If each agent is assumed to have a different random variable to represent the cost for each product, then problem involves solving a multidimensional mechanism design problem. This problem is recognized to be hard. Therefore, our approach is to assume that suppliers' costs can be parametrized by a single type, which can be interpreted as if the auctioneer knows the agents' cost structures but not their underlying cost parameter. This approach is commonly used in the literature to overcome the multidimensional mechanism design problem [17].

For $i \in N$, let P_i denote the set of products offered by supplier i. We assume that agent i has cost $c_{ip}(\theta_i)$ for product $p \in P_i$, where θ_i is agent is type. The utility function of supplier i is given by

$$u_i = t_i - \sum_{p \in P_i} c_{ip}(\theta_i) x_{ip},$$

where x_{ip} is the amount of product p allocated to i, t_i is the payment i receives in the auction, and θ_i is his type. Similarly, the interim utility for supplier i when he reports cost θ_i' and has true cost

θ_i is given by:

$$U_i(\theta_i'|\theta_i) = T_i(\theta_i') - \sum_{p \in P_i} c_{ip}(\theta_i) X_{ip}(\theta_i').$$

For each pair (i, p) with $i \in N$ and $p \in P_i$, we define the modified virtual cost as:

$$v_{ip}(\theta_i) = c_{ip}(\theta_i) + \frac{F_i(\rho(\theta_i))}{f_i(\theta_i)} \left(c_{ip}(\theta_i) - c_{ip}(\rho(\theta_i))\right).$$

As usual, we assume virtual costs to be increasing. Furthermore, we require that the function $h_i : \mathbb{R}^{|P_i|} \times \mathbb{R} \to \mathbb{R}$ defined as $h_i(x_i, \theta_i) = \sum_{p \in P_i} c_{ip}(\theta_i) x_{ip}$ satisfies the increasing differences property. Under these assumptions, the optimal solution to the relaxed problem is characterized by the following proposition.

Proposition C.1. *Suppose that $(\boldsymbol{x}, \boldsymbol{t})$ satisfy the following conditions:*

1. *The allocation function satisfies for all $\boldsymbol{\theta} \in \Theta$,*

$$\begin{aligned} x(\boldsymbol{\theta}) \quad \in \quad & \text{argmax} \sum_{i=1}^{n} \sum_{p \in P_i} k_{ip}(x(\boldsymbol{\theta})) - v_{ip}(\theta_i) x_{ip}(\boldsymbol{\theta}) \\ s.t. \quad & \sum_{i=1}^{N} \sum_{p \in P_i} x_{ip}(\boldsymbol{\theta}) = 1, \qquad x_{ip}(\boldsymbol{\theta}) \geq 0 \quad \forall i \in N, p \in P_i \ . \end{aligned}$$

2. *Interim expected transfers satisfy for all $i \in N$ and $\theta_i^j \in \Theta_i$:*

$$T_i(\theta_i^j) = \sum_{p \in P_i} c_{ip}(\theta_i^j) X_{ip}(\theta_i^j) + \sum_{k=j+1}^{|\Theta_i|} \sum_{p \in P_i} \left(c_{ip}(\theta_i^k) - c_{ip}(\theta_i^{k-1})\right) X_{ip}(\theta_i^k)$$

Then, $(\boldsymbol{x}, \boldsymbol{t})$ is an optimal mechanism for the relaxed problem.

Ideally, we would like to use the the characterization of the optimal solution to the relaxed problem to study the original problem. The optimal demands for the relaxed problem still have an intuitive form, similar to the single-product case. However, the expected transfers constraints differ. While demands depend on both the individual product and the cost realization, the expected transfers only depend on the cost realization. Therefore, for each cost realization, the expected transfers constraints involve terms for potentially many products. This introduces some additional complexities in the analysis, and the extension of Theorem C.1 to the multiproduct case is not straightforward.

Surprisingly, under sufficient conditions, we are able to show that our main result still holds. That is, there exists prices under which we have $OPT(P_0) = OPT(P_1)$. This is formalized by the following theorem.

Theorem C.1. *Consider the general setting in which agents have arbitrary costs distributions and offer any arbitrary number of products. Then, there exists $c^* \in \mathbb{N}$, $d^* \in \mathbb{R}_+$ such that, whenever the following conditions are simultaneously satisfied,*

1. *There exists a profile $\boldsymbol{\theta} \in \Theta$ such that $p_i \in Q(\boldsymbol{\theta})$ for all $p_i \in P_i$ and all $i \in N$. Furthermore, there exists a $d^* \in \mathbb{R}$ such that, for all $\boldsymbol{\theta}' \in \Theta$ with $|\boldsymbol{\theta} - \boldsymbol{\theta}'|_\infty \leq d^*$ we have $Q(\boldsymbol{\theta}') = \cup_{i \in N} P_i$.*

2. *$|\Theta_i| \geq c^*$ for all $i \in N$, and for every $i \in N$ and every $\theta^j \in \Theta_i$, we have $\max_{p \in P_i}\{v_{ip}(\theta_i^{j+1}) - v_{ip}(\theta_i^j)\} \leq d^*/3$.*

we have $OPT(P_0) = OPT(P_1)$.

The proof of Theorem C.1 can be found in the companion appendix. Although the intuition is similar to the single-product case, there are some fundamental differences. For example, the set $Q(\boldsymbol{\theta})$ now denotes the active products rather than the active suppliers. Note that a single supplier can simultaneously have many different products in the assortment, which will be reflected in the expected transfer constraints. In addition, as the cost realization of a supplier is simultaneously valid for all his products, we need to guarantee that the grid is thin enough for all products offered by the supplier.

C.2 Demand Elasticity

Throughout this work we have assumed that demand is inelastic; regardless the prices, exactly one unit is consumed across all substitute products. This is a natural constraint to impose when modeling some specific FAs such as dialysis supply, in which the aggregate demand is inelastic. In some FAs, however, it is not unreasonable to suppose that the actual quantity purchased will depend on the prices: for instance, a school seeking to renovate two computer labs might decide to renovate only one of them if the price of computers is too high. Therefore, one reasonable extension to our model would be to consider an elastic demand setting by relaxing the constraint that demands should add up to one.

The problem of auctions with endogenous quantities is not new, it was first introduced by Hansen [14]. In his paper, even though demand is elastic, the auction has a unique winner (the lowest-price bidder). Therefore, determining the allocation is easy and only the quantity is endogenous. This unique-winner assumption is usually common to all the literature in the area. In our problem instead, both winners (i.e., agents that are in menu) and quantities should be endogenous, which adds significant difficulties to the analysis.

To illustrate, consider the general affine demand model introduced in Section 5.3. By using the same arguments as in Proposition 4.1, we know that the optimal solution to the relaxed problem must satisfy:

$$x(\boldsymbol{\theta}) = \underset{y \geq 0}{\operatorname{argmax}}\, \boldsymbol{c}\boldsymbol{y} - \frac{1}{2}\boldsymbol{y}^T \boldsymbol{D}\boldsymbol{y} - \boldsymbol{v}(\boldsymbol{\theta})\boldsymbol{y},$$

where $\boldsymbol{v}(\boldsymbol{\theta})$ is the vector of virtual costs. Unfortunately, given our market primitives, in the general case this implies a gap between the optimal solutions of the relaxed and the original problem. As the consumer surplus function is assumed to be strictly concave, it has a unique optimal solution. Whenever both agents have positive allocations in the optimal solution to the relaxed problem, the only way to replicate those demands in original problem is by setting the prices equal to the virtual costs. However, this choice of prices generally violates the incentive compatibility constraints.

