

Matching Models for Preference-sensitive Group Purchasing

TYLER LU, University of Toronto
CRAIG BOUTILIER, University of Toronto

Matching buyers and sellers is one of the most fundamental problems in economics and market design. An interesting variant of the matching problem arises when self-interested buyers come together in order to induce sellers to offer *quantity or volume discounts*, as is common in buying consortia, and more recently in the consumer group couponing space (e.g., Groupon). We consider a general model of this problem in which a group or buying consortium is faced with volume discount offers from multiple vendors, but group members have *distinct preferences* for different vendor offerings. Unlike some recent formulations of matching games that involve quantity discounts, the combination of varying preferences and discounts can render the core of the matching game empty, in both the transferable and nontransferable utility sense. Thus, instead of coalitional stability, we propose several forms of Nash stability under various epistemic and transfer/payment assumptions. We investigate the computation of buyer-welfare maximizing matchings and show the existence of transfers (subsidized prices) of a particularly desirable form that support stable matchings. We also study a nontransferable utility model, showing that stable matchings exist; and we develop a variant of the problem in which buyers provide a simple preference ordering over “deals” rather than specific valuations—a model that is especially attractive in the consumer space—which also admits stable matchings. Computational experiments demonstrate the efficacy and value of our approach.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: stable matching, preferences, demand aggregation, group purchasing, volume discounts, daily deals, cooperative games.

1. INTRODUCTION

Matching buyers and sellers is one of the most fundamental problems in economics and market design. A wide variety of models and mechanisms have been developed that reflect different assumptions about the demands, valuations/preferences, and knowledge of the market participants and their ability to cooperate. Each leads to its own computational challenges when developing algorithms for computing stable (core) matchings, Nash equilibria, clearing prices or other solution concepts.

In this paper, we address the problem of *cooperative group buying*, in which a group of buyers coordinate their purchases to realize volume discounts, mitigate demand risk, or reduce inventory costs. Group buying has long been used for corporate procurement, via industry-specific buying consortia or broadly based group purchasing organizations (GPOs) [Chen and Roma 2010]. The advent of the Internet, in particular, has helped businesses with no prior affiliation more easily aggregate their demand [Anand and Aron 2003]. Consumer-oriented group purchasing has also been greatly facilitated by the web; and the recent popularity of volume-based couponing and “daily

Authors' address: Department of Computer Science, University of Toronto, {t1,cebly}@cs.toronto.edu.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
EC'12, June 4–8, 2012, Valencia, Spain. Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

deal” providers like Groupon and Living Social (and services that aggregate such deals) has propelled group discounts into the public consciousness.

Group buying and demand aggregation has been studied from several perspectives, and many models have been proposed for their analysis. However, we consider a vital ingredient of group buying that has received insufficient attention in the literature, namely, the fact that buyers often have *distinct preferences* for the offerings of different vendors. Most matching models with volume discounts assume that vendor offerings are indistinguishable to buyers, which significantly limits their applicability. For instance, suppose two buyers X and Y are (jointly) comparing the offers of two vendors for some item: A offers a price of 10 for one unit, but a discounted price of 8 if both buy from him; and B offers a single price of 9 per unit. If A and B are indistinguishable, X and Y should cooperate and buy from A . But suppose X prefers B (with valuation 11.5) to A (valuation 10). In this case, X would prefer to stick with B unless Y offers some payment to switch vendors (Y would gladly share some of her generated surplus with X for this purpose). Without the ability to express preferences over vendors, “group buying” would not emerge even in this trivial example.

While matching becomes much more subtle in such models, assigning buyers to vendors in a way that triggers volume discounts, *while remaining sensitive to buyer preferences*, offers flexibility and efficiency gains that greatly enhance the appeal of group buying. Consider a group of businesses or *buyers* working with a GPO to procure supplies within a specific product category (e.g., manufacturing materials, packaging, transportation, payroll services, etc.). The GPO is able to negotiate volume discounts from a handful of suppliers or *vendors*, possibly with multiple discount thresholds. Buyers generally have different valuations for the offerings of different vendors (e.g., buyers may have slightly different manufacturing specifications; or may prefer the contract, payment or delivery terms of certain vendors). A suitable matching of buyers to vendors must trade off these preferences with the triggered discount prices.

The same issues arise in consumer domains. Suppose a daily deal aggregator creates a “marketplace” for some product category, say, spas. Multiple spas offer deals that only trigger if a certain quantity is sold. Buyers are faced with a dilemma: they may want only one item, but are uncertain about which deal will trigger. If they only offer to buy (i.e., conditionally purchase) their most preferred spa, they may not get *any* deal if their preferred deal does not trigger. But if they offer on multiple spas to hedge that risk, they run the opposite risk of obtaining more items than they want. A matching model that allows consumers to specify preferences for items relative to their discounted prices provides flexibility that benefits both consumers and retailers.

Our model. In broad strokes, our model assumes a set of vendors offering products (e.g., within a specific product category). Interacting with some GPO or informal buying group, vendors offer (possibly multiple) volume discounts that trigger if the group collectively buys in a certain quantity. We assume these are proposed or negotiated in advance, and take them to be fixed, posted prices. For ease of exposition, we assume buyers have unit demand, hence treat items as partial substitutes. Each buyer has valuations for each item and quasilinear utility.

Since vendor prices are fixed, our aim is to find an allocation of items to buyers that maximizes social welfare (i.e., sum of buyers’ utilities) given the discounts that trigger, while ensuring *stability*, or buyer “satisfaction” with the resulting allocation at the triggered prices. We consider two main variants of this problem. In the *transferable utility (TU)* model, the gains due to demand aggregation can be transferred between buyers to ensure cooperation. In the *non-transferable utility (NTU)* model, each buyer pays the (triggered) price of her allocated item. Both models have a role to play in specific business and consumer applications. We also consider various forms of knowledge

and recourse on the part of the buyer (e.g., whether they know only which discounts triggered, or have knowledge of the entire allocation and discount schedule).

Our results. Since vendor prices are fixed given some demanded quantity, the model induces a *coalitional game* among the buyers, which we refer to as a *discount matching game*. Vendor discounts introduce significant externalities in the corresponding matching problem: this leads to the emptiness of *core* of such games in certain instances, both in the TU and the NTU sense. As a consequence, we consider unilateral deviations from the matching, and focus on the weaker notion of *Nash stability* under several different epistemic assumptions.

We focus first (and primarily) on TU games. We establish that stable matchings (under all epistemic assumptions) not only exist, but that they maximize social welfare. Moreover, they can be realized using transfers only between buyers that are matched to the same vendor. We then consider computation of social welfare maximizing matchings: we show that the corresponding decision problem is NP-complete, but that, given a (fixed) set of discount prices, computing an optimal allocation can be done in polynomial time. As a result, a mixed integer programming (MIP) model of the problem can be formulated in which binary matching variables can be relaxed (as is typical in matching/assignment problems [Roth et al. 1993]), leaving a MIP whose only integer variables represent the triggering of specific discount thresholds (which, in practice, are relatively few). Experiments demonstrate the efficacy of the formulation.

We then consider the NTU discount matching game, and show stable matchings exist. Finally, we consider *qualitative discount matching games*, a variant in which buyers do not specify valuations for items, but simply *rank* the deals offered (where a *deal* is any item and one of its discounted prices). This model is especially appealing in consumer domains, where buyers may be unable to articulate precise valuations for items, but can easily compare any two items at specific prices. As long as the rankings are *rationalizable* (i.e., correspond to quasi-linear preferences under some latent valuation), again stable matchings are guaranteed to exist.

We do not address incentive issues with respect to reporting of buyer preferences. This is an important part of the design of such markets, but one we leave to future research. Truthful reporting of valuations is commonly assumed in work on procurement and inventory management (see below), where parties interact repeatedly. Similarly, we assume that sellers simply post (base and discounted) prices without regard to strategic interaction with buyers. While interactions *between* sellers w.r.t. strategic price-setting is also of interest, the way in which “between-seller” equilibrium prices and discount schedules are set does not impact group buying decisions.

Related work. Assignment games and matching markets have a rich history, and the literature is rife with connections between various forms of (individual and coalitional) stability, competitive equilibrium prices, etc. [Shapley and Shubik 1971; Gale and Shapley 1962; Demange et al. 1986]. While a general discount market model would consider strategic behavior on the part of both buyers and sellers, we take seller prices as given and focus on the one-sided problem that results by considering only the strategic behavior of buyers. Of special relevance is work on assignment models, auctions, and procurement optimization that deals explicitly with quantity discounts, buyer/bidder cooperation, and externalities in assignments.

Within the context of auctions, Kothari et al. [2005] consider multi-unit (reverse) auctions with discount tiers, and use the VCG mechanism, but consider only a single buyer with no preferences over sellers.¹ Conversely, Matsuo et al. [2005] model the

¹They also consider forward auctions with decreasing marginal utilities (see also [Lehmann et al. 2006]).

problem of a single seller offering multiple items, each with discount schedules. Buyers with combinatorial preferences bid for items, and allocations/prices are set using VCG; unlike our model, the discounts are not “posted prices” in the usual sense, but are merely used as reserve prices. While the mechanism and assumptions are quite different, and computation is not considered, their motivations are similar to ours. Leyton-Brown and Shoham [2000] study bidding clubs which collude in auction mechanisms to lower prices, and devise payment schemes that induce participation.

