

On the Value of using Group Discounts under Price Competition

Reshef Meir¹ Tyler Lu² Moshe Tennenholtz³

Craig Boutilier²

¹Harvard University

rmeir@seas.harvard.edu

²University of Toronto

{cebly, tl}@cs.toronto.edu

³Technion-Israel Institute of Technology and Microsoft Research, Herzliya

moshet@microsoft.com

December 9, 2013

Abstract

The increasing use of group discounts has provided opportunities for buying groups with diverse preferences to coordinate their behavior in order to exploit the best offers from multiple vendors. We analyze this problem from the viewpoint of the vendors, asking under what conditions a vendor should adopt a volume-based price schedule rather than posting a fixed price. We consider both the case of monopolist vendors and cases where a vendor competes with other vendors. When vendors have uncertainty about buyers' valuations specified by a known distribution, we show that a vendor is always better off posting a fixed price, provided that buyer types (valuations) are i.i.d. and that other vendors also use fixed prices. We also show that these assumptions cannot be relaxed: if buyer types are not i.i.d., or other vendors post discount schedules, then posting a schedule may yield a higher profit for the vendor. We provide similar results under a distribution-free uncertainty model, where vendors minimize their maximum regret over all type realizations.

1 Introduction

Online services offering consumer group discounts represent an important and growing segment of online sales. Despite margin pressures, services such as Groupon, Living Social, and hundreds of others remain successful, offering consumers a choice of multiple, competing offers from vendors of identical or similar products. The move to offering large-scale group discounts is emerging in very sophisticated markets as well, as evidenced by programs in the UK, Australia, and other countries that encourage consumers to engage in coalitional buying activities for electricity, life insurance, and

other such products and services.¹ This abundance presents difficult decisions for the buyer, since the optimal purchase depends not only on her preferences, but also on the choices of other buyers (which determine the triggered price); intelligent search and coordination mechanisms allow mediators to assemble large group of buyers in order to take advantage of offered discounts (see Related Work). Vendors too face complex decisions in the face of strategic competitors and buyers (especially when the latter coordinate their purchases using online services): they must decide on a *complete pricing strategy*, setting volume-based prices instead of a single posted price.

In this paper we assess the value of offering group discounts from the perspective of the vendors, and take some initial steps towards delineating conditions under which such discounts may increase vendor revenue. Our starting point is the group buying model recently proposed by Lu and Boutilier [2012]. In this model (henceforth, the *LB model*), vendors of similar products each propose volume discounts for their product, and buyers each seek a single product from this set. Each buyer has preferences for the distinct products which, together with the final price—as triggered by purchase volume—determine her utility. Lu and Boutilier study various forms of stable assignment of buyers to specific vendors in this model, with and without transferable utility. They also develop algorithms for computing such stable assignments. However, they approach the problem from the perspective of the buying group or buying coalition, modeling the incentives for different buyers to coordinate their purchases with specific vendors. Critically, they take the (volume-discounted) vendor prices to be exogenously determined and fixed (independent of buyer behavior). As such, they do not address the incentives for vendors to offer such discounts in the first-place, nor strategic interactions involving the vendors.

Incentives for group discounts. The discussion above leads to a natural question: why would a vendor offer price discounts based on the number of buyers of its product? There are at least three potential reasons. First, if a vendor has economies of scale (i.e., when the marginal costs of producing a good or offering a service decreases with the number of units sold), at least some of the cost savings induced by unit-sales volume can be passed on to consumers to increase total sales and revenues. Second, there may be implicit or explicit network effects, such as word of mouth advertising, where initial discounts that bring in large volumes of buyers can result in greater additional sales among further buyers at a future point in time. When customer acquisition costs are high through traditional channels, group discounts may induce buyers—in order to extract a better price—to recruit their friends who may be unaware of the product or service being offered.² Finally, group discounts may allow for price discrimination between buyers with different preferences. This may be especially true in circumstances in which a vendor is uncertain about the valuations buyers have for her product or service and can be used against buyer types. While all three factors (and likely several others) play a role in driving volume discounts in real markets, in this work we

¹See for example, <https://www.which.co.uk/switch/> or <https://www.onebigswitch.com.au/>.

²Such discounts are sometimes motivated by considering businesses in which customer acquisition costs are high, but where initial discounts pay for themselves by retaining newly acquired customers for future sales [Edelman et al., 2011].

focus exclusively on the third factor in order to study its effect on pricing strategies in isolation. We do so by assuming that production costs are linear (i.e., there are no economies of scale) and that the set of potential buyers is fixed and aware of all vendor offerings (i.e., there are no advertizing or acquisition effects).

A two-stage competition. Several models have been proposed that examine vendor incentives for offering volume-based discounts under a variety of utility and informational assumptions. We discuss these in depth in Section 1.1. Most such proposals adopt a two-stage model of interaction, along the lines of Stackelberg games, in which one or more vendors first commit to a pricing strategy, or *discount schedule*, and then buyers make individual or coordinated purchasing decisions. The LB model reflects the *coordinated* behavior of buyers in the presence of multiple discount schedules, a behavior that has become increasingly feasible using online services (e.g., Groupon, Living Social, Which? Switch) to assess preferences and form suitable buying groups. As such, it is natural to assess its fit within the “standard” two-stage framework, and analyze how such coordinated buyer behavior impacts vendor pricing strategies. Specifically, we analyze the conditions under which a vendor can derive value by using a *discount schedule* rather than a *fixed price*.

Since the LB model does not account for the incentives or behavior of vendors, we extend it to include vendors as strategic entities to form a game that reflects the strategic interactions of all participating parties. To do so, we must specify the structure of vendor utility functions, vendor beliefs about buyer valuations, and the relation between the two. We consider three natural models that differ in vendor information structure, each of which is elaborated below. Our *complete information model* assumes that vendors know the valuations (or *types*) of all buyers for their products. In our *Bayesian model*, vendor beliefs take the form of a distribution over buyers’ types and the utility of their strategies is determined using the expected response of the buyers to their prices and/or discounts. Finally, we analyze a *distribution-free model*, where vendors know only the set of *possible* buyer types, but not the frequency with which they occur. Here they try to maximize worst-case utility over possible realizations of buyer types. Despite the key informational difference, in all models, vendor utility remains linear in the number of units sold. In each setting, we analyze conditions under which a vendor may profit by offering group discounts when competing with other vendors.

To summarize our model, we study a two-stage game with multiple vendors and multiple buyers. In the first stage, vendors propose price schedules, exploiting available information about buyer types, which varies across the three models. Buyers then coordinate their buying behavior, based on their realized types (if unknown to the vendors), using the LB (non-cooperative) mechanism. Under this non-cooperative model of buyer behavior, buyers have full information about the offers and the types of other buyers, but cannot transfer payments or make binding agreements. Buyers form stable groups (as defined below), where no *single buyer* can benefit by switching to a different vendor.

1.1 Related work

Volume-based pricing has been studied extensively, but often using motivations different than ours. One line of research focuses on the effect of volume discounts on purchase management and the induced efficiencies in supply chains [Monahan, 1984; Lal and Staelin, 1984; Wang, 2002]. For instance, quantity discounts can increase order quantities from a single or multiple buyers. Reduced setup, inventory and shipping costs can more than compensate suppliers for their reduced margins, while saving buyers money. The focus of such work is on optimizing pricing, though strategic elements are sometimes assessed. Edelman et al. [2011] provide an analysis of group discounts that focuses on the value of discounts for customer acquisition. However, as noted above, we ignore the benefits of non-linear production costs and buyer acquisition in our work.

A different model was suggested by Anand and Aron [2003], with motivations very similar to ours: buyer utility is quasi-linear in price, and vendor utility is linear in the number of units sold (as in our model). Volume discounts are used to attract buyers that would otherwise refrain from purchase. The main difference with our model is their assumption of weak buyer coordination: buyers are uncertain of the valuations of others and do not coordinate their choices. Anand and Aron further limit their analysis to a monopolist (single vendor) and several very specific classes of buyers. Under a variety of conditions, they prove that a monopolist with a fixed marginal production cost *cannot increase its profit* by posting a discount schedule rather than a fixed price. However, a schedule may be the best strategy for a monopolist, for example, when facing buyers whose types are correlated by some signal that provides some indication of the quality of the product.

Somewhat less related (but still within the two-stage framework) is the *group buying auction* model [Chen et al., 2007]. Here a vendor posts a discount schedule, then buyers arrive sequentially and can announce the price at which they are willing to buy (rather than just joining the group). These announcements, in turn, may affect the estimates of other buyers regarding the eventual price, and their decision to join the group. Chen *et al.* show that a monopolist facing i.i.d. buyers cannot benefit by offering volume discounts unless it is risk-seeking or has decreasing marginal costs. More recently, Chen *et al.* [2010] have shown how to derive the optimal discount schedule for a vendor facing a particular class of (non-i.i.d.) buyers, both as a monopolist and when competing against other (fixed-price) vendors. Other discount-based auction mechanisms have also been developed [Matsuo et al., 2005; Prashanth and Narahari, 2008].

The algorithmic aspects of finding a service that is most suitable for a group of buyers (including scenarios with group discounts) have been studied in the AI literature [Sarne and Kraus, 2005; Manisterski et al., 2008], where buyers are assumed to be fully cooperative.

Finally, various stable mechanisms for buyer coordination have been suggested assuming *transferable utility* [Yamamoto and Sycara, 2001; Li et al., 2005; Lu and Boutilier, 2012], which requires the possibility that buyers can construct and engage in binding agreements. In our models, as in [Anand and Aron, 2003; Chen et al., 2007] and in the non-transferable-utility version of [Lu and Boutilier, 2012], we assume a non-cooperative setting that excludes monetary transfers among buyers, and focus on

(one-shot) vendor revenue maximization.

