

On the Value of Using Group Discounts under Price Competition

Reshef Meir¹ and Tyler Lu² and Moshe Tennenholtz³ and Craig Boutilier²

¹The Hebrew University of Jerusalem, Israel and Microsoft Research, Herzliya
reshef.meir@mail.huji.ac.il

²University of Toronto
{cebly,tl}@cs.toronto.edu

³Technion-Israel Institute of Technology and Microsoft Research, Herzliya
moshet@microsoft.com

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, either as a monopolist or when competing 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 buyers' types are i.i.d. and that other vendors also use fixed prices. We also show that these assumptions cannot be relaxed: if buyers are not i.i.d., or other vendors post discount schedules, then posting a schedule may yield 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.

Introduction

Online services offering consumer group discounts represent an important and growing segment of online sales. Despite margin pressures, services such as Groupon, Google Offers and hundreds of others remain successful, offering consumers a choice of multiple, competing offers from vendors of identical or similar products. 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). 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, and suggest corresponding algorithms. However, they assume discounts to be fixed and given, modeling the interaction as a (both a cooperative and non-cooperative) game among the buyers themselves. As such, they do not address the incentives for vendors to offer such discounts in the first-place, nor strategic interactions involving the vendors.

Several models have been proposed that examine vendor incentives for offering discounts (see Related work section that follows) under a variety of utility and informational assumptions. Most adopt a two-stage model of interaction (along the lines of Stackelberg games), in which one or more vendors first commit to a *discount schedule* (or pricing strategy), 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, something that has become increasingly feasible using online services 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 behavior affects vendors. Specifically, we analyze the conditions under which a vendor can derive value by using a discount schedule rather than a fixed price.

To extend the LB model into a game that reflects the strategic interactions of both buyer and vendors, we must specify the vendors' utility structure, their beliefs about buyer valuations, and the relation between the two. We consider three natural models that differ in vendor information structure: (a) a *complete information model*, where vendors know buyer valuations (or types); (b) an *expected utility model*, where vendor beliefs take the form of a distribution over the buyers' types; and (c) a *distribution-free model*, where vendors know only the set of *possible* buyer types. In each model, vendor utility is linear in the number of units sold.

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 in each model. Types are then determined (if they were unknown) and buyers co-

ordinate their purchases using the LB mechanism. We assume a non-cooperative model of buyer behavior in which they have full information about offers and other buyers' types, but cannot transfer payments or make binding agreements. Thus buyers form stable partitions, where no *single buyer* can benefit by switching to a different vendor.

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.

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 monopolist, for example when facing buyers whose types are correlated by a signal on the quality of the product.

Somewhat less related (but still within the two-stage framework) is the *group buying auction* model (Chen, Chen, and Song 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 gain using 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, Ito, and Shintani 2005; Prashanth and Narahari 2008).

Different buyer coordination mechanisms have been suggested assuming *transferable utility* (Yamamoto and Sycara 2001; Li *et al.* 2005; Lu and Boutilier 2012), which requires the possibility of binding agreements. In our models, as in (Anand and Aron 2003; Chen, Chen, and Song 2007) and in the non-cooperative version of (Lu and Boutilier 2012), we assume a non-cooperative setting that excludes mone-

tary transfer among buyers, and focus on (one-shot) vendor revenue maximization.

One of the prime economic motivations for vendor discounts in Groupon-like models is customer acquisition, where (often steep) discounts incur a loss in the short-term, but longer-term repeat business justifies this cost (Edelman, Jaffe, and Kominers 2011). This is not reflected in any of the models discussed here, ours included.

Our contribution. Our main contribution is the analysis of the impact of buyer coordination (in the LB model) on vendor pricing, in particular, in the presence of competing vendors. We first show that with complete information there is no reason to use group discounts. In the Bayesian (expected utility) model, we prove that if buyer valuations are independent and identically distributed, and all other vendors use fixed prices, then a fixed price is optimal. However, if *any* of these conditions is relaxed, then a vendor may gain by posting a discount schedule rather than a fixed price. We provide a similar result in the distribution-free setting: a vendor facing buyers with the same set of possible valuations, with other vendors offering fixed prices, should also post a fixed price in order to minimize regret.