For an overview of the literature on group buying, see [Anand and Aron 2003; Chen and Roma 2010]. Discount schedules have received considerable attention in the operations research literature on procurement. Some work considers procurement optimization in the face of discounts, e.g., Goossens et al. [2007] deal with the problem of optimizing the procurement of multiple goods by a single buyer faced with suppliers offering total quantity discounts (which they show to be NP-hard and inapproximable). Most closely related to our problem is work on *cooperative procurement* and *inventory pooling* which models the problem of retailers coordinating their purchases from multiple suppliers. While this can be used to generate volume discounts, another motivation is to mitigate risk and lower restocking and holding costs. Of course, all buyers benefit from this cooperation, and the gains must be shared, leading to a cooperative game. van den Heuvel et al. [2007] develop a model of this problem, and show that the core of the resulting game is non-empty, but do not consider discounts or buyer preferences for different suppliers (see [Drechsel and Kimms 2010] for a procedure for computing the core of this game). Chen [2009] considers a multistage model of inventory pooling that incorporates volume discounts. The buying group places orders with suppliers—prior to buyers receiving demand signals for their products, much like news-vendor models—who immediately ship items to warehouses for holding. Once the demand is known, buyers draw on the stock from various warehouses. While buyers have no preferences over suppliers, transportation costs from each warehouse differ, so they do have preferences over which warehouses fulfill their demands. Chen shows that the core of the resulting game is non-empty when: supplier prices are linear (i.e., there are no discounts); or there are discounts, but all buyers draw from a single warehouse. Since buyer “preferences” are only for warehouses, his core existence result can be interpreted in our setting (conceptually at least) as applying when quantity discounts are offered, or when buyers have preferences over suppliers. Our results show that when *both factors* are considered, the core may be empty.

Recently, the “Groupon phenomenon” has attracted academic attention, leading to several interesting investigations of its economic value, but much of this comes from the perspective of retailers and their strategies, potential gains, etc. (see, e.g., [Edelman et al. 2011]), and sets aside the possibility of deal aggregation and optimal deal assignment to a buying group. Finally, considerable attention has been paid to matching markets with externalities (see, e.g., [Sasaki and Toda 1996; Echenique and Yenmez 2007; Hafalir 2008; Kominers 2010; Bodine-Baron et al. 2011]). As in our case, externalities cause difficulties for stability in many of these models.

2. BACKGROUND

We focus on the coalitional game induced among buyers, so we describe key concepts from cooperative game theory (see, e.g., [Myerson 1991]). Let N be a finite set of n agents. A *coalitional game with transferable utility (TU-game)* is specified by a *characteristic (or value) function* $v : 2^N \rightarrow \mathbb{R}$, which

defines the *value* $v(C)$ of each coalition $C \subseteq N$. Intuitively, $v(C)$ is the payoff members of C jointly receive by cooperating effectively: only the total payoff matters since it can be transferred freely among the members of C . A *payoff vector* $t = (t_1, \dots, t_n)$ assigns a payoff to each $i \in N$.

Definition 2.1. The *core* of a TU-game is the set of payoff vectors t s.t. $\sum_{i \in N} t_i = v(N)$ and $\sum_{i \in C} t_i \geq v(C)$ for all $C \subseteq N$.

A payoff vector in the core of a TU-game is sufficient to ensure that no coalition of agents deviates from some prescribed course of action.

A *coalitional game with non-transferable utility (NTU-game)* is given by a value function v that assigns to each $C \subseteq N$ a set of feasible (local) payoff vectors $v(C) \subseteq \mathbb{R}^{|C|}$. Intuitively, these are the possible vectors of payoffs that agents in C could (individually) receive if they act cooperatively. These payoffs cannot be transferred.

Definition 2.2. The *core* of an NTU-game is the set of payoff vectors $t \in v(N)$ s.t. there is no $t' \in v(C)$, for any $C \subseteq N$, where $t'_i > t_i$ for some $i \in C$, and $t'_i \geq t_i$ for each $i \in C$.

3. DISCOUNT MATCHING AND SOLUTION CONCEPTS

A *discount matching market* involves a collection of items with sufficient similarity that many buyers consider them to be partial substitutes. The elements of the market are as follows. We have a *buying group* of n buyers, $i \in N$, who have agreed to coordinate their purchases in a particular item category. We have m vendors, $j \in M$, each offering a single, differentiated item within the category. Vendor quantities are assumed to be sufficient to meet the needs of the buying group (i.e., unlimited for practical purposes), and each vendor j posts a *discount schedule* consisting of two vectors: a vector of D positive *discount thresholds* $[\tau_j^1, \tau_j^2, \dots, \tau_j^D]$, with $\tau_j^d < \tau_j^{d+1}$ for all $d < D$; and a vector of $D + 1$ positive *prices* $[p_j^0, p_j^1, \dots, p_j^D]$, with $p_j^d > p_j^{d+1}$ for all $d < D$. The interpretation is straightforward: if the total purchase volume committed by a buying group to that vendor is at least τ_j^d , but less than τ_j^{d+1} , the vendor sells the item for a unit price of p_j^d . Define $\tau_j^0 = 0$, which gives a base (undiscounted) price of p_j^0 .² Each buyer $i \in N$ has her own preferences over vendors/items given by a utility function or *valuation* $v_i : M \mapsto \mathbb{R}$. Here $v_i(j) = v_{ij}$ denotes i 's valuation for vendor j 's product. Buyers have quasi-linear payoffs: should i obtain j for price p , her net utility or *payoff* is $\pi_i = v_{ij} - p$. We focus on the case of *unit demand*: each buyer wants at most one item.³ We assume a *dummy item*, with price and valuation 0 for every buyer, which allows any buyer to remain “unmatched.”

Formally, a *discount matching market* takes the form $DM = (N, M, (\delta_j)_{j \in M}, (v_i)_{i \in N})$ with: N , the set of buyers; M , the set of vendors/items; a set of discount schedules, $\delta_j = (\tau_j, p_j)$ for each vendor j ; and a set of valuation functions, v_i for each buyer i .

Given the preferences of the buyers and the (discounted) prices posted by the vendors, our aim is to find an assignment or *matching* of buyers to vendors that satisfies some objective on the part of the buying group. Formally, a matching is simply a mapping $\mu : N \mapsto M$, with $\mu(i) = j$ indicating that buyer i has been assigned to purchase item j . Let $n(\mu, j) = |\mu^{-1}(j)|$ denote the number of buyers matched to vendor j under μ , and $p_j(q)$ the price of j that triggers if quantity q is demanded. We denote by $p_j(\mu) = p_j(n(\mu, j))$ the discounted price for item j triggered by matching μ , and $p(\mu)$ the corresponding price vector. The *payoff vector* for matching μ (ignoring possible transfer) is $\pi(\mu)$, where i 's payoff under μ is $\pi(\mu)_i = v_i(\mu(i)) - p_{\mu(i)}(\mu)$.

²The one to one correspondence between vendors and items, and fixed number of discount thresholds D across vendors is for ease of exposition only.

³This is for ease of exposition. The extension to multi-unit demand is trivial if demands are homogeneous, and the model requires only modest modifications if demand is heterogeneous. The extension to combinatorial preferences is more involved and is left as future work.

X: A:4, B:0	A: 0:3; 3:2; 8:1	X: A:5, B:4, C:1	A: 0:3; 2:1
Y: A:4, B:0	B: 0:1	Y: B:5, C:4, A:1	B: 0:3; 2:1
Z: A:2, B:2		Z: C:5, A:4, B:1	C: 0:3; 2:1
Buyers (Valuations)	Vendors (Discounts)	Buyers (Valuations)	Vendors (Discounts)
(a)		(b)	

Fig. 1. Two discount matching markets. Buyers X , Y and Z are shown with their valuations for specific vendors ($S:v$ indicates value v for vendor S). Vendors A , B and C are shown with the discount schedules (each $q:p$ indicates a threshold quantity q and associated unit price p).

Given matching μ , the buying group is responsible for paying the total price incurred at the demanded quantities. The responsibility can be passed on directly to the buyers—i.e., if $\mu(i) = j$ then i pays $p_j(\mu)$ —or the prices of certain buyers may be subsidized by increasing the payments of others. The latter case requires that utility be transferred between buyers in the form of payments, while the former presumes nontransferable utility. We elaborate on each of these models.