1.2 Outline and Contribution

Our main contribution in this work is an analysis of the impact of buyer coordination (in the LB model) on vendor pricing, in particular, in the presence of competing vendors. The remainder of the paper is organized as follows. In Section 2, we introduce our underlying model, elaborating on the range of buyer behaviors, reviewing the LB model and the behavior it induces within the buying group, and outline the pricing and discounting strategies available to vendors. We begin our analysis in Section 3, showing that with complete information there is no incentive for vendors to use group discounts. We outline the Bayesian (expected utility) model in Section 4, and prove that if buyer valuations are independent and identically distributed, and all other vendors use fixed prices, then fixed prices are optimal for any vendor. However, we show that if *any* of these conditions is relaxed, then a vendor may benefit by posting a discount schedule rather than a fixed price. With correlated buyers, this gain can grow linearly with the number of buyers. In Section 5, we provide a similar result in the distribution-free setting: a vendor facing buyers with identical sets of possible valuations—with other vendors offering fixed prices—should also post a fixed price in order to minimize regret. We conclude in Section 6 with some suggested future research directions, and offer some preliminary results on equilibrium analysis and ties to optimal auction theory.

2 Model and Notation

We use uppercase letters to denote row vectors of size m or sets, bold letters to denote column vectors of size n , and bold uppercase to denote matrices.

Assume a set N of n buyers and a set M of m vendors. Vendors each offer a single product for sale, of which they have an “unlimited” supply (enough to satisfy the demands of all buyers). Each buyer i has a *type* V_i , which is an m -vector of non-negative values, where v_{ij} represents i ’s valuation for j ’s product. We denote by $A_i \subseteq \mathbb{R}^m$ the set of possible types for buyer i . Each buyer has *unit demand*: she will buy at most one product from at most one vendor. Let \mathbf{v}_j denote the vector of values for vendor j (over all $i \in N$) and $\mathbf{V} = (v_{ij})_{i \in N, j \in M}$ be the matrix of all valuations. We can think of $(V_i)_{i \in N}$ as the *rows* of \mathbf{V} , and $(\mathbf{v}_j)_{j \in M}$ as its *columns*.³ The set of all possible type matrices is denoted by $\mathcal{A} \subseteq \mathbb{R}^{n \times m}$. A_i is thus the projection of \mathcal{A} over the dimensions relevant to buyer i . Each vendor has a fixed cost c_j for producing and/or distributing one unit of its product, which is common knowledge among vendors.

Two-stage interaction. Vendors and buyers engage as follows: in the first stage of the game, each vendor posts a *discount schedule*, a non-increasing vector $\mathbf{p}_j : [n] \rightarrow \mathbb{R}_+$, where $p_j(t)$ is the price offered if t buyers each purchase j ’s item [Anand and

³Throughout the paper, uppercase letters represent column vectors of length m (and also sets). Bold lowercase letters represent row vectors of length n . Bold uppercase letters represent $n \times m$ matrices. See Table 1 for an example.

Aron, 2003]. Let \mathcal{P} be the set of all possible discount schedules and $\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_m)$ a *profile* of schedules, one per vendor. We refer to a schedule with a single fixed price as a *trivial schedule*, which for simplicity is denoted by $p_j \in \mathbb{R}_+$.

In the second stage of the game, each buyer selects a single vendor from whom to purchase a product, or decides to abstain from a purchase. We will elaborate the mechanism by which buyers coordinate their purchases—the LB model—below. The selection of vendors can be modeled as an assignment $\mu : N \rightarrow M \cup \{0\}$ of buyers to vendors, where $\mu(i) = 0$ reflects the fact that i abstains from purchase. It will be convenient to represent this assignment as a partitioning of the buyers into $\mathbf{S} = (S_0, S_1, \dots, S_m)$, where $S_j = \mu^{-1}(j) \subseteq N$ is the set of buyers assigned to j (and S_0 are the abstainers). An *outcome* (\mathbf{P}, \mathbf{S}) of the game comprises a set of schedules $\mathbf{P} = (\mathbf{p}_j)_{j \in M}$, and a partitioning \mathbf{S} . Under a given outcome (\mathbf{P}, \mathbf{S}) , a buyer $i \in S_j$ pays price $p_j(|S_j|)$ for any $j \neq 0$.

2.1 The Lu-Boutilier model

We now define the LB model in more detail. Following Lu and Boutilier [2012], we assume buyers are strategic but vendors are not for the purposes of defining buyer behavior, and therefore assume that discount schedules posted by vendors are fixed. This is standard for a two-stage leader-responder game (or Stackelberg game). We will extend this basic setting in Section 2.2 by adding vendor utilities, strategies, and informational assumptions to model the strategic interactions of vendors.

Assume a profile of discount schedules \mathbf{P} has been fixed by the vendors. Once buyers are assigned to specific vendors, the item prices are set by (\mathbf{P}, \mathbf{S}) as defined above. Buyer utility is *quasi-linear* in price: the utility of $i \in S_j$ is $u_i(\mathbf{P}, \mathbf{S}) = v_{ij} - p_j(|S_j|)$. For ease of exposition we assume buyers are never indifferent between products (vendors); as such, we assume a predetermined ordering of vendors, specific to each buyer, that is part of its type V_i . This is used to break ties between vendors for whom the buyer otherwise has the same utility.

Buyer behavior. If \mathbf{P} consists of fixed prices, every buyer has a strongly dominant strategy (recall we assume strict preferences). However, if one or more vendors post non-trivial discount schedules, the optimal behavior of one buyer may depend on the decisions of *other* buyers. When considered as a *cooperative game* where the utilities are either transferable or non-transferable, Lu and Boutilier showed that the core of the game may be empty. That is, there are discount schedules where, given any partitioning of buyers, some *coalition* of buyers can do better by moving to a different vendor. Thus as in that work, we consider outcomes that are stable to unilateral deviations by buyers from a given outcome.

We assume that if buyer i switches from vendor $j = \mu(i)$ in outcome (\mathbf{P}, \mathbf{S}) to some other vendor j' , she enjoys the (potentially reduced) price $p_{j'}(|S_{j'}| + 1)$ induced by her deviation. *Strong stability* requires that no single buyer gains from such a deviation. For *any profile of discount schedules* \mathbf{P} and type matrix \mathbf{V} , there is some partition \mathbf{S} that is strongly stable [Lu and Boutilier, 2012]. We refer to such a partition as a *stable buyer partition (SBP)*. An SBP is, in fact, a pure Nash equilibrium in the second stage

Valuations of	v_1 v_2	Price for	p_1 p_2
Buyer a	$\begin{pmatrix} 7 & 3 \\ 7 & 3 \\ 3 & 4 \end{pmatrix}$	1 buyer	$\begin{pmatrix} 8 & 3 \\ 4 & 3 \\ 1 & 3 \end{pmatrix}$
Buyer b		2 buyers	
Buyer c		3 buyers	

Table 1: The left columns show the valuation matrix \mathbf{V} , where v_{ij} is the value in the i 'th row and the j 'th column. For example, the type of Buyer b is $V_b = (7, 3)$. The right columns show vendors and their current discount schedules \mathbf{P} . Note that vendor 2 uses a fixed price of $p_2 = 3$.

of our game. However, since our focus is on vendor behavior in this work, we reserve the term *equilibrium* for analysis of the first stage of the game (i.e., vendor play).

We note that there may be multiple SBPs in any game. We make no assumptions about the SBP that is chosen in the second stage of our game except that it is Pareto efficient (i.e., that is not Pareto-dominated by another SBP). From among efficient SBPs, we may select in an arbitrary, pre-defined way. For example, Lu and Boutilier [2012] describe a method for finding SBPs that maximize social welfare, which could readily be incorporated into our model. Nothing in our results depends on the means of selection. Thus for every schedule \mathbf{P} and type matrix \mathbf{V} there is a unique outcome (\mathbf{P}, \mathbf{S}) , where $\mathbf{S} = \mathbf{S}(\mathbf{P}, \mathbf{V})$ is a Pareto-optimal SBP induced by \mathbf{P}, \mathbf{V} .

Example. Consider the (second-stage) game among buyers described in Table 1. Suppose that every buyer selects her favorite vendor, ignoring prices and breaking ties lexicographically. Then we obtain the partition $\mathbf{S} = (\{a, b\}, \{c\})$. Under partition \mathbf{S} , we have $u_a(\mathbf{P}, \mathbf{S}) = u_b(\mathbf{P}, \mathbf{S}) = 7 - p_1(2) = 3$; $u_c(\mathbf{P}, \mathbf{S}) = 4 - p_2(1) = 1$. \mathbf{S} is *not* a SBP, as buyer c has an incentive to move to vendor 1. Indeed, in the partition $\mathbf{S}' = (\{a, b, c\}, \emptyset)$ we have $u_c(\mathbf{P}, \mathbf{S}') = 3 - p_1(3) = 2 > u_c(\mathbf{P}, \mathbf{S})$. The partition \mathbf{S}' is an SBP, as no buyer has an incentive to deviate from it.

Note that \mathbf{S}' is not the only SBP. The partition \mathbf{S}'' where all buyers go with vendor 2 is also stable. Indeed, if buyer a moves to vendor 1, her utility will drop from $3 - p_2(3) = 0$ to $7 - p_1(1) = -1$, and similarly for b and c . However, \mathbf{S}' is the only Pareto optimal SBP. Thus $\mathbf{S}' = \mathbf{S}(\mathbf{P}, \mathbf{V})$.

2.2 Vendors as agents

Our focus in this work is on the strategic interactions among vendors. Once they select their prices and/or discount schedules, we assume that buyer behavior is dictated by the (nontransferable utility version of the) LB model above. We now turn our attention to the vendors.

Given an outcome (\mathbf{P}, \mathbf{S}) , the utility of vendor $j \in M$ is simply the revenue derived from the buyers assigned to it less the cost of the products sold: $U_j(\mathbf{P}, \mathbf{S}) = |S_j| \cdot (p_j(|S_j|) - c_j)$, where c_j is the cost of a single product to j . Recall that $\mathbf{S}(\mathbf{P}, \mathbf{V})$ is the SBP induced by the prices \mathbf{P} . This allows us to write $U_j(\mathbf{P}, \mathbf{V}) \equiv U_j(\mathbf{P}, \mathbf{S}(\mathbf{P}, \mathbf{V}))$.