Even though the optimal relaxed solution cannot be mimicked, solving the relaxation still give us some useful information regarding the original problem. To that end, we consider the problem of two ex-ante identical agents and two possible types, θ_L and θ_H. We calculated the optimal solution to both problems for different combination of paremeters $\boldsymbol{c}$, $\boldsymbol{D}$, θ_L and θ_H and different distributions. In general, we considered own-price elasticities in the range $[-7, -0.3]$. We discovered that the optimal solution to the original problem generally imitates the market structure of the relaxed problem, i.e., the decision on how many suppliers to include in the assortment agrees in both problems. In addition, the same constraints bind in optimality: the IR constraint for the high type and the IC constraint for the low type. Whenever the high-type agents is never in the menu (high elasticity case), the optimal solution of the relaxed problem can be implemented in the original problem. This is straightforward; as the demand for the high-type is always zero, the low-type will not have an incentive to misreport if he is offered his own cost as price. However, in the general case a gap exists. In such cases, the price of the high-type is set at θ_H at optimality, and the prices of the low-type are higher than in the relaxed problem. As a result, when compared to the relaxation, the demand of low-type agents decreases in the original problem and the demand of the high-type increases. For all combination of parameters, the gap between the relaxed and original problem was less than 5%.

D Proof of Lemma 5.1

Proof. We start by stating the KKT conditions for problem (LD($\boldsymbol{p}$)):

$$\begin{aligned} \boldsymbol{c} - \boldsymbol{D}\boldsymbol{x} - \boldsymbol{p} + \lambda\boldsymbol{1} + \boldsymbol{q} &= \boldsymbol{0} \\ \boldsymbol{1}'\boldsymbol{x} &= 1 \\ \boldsymbol{x} &\geq \boldsymbol{0} \\ \boldsymbol{x}'\boldsymbol{q} &= \boldsymbol{0}, \end{aligned} \tag{21}$$

where λ is the multiplier associated to the equality constraint and $\boldsymbol{q}$ is the vector of multipliers associated to the non-negativity constraints. Define $\boldsymbol{v} = \boldsymbol{c} - \boldsymbol{D}\boldsymbol{x} - \boldsymbol{p} + \lambda\boldsymbol{1}$. By the KKT conditions we must have that $v_i = c_i - \boldsymbol{D}_{i,*}\boldsymbol{x} - p_i + \lambda = 0$, for all $i \in Q$. Therefore,

$$\boldsymbol{0} = \boldsymbol{v}_Q = \boldsymbol{c}_Q - \boldsymbol{D}_Q\boldsymbol{x}_Q - \boldsymbol{p}_Q + \lambda\boldsymbol{1}_Q.$$

As $\boldsymbol{D}$ is positive definite and $\boldsymbol{D}_Q$ is a principal submatrix of $\boldsymbol{D}$ we have that $(\boldsymbol{D}_Q)^{-1}$ exists and, furthermore,

$$\boldsymbol{x}_Q = (\boldsymbol{D}_Q)^{-1}\left(\boldsymbol{c}_Q - \boldsymbol{p}_Q + \lambda\boldsymbol{1}_Q\right)$$

In addition, by the feasibility constraint, we must have $\boldsymbol{1}'_Q\boldsymbol{x}_Q = 1$ and hence,

$$1 = \boldsymbol{1}'_Q\boldsymbol{x}_Q = \boldsymbol{1}'_Q(\boldsymbol{D}_Q)^{-1}\left(\boldsymbol{c}_Q - \boldsymbol{p}_Q + \lambda\boldsymbol{1}_Q\right)$$

which implies

$$\lambda = \frac{1 - \boldsymbol{1}'_Q(\boldsymbol{D}_Q)^{-1}\left(\boldsymbol{c}_Q - \boldsymbol{p}_Q\right)}{\boldsymbol{1}'_Q(\boldsymbol{D}_Q)^{-1}\boldsymbol{1}_Q}\ .$$

Hence,

$$\boldsymbol{x}_Q = (\boldsymbol{D}_Q)^{-1}\left(\boldsymbol{c}_Q - \boldsymbol{p}_Q + \left(\frac{1 - \boldsymbol{1}'_Q(\boldsymbol{D}_Q)^{-1}\left(\boldsymbol{c}_Q - \boldsymbol{p}_Q\right)}{\boldsymbol{1}'_Q(\boldsymbol{D}_Q)^{-1}\boldsymbol{1}_Q}\right)\boldsymbol{1}_Q\right)\ ,$$

as desired. $\square$

E Proof of Main Theorems

In this section we prove our main theorems. In particular, we prove a more general theorem (Theorem E.1), which generalized the statements of Theorem 5.2, Theorem B.1, and Theorem C.1. Throughout this section, we use several basic definitions and concepts from linear algebra. We refer the reader to [**?**].

E.1 The coefficient matrix

Abusing notation, let $\boldsymbol{M}$ and $\boldsymbol{m}$ be the coefficient matrix and the corresponding RHS respectively defined by linear equations in ($M_i(\boldsymbol{\theta})$) and ($T_i(\theta_i^j)$), where each column is associated with a price $p_i(\boldsymbol{\theta})$. We can safely discard the columns corresponding to prices $p_i(\boldsymbol{\theta})$ such that $i \notin Q(\boldsymbol{\theta})$, as all the coefficients of such columns are zero. The resulting matrix $\boldsymbol{M}$ will have $\sum_{\theta\in\Theta}|Q(\boldsymbol{\theta})|$ columns as we have one price variable per active supplier and per profile of costs. In addition, for each $\theta \in \Theta$, there will be $|Q(\boldsymbol{\theta})|-1$ rows given by the constraints in Eqs. ($M_i(\boldsymbol{\theta})$) and $\sum_{i\in N}\sum_{\theta_i\in\Theta_i}\mathbb{I}[\exists\ \boldsymbol{\theta}_{-i}:\ i \in Q(\theta_i,\boldsymbol{\theta}_{-i})] \leq |\Theta|$ rows given by the constraints in Eqs. ($T_i(\theta_i^j)$). The preceding observations are summarized by the following remark:

Remark E.1 (Dimension of the coefficient matrix)**.** *The coefficient matrix $\boldsymbol{M}$ has $\sum_{\theta\in\Theta}|Q(\boldsymbol{\theta})|$ columns and $\sum_{\theta\in\Theta}|Q(\boldsymbol{\theta})| - \Theta + \sum_{i\in N}\sum_{\theta_i\in\Theta_i}\mathbb{I}[\exists\ \boldsymbol{\theta}_{-i}:\ i \in Q(\theta_i,\boldsymbol{\theta}_{-i})]$ rows. Further, the number of columns is greater or equal than the number of rows.*

By the Rouché-Frobenius theorem, a system of linear equations $\boldsymbol{Mp}=\boldsymbol{m}$ is consistent (has a solution) if and only if the rank of its coefficient matrix $\boldsymbol{M}$ is equal to the rank of its augmented matrix $[\boldsymbol{M}|\boldsymbol{m}]$. To show whether the system of equations has a solution, we use an equivalent definition of consistency.