3.1. Transferable Utility

The *transferable utility (TU)* assumption allows side payments between buyers to induce cooperation. Some buyers may subsidize the prices paid by others if this can profitably trigger new discounts. For example, in Fig. 1(a), buyers X and Y are only willing to purchase item A ; and Z will only buy B (at the base price, with payoff of 1), even if the first discount triggers for A . This means X and Y pay the base price for A (payoff 1 each). But they would benefit if Z bought A : this would trigger a discount, reducing their costs by 1 each. If they agreed to subsidize Z 's purchase by each offering a side payment of some amount $0.5 < s < 1.0$, all three would benefit: Z 's payoff is now $2s > 1$, and X and Y have payoff $2 - s > 1$. Such side payments make sense for industrial buying consortia, where the demands of the buyers are aggregated by a central coordinator who finds deals on their behalf.⁴ This may also be appropriate in certain consumer buying groups (e.g., where there is sufficient familiarity so that buyers are willing to pay more than “posted” (discounted) prices). A key distinction between our model of demand aggregation/group buying and those considered more traditionally is that we explicitly allow buyers to express preferences over different product offerings.

The most natural way to match buyers and vendors under TU is to find a matching that maximizes total buyer payoff, and then determine payments that render this matching *stable* in the sense that that no buyer wants to deviate from the matching. We first consider stability w.r.t. coalitional deviations, defining *core matchings* in the discount matching market. We make use of the following concepts. The *sub-problem* $DM(S)$ of DM is the matching problem induced by restricting DM to the set of buyers $S \subseteq N$ (with discounts schedules and preferences of $i \in S$ unchanged). The *social welfare* of a matching μ in DM is the surplus realized by the buyers:

$$SW(\mu; DM) = \sum_i \pi(\mu)_i = \sum_i v_i(\mu(i)) - \sum_j n(\mu, j)p_j(\mu).$$

$MSW(DM)$ is the social welfare of the matching $\mu^*(DM)$ that maximizes total buyer surplus. A *transfer vector* $t \in \mathbb{R}^N$ for DM is any vector of payoffs, where t_i is interpreted as a (possibly negative) payment to buyer i for participating in the buying group. Transfer vector t is *feasible for matching* μ iff $\sum_i t_i \leq SW(DM)$. The *net payoff*

⁴In general, such consortia will often *negotiate* volume prices on behalf of the group. This introduces a strategic element to price setting that we deliberately ignore in this paper (see earlier comments).

vector for matching μ with transfers t is $\rho(\mu, t)$, where i 's net payoff under (μ, t) is:

$$\rho(\mu, t)_i = \pi(\mu)_i + t_i = v_i(\mu(i)) - p_j(\mu) + t_i.$$

Definition 3.1. Let $DM = (N, M, (\delta_j)_{j \in M}, (v_i)_{i \in N})$ be a discount matching market with transferable utility, μ some matching, and $t \in \mathbb{R}^N$ a feasible transfer for μ . Then (μ, t) is in the *core of DM* iff there is no coalition $S \subseteq N$ such that the submarket $DM(S)$ admits matching μ^S and feasible transfer t^S that makes all members of S better off. In other words, for any S and matching μ^S , there is no feasible transfer $t^S \in \mathbb{R}^{|S|}$ for μ^S such that $\rho(\mu^S, t^S)_i > \rho(\mu, t)_i$ for all $i \in S$.

This definition of a core allocation (matching-transfer pair) is strong in that it requires that any defecting coalition be able to obtain better value without considering the actions of buyers outside that coalition. However, the following observation regarding the potential emptiness of the core applies equally to weaker concepts (e.g., that anticipate discounts triggered by buyers outside the coalition). The discount matching game is clearly superadditive, so one can define the core using a characteristic function representation and focus purely on division of the total payoff $MSW(DM)$ under the socially optimal matching. Instead, we defined it explicitly in terms of matchings so that parallels to NTU markets and Nash stability are more evident below.

Unfortunately, the core is not a reliable solution concept under TU:

OBSERVATION 3.2. *The core of a discount matching market with TU can be empty.*

Fig. 1(b) shows a DM with an empty core. There are three surplus maximizing matchings, one corresponding to each item: the matching for an item assigns the pair of buyers with valuations $(5, 4)$ to that item (e.g., X and Y to B), and the third buyer to their most preferred undiscounted item (e.g., Z to C). Each has $SW(\mu) = 9 = MSW(DM)$. However, none of these is in the core. Consider matching $\mu^1(X) = \mu^1(Y) = B$, $\mu^1(Z) = C$. Any t that renders μ^1 stable must split the surplus of 7 generated by X and Y among them so that at least one of them attains payoff of no more than 3.5. If X 's payoff is less than 3.5, then $DM(\{X, Z\})$ renders μ^1 unstable, since X, Z can agree to purchase A , and share the generated surplus of 7 in a way that makes both better off (e.g., X takes 4 and Z takes 3). If Y 's payoff is less than 3.5, then $DM(\{Y, Z\})$ renders μ^1 unstable, since Y, Z would agree to purchase C . As a consequence we will instead consider matchings that are unilaterally or *Nash stable* below.

3.2. Non-transferable Utility

Transfers in a discount matching market are infeasible or undesirable in some settings. In consumer group buying, for example, a deal aggregator may collect a set of deals in a specific product category (e.g., spas, restaurants, etc.) with posted discount prices, and attract otherwise “unaffiliated” consumers considering these products. Consumers may be wary of paying “variable” prices that differ from posted discounts. In business settings, members of a GPO may be willing to aggregate their demand, but may be unwilling to explicitly subsidize the purchases of a competitor. Transfers can also greatly increase the opportunity for manipulation (see Sec. 8).

In a *nontransferable utility (NTU)* market, buyer i 's utility for a matching μ is exactly her payoff $\pi(\mu)_i = v_i(\mu(i)) - p_{\mu(i)}(\mu)$, so no transfers need to be considered. Core matchings in an NTU market are defined as follows:

Definition 3.3. Let $DM = (N, M, (\delta_j)_{j \in M}, (v_i)_{i \in N})$ be a discount matching market with nontransferable utility. Matching μ is in the *core of DM* iff there is no coalition $S \subseteq N$ such that submarket $DM(S)$ admits a matching μ^S that makes all members of S at least as well off, and at least one member strictly better off. In other words, for

any S and matching μ^S for $DM(S)$, we have either: $\pi(\mu^S)_i < \pi(\mu)_i$ for some $i \in S$; or $\pi(\mu^S)_i = \pi(\mu)_i$ for all $i \in S$.

As in TU games, the core is too strong a solution concept for NTU markets:

OBSERVATION 3.4. *The core of a discount matching market with non-transferable utility can be empty.*

The emptiness of the core in a TU market does not imply that the core of the corresponding NTU game is empty, since the inability to make transfers restricts the ability of certain coalitions to form. However, the example in Fig. 1(b) discussed above has no core matching even under the NTU assumption.

3.3. Nash Stability

Since core matchings may not exist in discount matching markets, we focus in the remainder of the paper on *stable matchings* in the Nash sense: a matching (together with an transfer vector if TU is assumed) is *stable* if no buyer can unilaterally achieve a higher payoff by choosing to buy a different item at the posted discounted price induced by the matching (accounting for her deviation).

If μ is a matching, let $\mu[i \leftarrow j]$ be the matching that is identical to μ with the exception that i is matched to j (if $\mu(i) = j$, then $\mu[i \leftarrow j] = \mu$). We define three notions of stability depending on what price a buyer can unilaterally demand. If buyers are able to select arbitrary vendors and benefit from the discount that results (including any additional discount they trigger), then i deviating from μ to a different vendor j gives her payoff $\pi(\mu[i \leftarrow j])_i$. If i can deviate to j , but is not permitted to trigger a new discount threshold, she must accept the current price for j as triggered under μ , realizing payoff $v_{ij} - p_j(\mu)$. Finally, if buyers deviate from their assignment, they can be “banished” from the buying group, leaving them to pursue their best outside option, which we take to be their preferred product at the base prices (i.e., when no discounts are offered). Let $\theta_i = \max_j v_{ij} - p_j^0$ be i 's *outside option (or reservation) value*.

Definition 3.5. Let $DM = (N, M, (\delta_j)_{j \in M}, (v_i)_{i \in N})$ be a discount matching market with non-transferable utility. Matching μ is *strongly stable* iff $\pi(\mu[i \leftarrow j])_i \leq \pi(\mu)_i$ for any buyer-vendor pair (i, j) . Matching μ is *myopically stable* iff $v_{ij} - p_j(\mu) \leq \pi(\mu)_i$ for any (i, j) . Matching μ is *weakly stable* iff $\theta_i \leq \pi(\mu)_i$ for any (i, j) .

Strong stability is the most robust of these solution concepts. It ensures coherence of the buying group in situations where buyers are unknown to one another, unable to easily negotiate with each other, or have little knowledge of each other's preferences (making coalitional deviation unlikely); but they do know item purchase volumes and which discounts triggered, hence the prices they could demand in the presence of other buyers. Myopic stability is suitable when triggered discounts are public, but buyers are unaware of either purchase volumes or discount thresholds (common in some consumer settings). Finally, weak stability, the solution concept with least applicability, might reflect a buying consortium that can “dictate” buyer behavior (with opting out of the group the only recourse available to a buyer). Clearly, strong stability implies myopic stability, which in turn implies weak stability.