Vendor behavior. Since the behavior of the buyers for any set of vendor discount schedules is well-defined, we can confine our analysis of the two-stage game to the first stage, where vendors announce prices. The incentives facing vendors in choosing their strategies depend critically on their knowledge of buyers’ types, as well as on their objective function. We consider three different models, which we briefly describe here (formal definitions appear in the sections that follow).

In the *full information model*, vendors know the precise buyer types and act to maximize their utility given the actions of other vendors. The *Bayesian (or expected utility) model* adopts a standard Bayesian game formulation: vendors have *partial information* in the form of a commonly-known distribution \mathcal{D} over (joint) buyer types, and act to maximize their *expected* utility, taking expectations over outcomes (SBPs) with respect to possible realizations of buyer types. The *strict uncertainty model* assumes that vendors possess even less information about the buyers: they know only the possible set of buyer types (i.e., only the *support* of the distribution is known). In this model, expected utility is ill-defined, so we instead adopt a common approach for such settings and assume vendors try to minimize their *worst-case regret* over all possible type realizations.

Best response and equilibrium. Informally, an equilibrium is a profile of vendor strategies (i.e., prices or discount schedules) such that no vendor prefers to use a different strategy, assuming that other vendors use their strategy in the profile, and the buyers behave as described above. Equivalently, a profile is *not in equilibrium* if some vendor has a *best response*, or alternative strategy, that it (strictly) prefers if other vendors use that profile. (The precise definition of “preferred strategy” depends on the model and how strategies are evaluated.)

Best responses are in some sense a more fundamental concept than equilibria, since analyzing equilibria depends on full understanding of available best responses. Furthermore, even in settings where we do not expect equilibria to emerge in the repeated interaction among vendors (or in settings where pure strategy equilibria may not exist, depending on the solution concept), an understanding of best-response dynamics provides natural insights into the likely outcomes of a game. Therefore, the main focus of this paper is the nature of vendor best responses to the actions of other vendors, and specifically the circumstances under which it is rational to respond with a non-trivial discount schedule rather than a fixed price. While not a focus of this work, all three models admit natural definitions of a vendor *equilibrium*, based on the corresponding best-response concept.

3 The Full Information Model

A game $G = \langle \mathbf{V}, C \rangle$ in the full information model is given by a buyer type matrix $\mathbf{V} = (v_{i,j})$ and vendors costs $C = (c_j)$. We will use this simple model to demonstrate a best-response by a vendor. Suppose that $c_1 = c_2 = 0$, and consider the example given in Table 1. Since by our assumption, under \mathbf{P} , \mathbf{V} buyers partition according to the SBP \mathbf{S}' (i.e., all buyers choose vendor 1), the utility of the vendors is $U_1 = 3 \cdot p_1(3) = 3$; $U_2 = 0$. This outcome is not an equilibrium for the vendors. For example, vendor 2

can change her discount schedule to $\mathbf{p}'_2 = (1, 1, 0)$, which would result in the SBP $\mathbf{S}(\mathbf{P}', \mathbf{V}) = (\{a, b\}, \{c\})$ with $U'_2 = 1$. Vendor 1 can also post a better schedule, by raising his prices to, say, $\mathbf{p}'_1 = (8, 6, 2)$. Thus, given the strategy of vendor 2 (or, more generally, all vendors except vendor 1), it is natural to ask what is the best discount schedule that vendor 1 can post.

As it turns out, the full information model is not especially interesting from our perspective.

Proposition 1. *In the full information model, there is always a fixed price p_j that is a best response for vendor j to any strategy profile \mathbf{P}_{-j} of the other vendors.*

Proof. If the vendors have full information, then they know exactly which buyer partitions will form given any profile of discounts \mathbf{P} . Suppose that vendor j expects to have t buyers under some nontrivial schedule \mathbf{p}_j . That is, $|S_j| = t$, where $\mathbf{S} = \mathbf{S}((\mathbf{p}_{-j}, \mathbf{p}_j), \mathbf{V})$. Then j can post a fixed price $p_j = p_j(t)$ and induce identical buyer behavior \mathbf{S} .

Let $\mathbf{P}^* = (\mathbf{p}_{-j}, p_j)$. To see why \mathbf{S} is still an SBP, assume toward a contradiction that some buyer i has a deviation from S_k to $S_{k'}$ under prices \mathbf{P}^* . First, $p_k(|S_k|) = p_k^*(|S_k|)$. If $k' \neq j$, then $p_{k'}(|S_{k'} \cup \{i\}|) = p_{k'}^*(|S_{k'} \cup \{i\}|)$. If $k' = j$, then

$$p_j(|S_j \cup \{i\}|) \leq p_j(|S_j|) = p_j = p_j^*(|S_j \cup \{i\}|).$$

In either case, buyer i would also deviate under prices \mathbf{P} , which is a contradiction. \square

This is not altogether surprising: the fact that some uncertainty is required to justify group discounts has previously been demonstrated, albeit in a somewhat different model, by Anand and Aron [2003]. As a consequence, volume discounts are of no value when buyer types are known (at least in the one-shot game with a fixed population of buyers).

4 The Bayesian Model

One reason for posting volume discounts rather than fixed prices is to hedge against uncertainty regarding the preferences (hence decisions) of the buyers. Specifically, a vendor can “insure” itself against the possibility that fewer buyers are attracted to its product than expected by using volume discounts. In the *Bayesian model* we assume that there is some joint distribution over buyers’ types $\mathcal{D} = \mathcal{D}(\mathcal{A})$ which is common knowledge among vendors. The game takes the form $G = \langle \mathcal{D}, C \rangle$. In the first stage of the game, vendors choose discount schedules, not knowing the buyers’ types. In the second stage, a type matrix $\mathbf{V} = (v_{ij})_{i \in N, j \in M}$ is drawn from \mathcal{D} . The goal of vendor j is to set a schedule \mathbf{p}_j that maximizes its expected utility:

$$U_j(\mathbf{P}, \mathcal{D}) = \mathbb{E}_{\mathbf{V} \sim \mathcal{D}}[U_j(\mathbf{P}, \mathbf{V})] = \mathbb{E}_{\mathbf{V} \sim \mathcal{D}}[U_j(\mathbf{P}, \mathbf{S}(\mathbf{P}, \mathbf{V}))].$$

A special case we consider is the case of *i.i.d. buyers*, in which $A_i = A$ for all $i \in N$. In other words, each buyer’s type is distributed according to a common distribution $\hat{\mathcal{D}}(A)$; and \mathcal{D} is the corresponding product distribution. We emphasize that within

$\widehat{\mathcal{D}}(A)$, buyer i 's preferences over different vendors may in fact be probabilistically dependent (i.e., v_{ij} and $v_{ij'}$ can be correlated).

4.1 A single vendor

First consider the case of a single vendor: suppose a monopolist is faced with distribution \mathcal{D} over buyer types. The simple example below demonstrates that a vendor can strictly increase its revenue, relative to any fixed price, using a non-trivial discount schedule. Assume two buyers, and a (discrete) type distribution that assigns probability 0.5 to each of two type matrices, $(3, 0)$ and $(2, 2)$. Note that buyers' valuations are correlated in \mathcal{D} . The optimal fixed price is $p = 2$, which guarantees revenue $U(p, \mathcal{D}) = 0.5 \cdot 2 \cdot 2 + 0.5 \cdot 2 = 3$. However, consider a discount schedule with a base price $p(1) = 3$, and a discounted price $p(2) = 2$. Its expected revenue, $U(\mathbf{p}, \mathcal{D}) = 0.5 \cdot 2 \cdot 2 + 0.5 \cdot 3 = 3.5$, is greater than that of the optimal fixed price. More detailed examples are given in Sections 4.3 and 4.4.

By contrast, if buyer valuations are i.i.d., the monopolist is always at least as well off using a fixed price:

Proposition 2. *Consider a single vendor facing n i.i.d. buyers with distribution \mathcal{D} . Let p^* be the optimal fixed price for the vendor. For any discount schedule \mathbf{p} , $U(\mathbf{p}, \mathcal{D}) \leq U(p^*, \mathcal{D})$.*

Proposition 2 in fact follows from a much more general result in auction design. Consider a vendor with an unlimited supply, facing i.i.d. bidders. It is known that *no truthful auction* yields more revenue, in expectation, than the optimal fixed price. It is not hard to see that any discount schedule is (equivalent to) a truthful auction mechanism (see also Section 6.2). For completeness, we provide a simple direct proof in of Proposition 2. Appendix A.

4.2 Multiple vendors

We now consider the best response of a vendor to the offers of other vendors. Suppose vendors other than j post discount schedules \mathbf{p}_{-j} . The best response of j is:

$$br_j^{EU}(\mathbf{p}_{-j}) = \operatorname{argmax}_{\mathbf{p}_j \in \mathcal{P}} U_j((\mathbf{p}_{-j}, \mathbf{p}_j), \mathcal{D}). \quad (1)$$

Our main result in the Bayesian model is that, assuming buyer types are independent and drawn from the same distribution, a vendor cannot benefit by using a discount schedule instead of a fixed price unless other vendors also use schedules. Below we show that these conditions are minimal: a non-trivial schedule can be of value if *any* of these three conditions is relaxed.

Theorem 3. *Let $G = \langle \mathcal{D}, C \rangle$ be a game with i.i.d. buyers. If all vendors except j post fixed prices, then the best response of vendor j is also a fixed price.*

Proof. W.l.o.g., we analyze the response of vendor 1 to fixed prices of the other vendors. Denote by q_2, \dots, q_m the (fixed) prices of all other vendors. Given distribution

$\widehat{\mathcal{D}}$ over a single buyer's types (a distribution over vectors of size m), we define the following single parameter distribution $\widehat{\mathcal{D}}'$:

$$Pr_{v \sim \widehat{\mathcal{D}}'}(v \geq x) \equiv Pr_{v \sim \widehat{\mathcal{D}}}(\forall j \in \{2, \dots, m\}, v_1 - v_j + q_j \gtrsim_j x), \quad (2)$$

where \gtrsim_j should be interpreted as $>$ if a buyer of type \mathbf{v} prefers vendor j over 1 in case of a tie, and otherwise as \geq .⁴ For instance, in the special case where the buyer always prefers vendor 1, we have $Pr_{v \sim \widehat{\mathcal{D}}'}(v \geq x) \equiv Pr_{v \sim \widehat{\mathcal{D}}}(v_1 - \max_{2 \leq j \leq m}(v_j - q_j) \geq x)$.