Omitted proofs can be found in the full version of this paper available online.¹

Model and Notation

We use uppercase 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 (an unlimited number of units of) a single product, while each buyer i has a *type* V_i , i.e., a vector of non-negative values v_{ij} for each vendor j 's product. Buyers have unit demand. We let $\mathbf{v}_j$ denote the vector of values for vendor j (over all $i \in N$) and $\mathbf{V}$ the full value matrix. Each vendor has a fixed cost c_j for producing one unit, 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 Aron 2003). Let $\mathcal{P}$ be the set of all discount schedules and $\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_m)$ a *profile* of schedules, one per vendor. A schedule with a single fixed price is a *trivial schedule*, and is denoted $p_j \in \mathbb{R}_+$. In the second stage, each buyer selects a single vendor (or abstains). An *outcome* $(\mathbf{P}, \mathbf{S})$ of the game is the set of schedules $\mathbf{P} = (\mathbf{p}_j)$, and an assignment $\mu : N \rightarrow M \cup \{0\}$ of buyers to vendors partitions them into $\mathbf{S} = (S_0, S_1, \dots, S_m)$, where S_j is the set of buyers assigned to j (S_0 are the abstainers). Given outcome $(\mathbf{P}, \mathbf{S})$, a buyer $i \in S_j$ pays $p_j(|S_j|)$.

We begin by defining the LB model, where buyers are strategic but vendors are not, and the schedules posted by vendors are fixed. We then extend this basic setting by

¹Available from <http://tinyurl.com/c3cerg3>.

adding vendor utilities, strategies, and informational assumptions to model the strategic interactions of vendors.

The LB model

We assume a profile of 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); we assume a predetermined vendor order for each buyer that is part of its type, and is used to break ties across vendors who have 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 there are non-trivial discount schedules, optimal buyer decisions may depend on the decisions of *other* buyers.

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)*.² There may be multiple SBPs in any game. We make no strong assumptions about the chosen SBP, but assume only that the buyers play a SBP that is efficient, i.e., that is not Pareto-dominated by another SBP. From these partitions, we may select arbitrarily in some pre-defined way. For example, Lu and Boutilier (2012) describe a method for finding SBPs that maximize social welfare, which could readily be adopted in our model.

Thus for every schedule $\mathbf{P}$ and type matrix $\mathbf{V}$ there is a unique outcome $(\mathbf{P}, \mathbf{S})$, where $\mathbf{S}$ is a SBP.

Vendors as agents

The utility of vendor $j \in M$ is simply the revenue derived from the buyers assigned to it: $U_j(\mathbf{P}, \mathbf{S}) = |S_j| \cdot (p_j(|S_j|) - c_j)$, where c_j is the cost of a single product to j . For any profile $\mathbf{P} \in \mathcal{P}^m$, let $\mathbf{S}(\mathbf{V}, \mathbf{P})$ be the SBP that is induced by the prices $\mathbf{P}$. This allows us to write $U_j(\mathbf{P}, \mathbf{V}) \equiv U_j(\mathbf{P}, \mathbf{S}(\mathbf{V}, \mathbf{P}))$.

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 (formal definitions appear in the sections that follow).

²An SBP is a pure Nash equilibrium in the second stage of our game. However, we reserve the term *equilibrium* for the first stage of vendor play.

In the *full information model*, vendors know the precise buyer types and try to maximize utility. 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 try to maximize expected utility. The *strict uncertainty model* assumes even less information: vendors only know 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 such that no vendor *prefers* to use a different strategy, assuming buyers and vendors behave as described above. Equivalently, a profile is *not in equilibrium* if some vendor has a *best response* that it (strictly) prefers when other vendors use that profile.

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 (or potentially when they do not exist, depending on the solution concept) best-response dynamics provide 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.

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 $C = (c_j)$. The full information model is not especially interesting from our perspective. If the vendors have full information, then they know exactly which buyer partitions will form given any profile of discounts. Thus if vendor j expects to have t buyers under some nontrivial schedule $\mathbf{p}_j$, it can post a fixed price $p_j = p_j(t)$ and induce identical buyer behavior. 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 (Anand and Aron 2003).