Things are more subtle in the TU model. The possibility of transfers means that a buyer i , matched to item j , could “agree” to help subsidize the price paid by another buyer (say k) for j , thereby triggering a discount that lowers i 's price. One possible deviation by i involves reneging on the subsidy agreement, by dropping her matched item j , then attempting to be rematched with j at its discounted price without paying her (part of the) subsidy to other buyers. No consortium would permit i to benefit this way, since without i 's subsidy, k might not agree to purchase item j . Let (μ, t) be a

matching-transfer pair and $s(j; \mu, t)$ be the set of *subsidized* buyers k matched to item j that receive a positive transfer (i.e., $\mu(k) = j$ and $t_k > 0$). We define i 's payoff in the new matching $\mu[i \leftarrow j]$ as follows: (a) if $\mu(i) \neq j$, then $\rho(\mu[i \leftarrow j], t[i \leftarrow j])_i = \pi(\mu[i \leftarrow j])_i$; (b) if $\mu(i) = j$, then $\rho(\mu[i \leftarrow j], t[i \leftarrow j])_i = v_i(\mu(i)) - p_j(n(\mu, j) - |s(j; \mu, t)|)$. In other words, i 's payoff is defined by assuming that any *subsidized* buyer matched to the same item j loses incentive to participate.⁵ We need not specify the “updated” transfer vector $t[i \leftarrow j]$; we require only that $t_i = 0$ after defection.

Definition 3.6. Let $DM = (N, M, (\delta_j)_{j \in M}, (v_i)_{i \in N})$ be a discount matching market with TU. Matching-transfer pair (μ, t) is *strongly stable* iff $\rho(\mu[i \leftarrow j], t[i \leftarrow j])_i \leq \rho(\mu, t)_i$ for any buyer-vendor pair (i, j) . The pair (μ, t) is *myopically stable* iff $v_{ij} - p_j(\mu) \leq \rho(\mu, t)_i$ for any (i, j) . Matching μ is *weakly stable* iff $\theta_i \leq \rho(\mu, t)_i$ for any (i, j) .

The motivation and application for these solution concepts are similar to those above. We now turn to the existence of stable matchings in both the TU and NTU models.

4. STABILITY AND EFFICIENCY UNDER TU

We first investigate stable matchings in the TU model, showing the existence of buyer-welfare maximizing strongly stable allocations (or matching-transfer pairs), despite the externalities created by volume discounts. This can be viewed as determining a form of personalized market clearing prices with the “consortium” or GPO acting as an intermediary. We then turn to the computation of welfare maximizing matchings and the required transfers.

4.1. Existence of Stable Matchings

Since the TU matching market is superadditive, the grand coalition maximizes social welfare. While the core is empty, we will show that any social welfare maximizing matching is in fact strongly stable by deriving appropriate side payments.

Given a market DM under the TU assumption, let μ be a *social-welfare maximizing (SWM)* matching (i.e., μ maximizes buyer surplus). Without side payments, μ may not be (even weakly) stable: Fig. 1(a) is an example where the SWM μ , which assigns all buyers to item A , cannot be sustained (even weakly) without a side payment to Z . Obviously, welfare maximization may require assigning a buyer to an item for which she has a negative payoff to obtain a larger discount benefitting other buyers.

However, stability can be realized in any SWM μ if transfers are allowed. We refer to any buyer i who is matched in μ to the item that maximizes her payoff at current prices and demanded quantities as an *ISM (individual surplus maximizing) buyer*; i.e., $\pi(\mu)_i = \max_j \pi(\mu[i \leftarrow j])_i$. Conversely, a buyer matched to an item that provides a lower payoff than some other item is a *non-ISM buyer*. Stability requires that some of the surplus generated by ISM buyers be transferred to non-ISM buyers. The transfer to any non-ISM k must be being sufficient to prevent k from preferring a product different from $\mu(k)$, and the transfer extracted from ISM buyer i must be small enough to maintain $\mu(i)$ as her preferred product. This can be seen as a form of personalized pricing, where ISM buyer i pays a price $p_{\mu(i)} + c_i$ for $\mu(i)$, where c_i is her *contribution* to the subsidies (i.e., a negative transfer t_i); and where any non-ISM k pays a price $p_{\mu(k)} - s_k$ for $\mu(k)$, where s_k is her *subsidy* (or positive transfer t_k).

First we show that such transfers exist, and furthermore, that the transfers can be limited so that ISM buyers matched to a specific item j subsidize only non-ISM buyers matched to the same item. This ensures that the allocation is budget balanced in a

⁵We could also consider defections of non-subsidized buyers matched to j , but these only lower i 's payoff further. Since stability holds w.r.t. “subsidized” defections (see below), it holds under this stronger notion too.

strong sense: the total paid by buyers matched to vendor j is exactly that required by j under μ .⁶ We introduce some notation: let $\pi_{-j}^*(i, \mu) = \max_{j' \neq j} \pi(\mu[i \leftarrow j'])_i$ denote i 's payoff if matched to the product $j' \neq j$ that maximizes her payoff at current prices and demanded quantities.

THEOREM 4.1. *Given any SWM matching μ , there exists a transfer vector t such that (μ, t) is strongly stable. Furthermore, the matching is per-vendor budget balanced: for any vendor j , $\sum \{t_i : i \in \mu^{-1}(j)\} = 0$.*

PROOF. We provide a proof sketch, showing that for each vendor j , balanced transfers can be made among only those buyers matched to j . Some notation:

- Let $N_0 = \mu^{-1}(j)$ be the set of buyers matched to item j . Let $N_0^+ = \{i \in N_0 : \text{ISM}(i)\}$ be the buyers in N_0 satisfying ISM. Let $N_0^- = N_0 \setminus N_0^+$ be the buyers not satisfying ISM. Let $p^0 = p_j(\mu)$ be the price of j in μ .
- For any $k \geq 1$, let $N_k = N_{k-1}^+$ and $p^k = p_j(|N_k|)$. Then define $N_k^+ = \{i \in N_k : \text{ISM}(i, j, p^k)\}$. These are the buyers in N_k^+ whose who prefer to remain matched to j even if its is increased to p^k (the price induced if the number of matched buyers is reduced to $|N_k|$).
- For any $k \geq 1$, define $N_k^- = N_k \setminus N_k^+$ to be those buyers who would prefer to be matched to a new item if the price of j is increased to p^k .

Intuitively, the set N_0^- are those buyers unwilling to be matched to j at its current price given current demanded quantities of other items, and require some subsidy; N_0^+ are those who prefer j and should be willing to subsidize buyers in N_0^- . Notice since μ is SWM, N_0^+ is empty only if N_0 is empty (otherwise welfare would increase by switching all buyers away from j in μ). However, not all buyers in $N_0^+ = N_1$ are willing to provide equal subsidies. N_1 can be further broken down into: those buyers N_1^+ who would prefer j even if the price were increased to p^1 (which would happen if buyers in N_0^- switched to their most preferred product); and those N_1^- who have a better option than j at the increased price p^1 . Those in N_1^+ would be willing to pay at least $p^1 - p^0$ to see the price of j remain at p^0 . Any $i \in N_1^-$ is willing to pay at most $v_{ij} - p^0 - \pi_{-j}^*(i, \mu)$.

This line of argument can be extended. Let K be the largest integer such that $N_K \neq N_{K-1}$: such a K must exist because at least one buyer must be removed (starting from N_0) at each iteration until either some N_k is empty, or the price p^k fails to increase. We then have must have: (a) For any buyer $i \in N_K$, $v_{ij} - p^{K-1} \geq \pi_{-j}^*(i, \mu)$ (note: N_K may be empty). Any such buyer is thus willing to pay (at least) up to p^{K-1} for j ; i.e., i will continue to maximize her surplus with product j even if she pays a premium of up to $p^{K-1} - p^0$ above the nominal price p^0 under μ . (b) For any buyer in $i \in N_k^-$ ($1 \leq k \leq K-1$), we have $v_{ij} - p^0 > v_{ij} - p^{k-1} \geq \pi_{-j}^*(i, \mu) > v_{ij} - p^k$. Thus any such buyer continues to maximize her surplus with product j even if she pays a premium of up to $v_{ij} - p^0 - \pi_{-j}^*(i, \mu) \geq p^{k-1} - p^0$ above the nominal price $p^0 = p_j$.

The total contributions available from each such buyer, while still maintaining ISM if they remain matched to j , are:

$$TC(j) = \sum_{i \in N_K} [p^{K-1} - p^0] + \sum_{1 \leq k < K} \sum_{i \in N_k^-} [v_{ij} - p^0 - \pi_{-j}^*(i, \mu)].$$

⁶The ability to ensure that the net transfer is zero among those buyers matched to a specific vendor j may be important in settings where a consortium *coordinates* purchases, but the transactions are *executed* directly by buyers and vendors (at computed personalized prices). Furthermore, the possibility that a vendor's customers are subsidizing purchases from one of its competitors may be viewed as undesirable.