Now, consider vendor 1 to be a monopoly, facing buyers that are sampled from distribution \mathcal{D}' (or, equivalently, each buyer is i.i.d. sampled from $\widehat{\mathcal{D}}'$). An indifferent buyer (with $v = p$) always chooses to buy a product from vendor 1 over abstaining from purchase. The vendor can attract k buyers at price p iff there are k buyers for which $v_{i,1} \geq p$. However, for every buyer i ,

$$\begin{aligned} Pr_{V \sim \mathcal{D}'}(\text{buyer } i \text{ buys from monopolist}) &= Pr_{V \sim \mathcal{D}'}(v_i \geq p) \\ &= Pr_{v_i \sim \widehat{\mathcal{D}}'}(v_i \geq p) && \text{(since } \mathcal{D}' = (\widehat{\mathcal{D}}')^n \text{)} \\ &= Pr_{v_i \sim \widehat{\mathcal{D}}}(\forall j \in \{2, \dots, m\}, v_{i,1} - v_{i,j} + q_j \gtrsim_j p) && \text{(by Eq. (2))} \\ &= Pr_{\mathbf{V} \sim \mathcal{D}}(\forall j \in \{2, \dots, m\}, v_{i,1} - v_{i,j} + q_j \gtrsim_j p) && \text{(since } \mathcal{D} = (\widehat{\mathcal{D}})^n \text{)} \\ &= Pr_{\mathbf{V} \sim \mathcal{D}}(\forall j \in \{2, \dots, m\}, v_{i,1} - p \gtrsim_j v_{i,j} - q_j) \\ &= Pr_{\mathbf{V} \sim \mathcal{D}}(\text{buyer } i \text{ prefers vendor 1 over all others}) \end{aligned}$$

That is, vendor 1, when viewed as a monopolist, will attract the same set of buyers at the same price as in the multi-vendor game. Thus the revenue of any discount schedule \mathbf{p} for the single vendor under distribution \mathcal{D}' is equal to the revenue it derives using \mathbf{p} when faced with vendors posting prices q_2, \dots, q_m under distribution \mathcal{D} . One subtle issue is that given pricing $\mathbf{p}_1 = \mathbf{p}$ in the multi-vendor game G , there may be several SBPs. However, the stable partition that extracts the best price (i.e., the one where the largest number of buyers select vendor 1, and is also stable) Pareto dominates all other SBPs. Thus $\mathbf{S}((\mathbf{p}_1, q_2, \dots, q_m), V)$ is uniquely defined, and S_1 coincides with the set of buyers that purchase in the single vendor setting.

Finally, by Proposition 2, the best strategy for the vendor is to post a fixed price p^* , which entails that $p^* = br_j^{EU}(q_2, \dots, q_m)$. \square

4.3 When discounts help

There are three main conditions underlying Theorem 3: (a) all buyers have the same marginal distribution of values; (b) buyer valuations are independent; and (c) all other vendors use fixed prices. We now show that these are, in a sense, *minimal* requirements for the optimality of fixed prices. Specifically, relaxing any of the three admits non-trivial schedules as best responses in some circumstances.

Proposition 4. *For any pair of conditions taken from (a), (b) or (c), there is a game with two vendors and two buyers where the best response of one vendor is a non-trivial discount schedule.*

⁴Recall that according to our assumption every type implies a strict order over vendors that breaks ties.

We prove this proposition by deriving examples of such games that violate each of these conditions in turn, and showing that discount schedules do in fact serve as best responses in each case.

Relaxing condition (a). We first assume conditions (b) and (c) hold, but allow buyers to have different marginal distributions over valuations. Consider a simple counterexample with two vendors $M = \{1, 2\}$ and two independent (but not i.i.d.) buyers $N = \{a, b\}$. Both vendors have zero cost. Buyer a prefers vendor 1: $v_{a1} = 10 + x$, where $x \sim U(0, 1]$; and $v_{a2} = 10$. Buyer b prefers vendor 2: $v_{b1} = 10$; and $v_{b2} = 10 + y$, where $y \sim U(0, 1]$.

Consider the fixed price profile $P^* = (1, 1)$. The expected revenue is $U_1(P^*) = U_2(P^*) = 1$ (in fact this occurs w.p. 1, as every vendor keeps exactly one buyer). We argue that if discounts are not allowed, then P^* is an equilibrium, i.e., that no vendor can earn more than 1 by posting a fixed price. Indeed, suppose that vendor 1 announces some price $q > 1$, then it keeps buyer a w.p. $(2 - q)$, and

$$U_1(q, 1) = (2 - q)q + (1 - q)0 = 2q - q^2.$$

Similarly, if $q < 1$, then the vendor keeps buyer a for sure, and gains buyer b w.p. $1 - q$. Thus

$$U_1(q, 1) = (1 - q)2q + q \cdot q = 2q - q^2.$$

In other words, in both cases $U_1(q, 1) = 2q - q^2$, which has a maximum at $q^* = 1 = p_1^*$. The argument for the second vendor is the same.

Nevertheless, if vendor 1 deviates to the non-trivial discount schedule $\mathbf{q}'_1 = (1, 3/4)$, then it can derive higher expected revenue. Vendor 1 keeps buyer a as before. W.p. $1/4$, buyer b has a preference of less than $1/4$ for vendor 2 (i.e., $y < 1/4$), and will select vendor 1 in the unique SBP $\mathbf{S}((\mathbf{q}'_1, p_2), V)$. Hence:

$$\begin{aligned} U_1(\mathbf{q}'_1, p_2) &= 1/4(2q'_2) + 3/4 \cdot q'_1 = 1/4(2 \cdot 3/4) + 3/4 \cdot 1 \\ &= 3/8 + 3/4 = 9/8 > 1 = U_1(P^*). \end{aligned}$$

Relaxing condition (b). Our next example shows that relaxing independence, but retaining conditions (a) and (c), also admits discounting as a best response. Consider the previous game, but with probability $1/2$, swap the preferences (types) of both buyers. In other words, either: (i) a prefers vendor 1 and b prefers vendor 2 (with the valuations outlined in the previous example); or (ii) b prefers vendor 1 and a prefers vendor 2; and each of (i) and (ii) occurs with probability $1/2$. This results in a symmetric distribution, but one that correlates their valuations of buyers a and b . The fixed profile $P = (1, 1)$ remains a *fixed price* equilibrium. Moreover, since the best response of vendor 1 to price 1 is $\mathbf{q}_1 = (1, 3/4)$ regardless of its type, it remains a best response in the new game.

Relaxing condition (c). Lastly, we describe a game with two i.i.d. buyers, maintaining conditions (a) and (b), but where the best response for vendor 1 to a discount

schedule posted by vendor 2 is itself a schedule (the full analysis of this example appears in Appendix B). Let $v_{a1} = v_{b1} = 10$, $v_{a2} = 10 + x_a$, and $v_{b2} = 10 + x_b$, where x_a and x_b are sampled i.i.d. from $\mathcal{D} = U[-1, 1]$. As long as prices are not too high (say, below 8) buyer i 's decision is determined only by the value difference x_i between her value for the two vendors. It is not hard to verify that the best response of every vendor j to the fixed price $p_{-j} = 1$ is $p_j = 1$, even if schedules are allowed (i.e., $P = (1, 1)$ is an equilibrium, see Section 6.1). However, suppose vendor 2 posts schedule $\mathbf{q} = (1, 0.8)$. Vendor 1's best response is not a fixed price: it can be shown that its optimal fixed price is $p^* \cong 0.922$, yielding revenue of 0.93656, while the schedule $(0.93, 0.914)$ yields slightly higher revenue of 0.93675.

4.4 Revenue bounds

As we have seen in the examples in Section 4.3, there are instances in which a vendor derives greater expected revenue by posting a discount schedule rather than a fixed price. It is thus natural to attempt quantify the potential benefit of using volume-based discounts. To simplify the approach, we limit our analysis to a monopolist vendor, and examine the potential increase in utility relative to the optimal fixed price. Our approach and bounds can easily be extended to any number of vendors.

When faced with n buyers, there are cases where a vendor can increase its profit by a factor of n by using group discounts; that is, the value of using discounts is unbounded as the number of buyers grows. The proof of the following proposition assumes that buyer valuations for the vendor's product are correlated. However, the types of correlations used can occur quite naturally when publicly observable signals (such as product quality, product reviews, or network effects and externalities) influence buyer valuations.

Proposition 5. *A vendor's expected utility for using a discount schedule, when faced with n buyers, can be no more than a factor of n greater than the expected utility of the optimal fixed price. Moreover, there are instances where this improvement is realized by a discount schedule.*

Proof. We first prove the upper bound, that no discount schedule improves on the optimal fixed price by more than a factor of n . Let \mathbf{p} be any discount schedule. Let B_j be the event that price $p(j)$ realized.

$$\begin{aligned}
n \max_q U(q, \mathcal{D}) &\geq \sum_{j \leq n} U(p(j), \mathcal{D}) = \sum_{j \leq n} p_j E_{\mathbf{V} \sim \mathcal{D}}[|\{i \in N : v_i \geq p(j)\}|] \\
&\geq \sum_{j \leq n} p(j) E_{\mathbf{V} \sim \mathcal{D}}[|\{i \in N : p(j+1) > v_i \geq p(j)\}|] \\
&\geq \sum_{j \leq n} p(j) Pr_{\mathcal{D}}(B_j) E_{\mathbf{V} \sim \mathcal{D}}[|\{i \in N : p(j+1) > v_i \geq p(j)\}| | B_j] \\
&\geq \sum_{j \leq n} Pr_{\mathcal{D}}(B_j) p(j) j = U(\mathbf{p}, \mathcal{D}).
\end{aligned}$$

Thus the expected revenue $\max_q U(q, \mathcal{D})$ of the optimal fixed price is at least $\frac{1}{n}$ of that of any discount schedule \mathbf{p} .