Lemma E.1 (Consistency of a system of linear equations)**.** *Consider the system of linear equations $\boldsymbol{Mp}=\boldsymbol{m}$. Let $\boldsymbol{M}_{i,*}$ denote the i^{th} row of $\boldsymbol{M}$. Then, the system is consistent (has a solution) if and only if for every vector $\boldsymbol{y}$ such that $\sum_i y_i\boldsymbol{M}_{i,*}=\boldsymbol{0}$, we have $\sum_i y_i m_i = 0$.*

For each row $\boldsymbol{M}_i(\boldsymbol{\theta})$, let $a^i_{\boldsymbol{\theta}}$ denote the associated coefficient. Similarly, we denote by $b^i_{\theta_i^j}$ the coefficient associated to row $T_i(\theta_i^j)$. Let $(\boldsymbol{a},\boldsymbol{b})$ be the vector of coefficients we just described. Then, for a system to be consistent we must have that for every vector $(\boldsymbol{a},\boldsymbol{b})$ such that:

$$\sum_{\boldsymbol{\theta}\in\Theta}\sum_{\substack{i\in Q(\boldsymbol{\theta})\\ i\neq\iota(Q(\boldsymbol{\theta}))}} a^i_{\boldsymbol{\theta}}\boldsymbol{M}_i(\boldsymbol{\theta}) + \sum_{i\in N}\sum_{\theta_i^j\in\Theta_i} b^i_{\theta_i^j}T_i(\theta_i^j) \;=\; 0 \tag{22}$$

the linear combination of the right hand side also equals zero, that is,

$$\sum_{\boldsymbol{\theta}\in\Theta}\sum_{\substack{i\in Q(\boldsymbol{\theta})\\ i\neq\iota(Q(\boldsymbol{\theta}))}} a^i_{\boldsymbol{\theta}}\left(\sum_{j\in Q(\boldsymbol{\theta})}\boldsymbol{A}(\boldsymbol{\theta})_{ij}(\boldsymbol{\theta})v_j(\theta_j)\right)+\sum_{i\in N}\sum_{\theta_i^j\in\Theta_i} b^i_{\theta_i^j}\left(\theta_i^j X_i(\theta_i^j)+\sum_{k=j+1}^{|\Theta_i|}(\theta_i^k-\theta_i^{k-1})X_i(\theta_i^k)\right)=0. \tag{23}$$

To conclude, we note that, whenever the rows of $\boldsymbol{M}$ are linearly independent, the only vector of coefficients satisfying Eq. (22) is $(\boldsymbol{a},\boldsymbol{b})=\boldsymbol{0}$ and therefore the system is trivially consistent.

For a given $\boldsymbol{\theta}\in\Theta$, we denote by $\boldsymbol{A}(\boldsymbol{\theta})$ the submatrix of $\boldsymbol{M}$ that contains the demand constraints for $\boldsymbol{\theta}$, that is, $\boldsymbol{A}(\boldsymbol{\theta})=(M_i(\boldsymbol{\theta}))_{i\in Q(\boldsymbol{\theta})\setminus\iota(\boldsymbol{\theta})}$. Recall from Section 5, that the demand constraints for both the Hotelling model and Vertical model can be expressed as:

$$p_{\vartheta_{\boldsymbol{\theta}}(i)}(\boldsymbol{\theta})-p_i(\boldsymbol{\theta})=v_{\vartheta_{\boldsymbol{\theta}}(i)}(\theta_{\vartheta_{\boldsymbol{\theta}}(i)})-v_i(\theta_i)\qquad\forall\boldsymbol{\theta}\in\Theta,\ i\in Q(\boldsymbol{\theta}),\ i\neq\iota(\boldsymbol{\theta}). \tag{24}$$

Therefore, we have that the i^{th} row of $\boldsymbol{A}(\boldsymbol{\theta})$ will consist of all zeros except for a 1 in column $\vartheta_{\boldsymbol{\theta}}(i)$ and a -1 in column i for $i=\eta(Q(\boldsymbol{p}))$ for all $i\in Q(\boldsymbol{\theta}),\ i\neq\iota(\boldsymbol{\theta})$.[29] The following claim characterizes the matrix $\boldsymbol{A}(\boldsymbol{\theta})$ for the general affine demand models defined in Section 5.3.

Claim E.1. *Let $\boldsymbol{F}=\boldsymbol{F}(\boldsymbol{\theta})=(\boldsymbol{D}_{Q(\boldsymbol{\theta})})^{-1}$. Then, for every $j\in Q(\boldsymbol{\theta})$ and every i such that $1\leq i\leq Q(\boldsymbol{\theta})$, the coefficient for $p_j(\boldsymbol{\theta})$ in equation i is given by:*

$$\boldsymbol{A}(\boldsymbol{\theta})_{ij}=-\boldsymbol{F}_{ij}+\frac{(\boldsymbol{1}'_{Q(\boldsymbol{\theta})}\cdot\boldsymbol{F}_{*,j})(\boldsymbol{F}_{i,*}\cdot\boldsymbol{1}_{Q(\boldsymbol{\theta})})}{\boldsymbol{1}'_{Q(\boldsymbol{\theta})}\boldsymbol{F}\boldsymbol{1}_{Q(\boldsymbol{\theta})}\boldsymbol{1}_{Q(\boldsymbol{\theta})}}. \tag{25}$$

The proof of the Claim is omitted, as it follows straightforward from the characterization of demand given in Lemma 5.1.

E.2 Definitions and notation

We now state some definitions that we will use to prove the main theorem. Recall that $\underline{\theta_i}$ and $\overline{\theta_i}$ denote the lowest and highest values in Θ_i. For each $j\in N$, let $\theta_j^u=\max\{\theta_j\in\Theta_j:\ \theta_j\in Q(\theta_j,\overline{\boldsymbol{\theta}_{-j}})\}$, that is, θ_j^u is the maximum θ_j under which there exists a profile $\boldsymbol{\theta}=(\theta_j,\boldsymbol{\theta}_{-j})$ such that $j\in Q(\boldsymbol{\theta})$. We may assume that $\underline{\theta_j}\leq\theta_j^u$ for all agents $j\in N$, as otherwise we can consider (w.l.o.g.) the reduced problem in which all agents for which the condition is violated are removed.

Two profiles $\boldsymbol{\theta},\boldsymbol{\theta}'\in\Theta$ are defined to be *adjacent* if and only if $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ only differ in one component and $Q(\boldsymbol{\theta})=Q(\boldsymbol{\theta}')$. To illustrate, consider Example 5.1. There, profiles $(2.5,2.3)$ and

[29] Alternatively, one could think of a matrix $\boldsymbol{A}(\boldsymbol{\theta})$ in which the i^{th} row has a 1 in column $\vartheta_{\boldsymbol{\theta}}(i)$ and a -1 in column i for $i=\eta(Q(\boldsymbol{p}))$, and a 1 in column $\vartheta_{\boldsymbol{\theta}}(i)$ and a -2 in column i and a 1 in column $\varrho_{\boldsymbol{\theta}}(i)$, for all other $i\in Q(\boldsymbol{\theta}),i\neq\iota(\boldsymbol{\theta})$. Note that both matrices will define the same solutions.