The only buyers matched to j that violate ISM are those in N_0^- . For each buyer $i \in N_0^-$, a subsidy of $\pi_{-j}^*(i, \mu) - v_{ij} - p^0$ will suffice to restore ISM while maintaining the match to j . Thus the total subsidy required is:

$$TS(j) = \sum_{i \in N_0^-} [\pi_{-j}^*(i, \mu) - v_{ij} - p^0].$$

Now suppose, by way of contradiction, that $TC(j) < TS(j)$. We show the social welfare of μ can be strictly improved as follows: assign each buyer in any set N_k^- for any $0 \leq k \leq K-1$ to its preferred outside option, i.e., the j' maximizing $\pi_{-j'}^*(i, \mu)$; leave all other buyers (including those in N_K) assigned as in μ . Notice that the price of any product $j' \neq j$ in the new matching is less than or equal to its price $p_{j'}$ in μ (since volume discounts can only improve for products other than j).

First, consider buyers that maintain or improve their payoff in the new matching. All buyers in N_0^- strictly improve their payoff by moving from non-ISM item j to one that has a greater payoff. Thus the increase in payoff is at least $TS(j)$ (it may be greater if new discounts trigger for some products). Note that all buyers not matched to j in μ have equal or greater payoff in the new matching (greater if new discounts trigger for their matched products).

Now consider all buyers that could possibly lose payoff. All buyers in N_k^- , $1 \leq k \leq K-1$, may lose payoff, but this loss is at most $v_{ij} - p^0 - \pi_{-j}^*(i, \mu)$ (it may be less if new volume discounts trigger). The loss in payoff to any buyer in N_K is exactly $p^{K-1} - p^0$ in this new matching. Thus the total loss in payoff is no more than $TC(j)$.

By assumption, $TC(j) < TS(j)$, implying that gain of payoff-improving buyers in the new matching outweighs the loss of payoff-losing buyers. This contradicts the fact that μ is SWM. Hence transfers can be determined among buyers matched to j that induce ISM. (The construction above shows exactly what maximal contributions can be extracted, and the minimal subsidies needed for each buyer.) \square

4.2. Maximizing Social Welfare

We now consider computation of social welfare maximizing matchings. This problem can be formulated as a mixed integer program (MIP) and solved using standard solvers. The formulation is straightforward: we describe key variables and constraints. *Assignment variables (binary)* μ_{ij} ($i \in N, j \in M$): Is buyer i matched to vendor j . Constraints:

$$\sum_{j \in M} \mu_{ij} \leq 1, \quad \forall i \in N. \quad (1)$$

Count variables N_j ($j \in M$): Number of buyers matched to j . Constraints:

$$N_j = \sum_{i \in N} \mu_{ij}, \quad \forall j \in M. \quad (2)$$

Threshold variables (binary) I_j^d ($d \leq D, j \in M$): Is j 's d th discount threshold met. Only one threshold indicator is non-zero, ensuring the correct price is selected during optimization. Constraints:

$$\tau_j^d I_j^d \leq N_j, \quad \forall d \leq D, j \in M \quad \text{and} \quad \sum_{d \leq D} I_j^d = 1, \quad \forall j \in M. \quad (3)$$

The objective of the optimization can be written in quadratic form:

$$\max \sum_{i \in N, j \in M} v_{ij} \mu_{ij} - \sum_{d \leq D, j \in M} p_j^d I_j^d N_j, \quad (4)$$

where the first sum indicates total buyer valuation under μ and the second total cost. The quadratic term $I_j^d N_j$ is linearized in a standard way, with auxiliary variable Z_j^d introduced to represent this product:

$$Z_j^d \leq I_j^d U \quad \forall d \leq D, j \in M \quad \text{and} \quad \sum_{d \leq D} Z_j^d = N_j \quad \forall j \in M, \quad (5)$$

where U is an upper bound on N_j (e.g., n , the number of buyers). The MIP is then:

$$\begin{aligned} \max \quad & \sum_{i,j} v_{ij} \mu_{ij} - \sum_{d,j} p_j^d Z_j^d \\ \text{s.t.} \quad & (1), (2), (3), \text{ and } (5). \end{aligned}$$

The computational complexity of MIPs grows exponentially with the number of integer variables, and the formulation above requires integer (0-1) matching variables μ_{ij} for each buyer-vendor pair. However, as is common in matching problems, one can in fact relax the matching variables to be continuous in $[0, 1]$. This leaves as the only 0-1 indicator variables those I_j^d indicating which discount thresholds trigger for each product. Suppose we fix the values of these variables (i.e., fix the discounts that are triggered for each product). We show that the relaxation—which, for fixed thresholds, is a linear program—must give rise to optimal solutions in which all variables μ_{ij} are integral. The proof formulates the optimization as a min-cost max-flow problem.⁷

THEOREM 4.2. *Let τ_1, \dots, τ_m be threshold values for vendors in M such that $\tau_1 + \dots + \tau_m \leq n$; and let p_1, \dots, p_m be the corresponding prices (i.e., where $p_j = p_j(\tau_j)$, for $j \in M$). We can find a matching that meets these thresholds, and maximizes social welfare at the threshold prices, $\sum_{i \in N} v_i(\mu(i)) - p_{\mu(i)}$, in polynomial time.*

Thm. 4.2 implies that the solution to the relaxed MIP will either assign 0-1 values to all μ_{ij} variables; or, if not, then the optimal threshold indicators I_j^d can be used to construct a matching using a (polytime) min-cost max-flow algorithm.

In the relaxed MIP, the only 0-1 indicator variables are the I_j^d variables denoting which discount thresholds trigger for each vendor. The number of such variables, mD , will be quite small in most discount matching applications, making the solution of this MIP computationally manageable. Despite this, a simple reduction from Knapsack shows that computing a SWM matching is theoretically intractable. Let DISCOUNT-MATCH be the following decision problem: Given $x \geq 0$, is there a matching μ whose total buyer welfare, $SW(\mu) = \sum_i u_i(\mu(i)) - \sum_j n(\mu, j) p_j(n(\mu, j))$, is at least x ?

THEOREM 4.3. *DISCOUNT-MATCH is NP-complete.*

4.3. Computing Payments

The results above show that transfers exist that support SWM matchings (i.e., render them strongly stable). Generally, a buying group will desire transfers that satisfy certain objectives. One reasonable objective is to minimize the total transfer needed to ensure stability. Others include minimizing the largest single contribution from any ISM buyer, or using other forms of fair division of the surplus.

Given SWM matching μ , computing payments that simultaneously minimize both the total subsidy and the largest individual contribution is straightforward. We know precisely which subsidies are needed for each non-ISM buyer; let TS denote the total

⁷Proofs of all results not included here can be found in a longer version of the paper at <http://www.cs.toronto.edu/~cebly/papers.html>.

subsidy required. Let $e_i = \pi(\mu)_i - \pi_{-\mu(i)}^*(i, \mu)$ be the surplus that can be extracted from buyer i without destabilizing μ , and $S = \{i \mid e_i > 0\}$ be the set of buyers that are *strictly* ISM. We first order buyers in S , b_1, \dots, b_k , sorted in ascending order by e_i (define $e_0 = 0$). To compute transfers, we consider each b_i in order, and test if there is a contribution c , $0 < c \leq e_{b_i} - e_{b_{i-1}}$ such that extracting c units of transfer from each $b_{i'}$, $i' \geq i$, meets the subsidy target of TS , i.e., with a total contribution of $(n - i + 1)c$. If so, we stop; otherwise, we increase the contributions of each $b_{i'}$, $i' \geq i$, by $e_{b_i} - e_{b_{i-1}}$; and continue to buyer b_{i+1} . This process terminates with a set of transfers that minimize the maximum subsidy paid by any single buyer (as well as the total subsidy).

5. NTU STABILITY

We now turn to the NTU model, demonstrating the existence of strongly stable matchings without transfers, and discussing the computation of such matchings.

THEOREM 5.1. *Let $DM = (N, M, (\delta_j)_{j \in M}, (u_i)_{i \in N})$ be a discount matching market with non-transferable utility. There exists a strongly stable matching for DM .*

PROOF. Let μ be a matching and i a buyer; let $\mu \uparrow i \in \arg \max_j \pi(\mu[i \leftarrow j])_i$ be a matching the results from i 's *best response* to μ . If $\mu(i) \in \arg \max_j \pi(\mu[i \leftarrow j])_i$, we require that $\mu \uparrow i = \mu$ (i.e., if i is ISM in μ , it will not change to some other item that is equally good). Let σ be an arbitrary ordering of buyers: a σ -*best reply path* is any sequence of matchings μ_1, \dots, μ_K in which each buyer makes a best response, in turn, to the prior matching, i.e., $\mu_2 = \mu_1 \uparrow \sigma_1$, $\mu_3 = \mu_2 \uparrow \sigma_2$, and more generally $\mu_{k+1} = \mu_k \uparrow \sigma_{k \bmod N}$. We will show that no best response path can be cyclic (apart from trivial cycles in which $\mu_{k+1} = \mu_k$). Since the number of matchings is finite, this suffices to show the existence of a strongly stable matching.