We now construct an example that achieves this factor n improvement. Consider n buyers who are uncertain about their valuation of the product, believing it to be of one of n different *product types* (e.g., the product has one of n different quality levels). Let $\varepsilon > 0$ be some constant.

A product of type $j \in [n]$ has j buyers who value it at $v_j = \frac{\varepsilon^{j-n}}{j}$, whereas all other buyers value it at 0. Note that when $\varepsilon = 1$, the average value of all buyers for every j is 1. We can think of event j as representing the level of detail in a product review: when $j = n$, no buyer knows anything about the product, so all buyers assign it an expected value of 1. As j decreases, the review becomes more detailed, providing additional information, so that $n - j$ buyers realize they do not want this product, while expected value of other buyers increase.

We define \mathcal{D}_ε so that for each $j < n$ the product is of type j with probability $q_j = \varepsilon^{n-j}$. With the remaining probability mass $q_n = 1 - \varepsilon - \varepsilon^2 - \dots - \varepsilon^{n-1}$, the product is of type n . In the best case, the vendor can extract the full surplus under any product type. Note that this can be achieved with the schedule $p^*(j) = v_j$, thus the optimal schedule \mathbf{p}^* extracts from any buyer her full expected value:

$$U(\mathbf{p}^*, \mathcal{D}_\varepsilon) = \sum_{j=1}^n q_j p^*(j) j = q_n + \sum_{j=1}^{n-1} \varepsilon^{n-j} \varepsilon^{j-n} = n - 1 + q_n,$$

which approaches n as $\varepsilon \rightarrow 0$.

Suppose instead that the vendor uses a fixed price. With any fixed price of the form $p_j = p^*(j) = v_j, j \in [n]$ —clearly there is no reason to use any other price—the vendor will achieve revenue of:

$$\begin{aligned} U(p_j, \mathcal{D}_\varepsilon) &= p_j \sum_{j' \leq j} q_{j'} \cdot j' = \sum_{j' < j} q_{j'} p_j j' + q_j p_j j \\ &\leq \sum_{j' < j} \varepsilon n + q_j p_j j \leq \varepsilon n^2 + 1 \xrightarrow{\varepsilon \rightarrow 0} 1. \end{aligned}$$

Thus for any n , we can set ε sufficiently small so that the revenue from any fixed price is arbitrarily close to 1, whereas the best discount schedule attracts revenue that approaches n . \square

What we see is that with such correlations, the vendor can effectively use group discounts to achieve price discrimination, which was impossible in the i.i.d. case. Another interesting question pertains to the maximum utility a vendor can gain by using discounts when facing buyers that are *distinct* but *independent*. The first part of Proposition 4 shows that a factor of *at least* $\frac{9}{8}$ improvement is achievable; but it remains an open question as to whether a constant upper bound can be proven.

5 The Strict Uncertainty Model

The assumption that vendors have distributional knowledge of buyers' types may not be viable in certain situations. In this section, we consider an alternative model of

uncertainty, the *strict uncertainty* model, where vendors know only the *possible types* that buyers may possess. The game is structured as in the Bayesian model, but rather than sampling buyer types from a distribution, *arbitrary* types from the type space $A_1 \times \cdots \times A_n$ are chosen. This type of game is known variously as a *game in informational form* [Holzman et al., 2004] or an *incomplete information game with strict type uncertainty* [Hyafil and Boutilier, 2004].

In such settings, one plausible objective for vendors is to select a strategy that, given the strategies of other vendors, maximizes its *worst-case utility* under all possible realizations of buyer types. However, such an approach is inappropriate in our setting. For example, if buyer valuations can fall below a vendor's cost, that vendor's worst-case utility is at most 0, regardless of its actions. We therefore consider a more natural objective, assuming each vendor selects a strategy that minimizes its *worst-case or maximum regret* over possible possible buyer types. The minimax regret approach has deep roots in decision making under strict uncertainty [Savage, 1972], has found use in robust optimization [Kouvelis and Yu, 1997] and decision support [Boutilier et al., 2006; Braziunas and Boutilier, 2010], and has been applied in various game-theoretic contexts [Hyafil and Boutilier, 2004; Ashlagi et al., 2006].

Notation. We adapt the definitions of *minimax regret* of Hyafil and Boutilier [2004] to our model. Recall that $\mathcal{A} \subseteq \mathbb{R}^{n \times m}$ is the set of possible types for buyers. The possible types $A_i \subset \mathbb{R}^m$ for buyer i take the form $(v_{ij})_{j \in M}$ and must occur as a column in some $\mathbf{V} \in \mathcal{A}$. It is sometimes useful to think of \mathcal{A} as the support of some distribution \mathcal{D} , in which case A_i would be the support of the marginal distribution over buyer i 's types.

Once vendors select strategies (prices) \mathbf{P} , suppose that the realized buyer types are given by \mathbf{V} . This results in the buyer partition $\mathbf{S} = \mathbf{S}(\mathbf{P}, \mathbf{V})$. We define the *regret* $Reg_j(\mathbf{P}, \mathbf{V})$ of vendor j relative to this outcome to be the difference between the maximum utility it could have achieved by choosing different prices given this type realization (taking other vendor strategies to be fixed) and its utility given its selected strategy:

$$Reg_j(\mathbf{P}, \mathbf{V}) = \max_{p'_j \in \mathbb{R}} U_j((p'_j, \mathbf{p}_{-j}), \mathbf{S}(\mathbf{P}', \mathbf{V})) - U_j(\mathbf{P}, \mathbf{S}(\mathbf{P}, \mathbf{V})),$$

where $\mathbf{P}' = (p'_j, \mathbf{p}_{-j})$. Note that w.l.o.g. the optimal schedule p'_j can be taken to be a fixed price (by Prop. 1).

Given their uncertainty about the buyer's types, and without a distribution over these types, we assume vendors assess the utility of a strategy by assuming a type profile realization that maximizes the degree to which that strategy is suboptimal. Specifically, let \mathbf{P} be a strategy (pricing) profile, where vendor j uses p_j . We define the *maximum regret* of vendor j to be:

$$MaxReg_j(\mathbf{P}) = \max_{\mathbf{V} \in \mathcal{A}} Reg_j(\mathbf{P}, \mathbf{V}).$$

The goal of each vendor is the selection of a strategy that *minimizes* its maximum regret given the strategies of other vendors. We define the *minimax best response* of vendor j

to strategy profile \mathbf{p}_{-j} of its competitors to be:

$$br_j^{MR}(\mathbf{p}_{-j}) = \operatorname{argmin}_{\mathbf{p}_j \in \mathcal{P}} \operatorname{MaxReg}_j(\mathbf{p}_j, \mathbf{p}_{-j}).$$

Note that maximum regret is minimized with respect to the types of the *buyers*, not the actions of other vendors, which are assumed to be known. Minimax regret equilibrium can be defined in the obvious way using this form of best response [Hyafil and Boutilier, 2004].

5.1 Independent types

We now assess the value of discounts in the strict uncertainty model, assuming vendors minimize max-regret.

As in the Bayesian model, we can identify restrictions on the type space \mathcal{A} that allow us to derive strong restrictions on the form of best responses. First, we say that the buyer types \mathcal{A} are *symmetric* if $A_i = A_{i'}$ for all $i, i' \in N$; in other words, all buyers have identical sets of possible preferences. We say that buyer types \mathcal{A} are *independent* if $\mathcal{A} = \times_{i \in N} A_i$; in other words, the realization of the type of a buyer does not restrict the possible realizations of any other buyer's type. These notions are related to the i.i.d. assumptions in the Bayesian model. Suppose that \mathcal{A} is the support of some unknown distribution \mathcal{D} . Clearly if buyers have identical marginals in \mathcal{D} then \mathcal{A} is symmetric; and if buyers are independent in \mathcal{D} then \mathcal{A} is also independent. However, the corresponding assumptions on \mathcal{A} are much weaker, and may hold even if \mathcal{D} does not have independence or identically distributed buyer types.

Lemma 6. *If all vendors use fixed prices $P = (p_1, \dots, p_m)$, and \mathcal{A} is symmetric and independent, then maximum regret for each vendor is realized when all buyers have the same type.*

Proof. We need to show that if $r = \operatorname{Reg}_j(P, \mathbf{V})$ is vendor j 's maximum regret, then there is a type $V^* \in A$ (which depends on the strategy profile) such that $\operatorname{Reg}_j(P, \mathbf{V}^*) \geq r$, where every row in \mathbf{V}^* equals to V^* .

We break our analysis into two cases. Suppose first that j 's maximum regret is realized by a buyer type profile in which its best response is to ask a *lower* price than its price p_j in P , which by definition must attract more buyers than p_j . That is, there is a $p'_j < p_j$ s.t. $|S'_j| > |S_j|$, where $\mathbf{S}' = \mathbf{S}(P', \mathbf{V})$, $P' = (p'_j, p_{-j})$. Then there is some buyer $i \in S'_j \setminus S_j$. Let $V^* = V_i$. We now compute the regret at profile P when all buyers are of type $V^* \in A$.

Denote $\mathbf{T} = \mathbf{S}(P, \mathbf{V}^*)$ and $\mathbf{T}' = \mathbf{S}(P', \mathbf{V}^*)$. Since i prefers vendor j under prices P' , we have that $v_{ij} - p'_j \geq v_{ij'} - p_{j'}$ for all $j' \neq j$. As $V^* = V_i$, in the partition \mathbf{T}' all buyers select j , and thus $|T'_j| = n \geq |S'_j|$.

Likewise, under prices P , buyer i prefers some other vendor j' over j , i.e. $v_{ij'} - p_{j'} \geq v_{ij} - p_j$. Thus in \mathbf{T} all buyers select $j' \neq j$, and $|T_j| = 0 \leq |S_j|$.