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. A vendor can “insure” itself against the possibility that fewer buyers than expected are drawn to its product. In the *Bayesian model* we assume each buyer i has a set of possible types $A_i \subseteq \mathbb{R}_+^m$, and there is some joint distribution over types $\mathcal{D} = \mathcal{D}(A_1 \times A_2 \times \dots \times A_n)$ which is common knowledge among vendors. A 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 $\mathbf{V} = (v_{ij})_{ij}$ is drawn from $\mathcal{D}$. The goal of vendor j is to set a schedule $\mathbf{p}_j$ that maximizes its expected utility:

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

A special case we consider is the case of *i.i.d. buyers*: $A_i = A$ for all $i \in N$, each buyer's type is distributed according to a common distribution $\widehat{\mathcal{D}}(A)$, and $\mathcal{D}$ is the corresponding product distribution.³

A single vendor

First consider the case of a single vendor: suppose a monopolist is faced with distribution $\mathcal{D}$. 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. Similar examples with continuous distributions are easily constructed.

By contrast, if buyers are *i.i.d.*, the monopolist is always better off using a fixed price.

Proposition 1. *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 $\mathbf{p}$, $U(\mathbf{p}, \mathcal{D}) \leq U(p^, \mathcal{D})$.*

Proof. 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). $\square$

Multiple vendors

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

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

where EU stands for Expected Utility. 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 2. *Let $G = \langle \mathcal{D}, C \rangle$ be a game with i.i.d. buyers. If all vendors except j are using fixed prices, then the best response of vendor j is also a fixed price.*

Proof sketch. W.l.o.g. we analyze vendor 1, and assume $q_2, \dots, q_m$ are the (fixed) prices of the other vendors. Given distribution $\mathcal{D}$ over buyers' types define a single parameter distribution $\mathcal{D}'$ s.t. for all $x \in \mathbb{R}$,

$$\Pr_{v \sim \mathcal{D}'}(v > x) \equiv \Pr_{\mathbf{v} \sim \mathcal{D}}(v_1 - \max_{2 \leq j \leq m} (v_j - q_j) > x).$$

When vendor 1 is a monopolist facing buyers sampled i.i.d. from $\mathcal{D}'$, it can attract k buyers at price p_1 iff there are k buyers for which $v_{i,1} > p_1$ (i.e., $\mathcal{D}'$ "simulates" the multi-vendor state in which vendor 1 finds himself).

The revenue of any schedule $\mathbf{p}$ for vendor 1 under $\mathcal{D}'$ is equal to the revenue it accrues using $\mathbf{p}$ when other vendors post prices $q_2, \dots, q_m$ under distribution $\mathcal{D}$ (our assumption that we select a Pareto-dominant SBP is required). By Prop. 1, the best strategy for vendor 1 is to post a fixed price p^* , i.e., $br_j(q_2, \dots, q_m) = p^*$. $\square$

There are three main conditions underlying Thm. 2: (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 3. *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.*

³Within $\widehat{\mathcal{D}}(A)$, any buyer i 's preferences over different vendors may be dependent (i.e., $v_{ij}, v_{ij'}$ can be correlated).

Relaxing condition (a). We first assume conditions (b) and (c) hold, but allow buyers to have different marginal distributions. 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 client 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 client a for sure, and gains client 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 schedule $\mathbf{q}'_1 = (1, 3/4)$, then it can do better: Vendor 1 always 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 $\mathbf{S}(V, (\mathbf{q}'_1, p_2))$. 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. This results in a symmetric distribution, but correlates their values. 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 (we omit the full analysis due to space constraints). 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 $\widehat{D} = U[-1, 1]$. 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 profile $P = (1, 1)$ is a Nash equilibrium even if schedules are allowed. However, suppose vendor 2 posts schedule $\mathbf{q}_2 = (1, 0.8)$. Vendor 1's best response is not a fixed price: it can be shown that its optimal fixed price is $p_1^* \cong 0.922$, yielding revenue of 0.93656, while the schedule $\mathbf{q}'_1(0.93, 0.914)$ yields slightly higher revenue of 0.93675.