(2.5, 2) are adjacent, but profiles (2.5, 2) and (2.5, 1) are not. We define two profiles $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$ to be *connected* if there exists a sequence of adjacent profiles such that one can go from $\boldsymbol{\theta}$ to $\boldsymbol{\theta}'$.

Definition E.1 (Acceptable set)**.** *We say a subset of profiles $\tilde{\Theta} \subseteq \Theta$ is an* acceptable set *if the following conditions are simultaneously satisfied:*

1. *$Q(\boldsymbol{\theta}) = N$ for every $\boldsymbol{\theta} \in \tilde{\Theta}$.*
2. *For each agent i, let $\tilde{\Theta}_i = \{\theta_i \in \Theta_i : \ \exists \boldsymbol{\theta}_{i-1}$ such that $(\theta_i, \theta_{-i}) \in \tilde{\Theta}\}$. Then, for every $\theta_i \in \Theta_i$ such that $\min \tilde{\Theta}_i \leq \theta_i \leq \max \tilde{\Theta}_i$ we must have $\theta_i \in \tilde{\Theta}i$. That is, each $\tilde{\Theta}_i$ must be an interval.*
3. *For every profile $\boldsymbol{\theta}$ such that $\theta_i \in \tilde{\Theta}_i$ for all $i \in N$, we must have $\boldsymbol{\theta} \in \tilde{\Theta}$. That is, any two profiles in $\tilde{\Theta}$ can be connected through profiles in $\tilde{\Theta}$ and $\tilde{\Theta}$ must be maximal.*

To illustrate, in Example 5.1 the set $\tilde{\Theta} = \{(1, 1), (2.5, 2)\}$ satisfies the first two conditions but violates the third one as the profiles are not connected. The above definition of acceptable set will help us characterize sufficient conditions under which the optima of the relaxed and original mechanisms agree. In particular, let a market be defined by the set of suppliers, their product characteristics and cost distributions, as well as the demand model. We define a *relaxation-is-optimal market*(RIOM) as follows.

Definition E.2 (RIOM)**.** *We say a market is RIOM if there exists an acceptable set $\tilde{\Theta}$ under which the following (additional) conditions are satisfied:*

(4) For every $i \in N$ we have $|\tilde{\Theta}_i| \geq 3$.

(5) Let $\boldsymbol{\theta} \in \Theta$ be a profile such that $\theta_i \geq \max \tilde{\Theta}_i$. Then, there exists a profile $\boldsymbol{\theta}' \in \tilde{\Theta}$ such that the profiles $\boldsymbol{\theta}, \boldsymbol{\theta}'$ are connected.

Intuitively, the above conditions can be satisfied when we require the difference in virtual costs between adjacent points in the support to be small enough. To illustrate, we show that the conditions of Theorem 5.2 imply that the market is RIOM. First, by condition (2) in the statement of the theorem, a profile $\boldsymbol{\theta}$ in which $Q(\boldsymbol{\theta}) = N$ must exist. Furthermore, $|v_{i+1}(\theta_i) - v_i(\theta_i)| \leq \delta(\ell_{i+1} - \ell_i)/2$ for all $i \in N$. As $v_i(\theta_i^{j+1}) - v_i(\theta_i^j) \leq \frac{\delta c^*}{4}$ for all $i \in N$, and $\theta_i^j \in \Theta_i$, it follows that by letting θ_i^k denote θ_i we have $Q(\theta_i^{k+2}, \boldsymbol{\theta}_{-i}) = Q(\theta_i^{k-2}, \boldsymbol{\theta}_{-i}) = N$, provided these exist. As $|\Theta_i| \geq 3$ for all $i \in N$, we must have than an acceptable $\tilde{\Theta}$ exists and $|\tilde{\Theta}_i| \geq 3$. Finally, the connectivity requirement follows from the fact that $v_i(\theta_i^{j+1}) - v_i(\theta_i^j) \leq \frac{\delta c^*}{4}$ for all $i \in N$, and $\theta_i^j \in \Theta_i$. Using the same arguments, it can be see that the conditions of Theorem B.1 imply the market is RIOM.

E.3 Auxiliary Lemmas and Properties

We first state the following remark.

Remark E.2. *Suppose that vector of coefficients* $(\boldsymbol{a}, \boldsymbol{b})$ *is such that the equality given by Eq.* (22)) *holds. If there exists* $\boldsymbol{\theta}_{-i}$ *such that* $Q(\theta_i, \boldsymbol{\theta}_{-i}) = \{i\}$, *then* $b^i_{\theta_i}$ *(the coefficient associated with row* $T_i(\theta_i)$*) must be zero.*

Note that the column corresponding to $p_i(\theta_i, \boldsymbol{\theta}_{-i})$ will have exactly one non-zero element located in row $T_i(\theta_i)$. Therefore, equality (22) will not hold unless the coefficient $b^i_{\theta_i}$ is zero. Next, we state and prove the following proposition.

Proposition E.1. *Suppose the coefficients* $(\boldsymbol{a}, \boldsymbol{b})$ *are such that equality in Eq.* (22) *holds. For each* $i \in N$ *and each* $\theta_i \in \Theta_i$, *let* $g_i(\theta_i)$ *be defined as* $g_i(\theta_i) = \frac{b^i_{\theta_i}}{f_i(\theta_i)}$. *Then for each* $\boldsymbol{\theta} \in \Theta$, *we must have*

$$\sum_{i \in Q(\boldsymbol{\theta})} g_i(\theta_j) x_i(\boldsymbol{\theta}) = 0 \tag{26}$$

Proof. Fix $\boldsymbol{\theta} \in \Theta$. We first show the result for the general affine demand model as described in Section 5.3. Recall that the coefficients of the matrix corresponding to the demand equations (that is, Eqs. ($M_i(\boldsymbol{\theta})$)) are as defined by Eq. (25). As the equality in Eq. (22) holds, for each $j \in Q(\boldsymbol{\theta})$ we must have:

$$b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta}) + \sum_{i=1}^{Q(\boldsymbol{\theta})-1} a^i_{\boldsymbol{\theta}} \left(-F_{ij} + (\mathbf{1}'_{Q(\boldsymbol{\theta})} \cdot \boldsymbol{F}_{*,j})(\boldsymbol{F}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})}) \right) = 0.$$