Bringing these inequalities together, we have:

$$\begin{aligned}
Reg_j(P, \mathbf{V}^*) &\geq U_j(P', \mathbf{T}') - U_j(P, \mathbf{T}) \\
&= |T'_j|(p'_j - c_j) - |T_j|(p_j - c_j) \geq |S'_j|(p'_j - c_j) - |S_j|(p_j - c_j) \\
&= U_j(P', \mathbf{V}) - U_j(P, \mathbf{V}) = r.
\end{aligned}$$

In other words, regret is maximized in a type profile in which all buyers have identical valuations.

The remaining case is when vendor j 's regret is maximized under some profile in which its best response p'_j is to *increase* its price, potentially losing some some buyers, but increasing its margins on those it keeps. Thus $p'_j > p_j$, and $0 < |S'_j| \leq |S_j|$. (In contrast to the previous case, here we require that p'_j is the optimal price in retrospect). In this case, there must be a buyer $i \in S'_j$ of type $V^* = V_i$. Again we consider the regret given type matrix \mathbf{V}^* (where every row equals to V^*). Since i still prefers j in P' , we have that $v_{ij} - p'_j \geq v_{ij'} - p_{j'}$ for all $j' \neq j$. Thus in \mathbf{T}' all buyers still prefer j , and $|T'_j| = n \geq |S'_j|$. Since $p_j < p'_j$, then clearly $|T_j| = n$ as well. As a result:

$$\begin{aligned}
r &= U_j(P', \mathbf{V}^*) - U_j(P, \mathbf{V}^*) = |S'_j|(p'_j - c_j) - |S_j|(p_j - c_j) \\
&\leq |S_j|(p'_j - c_j) - |S_j|(p_j - c_j) = |S_j|(p'_j - p_j) \leq n(p'_j - p_j)
\end{aligned}$$

On the other hand,

$$\begin{aligned}
Reg_j(P, \mathbf{V}^*) &\geq U_j(P', \mathbf{T}') - U_j(P, \mathbf{T}) \\
&= |T'_j|(p'_j - c_j) - |T_j|(p_j - c_j) = n(p'_j - p_j) \geq r.
\end{aligned}$$

Again regret is maximized when all buyers have the same type.

We emphasize that the type V^* depends on the profile P , and for every profile there may be a different “worst-case” type. \square

Our main result in analyzing the strict uncertainty model is similar in spirit to Theorem 3.

Theorem 7. *If all vendors except j use fixed prices, and \mathcal{A} is symmetric and independent, then $br_j(p_{-j})$ is a fixed price.*

Proof. Let \mathbf{q}_j be the schedule that is the best response to p_{-j} , i.e., $MaxReg_j(\mathbf{q}_j, p_{-j})$ is minimal. Let $p_j = q_j(n)$, i.e., the price induced by that schedule when j attracts n buyers, and let $P = (p_j, p_{-j})$. We show that $MaxReg_j(P) = r$ is also minimal. Intuitively, the proof shows that the only part of j 's strategy that impacts its utility (in the worst case) is the price set in the schedule for all n buyers. Thus, the fixed price $p_j = q_j(n)$ is as good as schedule \mathbf{q}_j .

Consider $MaxReg_j(p_j, p_{-j})$ as a function of p_j . For any p_j , there is some type matrix \mathbf{V}^* where maximum regret under P is realized, i.e., $Reg_j(P, \mathbf{V}^*) = MaxReg_j(P) = r$. There is an optimal price p'_j for \mathbf{V}^* such that $Reg_j(P, \mathbf{V}^*) = U_j((p'_j, p_{-j}), \mathbf{V}^*) - U_j(P, \mathbf{V}^*)$. By Lemma 6, w.l.o.g. all buyers have the same type in \mathbf{V}^* , and thus either $S_j = S_j(P, \mathbf{V}^*)$ includes all buyers or S_j is empty.

Suppose $MaxReg_j(\mathbf{q}_j, p_{-j}) < r$. By definition $Reg_j((\mathbf{q}_j, p_{-j}), \mathbf{V}) < r$ for any \mathbf{V} , in particular for the profile \mathbf{V}^* . However, in \mathbf{V}^* either $|S_j| = n$ or $|S_j| = 0$ for any prices. Recall that $\mathbf{S} = \mathbf{S}(P, \mathbf{V}^*)$ and denote $\mathbf{S}' = \mathbf{S}(p'_j, p_{-j}, \mathbf{V}^*)$; $\mathbf{T} = \mathbf{S}((\mathbf{q}_j, p_{-j}), \mathbf{V}^*)$. In particular, $|T_j| \in \{0, n\}$.

If $|T_j| = 0$, then $|S_j| = 0$ as well, since at price $p_j = q_j(n)$ vendor j does not attract any buyer of type V^* . If $|T_j| = n$, then the vendor attracts all buyers of type V^* at price p_j and thus $|S_j| = n = |T_j|$, and

$$|T_j|(q_j(|T_j|) - c_j) = |S_j|(q_j(|n|) - c_j) = |S_j|(p_j - c_j).$$

Note that in either case $|T_j|(q_j(|T_j|) - c_j) = |S_j|(p_j - c_j)$. Thus for some p'_j ,

$$\begin{aligned} Reg_j((\mathbf{q}_j, p_{-j}), \mathbf{V}^*) &= U_j(p'_j, p_{-j}, \mathbf{V}^*) - U_j((\mathbf{q}_j, p_{-j}), \mathbf{V}^*) \\ &= |S'_j|(p'_j - c_j) - |T_j|(q_j(|T_j|) - c_j) \\ &= |S'_j|(p'_j - c_j) - |S_j|(p_j - c_j) \\ &= U_j(p'_j, p_{-j}, \mathbf{V}^*) - U_j(P, \mathbf{V}^*) \\ &= Reg_j(P, \mathbf{V}^*). \end{aligned}$$

Therefore,

$$r = MaxReg_j(P) = Reg_j(P, \mathbf{V}^*) \leq MaxReg_j(\mathbf{q}_j, p_{-j}).$$

In other words, $p_j \in br_j^{MR}(p_{-j})$, as required. \square

5.2 Distinct and correlated types

With identical types spaces, we see that discounts provide no value to a vendor if other vendors use fixed prices. However, analogous to the Bayesian model, if the type spaces of some buyers differ from those of others, then the maximum regret of a vendor *can* be (strictly) minimized by posting a non-trivial schedule. More specifically, as in the expected utility model, if we relax either symmetry or independence, the result above no longer holds.

Proposition 8. *There are discount games in which a vendor minimizes its maximum regret by posting a non-trivial schedule if either (a) \mathcal{A} is asymmetric and independent; or (b) \mathcal{A} is symmetric and non-independent.*

Proof. We begin with part (a). Consider a game with a single vendor having zero cost and with three buyers. The possible valuations (types spaces) for the buyers are $A_1 = [6, 12]$; $A_2 = A_3 = [0, 6]$, and $\mathcal{A} = A_1 \times A_2 \times A_3$.

We first show that the optimal fixed price is $p^* = 4$, and that $MaxReg(4) = 8$. Clearly $p^* \leq 6$, since for every $p > 6$ the types $(6, 6, 6)$ lead to a utility of 0, whereas a price of $p' = 6$ would give the maximum possible utility of 18; hence the maximum regret of any $p > 6$ is 18.

For every $p \leq 6$, the maximum regret is attained for the types $\mathbf{v} = (p - \varepsilon, p - \varepsilon, 12)$. Note that $U(p, \mathbf{v}) = p$, whereas $U(p - \varepsilon, \mathbf{v}) = 3p - 2\varepsilon$ and $U(12, \mathbf{v}) = 12$. Thus

$MaxReg(p) \geq Reg(p, \mathbf{v}) = \max\{12 - p, 3p - 2\varepsilon - p\} \cong \max\{12 - p, 2p\}$. The latter term is minimized when $12 - p = 2p$, hence for $p^* = 4$, maximum regret is 8.

Next, we show that the schedule $\mathbf{p} = (6, 4, 4)$ has maximum regret of at 6. We consider three cases:

- If $v_2, v_3 \geq 4$, then the realized price is 4, and $U(4, \mathbf{v}) \geq 3 \cdot 4 = 12$. On the other hand, the maximum utility is 18, thus $Reg(4, \mathbf{v}) \leq 18 - 12 = 6$.
- If $v_2 \geq 4 > v_3$ (or vice versa), then the realized price is 4, and $U(4, \mathbf{v}) \geq 2 \cdot 4 = 8$. However, the maximum utility in this case is obtained by either selling one item at price 12, or two items at price 6 (since $v_3 < 6$), or three items at price $4 - \varepsilon$. In each case the optimal utility is no more than 12, and thus $Reg(4, \mathbf{v}) \leq 12 - 8 = 4$.
- If $v_2, v_3 < 4$, then the realized price is 6, and $U(6, \mathbf{v}) \geq 6$. The optimal strategy under this realization is to either sell a single item at price 12, or all three items at price $4 - \varepsilon$. Therefore, $Reg(6, \mathbf{v}) \leq 12 - 6 = 6$. The maximum regret of this discount schedule is less than that of the fixed price with minimax regret.

For the proof of part (b), note that for any number of buyers and vendors we can easily impose symmetry (at the expense of independence), similarly to the randomized swap in the Bayesian model. Formally, we define $\mathcal{A}' = \{\pi(\mathbf{V}) | \mathbf{V} \in \mathcal{A}, \pi \in S_n\}$ (i.e., all possible buyer permutations over the allowed type matrices). Since vendors experience the same regret regardless of the chosen permutation, a vendor's utility under \mathcal{A} and \mathcal{A}' is always the same. By applying this construction to the above example, we get the proof for (b). \square

In contrast to the Bayesian model, it is an open question whether fixed prices are still dominant when the restriction on other vendors is relaxed.