Suppose by way of contradiction that there is a (nontrivial) cyclic best response path $\mu_1, \mu_2, \dots, \mu_K$, with $K \geq 3$ and $\mu_1 = \mu_K$ (assume all other duplicate matchings have been deleted.) Each move from μ_k to μ_{k+1} corresponds to some buyer i_k switching from vendor j_k^- at its existing price p_k^- to a different vendor j_k^+ at (existing or newly triggered) price p_k^+ . Since each move is a best response, we must have

$$v_{i_k}(j_k^+) - p_k^+ > v_{i_k}(j_k^-) - p_k^-$$

(the inequality must be strict). Summing these such moves and rearranging, we obtain:

$$\sum_{k < K} v_{i_k}(j_k^+) - v_{i_k}(j_k^-) > \sum_{k < K} p_k^+ - p_k^-. \quad (6)$$

Since the best response path is cyclic, any buyer i active in the path first leaves its initially matched vendor at some point k , joins/leaves other vendors, and finally rejoins its initially matched vendor at some point $k' > k$. So for any fixed buyer i and item j any positive occurrence of v_{ij} (when i joins j) on the LHS of Eq. 6 is matched by a negative occurrence (when i leaves j). Hence the LHS of Eq. 6 sums to zero.

Since each item j is matched to the same number of buyers in μ_1 and in μ_K , the sequence of prices on the RHS corresponding to any fixed product j must be balanced by similar reasoning (each price corresponding to some joining buyer is matched by the price corresponding to some leaving buyer). Hence the RHS of Eq. 6 also sums to zero, contradicting the requirement that the inequality be strict. \square

With the existence of strongly stable matchings established, we now turn attention to their computation. We first consider computing SWM matchings subject to strong stability constraints. We can adapt the MIP from Sec. 4.2, but because transfers are not permitted, we modify it by introducing the following indicator variables and strong stability constraints.

Threshold variables (binary) \tilde{I}_j^d ($d \leq D, j \in M$): Is j 's d th discount threshold triggered if one additional buyer is assigned to j . Constraints:

$$\tau_j^d \tilde{I}_j^d \leq N_j + 1, \quad \forall j \in M, d \leq D, \quad \text{and} \quad \sum_{d \leq D} \tilde{I}_j^d = 1, \quad \forall j \in M. \quad (7)$$

$$(n+2)\tilde{I}_j^D + \sum_{d < D} \tau_j^{d+1} \tilde{I}_j^d \geq N_j + 2 \quad \forall j \in M. \quad (8)$$

Stability constraints: Ensure that i 's payoff is no higher when assigned to a different vendor:

$$\sum_{j \in M} \sum_{d \leq D} \mu_{ijd} (v_{ij} - p_j^d) \geq v_{ij'} - \sum_{d \leq D} \tilde{I}_j^d p_j^d - L\mu_{ij'}, \quad \forall i \in N, j \in M, \quad (9)$$

where $\mu_{ijd} \in \{0, 1\}$, $\sum_{d \leq D} \mu_{ijd} = \mu_{ij}$, L is an upper bound on valuations, and the term $L\mu_{ij'}$ ensures the constraints are trivially satisfied if $j = j'$.

Note that we can no longer relax the variables μ_{ij} and obtain an integral optimal solution (We can relax the new indicators μ_{ijd} .) This makes the problem *prima facie* more difficult to solve than SWM matching with transferable utility. We are currently investigating its complexity. Another important question is whether iterative algorithms in the style of Gale-Shapley [1962] can be developed for this problem (either for SWM matchings or arbitrary stable matchings). Notice that the best response ‘‘algorithm’’ used in the proof of Thm. 5.1 can be used directly to compute stable matchings (not necessarily SWM), though it is unlikely to be practical.

6. QUALITATIVE PREFERENCES

The models examined so far have assumed that buyers specify valuations for items, allowing the computation of the precise payoff to any buyer under a given matching, as well as providing the ability to determine personalized prices/transfers when TU is assumed. In many applications, such as corporate demand aggregation, buyers will be relatively sophisticated, able to distinguish vendor offerings clearly, and willing to specify valuations with some degree of precision. In other settings, this may not be the case. For example, in consumer group buying, it is unreasonable to expect buyers to specify precise valuations for items.

In the NTU model, there is in fact no need to have consumers specify valuations for items. We refer to any pair (j, p_j) consisting of an item and its price as a *deal*. The discount schedules provided by vendors correspond to a finite set of deals, with each item occurring in one deal for each of its possible (base or discounted) prices. To determine a stable matching in the NTU model, it is sufficient to simply have buyers *rank* the deals that have been proposed rather than specify valuations.

Definition 6.1. Let $(M, (\delta_j)_{j \in M})$ be a set of vendors and associated discount schedules. The set of *deals* is

$$L = \{(j, p_j) : j \in M, (\tau, p_j) \in \delta_j \text{ for some } \tau\}.$$

A *deal ranking* \succeq is any total preorder over L .

Each buyer i specifies a deal ranking \succeq_i indicating her relative preference for each item at that each of its discounted prices. Let \succ_i denote strict preference. We often assume that buyer rankings are consistent with some underlying valuation:

Definition 6.2. A deal ranking \succeq over L is *rationalizable* iff there is some valuation function v s.t., for any pair of deals in L , $(j, p_j) \succeq (k, p_k)$ iff $v(j) - p_j \geq v(k) - p_k$.

Allowing the specification of deal rankings can ease the burden on unsophisticated buyers considerably. For example, suppose three vendors (say, spas), $j \in \{A, B, C\}$, each offer a base price b_j and a single discounted price d_j . A buyer need only rank the six induced deals. Moreover, buyers need not rank unacceptable deals; e.g., a buyer might only rank $(A, d_A) \succ (B, d_B) \succ (B, b_B)$, thereby deeming A at its base price, and C at either price, to be unacceptable. (This can be treated in the model as ranking these items below some dummy item with price and valuation zero.)⁸

A *qualitative discount matching market* is $QDM = (N, M, (\delta_j)_{j \in M}, (\succeq_i)_{i \in N})$ where instead of valuations, buyers specify a ranking over deals. Stability is defined in an analogous way to the case of NTU markets with valuations:

Definition 6.3. Let $QDM = (N, M, (\delta_j)_{j \in M}, (\succeq_i)_{i \in N})$ be a qualitative discount matching market. A matching μ is *strongly stable* iff

$$(\mu(i), p_{\mu(i)}(\mu)) \succeq_i (\mu[i \leftarrow j](i), p_{\mu[i \leftarrow j]}(\mu[i \leftarrow j]))$$

for any buyer-vendor pair (i, j) .

It follows directly from Theorem 5.1 that stable matchings exist if the deal ranking of every buyer is rationalizable.

COROLLARY 6.4. *Let $QDM = (N, M, (\delta_j)_{j \in M}, (\succeq_i)_{i \in N})$ be a qualitative discount matching market. If each deal ranking \succeq_i is rationalizable, then there exists a strongly stable matching for QDM .*

Computation of stable matchings for qualitative markets can be accomplished using a MIP that is conceptually similar to the one developed for the NTU model, though with some technical differences, since welfare maximization is not possible without valuations. In addition, even if all deal rankings are rationalizable, the stability of a matching in this setting does not come with any guarantees regarding maximum social welfare w.r.t. the latent valuations. Development of a matching algorithm in the style of Gale-Shapley would certainly be desirable in this context.

7. EXPERIMENTS

We experiment with the discount matching model to develop insights into the value of cooperative buying (relative to a model of naive consumer deals), the relative level of subsidy needed to support social welfare maximization in the TU model, the “price of stability,” and the performance of our MIP formulation.

We generate random problem instances using a synthetic model intended to reflect utilities and discounts in a setting involving a consumer “daily deal” aggregator⁹ Roughly, the model assumes each vendor’s product has underlying attributes which determine overall quality, while vendors themselves have a *brand value*. Base prices are proportional to quality (with additive noise). The discounted price at threshold level t is a random percentage (chosen from a plausible range) of the price at the prior level $t - 1$. Thresholds are based on the “expected” number of buyers per vendor n/m , plus some small random fraction of n/m . User valuations are linear in product attributes and brand value, with weights generated uniformly at random. Valuations are scaled relative to vendor prices.

We first consider the degree of subsidy needed to support SWM matchings in the TU model. We consider problems with $m = 10$ vendors and vary the number of buyers

⁸If two deals have very similar payoff to a buyer, one might observe violations of rationalizability. Detecting such violations is straightforward, and can be handled by using approximate rationalizability, which will lead to approximately stable matchings.

⁹Details are provided in the longer version of the paper at <http://www.cs.toronto.edu/~cebly/papers.html>.




Fig. 2. Subsidies required to achieve stability (error bars of one (sample) std. dev. shown in line plots).




Fig. 3. MIP solution times as n (left) and m (center) vary: plots show 25-75 quantile range, median, and average. Right: Price (social welfare) of achieving NTU stability.