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 $A_1 \times \dots \times A_n$ are chosen. One plausible vendor objective is to maximize *worst-case utility*, but such an approach is inappropriate in our setting. For example, if buyer valuations can lie 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*. The minimax regret approach has deep roots in decision making (Savage 1972), and it has been applied in various game-theoretic contexts (Hyafil and Boutilier 2004; Ashlagi, Monderer, and Tennenholtz 2006).

Notation. We adapt the definitions of *minimax regret* from (Hyafil and Boutilier 2004) to our model. Let $A_i \subset \mathbb{R}^m$ be the set of possible types for buyer i , and $\mathbf{A} = \times_{i \in N} A_i$. Once vendors select strategies (prices) $\mathbf{P}$, suppose realized buyers' types are $\mathbf{V}$, resulting in the buyer partition $\mathbf{S} = \mathbf{S}(\mathbf{V}, \mathbf{P})$. The *regret* $Reg_j(\mathbf{P}, \mathbf{V})$ of vendor j in this outcome is the difference between its maximal profit in retrospect, and its actual profit:

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

where $\mathbf{P}' = (p'_j, \mathbf{p}_{-j})$. Note that w.l.o.g. p'_j is a fixed price and not a schedule.

Without a type distribution, vendors assume the worst-case realization of types. The *maximum regret* over all possible types is:

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

The goal of each vendor is therefore the selection of a strategy that minimizes its maximum regret. The best response to strategy profile $\mathbf{p}_{-j}$ is:

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

Note that regret is minimized w.r.t. the types of the *buyers*, not the actions of other vendors, which are assumed to be known.⁴

Discounts and regret

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

Lemma 4. *If all vendors use fixed prices, and buyer type spaces are symmetric (i.e., $A_i = A$ for all i), then maximum regret for each vendor is realized when all buyers have the same type.*

⁴Minimax regret equilibrium can be naturally defined, as a profile where the best response of every agent is its current action.

Our main result in the strict uncertainty model is similar in spirit to Thm. 2.

Theorem 5. *If all vendors except j use fixed prices, and buyer type spaces are symmetric, then $br_j^{MR}(p_{-j})$ 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 for n buyers, and $P = (p_j, p_{-j})$. We will show that $MaxReg_j(P) = R$ is also minimal. Intuitively, the proof shows that the only part of j 's strategy that is being used in practice (in the worst case) is the price for the complete set N . Thus, 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 max-regret under P is realized, i.e., $Reg_j(P, \mathbf{V}^*) = MaxReg_j(P) = R$. There is an optimal price p'_j for $\mathbf{V}^*$ s.t. $Reg_j(P, \mathbf{V}^*) = U_j((p'_j, p_{-j}), \mathbf{V}^*) - U_j(P, \mathbf{V}^*)$. By Lemma 4, w.l.o.g. all buyers have the same type in $\mathbf{V}^*$, denoted by $V^* \in A$. Thus either $S_j = S_j(P, \mathbf{V}^*)$ has 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 uniform profile $\mathbf{V}^* = (V^*, \dots, V^*)$. However, in $\mathbf{V}^*$ either $|S_j| = n$ or $|S_j| = 0$ for any prices. Recall that $\mathbf{S} = \mathbf{S}(\mathbf{V}^*, P)$ and denote $\mathbf{S}' = \mathbf{S}(\mathbf{V}^*, (p'_j, p_{-j}))$; $\mathbf{T} = \mathbf{S}(\mathbf{V}^*, (\mathbf{q}_j, p_{-j}))$. 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 ,

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

i.e., a contradiction. Therefore

$$MaxReg_j(p_j, p_{-j}) = R \leq MaxReg_j(q_j, p_{-j}),$$

i.e. $p_j \in br_j^{MR}(p_{-j})$, as required. $\square$

Non-identical type spaces. 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 type spaces are distinct, then a single vendor *can* derive value by posting a non-trivial schedule.