6 Discussion

We have investigated conditions under which vendors may benefit from posting group or volume discounts for groups of buyers—assuming that buyers can coordinate their purchasing activities—relative to the posting of fixed prices. We showed that, when facing i.i.d. buyers that use the coordination mechanism of Lu and Boutilier [2012], complex discount schedules cannot yield greater revenue than that generated using the optimal fixed price. This holds whether vendors know the distribution of buyer types or simply the support of this distribution. This is consistent with similar findings in other models of group buying (see Section 1.1). This robust result highlights the fact that the design of effective pricing schemes for group buying should focus on settings where group discounts provide vendor value, including domains where buyer valuations are correlated by unobservable factors (such as perceived quality or advertising impact), marginal production costs are decreasing, vendors are risk-seeking, or where discounts have viral or long-term acquisition benefits.

We now outline several directions for future research, with an emphasis on equilibrium analysis and the connection to optimal auctions.

6.1 Vendor Equilibria

While this paper focused on the structure of a vendor's *best-response*, natural questions arise regarding to the existence and properties of (pure) Nash equilibria in our model. Since our game has two steps, the appropriate solution concept is a subgame perfect equilibrium.

Equilibrium concepts. Formally, a (pure) *subgame perfect equilibrium* (SPE) of a complete information game $G = \langle \mathbf{V}, C \rangle$ is a profile of discount schedules $\mathbf{P} \in \mathcal{P}^m$ such that no vendor j prefers $\mathbf{S} = \mathbf{S}(\mathbf{P}, V)$ over $\mathbf{S}' = \mathbf{S}((\mathbf{p}_{-j}, \mathbf{p}'_j), V)$, for any $\mathbf{p}'_j \in \mathcal{P}$. A *fixed subgame perfect equilibrium* (FSPE) in G is defined similarly, with the exception that only trivial (fixed price) schedules are allowed. Note that an FSPE is neither a special case nor more general than an SPE since the restriction applies both to the profile \mathbf{P} and to the deviation \mathbf{p}'_j . However, by Prop. 1, the existence of an SPE in the full information model entails the existence of an FSPE with the same utilities for each vendor; and any FSPE is also an SPE in the full information model.

To illustrate, consider a game with two vendors (with zero costs), and two buyers. Buyer a has preferences $v_a = (3, 1)$ (i.e., a prefers vendor 1), whereas $v_b = (1, 3)$. The fixed profile $P = (3, 3)$ (i.e., $p_1 = 3, p_2 = 3$) is a FSPE. Each vendor has utility $U_i = 3$. Any $p'_i > 3$ will result in a utility of 0, whereas setting the price low enough to attract both buyers will result in a utility of at most 2.

When we consider the more complex models of partial information, we can construct similar definitions of SPE and FSPE using the already-defined concepts of best response. For example, in the Bayesian model \mathbf{P} is an SPE in the game $G = \langle \mathcal{D}, C \rangle$ if $\mathbf{p}_j \in br_j^{EU}(\mathbf{p}_j, G)$ for all $j \in M$. Due to Theorems 3 and 7, we know that under an i.i.d. distribution any FSPE in the Bayesian/Distribution-free model is also an SPE. However, this does not hold when the distribution is not i.i.d.: in the examples described in proof of Proposition 4(a) and (b), the fixed profile $P^* = (1, 1)$ is an FSPE but not an SPE. Also note that the existence of an SPE does not entail the existence of an FSPE, since the best response to a discount schedule might not be a fixed price (see Proposition 4(c)).

A game with no pure equilibrium. It is interesting to note that, in general, a (pure) SPE/FSPE may not exist at all—even in the complete information model and even when discounts are forbidden.

We now describe an example with two vendors and four buyers with known types that does not admit an FSPE (and by Proposition 1, no SPE). Buyer types are: $a_1 = (9, 0); a_2 = (0, 9); b_1 = b_2 = (5, 5)$. Buyer $b_i, i \in \{1, 2\}$, has a weak preference to vendor i in the case of a tie. Both vendors have zero cost. It can be verified by detailed case analysis that given any fixed price profile $P = (p_1, p_2)$ there is at least one vendor who would raise or lower her price to improve her utility.

Let x, y be the fixed prices offered by vendors 1 and 2, respectively. We break our analysis into several cases, and show in each that at least one vendor can increase her utility by changing her price.

- If $x > 9$ or $y > 9$ then the relevant vendor has no buyers and will lower his price to 9.
- If $9 > x > 5$ then $x' = 9$ is an improvement. Likewise for $9 > y > 5$.
- If $x = y = 9$ then $y' = 5$ brings the utility of vendor 2 from $9 \cdot 1$ to $5 \cdot 3 = 15$.
- If $x = 9, y < 5$, then $y' = 5$ is an improvement. Likewise for $y = 9, x < 5$.
- If $x = 9, y = 5$, then $x' = 5$ increases the utility of vendor 1 to $5 \cdot 2 = 10 > 9$. Similarly for $x = 5, y = 9$.
- If $x = y \in (4, 5]$ then $x' = 4$ improves the utility of vendor 1 from $2x \leq 10$ to $3(x') = 12$. Similarly, the higher vendor will deviate to 4 for any $x, y \in (4, 5]$.
- If $y \leq x \leq 4$ then $x' = 9$ improves the utility of vendor 1 from (at most) $2x \leq 8$ to 9. Similarly for vendor 2 if $x \leq y \leq 4$.
- The last case is $y \leq 4, x \in (4, 5]$. Then $x' = 9$ improves the utility of vendor 1 from $x \leq 5$ to 9. Similarly for vendor 2 if $x \leq 4, y \in (4, 5]$.

Intuitively, the market can be understood as follows. Each vendor has some “core buyers” (a_1, a_2 in our example), and “flexible buyers” that have no vendor preference and are influenced only by price (b_1, b_2). Suppose we start from a state where both prices are high. The vendors compete for the flexible buyers by lowering prices, until at some point it is more worthwhile to one of the vendors (say 1) to opt out of the competition. Then vendor 1 increases the price, keeping only its core buyers. Now vendor 2 can also raise its price since there is no competition, and we are back in the start state.

Developing conditions under which pure SPE and FPSE exist is of great interest, especially in cases where all vendors use group discounts.

6.2 Group discounts and optimal auctions

In his seminal study on optimal auctions, Myerson [1981] characterized truthful auction mechanisms, showing that when a seller with a single good faces i.i.d. buyers, the optimal mechanism is effectively a second-price auction with a *fixed* reserve price. In a related analysis of a model closer to ours, to our model, Goldberg et al. [2001] show that the worst-case revenue of using a fixed price to sell digital goods (i.e., unlimited supply) is at least as great as that any other *truthful* “auction” mechanism. They also provided bounds on the revenue attainable by such a mechanism when compared to a variety of benchmarks, which later became standard in the literature.

Consider the following auction rule that sets a price for bidder i , when submitted bids are (b_1, \dots, b_n) : “if there are at least $k - 1$ other bidders s.t. $b_j \geq p(k)$, then allocate an item to i for a price of $p_i = q(k)$ ”. Clearly the outcome of such an auction is equivalent to the outcome of posting the discount schedule \mathbf{q} . Moreover, the allocation rule is monotone in the bid, and the price p_i is independent of b_i . Thus according to Myerson, this auction is truthful. While the Goldberg et al. do not deal with a Bayesian

setting, it is known that the result still holds when bidder valuations are sampled i.i.d. from some distribution.

It would be interesting to study the relationship between optimal truthful mechanisms and group discount mechanisms in more depth in various settings.⁵

6.3 Other directions

Within our current model, further research is needed to understand the full impact of group discounts when buyer valuations are correlated by signals—such as product quality, vendor reputation, or advertising.

Even for a monopolist, the question of computing optimal discount schedules (given a correlated distribution over buyer types) is non-trivial. When facing competing vendors the problem becomes even more complex. It is therefore important to identify natural classes of distributions where optimal or near-optimal schedules can be computed efficiently.

Another interesting question is what to expect when buyers use stronger coordination mechanisms, such as those that allow *transferable utility* [Yamamoto and Sycara, 2001; Lu and Boutilier, 2012]. We conjecture that our main result about the optimality of a fixed price when buyers are i.i.d. still holds under different variations such as transferable utility.

Network externalities is the common name for the phenomenon where buyers assign utility to sharing a vendor or a service (such as a cellular provider or online social network) [Katz and Shapiro, 1985; Banerji and Dutta, 2009; Feldman et al., 2013]. One way to think of group discounts is in terms of adding *indirect* network externalities to the interaction among buyers—the interactions are indirect in the sense that they mediated by the vendors (or by arbitrators of buying groups) and depend on vendor strategies, with the buyers themselves potentially even unaware of the choices of other buyers. Incorporating network structure into our model of group discounts may also be very relevant, as such structure may determine which groups of buyers are able or are likely to coordinate—especially in the context of transferable utility.

⁵We thank Brendan Lucier for highlighting this connection.

References

- Anand, K. and R. Aron (2003). Group buying on the web: A comparison of price-discovery mechanisms. *Management Science* 49(11), 1546–1562.
- Ashlagi, I., D. Monderer, and M. Tennenholtz (2006). Resource selection games with unknown number of players. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-06)*, pp. 819–825. ACM.
- Banerji, A. and B. Dutta (2009). Local network externalities and market segmentation. *International Journal of Industrial Organization* 27(5), 605 – 614.
- Boutilier, C., R. Patrascu, P. Poupart, and D. Schuurmans (2006). Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence* 170(8–9), 686–713.
- Braziunas, D. and C. Boutilier (2010). Assessing regret-based preference elicitation with the UTPREF recommendation system. In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC’10)*, Cambridge, MA, pp. 219–228.
- Chen, J., X. Chen, and X. Song (2007). Comparison of the group-buying auction and the fixed pricing mechanism. *Decision Support Systems* 43(2), 445–459.
- Chen, J., R. Kauffman, Y. Liu, and X. Song (2010). Segmenting uncertain demand in group-buying auctions. *Electronic Commerce Research and Applications* 9(2), 126–147.
- Edelman, B., S. Jaffe, and S. D. Kominers (2011). To Groupon or not to Groupon: The profitability of deep discounts. Technical Report Working Paper Number 11-063, Harvard Business School, Cambridge, MA.
- Feldman, M., R. Meir, and M. Tennenholtz (2013). Competition in the presence of social networks: How many service providers maximize welfare? In *Proceedings of the 9th Workshop on Internet & Network Economics*. to appear.
- Goldberg, A. V., J. D. Hartline, and A. Wright (2001). Competitive auctions and digital goods. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-01)*, pp. 735–744. Society for Industrial and Applied Mathematics.
- Holzman, R., N. Kfir-Dahav, D. Monderer, and M. Tennenholtz (2004). Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior* 47(1), 104–123.
- Hyafil, N. and C. Boutilier (2004). Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (UAI-04)*, pp. 268–277. AUAI Press.
- Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *The American Economic Review* 75(3), 424–440.

