Εθνικό Μετσόβιο Πολυτεχνείο
Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Χρηματικά Ανταλλάξιμη Ωφέλεια σε Συνδυαστικές Δημοπρασίες

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Κυριάκος Λωτίδης

Επιβλέπων: Δημήτρης Φωτάκης
Αναπληρωτής Καθηγητής Ε.Μ.Π.

Εργαστήριο Λογικής και Επιστήμης Υπολογισμών
Αθήνα, Οκτώβριος 2018

Εθνικό Μετσόβιο Πολυτεχνείο
Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Χρηματικά Ανταλλάξιμη Ωφέλεια σε Συνδυαστικές Δημοπρασίες

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Κυριάκος Λωτίδης

Επιβλέπων: Δημήτρης Φωτάκης
Αναπληρωτής Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 31η Οκτωβρίου 2018.

....	....	....
Δημήτρης Φωτάκης	Άρης Παγουρτζής	Νικόλαος Παπασπύρου
Αν. Καθηγητής Ε.Μ.Π.	Αν. Καθηγητής Ε.Μ.Π.	Αν. Καθηγητής Ε.Μ.Π.

Εργαστήριο Λογικής και Επιστήμης Υπολογισμών
Αθήνα, Οκτώβριος 2018

....
Κυριάκος Λωτίδης
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Περίληψη

Στην παρούσα διπλωματική εργασία μελετούμε φιλαλήθεις μηχανισμούς σε Συνδυαστικές Δημοπρασίες με budgets. Για την περίπτωση που δεν υπάρχει περιορισμός στα budgets, παρόλο που υπάρχει μηχανισμός που βρίσκει τη βέλτιστη λύση για το Social Welfare, η εφαρμογή του είναι υπολογιστικά αδύνατη. Συνεπώς, πρέπει να αναζητήσουμε προσεγγιστικούς μηχανισμούς που να εκτελούνται αποδοτικά. Στο πλαίσιο αυτό, παρουσιάζουμε ένα διαφορετικό είδος μηχανισμών, τις δημοπρασίες αυξανόμενης τιμής, και τις ιδέες των clearing prices και Walrasian Equilibrium. Επιπλέον, παρουσιάζουμε μηχανισμούς για συνδυαστικές δημοπρασίες με χρήση demand query oracles, που εμπνεύστηκαν από την ιδέα των clearing prices, και επιτυγχάνουν τους καλύτερους λόγους προσέγγισης, μέχρι σήμερα. Αναλύοντας τα βασικά συστατικά των μηχανισμών αυτών, ερευνούμε τις προϋποθέσεις, κάτω από τις οποίες μπορούμε να επεκτείνουμε τα αποτελέσματα αυτά για το Liquid Welfare, μία μετρική μέτρησης της αποτελεσματικότητας των μηχανισμών με budget-restricted παίχτες, η οποία προτάθηκε από τους Dobzinski και Leme.

Έπειτα, δείχνουμε ότι κάποια από τα πιο γνωστά αποτελέσματα μηχανισμών για τη προσέγγιση του Social Welfare με submodular (ή XOS) παίχτες μπορούν να προσαρμοστούν για τη μετρική του Liquid Welfare. Πιο συγκεκριμένα, για τη βελτιστοποίηση του Liquid Welfare σε συνδυαστικές δημοπρασίες με submodular παίχτες, παίρνουμε ένα φιλαλήθη $O(\log m)$-προσεγγιστικό μηχανισμό, όπου m ο αριθμός των αντικειμένων, προσαρμόζοντας το μηχανισμό των Krysta και Vöcking.

Στη συνέχεια, με βάση την ιδέα του large market assumption, παρουσιάζουμε μία νέα έννοια για competitive markets και δείχνουμε ότι σε τέτοιες αγορές, το Liquid Welfare μπορεί να προσεγγιστεί με ένα σταθερό παράγοντα. Τέλος, για το Bayesian setting, προσαρμόζοντας τα αποτελέσματα των Feldman et al., παίρνουμε έναν φιλαλήθη $O(1)$-προσεγγιστικό μηχανισμό για την περίπτωση που τα valuation των παιχτών παράγονται σαν ανεξάρτητα δείγματα από γνωστές κατανομές.