(up to $n = 1000$) and generate 10 instances of each problem size. For each instance, we first compute the SWM matching, and measure the *total subsidy* required to keep non-ISM buyers satisfied with their match as a fraction of the *total ISM surplus*, i.e., the contributions (transfers) ISM buyers would be willing to make while maintaining satisfaction with their matches. Fig. 2 (top) shows these results. We see (left) that the average contribution is around 3% of available surplus and is stable as n increases (the histogram on the right shows the distribution of these fractions over all instances). Similarly, the bottom plot/histogram of Fig. 2 show the ratio of the *average contribution* of an ISM buyer to the *maximum contribution* of any ISM buyer (when contributions are computed to minimize the maximum transfer as in Sec. 4.3). We see that the average contribution is usually very close to the maximum—on average within 99%—indicating a reasonable degree “fairness” in the provision of subsidies.

We next examine performance of our (relaxed) MIP formulation for computing SWM matchings. Fig. 3 (left and middle) shows wall-clock solution time—using CPLEX 12.2, 1GHz desktop, no tuning—for 10 vendors and a variable number of buyers, and for 200 buyers and a variable number of vendors (averaged over 10 instances each). Solution time, even for 1000 buyers and 10 vendors, is reasonable (avg. 2min). When fixing $m = 10$, the increase in solution time as n grows appears to be polynomial (as suggested by Thm. 4.2). Solution time for fixed $n = 200$ for $m \leq 25$ is also good. Since the number of vendors m will be relatively small in many settings, this is encouraging.

Fig. 3 (right) illustrates the ratio of the maximum social welfare achievable under the NTU model (i.e., when imposing stability constraints with no possibility of transfers) with the true maximum social welfare (which can be achieved under TU). This can be viewed as a *price of stability*. We see this ratio is very close to 1 (0.98 on avg.,




Fig. 4. Buyer satisfaction and social welfare approximation quality under the online matching (error bars show one (sample) std. dev.).

never below 0.92). This suggests that imposing stability under NTU does not significantly distort social welfare. Equivalently, it means that disallowing transfers in the buying model, while still insisting on stability, actually imposes a relatively small cost in terms of social welfare.

Finally, we evaluate the *value* of cooperative buying in a consumer “daily deal” scenario by comparing our centralized matching model with an “online” model in which customers arrive randomly, search for deals, and greedily accept the best deal at the current triggered discounts (i.e., choose based on the volume and choices of *prior* customers, but not anticipating future arrivals). We (a) fix the number of vendors at $m = 10$ and vary the number of buyers (up to $n = 1000$); and (b) fix $n = 500$ and vary m (up to 25). At the end of the process, we measure the quality of this self-directed “matching” in various ways. First we measure the fraction of non-ISM buyers (i.e., who would prefer a different product at the final discount prices than that they accepted upon arrival). Fig. 4 (top) shows these results, indicating that a significant fraction of the buyers—on average 14%—are unsatisfied with their “online” matched item—this varies little as n increases, but increase marginally with m . We also measure *average normalized regret (ANR)* over the non-ISM buyers: normalized regret for i is $(\pi_i^* - \pi_i)/\pi_i^*$, where π_i is her actual payoff and π_i^* is the payoff for her preferred vendor. Fig. 4 (middle) shows that ANR is about 11%, suggesting non-ISM buyers lose significant value with the inability to coordinate their activities. Finally, Fig. 4 (bottom) compares the social welfare of the online matching with the optimum; here we see reasonable performance, as it attains, on average 93% of optimal social welfare.¹⁰

These results suggest that coordinated purchasing through cooperative buying in the presence of multiple vendor discounts can offer reasonable increases in social welfare and, more importantly, increase buyer satisfaction significantly.

8. CONCLUDING REMARKS

We presented a matching model for cooperative group buying in markets where multiple vendors offer discounts for their products, but buyers have distinct preferences over the vendors. We considered both transferable and nontransferable utility settings; and while the induced cooperative game among buyers may have an empty core, we showed that various forms of Nash stable matchings exist in both settings. While computing a social welfare maximizing matching was shown to be NP-hard, we developed a MIP formulation that admits a partial relaxation that can be solved to optimality effectively. Our empirical results suggest that cooperative buying can offer gains in buyer

¹⁰This is in part a function of the utility model, and that we don’t model customers abandoning a search.

welfare—and, we believe, help discount vendors attract and retain customers—even when stability is required in the presence of distinct buyer preferences for vendors.

Our model and analysis leave a number of important questions unaddressed. From a computational perspective, iterative algorithms in the style of Gale-Shapley for finding optimal stable matchings, especially with ordinal “deal” preferences, are of great interest, as are auction-based approaches. With respect to incentives, one of our primary goals is the design of mechanisms that account for the strategic reporting of buyer valuations, as well as the price-setting behavior of vendors in a two-sided market. There are many important extensions to our model that we are also investigating, including: vendor quantity limits; buyers with heterogeneous, combinatorial, multi-unit valuations; and new solution concepts and algorithms that reduce buyer burden by intelligently querying for partial preferences (e.g., partial rankings or imprecise valuations) and compute robust matchings.

Acknowledgements. We thank Darius Braziunas for useful discussions, and the anonymous reviewers for their helpful comments. This work was supported by NSERC.

REFERENCES

- ANAND, K. S. AND ARON, R. 2003. Group buying on the web: A comparison of price-discovery mechanisms. *Management Science* 49, 11, 1546–1562.
- BODINE-BARON, E., LEE, C., CHONG, A., HASSIBI, B., AND WIERMAN, A. 2011. Peer effects and stability in matching markets. In *Proc. 4th Symposium on Algorithmic Game Theory*. 117–129.
- CHEN, R. R. AND ROMA, P. 2010. Group buying of competing retailers. *Prod. and Oper. Mgmt.* 20, 181–197.
- CHEN, X. 2009. Inventory centralization games with price-dependent demand and quantity discount. *Oper. Res.* 57, 6, 1394–1406.
- DEMANGE, G., GALE, D., AND SOTOMAYOR, M. 1986. Multi-item auctions. *J. Pol. Econ.* 94, 863–872.
- DRECHSEL, J. AND KIMMS, A. 2010. Computing core allocations in cooperative games with an application to cooperative procurement. *Intl. J. Prod. Econ.* 128, 1, 310–321.
- ECHENIQUE, F. AND YENMEZ, M. B. 2007. A solution to matching with preferences over colleagues. *Games and Economic Behavior* 59, 1, 46–71.
- EDELMAN, B., JAFFE, S., AND KOMINERS, S. D. 2011. Togroupon or not togroupon: The profitability of deep discounts. Tech. Rep. Working Paper Number 11-063, Harvard Business School, Cambridge, MA.
- GALE, D. AND SHAPLEY, L. S. 1962. College admissions and the stability of marriage. *American Mathematical Monthly* 69, 1, 9–15.
- GOOSSENS, D. R., MAAS, A. J. T., SPIEKSMAS, F. C. R., AND VAN DE KLUNDERT, J. J. 2007. Exact algorithms for procurement problems under a total quantity discount structure. *Eur. J. Oper. Res.* 178, 603–626.
- HAFALIR, I. E. 2008. Stability of marriage with externalities. *Intl. J. Game Theory* 37, 3, 353–369.
- HEUVEL, W., BORM, P., AND HAMERS, H. 2007. Economic lot-sizing games. *Eur. J. Oper. Res.* 176, 2, 1117–1130.
- KOMINERS, S. D. 2010. Matching with preferences over colleagues solves classical matching. *Games and Economic Behavior* 68, 2, 773–780.
- KOTHARI, A., PARKES, D. C., AND SURI, S. 2005. Approximately-strategyproof and tractable multi-unit auctions. *Decision Support Systems* 39, 1, 105–121.
- LEHMANN, B., LEHMANN, D., AND NISAN, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2, 270–296.
- LEYTON-BROWN, K., SHOHAM, Y., AND TENNENHOLTZ, M. 2000. Bidding clubs: Institutionalized collusion in auctions. In *Proc. Second ACM Conference on Electronic Commerce (EC’00)*. 253–259.
- MATSUO, T., ITO, T., AND SHINTANI, T. 2005. A volume discount-based allocation mechanism in group buying. In *Proc. Intl. Workshop on Data Engineering Issues in E-commerce*. 59–70.
- MYERSON, R. B. 1991. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge.
- ROTH, A. E., ROTHBLUM, U. G., AND VANDE VATE, J. H. 1993. Stable matchings, optimal assignments, and linear programming. *Mathematics of Oper. Res.* 18, 4, 803–828.
- SASAKI, H. AND TODA, M. 1996. Two-sided matching problems with externalities. *J. Economic Theory* 70, 1, 93–108.
- SHAPLEY, L. S. AND SHUBIK, M. 1971. The assignment game I: The core. *Intl. J. Game Th.* 1, 111–130.

Online Appendix to: Matching Models for Preference-sensitive Group Purchasing

TYLER LU, University of Toronto
CRAIG BOUTILIER, University of Toronto

A. MODEL OF RANDOM DISCOUNT MATCHING INSTANCES

This appendix describes the model of random discount matching instances generated for the experiments in Sec. 7. While this model may not be realistic in some ways, it represents what the authors believe to be plausible.