Consider a game with a single vendor having zero cost and with three buyers. The values for the buyers are $v_1 \in [6, 12]$; $v_2, v_3 \in [0, 6]$. In this game, the best fixed price for the vendor is $p^* = 4$, and $MaxReg(p^*) = 8$. However by posting the discount schedule $\mathbf{p} = (6, 4, 4)$, the vendor attains regret at most 6. This can be shown

by splitting the possible valuations into cases, and deriving the maximum regret for each case separately. For example, if $v_2, v_3 \geq 4$, then the realized price is 4, and $U(4, V) \geq 3 \cdot 4 = 12$. On the other hand, maximal utility is 18, thus $Reg(4, V) \leq 18 - 12 = 6$. The other two cases, where either one or both values are less than 4, are treated similarly.

We conjecture that in the strict uncertainty model, fixed prices are dominant even if the restriction on other vendors is relaxed (in contrast to the Bayesian model).

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; a finding that is consistent with similar findings in other models of group buying (see the Related work section). 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 buyers' 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.

Future work. A number of interesting directions for future research remain. One interesting question is whether similar results hold when buyers use stronger coordination mechanisms, such as those that allow transferable utility (Yamamoto and Sycara 2001; Lu and Boutilier 2012). 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—and to develop algorithms that compute optimal discounts for such settings.

Other important questions relate to the existence and properties of equilibria in our model. We have derived some preliminary results showing that *pure* vendor equilibria may not exist in our model, either with or without discounts, even in the complete information model (and even with a single buyer!). Developing conditions under which such equilibria exist is of great interest, especially in cases where all vendors use group discounts.

Acknowledgments

The authors thank Uri Feige, Asaf Horev and Liad Blumrosen for their comments and the anonymous reviewers for helpful suggestions. Boutilier and Lu acknowledge the support of the Natural Sciences and Engineering Research Council (NSERC). Lu was also supported by a Google Ph.D. Fellowship.

References

- Anand, K., and Aron, R. 2003. Group buying on the web: A comparison of price-discovery mechanisms. *Management Science* 49(11):1546–1562.
- Ashlagi, I.; Monderer, D.; and Tennenholtz, M. 2006. Resource selection games with unknown number of players. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, 819–825. ACM.
- Chen, J.; Kauffman, R.; Liu, Y.; and Song, X. 2010. Segmenting uncertain demand in group-buying auctions. *Electronic Commerce Research and Applications* 9(2):126–147.
- Chen, J.; Chen, X.; and Song, X. 2007. Comparison of the group-buying auction and the fixed pricing mechanism. *Decision Support Systems* 43(2):445–459.
- Edelman, B.; Jaffe, S.; and Kominers, S. D. 2011. Togroupon or not togroupon: The profitability of deep discounts. Technical Report Working Paper Number 11-063, Harvard Business School, Cambridge, MA.
- Hyafil, N., and Boutilier, C. 2004. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, 268–277. AUAI Press.
- Lal, R., and Staelin, R. 1984. An approach for developing an optimal discount pricing policy. *Management Science* 30(12):1524–1539.
- Li, C.; Chawla, S.; Rajan, U.; and Sycara, K. 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 Boutilier, C. 2012. Matching models for preference-sensitive group purchasing. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 723–740. ACM.
- Matsuo, T.; Ito, T.; and Shintani, T. 2005. A volume discount-based allocation mechanism in group buying. In *Proceedings of the International Workshop on Data Engineering Issues in E-commerce*, 59–70.
- Monahan, J. 1984. A quantity discount pricing model to increase vendor profits. *Management science* 30(6):720–726.
- Prashanth, B., and Narahari, Y. 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*, 995–1000.
- Savage, L. 1972. *The foundation of statistics*. Dover publications.
- Wang, Q. 2002. Determination of suppliers’ optimal quantity discount schedules with heterogeneous buyers. *Naval Research Logistics (NRL)* 49(1):46–59.
- Yamamoto, J., and Sycara, K. 2001. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proceedings of the fifth international conference on Autonomous agents*, 576–583. ACM.