- Kouvelis, P. and G. Yu (1997). *Robust Discrete Optimization and Its Applications*. Dordrecht: Kluwer.
- Lal, R. and R. Staelin (1984). An approach for developing an optimal discount pricing policy. *Management Science* 30(12), 1524–1539.
- Li, C., S. Chawla, U. Rajan, and K. Sycara (2005). Mechanism design for coalition formation and cost sharing in group-buying markets. *Electronic Commerce Research and Applications* 3(4), 341–354.
- Lu, T. and C. Boutilier (2012). Matching models for preference-sensitive group purchasing. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 723–740. ACM.
- Manisterski, E., D. Sarne, and S. Kraus (2008). Enhancing cooperative search with concurrent interactions. *Journal of Artificial Intelligence Research* 32(1), 1–36.
- Matsuo, T., T. Ito, and T. Shintani (2005). A volume discount-based allocation mechanism in group buying. In *Proceedings of the International Workshop on Data Engineering Issues in E-commerce*, pp. 59–70.
- Monahan, J. (1984). A quantity discount pricing model to increase vendor profits. *Management Science* 30(6), 720–726.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73.
- Prashanth, B. and Y. Narahari (2008). Efficient algorithms for combinatorial auctions with volume discounts arising in web service composition. In *Proceedings of the 4th IEEE Conference on Automation Science and Engineering*, Washington, DC, pp. 995–1000.
- Sarne, D. and S. Kraus (2005). Cooperative exploration in the electronic marketplace. In *AAAI*, Volume 5, pp. 158–163.
- Savage, L. J. (1972). *The Foundation of Statistics*. Dover.
- Wang, Q. (2002). Determination of suppliers’ optimal quantity discount schedules with heterogeneous buyers. *Naval Research Logistics (NRL)* 49(1), 46–59.
- Yamamoto, J. and K. Sycara (2001). A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proceedings of the Fifth International Conference on Autonomous Agents*, pp. 576–583. ACM.

A Proof of Proposition 2

W.l.o.g., the optimal fixed price p^* can be set deterministically (i.e., randomized pricing cannot do better). Let $r^* = p^* \Pr_{\mathcal{D}}(v > p^*)$ be the optimal expected revenue that can be extracted from a single buyer. Applying the optimal fixed price p^* to all n buyers gives an expected revenue of nr^* .

Assume, by way of contradiction, that some discount schedule $\mathbf{p} = (p(1), \dots, p(n))$ yields strictly greater revenue than nr^* . Let r_i be the expected revenue extracted from buyer i using \mathbf{p} . Then $\sum_i r_i > nr^*$, i.e., there is at least one buyer (w.l.o.g. assume buyer n) s.t. $r_n > r^*$. We now construct a pricing strategy that yields revenue r_n from buyer n . Independently sample $n - 1$ values from \mathcal{D} , simulating the first $n - 1$ buyers, and sort values so that $v_1 \geq \dots \geq v_{n-1}$. Now select price $p(1)$ iff $v_1 < p(1)$, $p(2)$ iff $v_2 < p(2) \leq v_1$, and more generally $p(k)$ iff $v_k < p(k) \leq v_{k-1}$. These events are pairwise disjoint and cover the entire event space (since the union of events 1 to k holds iff least $n - k$ buyers have values below $p(k)$).

Let A_k denote the k 'th event, and B_k the corresponding event when actual buyer values are drawn from \mathcal{D}^{n-1} . Clearly $\Pr(A_k) = \Pr(B_k)$. Moreover, when B_k occurs, *exactly* $k - 1$ buyers have value at least $p(k)$. Thus buyer n purchases iff $v_n \geq p(k)$ as well, and pays $p(k)$ if so. However, this is exactly the purchase probability and price paid by a single buyer when the proposed price is $p(k)$. Thus the revenue is $\sum_{k=1}^n \Pr(A_k) \Pr(v \geq p(k) | A_k) p(k)$ (from the single buyer), i.e.,

$$\sum_{k=1}^n \Pr(B_k) \Pr(v_n \geq p(k) | B_k) p(k) = r_n > r^*.$$

Thus \mathbf{p} extracts more than r^* from a single buyer (a contradiction).

B Proof of Proposition 4

Although this is not required for the proof, we show that in all three examples there is a FSPE. The value of discounts in the first two cases (relaxing (a) and relaxing (b)) has been shown in full in the main text, so we only show FSPE.

Lemma 9. *Consider the game described in Proposition 4(a). $P = (p_1 = 1, p_2 = 1)$ is a FSPE.*

Proof. Note that in profile P , w.p. 1 buyer a will go to vendor 1, and buyer b will go to vendor 2, thus $u_1(1) = u_2(1) = 1$.

Suppose that vendor 1 announces some price $q > 1$, then it keeps client a w.p. $(2 - q)$, and

$$u_1(q) = (2 - q)q + (1 - q)0 = 2q - q^2.$$

Similarly, if $q < 1$, then the vendor keeps client a for sure, and gains client b w.p. $1 - q$, Thus

$$u_1(q) = (1 - q)2q + q \cdot q = 2q - q^2.$$

I.e. in both cases $u_1(q) = 2q - q^2$, which has a maximum in $q^* = 1 = p_1$. The argument for the second vendor is the same. \square

-1 ← x_a → 1			
-1	a always prefers vendor 1	a always prefers vendor 1	a always prefers vendor 1
↑	b always prefers vendor 1	b prefers vendor 1 given that a does	b always prefers vendor 2
x_b	b always prefers vendor 1	Each buyer prefers vendor i	b always prefers vendor 2
↓	a prefers vendor 2 given that b does	if the other buyer does	a prefers vendor 2 given that b does
1	a always prefers vendor 2	a always prefers vendor 2	a always prefers vendor 2
	b always prefers vendor 1	b prefers vendor 2 given that a does	b always prefers vendor 2

By using the same argument as in the main text, $P = (1, 1)$ is also an FSPE in the game described in Proposition 4(b).

Relaxing condition (c) Lastly, we describe a game with two i.i.d. buyers (i.e. holding conditions (a,b)), where the best response to a schedule posted by vendor 2 is also a schedule.

We recall the definitions from the main text. Both buyers have $v_{a1} = v_{b1} = 10$. The preference for vendor 2 is $v_{a2} = 10 + x_a; v_{b2} = 10 + x_b$, where x_a and x_b are sampled i.i.d from $\mathcal{D} = U[-1, 1]$. Note that as long as prices are not too high (say, below 8) the decision of buyers is determined only by the difference between preference to vendor 1 and to vendor 2, i.e. by the values x_a and x_b .

Lemma 10. *The profile $P = (1, 1)$ is an FSPE. The expected revenue in P is 1 to each vendor.*

Proof. Clearly in P each vendor gets every buyer w.p. $1/2$. Thus $U_i(P, \mathcal{D}) = 1/2 \cdot 1 + 1/2 \cdot 1 = 1$. Suppose that vendor 2 switches to $p'_2 \neq p_2 = 1$. For any $2 > p'_2 > 0$, the revenue from each buyer is $p'_2 pr(x_{ib} > p'_2 - 1)$. Thus

$$U_2(p_1, p'_2) = 2p'_2 pr(x_{ib} > p'_2 - 1) = 2p'_2(2 - p'_2)/2 = 2p'_2 - (p'_2)^2 < 2 = U_2(P).$$

For $p'_2 \notin [0, 2]$ the revenue is even lower, thus $p_2 = br_2^{EU}(p_1)$. The same analysis holds for vendor 1, thus $P = (1, 1)$ is a FSPE. \square

Lemma 11. *Suppose that vendor 2 posts the schedule $\mathbf{q} = (1, 0.8)$. Then the best response of vendor 1 is not a fixed price.*

Proof. Denote the strategies of vendors 1 and 2 by $\mathbf{p} = (p, p')$ and $\mathbf{q} = (q, q')$, respectively. The revenue of vendor 1 can be written as a function of \mathbf{p} and \mathbf{q} . The prices divide the type space (and thus the probability space) to 9 regions as follows.

In each cell we know exactly how many buyers bought from vendor 1 and at what price. It remains to compute the probability of each cell. We denote the cells by T=top, M=med, B=bottom, L=left, R=right. We assume that in the middle cell MM both buyers select the same vendor, with equal probability to each vendor.

When the maximal distance between prices is no more than 1,

$$\begin{aligned} U_1(\mathbf{P}, \mathcal{D}) &= 2p' pr(TL) + 2p' pr(TM) + 2p' pr(ML) \\ &\quad + p \cdot pr(TR) + p \cdot pr(BL) + 1/2 \cdot 2p' pr(MM) \\ &= \frac{1}{4}[2p'(q' - p + 1)^2 + 4p'(q' - p + 1)(q - q' + p - p') \\ &\quad + 2p(q' - p + 1)(p' - q + 1) + p'(q - q' + p - p')^2]. \end{aligned}$$

Now, suppose that vendor 2 posts the schedule $\mathbf{q} = (1, 0.8)$. Using the formula, it can be verified that the best fixed response to \mathbf{q} (i.e. under the constraint $p' = p$) is $p^* \cong 0.922$, which yields a revenue of $U_1((p^*, \mathbf{q}), \mathcal{D}) \cong 0.93656$. However, the schedule $p = 0.93, p' = 0.914$ yields a slightly higher revenue of 0.93675. While this is not a large improvement, it still indicates that condition (a) is necessary for Theorem 3. \square