Λέξεις Κλειδιά: Σχεδιασμός Μηχανισμών, Συνδυαστικές Δημοπρασίες, Walrasian Equilibrium, Demand Queries, Περιορισμοί στο Budget, Liquid Welfare, Competitive Markets.

Abstract

In this thesis, we study truthful mechanisms in *Combinatorial Auctions* with budgets. For the budget-unrestricted case, although there is a mechanism that provides a welfare maximizing solution in a truthful way, its implementation is computationally intractable in most cases. Therefore, we have to design approximation mechanisms that can be executed efficiently. In this sense, we present a different kind of mechanisms, the *Ascending Price* auctions, and the notion of *clearing prices* and *Walrasian Equilibirum*. Furthermore, we present mechanisms for Combinatorial Auctions through demand query oracles that are motivated by the clearing prices and succeed the best approximation ratios, until now. Analyzing the basic components of such mechanisms, we investigate the conditions under which we can extend these results for the *Liquid Welfare*, a notion of efficiency for budget-constrained bidders introduced by Dobzinksi and Leme.

In this work, we show that some of the best known truthful mechanisms that approximate the Social Welfare for Combinatorial Auctions with submodular (or XOS) bidders through demand query oracles can be adapted so that they retain truthfulness and achieve asymptotically the same approximation guarantees for the Liquid Welfare. More specifically, for the problem of optimizing the Liquid Welfare in Combinatorial Auctions with submodular bidders, we obtain a universally truthful randomized $O(\log m)$-approximate mechanism, where m is the number of items, by adapting the mechanism of Krysta and Vöcking.

Additionally, motivated by large market assumptions often used in mechanism design, we introduce a notion of competitive markets and show that in such markets, Liquid Welfare can be approximated within a constant factor by a randomized universally truthful mechanism. Finally, in the Bayesian setting, we obtain a truthful $O(1)$-approximate mechanism for the case where bidder valuations are generated as independent samples from a known distribution, by adapting the results of Feldman et al. .

Keywords: Mechanism Design, Combinatorial Auctions, Walrasian Equilibrium, Posted-Price mechanisms, Demand Queries, Budget Constraints, Liquid Welfare, Competitive Markets.

Ευχαριστίες

Με την ολοκλήρωση της παρούσας διπλωματικής εργασίας θα ήθελα να ευχαριστήσω θερμά τον κ. Φωτάκη που, ως καθηγητής και ως άνθρωπος, με στήριξε με τις πολύτιμες συμβουλές και την καθοδήγησή του και με εισήγαγε στον κόσμο της Θεωρητικής Πληροφορικής και της έρευνας. Ακόμα, οφείλω ένα μεγάλο ευχαριστώ στην Χαρά Ποδηματά, που μέσα από πολύωρες συζητήσεις και την διαρκή ενασχόληση και βοήθεια της, καταφέραμε να δημοσιεύσουμε τα αποτελέσματά μας. Επίσης, θα ήθελα να ευχαριστήσω ιδιαίτερα την οικογένειά μου και τους φίλους μου Γρηγόρη, Παναγιώτη, Κώστα και Γιάννη για τη στήριξη και τη βοήθεια που μου παρείχαν. Τέλος, οφείλω ένα τεράστιο ευχαριστώ στην Ιωάννα που με στήριξε σε όλες μου τις δυσκολίες.

Κυριάκος Λωτίδης

Contents

List of Algorithms

Chapter 1

Εκτεταμένη Ελληνική Περίληψη

Στο σημείο αυτό, θα συνοψίσουμε το περιεχόμενο της παρούσας διπλωματικής, δίνοντας βασικούς ορισμούς και θεωρήματα, χωρίς αποδείξεις.