Therefore

$$\begin{aligned}
\sum_{j \in Q(\boldsymbol{\theta})} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta}) &= -\sum_{j \in Q(\boldsymbol{\theta})} \sum_{i=1}^{Q(\boldsymbol{\theta})-1} a^i_{\boldsymbol{\theta}} \left(-\boldsymbol{F}_{ij} + \frac{(\mathbf{1}'_{Q(\boldsymbol{\theta})} \cdot \boldsymbol{F}_{*,j})(\boldsymbol{F}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})})}{\mathbf{1}'_{Q(\boldsymbol{\theta})} \boldsymbol{F} \mathbf{1}_{Q(\boldsymbol{\theta})} \mathbf{1}_{Q(\boldsymbol{\theta})}} \right) \\
&= -\sum_{i=1}^{Q(\boldsymbol{\theta})-1} a^i_{\boldsymbol{\theta}} \left(\sum_{j \in Q(\boldsymbol{\theta})} \left(-\boldsymbol{F}_{ij} + \frac{(\mathbf{1}'_{Q(\boldsymbol{\theta})} \cdot \boldsymbol{F}_{*,j})(\boldsymbol{F}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})})}{\mathbf{1}'_{Q(\boldsymbol{\theta})} \boldsymbol{F} \mathbf{1}_{Q(\boldsymbol{\theta})} \mathbf{1}_{Q(\boldsymbol{\theta})}} \right) \right) \\
&= -\sum_{i=1}^{Q(\boldsymbol{\theta})-1} a^i_{\boldsymbol{\theta}} \left(-\boldsymbol{F}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})} + \boldsymbol{F}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})} \left(\sum_{j \in Q(\boldsymbol{\theta})} \frac{(\mathbf{1}'_{Q(\boldsymbol{\theta})} \cdot \boldsymbol{F}_{*,j})}{\mathbf{1}'_{Q(\boldsymbol{\theta})} \boldsymbol{F} \mathbf{1}_{Q(\boldsymbol{\theta})} \mathbf{1}_{Q(\boldsymbol{\theta})}} \right) \right) \\
&= -\sum_{i=1}^{Q(\boldsymbol{\theta})-1} a^i_{\boldsymbol{\theta}} \left(-\boldsymbol{F}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})} + \boldsymbol{F}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})} \right) \\
&= 0
\end{aligned}$$

To complete the proof, note that $\sum_{j \in Q(\boldsymbol{\theta})} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta}) = f(\boldsymbol{\theta}) \left(\sum_{j \in Q(\boldsymbol{\theta})} g_j(\theta_j) x_j(\boldsymbol{\theta}) \right) = 0.$

Hence, $\sum_{j \in Q(\boldsymbol{\theta})} g_j(\theta_j) x_j(\boldsymbol{\theta}) = 0$ as desired.

Next, we establish the result for the Hotelling and vertical models. In particular, we show that whenever the coefficients (a, b) are such that equality (22) holds, then for each $\boldsymbol{\theta} \in \Theta$, we must have:

$$a^i_{\boldsymbol{\theta}} = \sum_{\{j \in Q(\boldsymbol{\theta}):\ j \leq i\}} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta}) \qquad \forall\ i \in Q(\boldsymbol{\theta}), i \neq \iota(\boldsymbol{\theta}),$$

and,

$$\sum_{j \in Q(\boldsymbol{\theta})} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta}) = 0$$

which implies the result.

Fix $\boldsymbol{\theta} \in \Theta$. We show that $a^i_{\boldsymbol{\theta}} = \sum_{\{j \in Q(\boldsymbol{\theta}):\ j \leq i\}} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta})$ by induction in the agents' number. Consider the coefficients $(\boldsymbol{a}, \boldsymbol{b})$ involving $i = \eta(Q(\boldsymbol{\theta}))$, i.e., i is the leftmost vertex agent in $Q(\boldsymbol{\theta})$. If $Q(\boldsymbol{\theta}) = \{i\}$ is the leftmost active vertex, then $b^i_{\theta_i} = 0$, there is no such coefficient $a^i_{\boldsymbol{\theta}}$ and the result vacuously holds. Otherwise, we have that $a^i_{\boldsymbol{\theta}} = b^i_{\theta_i} f(\boldsymbol{\theta}_{-i}) x_i(\boldsymbol{\theta})$, which establishes the basis for the induction.

Suppose that the claim holds for every coefficient associated to the columns $p_j(\boldsymbol{\theta})$ with $j \in Q(\boldsymbol{\theta})$ and $j < i$. We show that it holds for the coefficients associated with $p_i(\boldsymbol{\theta})$ with $i \in Q(\boldsymbol{\theta})$. Consider the column associated to $p_i(\boldsymbol{\theta})$. If $i \neq \iota(Q(\boldsymbol{\theta}))$, then we need $a^{\varrho_{\boldsymbol{\theta}}(i)}_{\boldsymbol{\theta}} - a^i_{\boldsymbol{\theta}} + b^i_{\theta_i} f(\boldsymbol{\theta}_{-i}) x_i(\boldsymbol{\theta}) = 0$. By inductive hypothesis, $a^{\varrho_{\boldsymbol{\theta}}(i)}_{\boldsymbol{\theta}} = \sum_{\{j \in Q(\boldsymbol{\theta}):\ j \leq \varrho_{\boldsymbol{\theta}}(i)\}} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta})$, and therefore $a^i_{\boldsymbol{\theta}} = \sum_{\{j \in Q(\boldsymbol{\theta}):\ j \leq i\}} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta})$ as desired. Finally, if $i = \iota(\boldsymbol{\theta})$, then $a^{\varrho_{\boldsymbol{\theta}}(i)} + b^i_{\theta_i} f(\boldsymbol{\theta}_{-i}) x_i(\boldsymbol{\theta}) = 0$ together with the inductive hypothesis imply $\sum_{j \in Q(\boldsymbol{\theta})} b^j_{\theta_j} f(\boldsymbol{\theta}_{-j}) x_j(\boldsymbol{\theta}) = 0$ as desired. To conclude, we note that the same result can be similarly obtained for the alternative definition of $\boldsymbol{A}$ for the hotelling and vertical cases. $\square$

Let $\boldsymbol{A} = \boldsymbol{A}(\boldsymbol{\theta})$ for any $\boldsymbol{\theta} \in \Theta$ such that $Q(\boldsymbol{\theta}) = N$ be as defined by Claim E.1. We are now going to prove two useful properties of $\boldsymbol{A}$.

Remark E.3. $\boldsymbol{A}$ *is symmetric.*

Note that, whenever $Q(\boldsymbol{\theta}) = N$, we have $\boldsymbol{F} = \boldsymbol{D}^{-1}$ where $\boldsymbol{F}$ is as defined in Claim E.1. By assumption, $\boldsymbol{D}^{-1}$ is symmetric and positive definite. Therefore, $\boldsymbol{A}$ is also symmetric by definition.

The second property is related to the rank of $\boldsymbol{A}$. Note that we want to find prices $\boldsymbol{p}$ such that $\boldsymbol{x}(\boldsymbol{p}) = \boldsymbol{x}(v(\boldsymbol{\theta}))$, where $v(\boldsymbol{\theta}) = (v_1(\boldsymbol{\theta}), \ldots, v_n(\boldsymbol{\theta}))$ is defined as the vector of virtual costs. That is, we must have $\boldsymbol{A}\boldsymbol{p} = \boldsymbol{A}v(\boldsymbol{\theta})$. We now show that the dimension of prices satisfying that is exactly one. In particular, we show that $\boldsymbol{A}$ has rank $n-1$.

Claim E.2. $\boldsymbol{A}$ *has rank* $n-1$.