Product (Vendor) Model

The key idea here is that different vendors will offer goods of varying “quality”, where quality can vary on several different attributes (including the special case of one quality measure). We use two attributes.

- For each vendor j , draw random quality $Q_a \sim \text{Beta}(4, 4)$ for each attribute a .
- Let overall quality Q_j for vendor j be $f(Q_{a_1}, Q_{a_2}, \dots) = 20 \prod_a Q_a$.
- Set the initial price $p_j^0 = Q_j + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, 0.05Q_j)$ is a small noise parameter

Discount Model

We keep the discount level to be fixed at 4 (including the default base threshold) for the experiments, although a straightforward extension of the description below can allow for arbitrary discount levels.

- Draw first discount percentage d_1 from the multinomial [5% 0.2; 10% 0.6; 15% 0.2] then set the discounted price $p_j^1 = (100 - d_1)\%$ of p_j^0 .
- Draw discount percentage d_2 from multinomial [5% 0.25; 10% 0.5; 15% 0.25] then set the discount price $p_j^2 = (100 - d_1 - d_2)\%$ of p_j^0 .
- Draw discount percentage d_3 from multinomial [5% 0.15; 10% 0.5; 15% 0.35] then set the discount price $p_j^3 = (100 - d_1 - d_2 - d_3)\%$ of p_j^0

To set the discount threshold, we note that each vendor would expect roughly $|M|/|N|$ buyers (ignoring variation in preferences, product quality, etc.) So we set thresholds according to the fraction of this expectation. First let $\varepsilon \sim \mathcal{N}(0, 0.02|M|/|N|)$.

- Set first threshold $\tau_j^1 = 0.75|M|/|N| + \varepsilon$.
- Set second threshold $\tau_j^2 = \tau_j^1 + 0.5|M|/|N| + \varepsilon$.
- Set third threshold $\tau_j^3 = \tau_j^2 + 0.75|M|/|N| + \varepsilon$.

User Preference Models

The basic intuition here is that different users care about different quality attributes, and that they may also care about the “brand” (vendor identity) independent of attributes, and that they may be more or less be price sensitive. So we might have some consumers that care mostly about price (care very little for quality); some who care mostly for specific attributes (possibly including brand), so are more price insensitive; and users who lie somewhere in between. Note: some consumers are clearly “discount” sensitive, and are inherently attracted to the discount level itself, but we do not model this. We will ignore price sensitivity, and instead focus on inherent preference.

Suppose that there are k quality attributes. Each user i has an inherent utility for a vendor j on a $[0, 1]$ scale generated by an additive utility function (leaving out i superscripts):

$$\left[\sum_{\ell=1}^k w_{\ell} a_{\ell j} \right] + w_b \cdot \text{BrandValue}_{ij}$$

where attribute/quality levels are normalized on a $[0, 1]$ scale (best value set to one, worst value set to zero), weights w_1, w_2, \dots, w_k plus brand weight w_b sum to one, and BrandValue_{ij} represents a component of a local value function over brands specific to buyer i .

We draw weights (including brand values) uniformly at random from $[0, 1]$ and normalize their sum to one. Users who are more brand sensitive (with high w_b) are likely skewed toward specific brands, hence the local value function should reflect this. Generally speaking such users will have more brand loyalty to higher quality products. However, we will ignore this for now.

We now have a utility for each vendor j on a $[0, 1]$ scale. We translate this to the willingness-to-pay as follows:

- Let j^* be most preferred product for buyer i : set willingness-to-pay for j^* to be $v_{ij^*} = p_{j^*}^0 + \varepsilon_{j^*}$, where $\varepsilon_{j^*} \sim \mathcal{N}(0.5p_{j^*}^0/2, 0.1p_{j^*}^0)$. The intuition here is that i is willing to pay roughly 50% more than j^* 's undiscounted base price, but with some (high) variance.
- Let j' be least preferred product for buyer i : set willingness-to-pay for j' to be $v_{ij'} = p_{j'}^0 + \varepsilon_{j'}$, where $\varepsilon_{j'} \sim \mathcal{N}(0.15p_{j'}^0, 0.1p_{j'}^0)$. Again the intuition is that i is willing to pay roughly 15% less than j' 's undiscounted base price, but with some high variance.
- Then set the willingness-to-pay for all other products to be proportional to their utility in the range $[v_{ij'}, v_{ij^*}]$.

B. PROOFS

PROOF OF THEOREM 4.2. We can reduce the optimal allocation problem given threshold values to an instance of the minimum-cost maximum-flow problem. Construct a network consisting of a source node s and sink node t . Let nodes a_1, \dots, a_n represent buyers and add arcs (s, a_i) with capacity 1 and zero cost. Include nodes r_1, \dots, r_m , and r'_1, \dots, r'_m and x . For all i, j , add arcs (a_i, r_j) and (a_i, r'_j) , each with capacity 1 and cost $L - (v_{ij} - p_j)$, where L is the largest buyer valuation. Also add arcs (a_i, x) with capacity 1 and cost L . Flow from a_i to r_j represents an assignment of buyer i to vendor j s.t. i is one of the buyers needed to meet threshold τ_j ; and flow from a_i to r'_j represents the assignment of i to j , but where i 's demand is considered in excess of τ_j . Flow on arcs to x represents unassigned buyers.

For each node r_j , there must be at most τ_j buyers assigned to it. We enforce this by adding arcs (r_j, t) with capacity τ_j and zero cost. There must be, in total, $n - \tau_1 - \dots - \tau_m$ buyers assigned to the "excess" nodes r'_j and x . We add a node e with arcs (r'_j, e) for $j \in M$, and arc (x, e) , each with infinite capacity and zero cost, and arc (e, t) with capacity $n - \tau_1 - \dots - \tau_m$ and zero cost.

A maximum flow of n can be achieved by routing the first $\tau_1 + \dots + \tau_m$ buyers to the r_j 's subject to capacity constraints, and remaining buyers to x . Recall the assignment of buyer i to vendor j corresponds to (integer) flows from a_i to either r_j or r'_j . In any max flow, there must be total flow of τ_j to r_j ; otherwise total flow to the r_j 's is less than $\tau_1 + \dots + \tau_m$, hence total flow to t is less than n given the capacity of arc (e, t) . This also means a max flow must reach full capacity at (e, t) .

Because the cost of arcs (a_i, r_j) , (a_i, r'_j) , and (a_i, x) is proportional to the negative surplus of buyer i when assigned to vendor j at price $p_j(\tau_j)$, the SWM matching (fixing threshold requirements) corresponds to an integer optimal min-cost max-flow solution where flows on the aforementioned arcs correspond to matchings. Standard results for the min-cost max-flow problem state that there is always an optimal *integer* flow, which can be found in polynomial time using, e.g., a modification of the Ford-Fulkerson algorithm. \square

PROOF OF THEOREM 4.3. The problem is in NP, since computing $SW(\mu)$ for a given matching μ is straightforward. For hardness, consider a reduction from the Knapsack problem where given integer capacity C , integer item weights $\{w_1, \dots, w_T\}$ and values $\{v_1, \dots, v_T\}$, and a number $y \geq 0$, we wish to determine whether there exists a subset of items $S \subseteq \{1, \dots, T\}$ such that $\sum_{i \in S} v_i \geq y$ and $\sum_{i \in S} w_i \leq C$. In the reduction, construct:

- T vendors each with one threshold level, i.e. $D = 1$. They have thresholds $\tau_i^0 = 0, \tau_i^1 = w_i$ with corresponding prices $p_i^0 = v_i, p_i^1 = 0$.
- $n = C$ buyers z_1, \dots, z_n each of whom have zero valuation for all vendors, and
- T additional buyers s_1, \dots, s_T where $v_{s_i}(i) = v_i$ and $v_{s_i}(j) = 0$ for all $i, j \in \{1, \dots, T\}, i \neq j$.
- Set $x = y$.

Suppose we have a “yes” instance of Knapsack, and that S is the set of feasible items such that its total value exceeds y . In the corresponding segmentation problem, we can allocate τ_i^2 of the z buyers, and an additional buyer s_i to each vendor $i \in S$. In such an allocation each vendor in S will set prices to zero. User $s_i, i \in S$, will obtain a payoff of v_i by being assigned to i . Thus the allocation has a payoff of $\sum_{i \in S} v_i$ which is at least $x (= y)$.

Suppose we have a “no” instance of Knapsack. Consider any allocation μ of buyers to vendors. For any vendor i in which the threshold τ_i^1 is not met, then the payoff of any buyer assigned to i cannot be positive. Consider the set S of vendors with at least τ_i^1 assigned buyers. Then the payoff at any vendor $i \in S$ is at most v_i because s_i has valuation v_i and all other buyers have zero valuation, while price of any vendor is zero. Thus the payoff of μ is at most $\sum_{i \in S} v_i$ which is strictly less than x , otherwise the corresponding item set S for Knapsack would be a feasible “yes” solution (since the number of buyers allocated to vendors in S satisfies $\sum_{i \in S} \tau_i^1 = \sum_{i \in S} w_i \leq n = C$ which is impossible by assumption). \square