1.1 Εισαγωγή

Με τον όρο 'Σχεδιασμός Μηχανισμών' (Mechanism Design) εννοούμε το σχεδιασμό μιας διαδικασίας σε ένα παίγνιο με στρατηγικούς παίχτες, με στόχο την εύρεση της βέλτιστης λύσης, σύμφωνα με κάποια αντικειμενική συνάρτηση. Στα παίγνια αυτά, κάθε παίχτης έχει την προσωπική του στρατηγική και δρα με γνώμονα αποκλειστικά την δική του ικανοποίηση. Σκοπός του μηχανισμού είναι να μη δώσει κίνητρο στους παίχτες να δηλώσουν ψευδείς πληροφορίες, ώστε να έχει εγγυήσεις για το αποτέλεσμα. Οι μηχανισμοί χωρίζονται σε direct revelation, όπου οι παίχτες ανακοινώνουν στο μηχανισμό τις προτιμήσεις τους, και σε indirect, όπου ο μηχανισμός εξελίσσεται ανάλογα με τις πράξεις των παιχτών. Ένας direct revelation μηχανισμός με n παίχτες και ένα σύνολο $\mathcal{O}$ από δυνατά αποτελέσματα λειτουργεί ως εξής: Κάθε παίχτης i έχει ένα valuation $v_i : \mathcal{O} \to R_{\geq 0}$ και ανακοινώνει στον μηχανισμό το *bid* του b_i. Ο μηχανισμός αντιστοιχεί το $\vec{b} = (b_1, \ldots, b_n)$ μέσω μιας συνάρτησης $f : \vec{b} \to \mathcal{O}$ σε ένα αποτέλεσμα και υπολογίζει ένα διάνυσμα πληρωμών $\vec{p} = (p_1, \ldots, p_n)$. Single-parameter περιβάλλον θεωρούμε όταν κάθε παίχτης πρέπει να ανακοινώσει μία μόνο τιμή, ενώ multi-parameter αν πρέπει πολλαπλές. Το single-parameter είναι αρκετά ευκολότερο και γι΄ αυτό θα ασχοληθούμε με το multi-parameter.

Ένα **multi-parameter** περιβάλλον ορίζεται στην γενική περίπτωση ως εξής:

- n στρατηγικοί παίχτες.
- Ένα πεπερασμένο σύνολο Ω από τα εφικτά αποτελέσματα.
- Κάθε παίχτης έχει ένα προσωπικό valuation $v_i(\omega)$, $\forall \omega \in \Omega$.

Η **ωφέλεια** κάθε παίχτη i ορίζεται στις περισσότερες περιπτώσεις ως $u_i = v_i(f(\vec{b})) - p_i(\vec{b})$ και ονομάζουμε τη στρατηγική b_i ενός παίχτης **κυρίαρχη στρατηγική**, εάν μεγιστοποιεί την ωφέλειά του, ανεξάρτητα με το τι θα παίξουν οι άλλοι. Αν η **κυρίαρχη στρατηγική** ενός παίχτη είναι να ανακοινώσει το πραγματικό του valuation, τότε ο μηχανισμός είναι **φιλαλήθης**.

Συνδυαστική Δημοπρασία είναι μία ειδική κατηγορία του multi-parameter περιβάλλοντος. Πιο συγκεκριμένα, αποτελείται από ένα σύνολο m αντικειμένων και n στρατηγικούς παίχτες. Κάθε παίχτης i έχει ένα valuation $v_i : 2^m \to \mathbb{R}_{\geq 0}$. Σκοπός είναι να υπολογίσουμε μία ανάθεση $S = (S_1, \ldots, S_n)$ των m αντικειμένων ώστε να μεγιστοποιήσουμε τη συνάρτηση αξιολόγησης, η οποία στις περισσότερες περιπτώσεις είναι η συνάρτηση **Κοινωνικής Ευημερίας** και ορίζεται ως $\sum_{i=1}^{n} v_i(S_i)$. Οι πιο βασικές κλάσεις valuation είναι οι additive, gross substitutes, submodular,XOS, subadditive και περιγράφονται αναλυτικά στο section 3.3.2