Proof. Let $\boldsymbol{I}$ denote the identity matrix of size n. Note that $\boldsymbol{A} = \boldsymbol{D}^1 \left(-\boldsymbol{I} + \mathbf{1}\frac{\mathbf{1}\boldsymbol{D}^{-1}}{\mathbf{1}\boldsymbol{D}^{-1}\mathbf{1}}\right)$. Therefore,

$$rank(\boldsymbol{A}) \geq rank(\boldsymbol{D}^1) + rank\left(-\boldsymbol{I} + \mathbf{1}\frac{\mathbf{1}\boldsymbol{D}^{-1}}{\mathbf{1}\boldsymbol{D}^{-1}\mathbf{1}}\right) - n = rank\left(-\boldsymbol{I} + \mathbf{1}\frac{\mathbf{1}\boldsymbol{D}^{-1}}{\mathbf{1}\boldsymbol{D}^{-1}\mathbf{1}}\right),$$

as $\boldsymbol{D}^{-1}$ has full rank. In addition, we have[30]

$$rank\left(-\boldsymbol{I} + \mathbf{1}\frac{\mathbf{1}\boldsymbol{D}^{-1}}{\mathbf{1}\boldsymbol{D}^{-1}\mathbf{1}}\right) \geq \left|n - rank\left(\mathbf{1}\frac{\mathbf{1}\boldsymbol{D}^{-1}}{\mathbf{1}\boldsymbol{D}^{-1}\mathbf{1}}\right)\right| \geq n-1,$$

as the matrix $\mathbf{1}\frac{\mathbf{1}\boldsymbol{D}^{-1}}{\mathbf{1}\boldsymbol{D}^{-1}\mathbf{1}}$ has rank exactly one. The converse follows just from the definition of $\boldsymbol{A}$, as we know that one row must be redundant as all demands must some up to one. □

We conclude this section by noting that Claim E.2 trivially holds for the hotelling and vertical cases. Also, note that Remark E.3 does not hold for the original definition of $\boldsymbol{A}(\boldsymbol{\theta})$ for the hotelling and vertical models, but it does hold for the alternative definition. We highlight that this will not affect the proof: essentially, we require that for every $i, j \in N$, the coefficient of p_j in the demand equation of i must be equation to the coefficient of p_i in the demand equation for j.

E.4 Main Theorem

We can now state and prove our main theorem. To avoid excessive notation, we assume that we are working with the general affine demand model as defined in Section 5.3 but all steps and calculations are also valid for the hotelling and vertical models, when the alternative definition of the matrix is assumed. For completeness, we clarify using a footnote when the validity of a step is not immediate.

Theorem E.1. *Consider the general setting in which agents have arbitrary costs distributions. If the market is RIOM, then $OPT(P_0) = OPT(P_1)$.*

Proof. To show $OPT(P_0) = OPT(P_1)$, we show that the system of equations is consistent. Let $(\boldsymbol{a}, \boldsymbol{b})$ be a vector of coefficients satisfying Eq. (22). Let $g_i(\theta_i)$ be as defined in the statement of Proposition E.1. As the market is RIOM, we know that there exists a subset of profiles $\tilde{\Theta} \subseteq \Theta$ that satisfies conditions (1)-(5). The idea of the proof is as follows. First, we show that if the market is RIOM all $g_i(\theta_i)$ must be zero. To do so, we start by proving that $g_i(\theta_i) = 0$ for all $\theta_i \in \tilde{\Theta}_i$, where $\tilde{\Theta}_i$ is as defined by condition (2). Then, we show that this implies $g_i(\theta_i) = 0$ for all $\theta_i \in \Theta_i$. We conclude the proof by showing that the fact that $g_i(\theta_i) = 0$ for all $\theta_i \in \Theta_i$ implies that the system is consistent.

[30] Matrix property: $rank(A - B) \geq |rank(A) - rank(B)|$

We now show that $g_i(\theta_i) = 0$ for all $\theta_i \in \tilde{\Theta}_i$. By assumption, $\tilde{\Theta}$ satisfies conditions (1)-(5). Therefore, for every $\boldsymbol{\theta} \in \tilde{\Theta}$ for all $i \in N$ we must have $Q(\boldsymbol{\theta}) = N$. Consider two profiles $\boldsymbol{\theta} = (\theta_i, \boldsymbol{\theta}_{-i})$ and $\boldsymbol{\theta'} = (\theta'_i, \boldsymbol{\theta}_{-i})$ which only differ in agent i's cost and such that $\boldsymbol{\theta}, \boldsymbol{\theta'} \in \tilde{\Theta}$. By the definition of $\tilde{\Theta}$, such pair of profiles exists (condition (4)). By Eq. (26), we must have $g_i(\theta_i)x_i(\boldsymbol{\theta}) + \sum_{j\neq i} g_j(\theta_j)x_j(\boldsymbol{\theta}) = 0$ and $g_i(\theta'_i)x_i(\boldsymbol{\theta'}) + \sum_{j\neq i} g_j(\theta_j)x_j(\boldsymbol{\theta'}) = 0$. Hence, by subtracting the second equality from the first one we obtain

$$g_i(\theta_i)x_i(\boldsymbol{\theta}) - g_i(\theta'_1)x_i(\boldsymbol{\theta'}) = \sum_{j\neq i} g_j(\theta_j)\left[x_j(\boldsymbol{\theta'}) - x_j(\boldsymbol{\theta})\right].$$

For each $j \in N$, we must have $x_j(\boldsymbol{\theta'}) - x_j(\boldsymbol{\theta}) = \boldsymbol{A}(\boldsymbol{\theta})_{j,i}\left(v_i(\theta'_i) - v_i(\theta_i)\right)$, where we used the fact that $A(\boldsymbol{\theta}) = A(\boldsymbol{\theta'})$ by definition (see Claim E.1). Let $\boldsymbol{A} = \boldsymbol{A}(\boldsymbol{\theta})$, and note that this $\boldsymbol{A}$ agrees with the one in Remark E.3 and Claim E.2. Hence, we can re-write the above equality as:

$$g_i(\theta_i)x_i(\boldsymbol{\theta}) - g_1(\theta'_i)x_i(\boldsymbol{\theta'}) = \left(v_i(\theta'_i) - v_i(\theta_i)\right)\left(\sum_{j\neq i} g_j(\theta_j)\boldsymbol{A}_{j,i}\right),$$

and therefore,

$$g_i(\theta_i)x_i(\boldsymbol{\theta}) - g_i(\theta'_i)x_i(\boldsymbol{\theta'})v_i(\theta'_i) - v_i(\theta_i) = \left(\sum_{j\neq i} g_j(\theta_j)\boldsymbol{A}_{j,i}\right). \tag{27}$$

Fix an arbitrary $j \in N$ with $j \neq i$ and $\boldsymbol{A}_{ij} \neq 0$.[31] Assume that j has cost θ_j in both $\boldsymbol{\theta}$ and $\boldsymbol{\theta'}$ as defined above. Let $\theta'_j \in \Theta_j$ be such that $\theta'_j \neq \theta_j$ and $\theta'_j \in \tilde{\Theta}_j$. Define $\tilde{\boldsymbol{\theta}} = (\theta_i, \theta'_j, \boldsymbol{\theta}_{-i,j})$ and $\tilde{\boldsymbol{\theta}}' = (\theta'_i, \theta'_j, \boldsymbol{\theta}_{-i,j})$. The only thing we assumed about θ_j was $\theta_j \in \tilde{\Theta}_j$. Therefore, the above equality must also hold for any $\tilde{\Theta}_j$. That is,

$$\frac{g_1(\theta_i)x_1(\tilde{\boldsymbol{\theta}}) - g_1(\theta'_i)x_1(\tilde{\boldsymbol{\theta}}')}{v_i(\theta'_i) - v_i(\theta_i)} = g_j(\theta'_j) + \sum_{k\neq i,j} g_k(\theta_k)\boldsymbol{A}_{k,i}.$$

By subtracting the inequality when j has cost θ_j from the one when his cost is θ'_j we get

$$\frac{g_i(\theta_1)\left(x_i(\tilde{\boldsymbol{\theta}}) - x_i(\boldsymbol{\theta})\right) - g_i(\theta'_1)\left(x_i(\tilde{\boldsymbol{\theta}}') - x_i(\boldsymbol{\theta'})\right)}{v_i(\theta'_i) - v_i(\theta_i)} = \boldsymbol{A}_{j,i}\left(g_j(\theta'_j) - g_j(\theta_j)\right).$$

[31] In the hotelling and vertical models, this implies that $j = i - 1$ or $j = i + 1$.