Ο μηχανισμός **VCG** , ο οποίος παρουσιάζεται αναλυτικά στο section 3.3, υπολογίζει μία ανάθεση που μεγιστοποιεί τη συνάρτηση κοινωνικής ευημερίας με φιλαλήθη τρόπο, ωστόσο είναι αδύνατη η εφαρμογή του, στις περισσότερες περιπτώσεις, λόγω της υπολογιστικής πολυπλοκότητας. Για το λόγο αυτό, πρέπει να σχεδιάσουμε μηχανισμούς που προσεγγίζουν τη βέλτιστη λύση. Μία κατηγορία τέτοιων μηχανισμών είναι οι *Posted-Price* μηχανισμοί.

Σε έναν **posted-price** μηχανισμό, δίνουμε τιμές στα αντικείμενα και στη συνέχεια, οι παίχτες έρχονται με μία σειρά και διαλέγουν το σύνολο από αντικείμενα που μεγιστοποιεί την ωφέλειά τους στις τιμές αυτές, $\vec{p}$. Πολλές φορές, εφαρμόζεται ένας κανόνας ενημέρωσης στις τιμές των αντικειμένων που επιλέγονται. Στο σημείο αυτό, είναι εμφανές πως ο λόγος προσέγγισης εξαρτάται από τον τρόπο που υπολογίζονται και ενημερώνονται οι τιμές των αντικειμένων. Το σύνολο των αντικειμένων που μεγιστοποιεί την ωφέλειά τους δίνεται μέσω **Demand Queries (DQ).** Ένα DQ$(v_i, U, \vec{p})$ επιστρέφει, δηλαδή, το:

$$S_i = \arg\max_{S \in U} \{v_i(S) - p(S)\}$$

Επίσης, λέμε ότι ένας τυχαιοποιημένος μηχανισμός έχει λόγο προσέγγισης ρ, εάν $\mathbb{E}[ALG] \geq \rho \cdot OPT$, όπου $\rho \leq 1$, ALG η λύση του μηχανισμού και OPT η βέλτιστη λύση.

1.2 Δημοπρασίες αυξανόμενης τιμής

Μια κατηγορία δημοπρασιών είναι οι **δημοπρασίες αυξανόμενης τιμής**. Σε αυτές τις δημοπρασίες, αρχικοποιούνται οι τιμές όλων των αντικειμένων στο 0, και κάθε παίχτης καλείται να διαλέξει το σύνολο των αντικειμένων που μεγιστοποιούν την ωφέλειά του στις υπάρχουσες τιμές, μέσα από μια επαναληπτική διαδικασία. Όταν ένας παίχτης διαλέγει ένα σύνολο, τότε κάθε αντικείμενο σε αυτό αυξάνει την τιμή του. Οι επόμενοι παίχτες μπορούν φυσικά να πάρουν κάποια από τα αντικείμενά του, τα οποία αν θέλει μπορεί να ξαναπάρει στον επόμενο γύρο σε κάποια μεγαλύτερη τιμή. Όταν κανένας παίχτης δεν επιθυμεί κάποιο άλλο αντικείμενο, πέρα από αυτά που του έχουν απομείνει, η δημοπρασία τερματίζει. Ένας τέτοιος μηχανισμός περιγράφεται στον αλγόριθμο 5. Το ζήτημα είναι να αναλύσουμε την αποτελεσματικότητα του μηχανισμού αυτού. Για το λόγο αυτό, πρώτα θα ορίσουμε κάποιες απαραίτητες έννοιες για την ανάλυσή μας.

Έστω n στρατηγικοί παίχτες και ένα σετ U από m διαφορετικά αντικείμενα. Ως **Walrasian Equilibrium** ορίζουμε ένα διάνυσμα τιμών $\vec{p}$ και μία ανάθεση $S = (S_1, \ldots, S_n)$, και λέμε ότι το $(S, \vec{p})$ είναι μία Walrasian Equilibrium, εάν ικανοποιούν τις εξής συνθήκες:

- Κάθε παίχτης i παίρνει το σύνολο που μεγιστοποιεί την ωφέλειά του (ή το $\emptyset$), δηλαδή:

$$S_i \in \arg\max_{T \subseteq U} \left\{ v_i(T) - \sum_{j \in T} p_j \right\}$$

- Ένα αντικείμενο $j \in U$ είναι απούλητο, μόνο εάν $p_j = 0$.

Η έννοια της Walrasian Equilibrium είναι πολύ σημαντική, καθώς εξασφαλίζει υψηλή Κοινωνική Ευημερία. Πιο συγκεκριμένα, σύμφωνα με το παρακάτω Θεώρημα, αν το $(S, \vec{p})$ είναι μία Walrasian Equilibrium, τότε η ανάθεση S μεγιστοποιεί την Κοινωνική Ευημερία.

Θεώρημα: Έστω μηχανισμός με n παίχτες, ένα σύνολο U από m διαφορετικά αντικείμενα, μία ανάθεση $S = (S_1, \ldots, S_n)$ κι ένα διάνυσμα τιμών $\vec{p}$. Αν $(S, \vec{p})$ είναι μία Walrasian ισορροπία, τότε η S μεγιστοποιεί την κοινωνικά ευημερία.

Απόδειξη: Έστω $O = (O_1, \ldots, O_n)$ μία ανάθεση που μεγιστοποιεί την κοινωνικά ευημερία και $P = \sum_{j \in U} p_j$. Αφού $(S, \vec{p})$ είναι μια Walrasian ισορροπία, κάθε παίχτης παίρνει το 'αγαπημένο' του σύνολο, δηλαδή για κάθε άλλο σύνολο $T \subseteq U$, έχουμε:

$$S_i \in \arg\max_{T \in U} \left\{ v_i(T) - \sum_{j \in T} p_j \right\}$$

Συνεπώς, για κάθε παίχτη i ισχύει:

$$v_i(S_i) - \sum_{j \in S_i} p_j \geq v_i(O_i) - \sum_{j \in O_i} p_j$$

Αθροίζοντας τις εξισώσεις για κάθε παίχτη, έχουμε:

$$v(S) - \sum_{i \in [n]} \sum_{j \in S_i} p_j \geq v(O) - \sum_{i \in [n]} \sum_{j \in O_i} p_j \tag{1.1}$$

Ωστόσο, επειδή στην ανάθεση S ένα αντικείμενο j είναι απούλητο μόνο αν $p_j = 0$, ο αρνητικός όρος του αριστερού μέλους της παραπάνω εξίσωσης αθροίζει σε όλα τα αντικείμενα που έχουν μη μηδενικές τιμές και άρα έχουμε:

$$\sum_{i \in [n]} \sum_{j \in S_i} p_j \geq \sum_{i \in [n]} \sum_{j \in O_i} p_j \tag{1.2}$$

Συνδυάζοντας τις εξισώσεις (1.1) και (1.2), παίρνουμε:

$$v(S) \geq v(O)$$

και άρα η S μεγιστοποιεί την κοινωνική ευημερία.

Οι τιμές σε ένα Walrasian διάνυσμα τιμών $\vec{p}$ ονομάζονται clearing prices, καθώς, υπό αυτές τις τιμές, η προσφορά ισούται με τη ζήτηση και άρα υπάρχει μια ισορροποία στην αγορα. Επίσης, αν έχουμε υπολογίσει clearing prices, τρέχοντας έναν posted-price μηχανισμό, καταλήγουμε σε βέλτιστη λύση.