However, note that $x_i(\tilde{\boldsymbol{\theta}}) - x_i(\boldsymbol{\theta}) = \boldsymbol{A}_{i,j}\left(v_j(\theta'_j) - v_j(\theta_j)\right)$. Therefore,

$$\boldsymbol{A}_{i,j}\frac{g_i(\theta_i) - g_i(\theta'_i)}{v_i(\theta'_i) - v_i(\theta_i)} = \boldsymbol{A}_{j,i}\frac{g_j(\theta'_j) - g_j(\theta_j)}{v_j(\theta'_j) - v_j(\theta_j)}.$$

Recall that $\boldsymbol{A}$ is symmetric (Remark E.3).[32] Therefore, whenever $\boldsymbol{A}_{i,j} \neq 0$ we must have:

$$\frac{g_i(\theta_i) - g_i(\theta'_i)}{v_i(\theta'_i) - v_i(\theta_i)} = \frac{g_j(\theta'_j) - g_j(\theta_j)}{v_j(\theta'_j) - v_j(\theta_j)}, \qquad \forall \theta_i \in \tilde{\Theta}_i,\ \forall \theta_j \in \tilde{\Theta}_j.$$

Furthermore, the above equality should hold for every $i, j \in N$ as we can find a sequence of agents $\{l_0 = i, \ldots, l_K = j\}$ such that $\boldsymbol{A}_{l_k,l_{k+1}} \neq 0$ for all $0 \leq k < K$.[33]

We now show that $g_i(\theta_i) = 0$ for all $\theta_i \in \tilde{\Theta}_i$. Suppose the numerator is zero for at least one pair of $g_i(\theta_i), g_i(\theta'_i)$. Then, $g_j(\theta_j) - g_j(\theta'_j)$ must be zero for every $j \in N$ and all pairs $\theta_j, \theta'_j \in \tilde{\Theta}_j$. We now show that $g_i(\theta_i) = g_j(\theta_j)$ must hold for every $\theta_i \in \tilde{\Theta}_i$ and $\theta_j \in \tilde{\Theta}_j$ and $i, j \in N$. This is trivial if $i = j$, as $g_i(\theta_i) - g_i(\theta'_i)$ must be zero for every $i \in N$ and all pairs $\theta_i, \theta'_i \in \tilde{\Theta}_i$. Otherwise, note that when $g_i(\theta_i) = g_i(\theta'_i)$, we have $g_i(\theta_i)x_i(\boldsymbol{\theta}) - g_i(\theta'_i)x_i(\boldsymbol{\theta}') = g_i(\theta_i)\boldsymbol{A}_{i,i}\left(v_i(\theta_i) - v_i(\theta'_i)\right)$. By Eq. (27) the above equality reduces to

$$\sum_{j \in N} g_j(\theta_j)\boldsymbol{A}_{i,j} = 0, \tag{28}$$

and this must be true for any $i \in N$. Let $\boldsymbol{A}_R$ denote the submatrix of $\boldsymbol{A}$ consisting of $(n-1)$ linearly independent rows. By Claim E.2, we know such matrix exists. Furthermore, we can assume that those are the $n-1$ demand equations that appear in the matrix coefficient $\boldsymbol{M}$. Let $\mathbf{g} = (g_1, \ldots, g_n)$ denote the vector of coefficients $g_i = g_i(\theta_i)$ for $\boldsymbol{\theta} \in \Theta$. By Eq. (28), the vector $\mathbf{g}$ must be in the nullspace of $\boldsymbol{A}_R$. However, as $\boldsymbol{A}_R \in \mathbb{R}^{(n-1)\times n}$ has dimension $(n-1)$ the dimension of its nullspace is at most 1. We will show that $\mathbf{1}$ is in Null$(\tilde{\boldsymbol{A}})$, which implies that *all* g_i with $i \in N$ must be equal.

Consider $\tilde{\boldsymbol{A}}_{i,*}$, that is, row i of the coefficient matrix $\tilde{\boldsymbol{A}}$. We will show that $\tilde{\boldsymbol{A}}_{i,*} \cdot \mathbf{1} = 0$. Note that

$$\tilde{\boldsymbol{A}}_{i,*} \cdot \mathbf{1} = \sum_j \left(-\boldsymbol{A}_{ij} + \frac{(\mathbf{1}'_{Q(\boldsymbol{\theta})} \cdot \boldsymbol{A}_{*,j})(\boldsymbol{A}_{i,*} \cdot \mathbf{1}_{Q(\boldsymbol{\theta})})}{\mathbf{1}'_{Q(\boldsymbol{\theta})}\boldsymbol{A}\mathbf{1}_{Q(\boldsymbol{\theta})}\mathbf{1}_{Q(\boldsymbol{\theta})}} \right) = -\boldsymbol{A}_{i,*} \cdot \mathbf{1} + \boldsymbol{A}_{i,*} \cdot \mathbf{1} = 0,$$

as desired. Therefore, $\mathbf{1}$ is in Null$(\tilde{\boldsymbol{A}})$ and $g_i(\theta_i) = g_j(\theta_j)$ for all $i, j \in N$, $\theta_i \in \tilde{\Theta}_i$, $\theta_j \in \tilde{\Theta}_j$.

Using that $g_i(\theta_i) = g_j(\theta_j)$ for all $\theta_i \in \tilde{\Theta}_i$ and $\theta_j \in \tilde{\Theta}_j$, we now show that $g_i(\theta_i) = 0$ for all $i \in N$ and all $\theta_i \in \Theta_i$ which implies $b^i_{\theta_i} = 0$ for all $\theta_i \in \tilde{\Theta}_i$. If $g_i(\theta_i) = 0$, for some $i \in N$ and $\theta_i \in \Theta_i$, we

[32] Note that this also holds for the alternative definition in the hotelling and vertical cases.

[33] Here we are implicitly assuming that matrix $\boldsymbol{A}$ has only one block. If $\boldsymbol{A}$ has more than one block, then we can use the same argument for each block.

are done. Otherwise, suppose that $g_i(\theta_i) = k \neq 0$ for all $i \in N$ and all $\theta_i \in \Theta_i$. By Proposition E.1 we have:

$$0 \;=\; \sum_{j\in Q(\boldsymbol{\theta})} g_j(\theta_j)x_j(\boldsymbol{\theta}) \;=\; k\left(\sum_{j\in Q(\boldsymbol{\theta})} x_j(\boldsymbol{\theta})\right) = k,$$

which is a contradiction.