Ωστόσο, η υπάρξη μιας Walrasian Equilibrium δεν είναι δεδομένη και εξαρτάται από την κλάση των valuation των παιχτών. Η κλάση Gross Substitutes, που ορίζεται αναλυτικά στο section 4.5 είναι η μεγαλύτερη κλάση συναρτήσεων για τις οποίες υπάρχει πάντα μία Walrasian Equilibrium. Με λίγα λόγια, η κλάση αυτή περιγράφεται από την εξής ιδιότητα: Έστω κάποιος παίχτης που αγόρασε ένα σύνολο αντικειμένων T, το οποίο μεγιστοποιούσε την ωφέλειά του, υπό ένα διάνυσμα τιμών $\vec{p}$ και $\vec{p}\,'$ ένα διάνυσμα τιμών με $p_j \leq p'_j$, για κάθε αντικείμενο j. Τότε, τα αντικείμενα $R \subseteq T$ των οποίων η τιμή παρέμεινε σταθερή πρέπει να ανήκουν στο σετ που μεγιστοποιεί την ωφέλειά του υπό το δίανυσμα τιμών $\vec{p}\,'$. Υπό αυτή την ιδιότητα, ο αλγόριθμος 5 τερματίζει σε μία Walrasian Equilibria.

1.3 Συνδυαστικές Δημοπρασίες με Demand Queries

Στο κεφάλαιο αυτό παρουσιάζουμε τα πιο σημαντικά αποτελέσματα των συνδυαστικών δημοπρασιών, που προκύπτουν από posted-price μηχανισμούς. Πιο συγκεκριμένα, αναλύουμε τα δομικά συστατικά τους και τις ιδέες, συνδυάζοντας τεχνική ανάλυση και διαίσθηση. Η ιδέα πίσω από τους posted-price μηχανισμούς είναι να δώσουμε τιμές στα αντικείμενα, τέτοιες ώστε, αν οι παίχτες έρχονται σε μία σειρά και παίρνουν το σύνολο που μεγιστοποιεί την ωφέλειά τους, να εξασφαλίσουμε ότι η τελική ανάθεση είναι κοντά στη βέλτιστη λύση. Για valuation κλάσεις, όπου η ύπαρξη clearing prices είναι δεδομένη, μπορούμε να βρούμε τιμές που ένας posted-price μηχανισμός, μεγιστοποιεί την κοινωνική ευημερία. Ωστόσο, όπως τονίσαμε στο προηγούμενο κεφάλαιο, κάτι τέτοιο δεν ισχύει για submodular, XOS και subadditive κλάσεις. Συνεπώς, αυτό που μπορούμε να κάνουμε είναι να βρούμε approximate clearing prices. Από το Παράδειγμα 5.1 καθίσταται σαφές ότι δεν αρκεί μόνο αυτό, αλλά πρέπει το revenue του μηχανισμού να είναι συγκρίσιμο με την αξία της βέλτιστης λύσης. Συνεπώς, αν ο μηχανισμός τερματίσει με την ανάθεση S, θέλουμε για τις τιμές των αντικειμένων να ισχύει:

$$\sum_{j \in S} p_j \geq a \cdot v(OPT)$$

όπου a ο λόγος προσέγγισης.

Στο worst-case setting, όπου δεν έχουμε καμία πληροφορία για τα valuation των παιχτών, οι Dobzinski et al. [21] εισήγαγαν έναν $O(\log^2 m)$- approximation posted-price μηχανισμό για submodular valuations. Η ιδέα πίσω από το μηχανισμό αυτό είναι η εύρεση μίας τιμής, ίδιας για όλα τα αντικείμενα, όχι πολύ υψηλής ούτε πολύ χαμηλής, η οποία να εξασφαλίζει ότι τα αντικείμενα που θα αγοραστούν, προσεγγίζουν τη βέλτιστη λύση με ένα παράγοντα $O(\log^2 m)$. Ο μηχανισμός αυτός παρουσιάζεται πιο αναλυτικά στο section 5.3.1.

Έπειτα, οι Krysta και Vöcking [41] κατάφεραν να προσεγγίσουν την βέλτιστη Κοινωνική Ευημερία με ένα παράγοντα $O(\log m)$ για submodular valuations. Η ιδέα του μηχανισμού προήλθε από τις δημοπρασίες αυξανόμενης τιμής, όπου η τιμή κάθε αντικειμένου αυξανόταν ανάλογα με τη ζήτησή του. Έτσι, ορίζοντας έναν χώρο αναζήτησης τιμών μεγέθους $\log m$, όταν ένας παίχτης αγοράζει κάποιο αγαθό η τιμή του αγαθού διπλασιάζεται. Για να διατηρηθεί η εφικτότητα της λύσης, ο μηχανισμός αναθέτει 'εικονικά' αντίγραφα των αντικειμένων, μαθαίνοντας έτσι τη 'σωστή' τιμή του κάθε αντικειμένου και τελικά αποφασίζει για την ανάθεσή τους ή μη μέσω ανεξάρτητων δοκιμών Bernoulli. Ο μηχανισμός αυτός παρουσιάζεται πιο αναλυτικά στο section 5.3.2.

Το πιο πρόσφατο αποτέλεσμα για submodular valuations είναι ένας $O(\sqrt{\log m})$ μηχανισμός από τον Dobzisnki [16], όπου συνδυάζει τις δύο προηγούμενες ιδέες, δηλαδή ίδιων και διαφορετικών τιμών. Ο τρόπος που το επιτυγχάνει αυτό είναι μέσω επαναλαμβανόμενων δημοπρασιών, οι οποίες γίνονται τελικές με μία μικρή πιθανότητα, με ίδια τιμή σε όλα τα αντικείμενα, μέσω των οποίων μαθαίνει τις διαφορετικές τιμές των αντικειμένων, ανάλογα με τον αν αγοράστηκαν ή όχι στις εκάστοτε τιμές. Ο μηχανισμός αυτός παρουσιάζεται πιο αναλυτικά στο section 5.3.3.

Τέλος, για το Bayesian setting, όπου τα valuations των παιχτών προέρχονται από κατανομές, τις οποίες γνωρίζει ο μηχανισμός, παρουσιάζουμε ένα $O(1)$-προσεγγιστικό μηχανισμό, τον οποίο εισήγαγαν οι Feldman et al. [29]. Ο τρόπος που υπολογίζει ο μηχανισμός αυτός τις τιμές είναι μέσω ghost samples από τις κατανομές. Πιο συγκεκριμένα, μέσω των samples υπολογίζει μία σχεδόν βέλτιστη λύση, και θέτει την τιμή κάθε αντικειμένου ίση με το μισό της συνεισφοράς του στη λύση αυτή. Ο μηχανισμός αυτός παρουσιάζεται πιο αναλυτικά στο section 5.4.

1.4 Liquid Welfare σε Συνδυαστικές Δημοπρασίες

Μέχρι στιγμής, όλοι οι μηχανισμοί που αναλύσαμε αγνοούσαν μία σημαντική παράμετρο: τον περιορισμό στα budgets των παιχτών. Επειδή η αντικειμενική της Κοινωνικής Ευημερίας (SW) δε μπορεί να προσεγγιστεί από κάποιο παράγοντα μικρότερο του n με φιλαλήθη τρόπο, όταν οι παίχτες έχουν περιορισμούς ρευστότητας, οι Dobzinski και Leme [19] πρότειναν σαν αντικειμενική συνάρτηση τη **Ρευστή Ευημερία (LW)**, η οποία ορίζεται ως το ελάχιστο του valuation και του budget κάθε παίχτη. Πιο συγκεκριμένα, το LW μιας ανάθεσης $S = (S_1, \ldots, S_n)$ ορίζεται ως:

$$LW(S) = \sum_{i=1}^{n} \min\{v_i(S_i), B_i\}$$

όπου B_i το budget του παίχτη i. Ως liquid valuation ορίζουμε το $\bar{v}_i(S_i) = \min\{v_i(S_i), B_i\}$.