Now suppose that there exists a pair $g_i(\theta_i), g_i(\theta_i')$ such that $\frac{g_i(\theta_i)-g_i(\theta_i')}{v_i(\theta_i')-v_i(\theta_i)} = k \neq 0$, and rewrite $g_i(\theta_i) = g_i(\theta_i') + k[v_i(\theta_i') - v_i(\theta_i)]$. Let $\theta_i, \theta_i', \theta_i'' \in \tilde{\Theta}_i$ and let $\theta_{-i} \in \tilde{\Theta}_{-i}$. Then, we must have

$$\begin{aligned}
\big(v_i(\theta_i') - v_i(\theta_i)\big)\sum_{j\neq i}\boldsymbol{A}_{ji}g_j(\theta_j) \;&=\; g_i(\theta_i)x_i(\boldsymbol{\theta}) - g_i(\theta_i')x_i(\boldsymbol{\theta}') \\
&=\; \big(g_i(\theta_i') + k[v_i(\theta_i') - v_i(\theta_i)]\big)\,x_i(\boldsymbol{\theta}) - g_i(\theta_i')x_i(\boldsymbol{\theta}') \\
&=\; g_i(\theta_i')\big(x_i(\boldsymbol{\theta}) - x_i(\boldsymbol{\theta}')\big) + k[v_i(\theta_1') - v_i(\theta_i)]x_i(\boldsymbol{\theta}) \\
&=\; g_i(\theta_i')\boldsymbol{A}_{ii}\big(v_i(\theta_i) - v_i(\theta_i')\big) + k[v_i(\theta_i') - v_i(\theta_i)]x_i(\boldsymbol{\theta})
\end{aligned}$$

By dividing on both sides by $v_i(\theta_i') - v_i(\theta)$ we obtain:

$$\sum_{j\neq i}\boldsymbol{A}_{ji}g_j(\theta_j) \;=\; -g_i(\theta_i')\boldsymbol{A}_{ii} + kx_i(\boldsymbol{\theta})$$

In addition, since $\theta_i'' \in \tilde{\Theta}_i$, we have $\frac{g_i(\theta_i'')-g_i(\theta_i')}{v_i(\theta_i')-v_i(\theta_i'')} = k$ and thus:

$$\sum_{j\neq i}\boldsymbol{A}_{ji}g_j(\theta_j) \;=\; -g_i(\theta_i')\boldsymbol{A}_{ii} + kx_i(\boldsymbol{\theta}'')$$

which is a contradiction as the virtual costs are strictly increasing and therefore $x_i(\boldsymbol{\theta}) \neq x_i(\boldsymbol{\theta}'')$.

Next, we show that $g_j(\theta_j) = 0$ for the remaining cases, that if, whenever $\theta_j < \min\tilde{\Theta}_i$ or $\theta_j > \max\tilde{\Theta}_j$. For $\theta_j < \min\tilde{\Theta}_j$ consider a profile $\boldsymbol{\theta} = (\theta_j, \boldsymbol{\theta}_{-j})$ such that $\theta_i \in \tilde{\Theta}_i$ for all $i \neq j$. By the definition of $\tilde{\Theta}_j$, we must have have $x_j(\boldsymbol{\theta}) > 0$. By Proposition **??** we have

$$0 \;=\; \sum_{i\in Q(\boldsymbol{\theta})} g_i(\theta_i)x_i(\boldsymbol{\theta}) \;=\; g_j(\theta_j)x_j(\boldsymbol{\theta}).$$

and therefore $g_j(\theta_j) = 0$ for all $\theta_j < \min\tilde{\Theta}_j$ and all $j \in N$. For $\theta_j > \max\tilde{\Theta}_j$, let $\boldsymbol{\theta} = (\theta_j, \boldsymbol{\theta}_{-j})$ be a profile such that $j \in Q(\boldsymbol{\theta})$. We may assume that $\boldsymbol{\theta}$ is such that $\theta_i \geq \min\Theta_i$ for all $i \in N$, as otherwise we can increase the $\theta_i < \min\Theta_i$ to satisfy this condition and j will still be active. By the definition of $\tilde{\Theta}$, $\boldsymbol{\theta} = (\theta_j, \boldsymbol{\theta}_{-j})$ must be connected to a profile $\boldsymbol{\theta}' \in \tilde{\Theta}$. That means, that there exists a sequence of adjacent profiles $\{\boldsymbol{\theta}_0 = \boldsymbol{\theta}', \ldots, \boldsymbol{\theta}_K = \boldsymbol{\theta}\}$. Given that $\boldsymbol{\theta}' \in \tilde{\Theta}$, we must have that $g_i(\theta'_i) = 0$ for all $i \in N$. Let k be the component in which $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$ differ. By Proposition E.1 we have $\sum_{i\in N} g_i((\theta_1)_i)x_i(\boldsymbol{\theta}_1) = 0$. As $\boldsymbol{\theta}'$ and $\boldsymbol{\theta}_1$ only differ in the k^{th} component, we must have $g_k((\boldsymbol{\theta}_1)_k) = 0$. We can inductively repeat this argument to show that all the g's corresponding to a profile in the path between $\boldsymbol{\theta}'$ and $\boldsymbol{\theta}$ must be zero, which implies $g_j(\theta_j) = 0$. Therefore, we have $g_j(\theta_j) = 0$ for all $i \in N$ and all $\theta_i \in \Theta_i$ which implies $b^i_{\theta_i} = 0$ for all $i \in N$ and all $\theta_i \in \Theta_i$.

To conclude the proof, we show that $b^i_{\theta_i} = 0$ for all $i \in N$ and all $\theta_i \in \Theta_i$ implies that the system is consistent. To that end, consider a vector $(\boldsymbol{a}, \boldsymbol{0})$ satisfying Eq. (22). For each $\boldsymbol{\theta} \in \tilde{\Theta}$, we have

$$\sum_{i=1}^{|Q(\boldsymbol{\theta})|-1} a^i_{\boldsymbol{\theta}} \left(\sum_{j\in Q(\boldsymbol{\theta})} \boldsymbol{A}(\boldsymbol{\theta})_{i,j} v_j(\theta_j) \right) = \sum_{j\in Q(\boldsymbol{\theta})} v_j(\theta_j) \left(\sum_{i=1}^{|Q(\boldsymbol{\theta})|-1} a^i_{\boldsymbol{\theta}} \boldsymbol{A}(\boldsymbol{\theta})_{i,j} \right) = 0,$$

as $(\boldsymbol{a}, \boldsymbol{0})$ satisfying Eq. (22) implies $\sum_{i=1}^{|Q(\boldsymbol{\theta})|-1} a^i_{\boldsymbol{\theta}} \boldsymbol{A}(\boldsymbol{\theta})_{i,j} = 0$. Hence, we have shown that $(\boldsymbol{a}, \boldsymbol{0})$ also satisfies Eq. (23). Therefore, the system is consistent and $OPT(P_1) = OPT(P_0)$ as desired. $\square$