Σε συνδυαστικές δημοπρασίες, δεν υπήρχε κανένα αποτέλεσμα μέχρι τώρα στην προσέγγιση του LW, εκτός από την περίπτωση με διαιρέσιμα αντικείμενα και additive valuations υπό ένα large market assumption. Ένα εύλογο ερώτημα είναι, γιατί δεν εφαρμόζουμε τους posted-price αλγόριθμους για SW με χρήση του $\text{DQ}(\min\{v, B\}, U, \vec{p})$; Αρχικά, πρέπει να αποδείξουμε ότι τα v και $\bar{v}$ ανήκουν στην ίδια κλάση, για v submodular, XOS ή subadditive. Η απόδειξη δίνεται στο παρακάτω Λήμμα.

Λήμμα 1: Έστω v μία submodular (αντίστοιχα XOS, subadditive) συνάρτηση. Τότε, για οποιοδήποτε $B \in \mathbb{R}_{\geq 0}$, $\bar{v} = \min\{v, B\}$ είναι επίσης sumodular (αντίστοιχα XOS, subadditive).

Απόδειξη: Αρχικά, η $\bar{v}$ διατηρεί τη μονοτονία. Κάνουμε την απόδειξη για κάθε περίπτωση ξεχωριστά.

- (submodular) Έστω v submodular συνάρτηση. Τότε, από τον ορισμό του submodularity, για sets $T \subseteq S$ και $j \notin S$ έχουμε:

$$v\left(S \cup \{j\}\right) - v\left(S\right) \leq v\left(T \cup \{j\}\right) - v\left(T\right)$$

Επίσης, επειδή v είναι μονότονη: $v(T) \leq v(S)$, το οποίο συνεπάγεται $\bar{v}(T) \leq \bar{v}(S)$. Έχουμε, λοιπόν τις εξής περιπτώσεις:

1. Αν $B \leq v(T \cup \{j\}) \leq v(S \cup \{j\})$. Τότε, για τα liquid valuations έχουμε: $\bar{v}(S \cup \{j\}) - \bar{v}(S) = B - \bar{v}(S) \leq B - \bar{v}(T) \leq \bar{v}(T \cup \{j\}) - \bar{v}(T)$, όπου η πρώτη ανισότητα ισχύει από τη μονοτονία.

2. Αν $\bar{v}(T \cup \{j\}) \leq \bar{v}(S \cup \{j\}) \leq B$. Τότε, $\bar{v}(S \cup \{j\}) - \bar{v}(S) = v(S \cup \{j\}) - v(S) \leq v(T \cup \{j\}) - v(T) = \bar{v}(T \cup \{j\}) - \bar{v}(T)$.

3. Αν $v(T \cup \{j\}) \leq B \leq v(S \cup \{j\})$. Τότε, έχουμε τις ακόλουθες περιπτώσεις: Αν $v(S) \geq B$ τότε, $\bar{v}(S \cup \{j\}) - \bar{v}(S) = 0 \leq v(T \cup \{j\}) - v(T) = \bar{v}(T \cup \{j\}) - \bar{v}(T)$. Αν $v(S) < B$, τότε $\bar{v}(S \cup \{j\}) - \bar{v}(S) = B - v(S) \leq v(S \cup \{j\}) - v(S) \leq v(T \cup \{j\}) - v(T) = \bar{v}(T \cup \{j\}) - \bar{v}(T)$. Τέλος, λόγω της μονοτονίας αυτές είναι οι μοναδικές περιπτώσεις.

- (XOS) Έστω v μία XOS συνάρτηση: Τότε, υπάρχουν αθροιστικές συναρτήσεις $\alpha_1, \ldots, \alpha_l$, τέτοιες ώστε $v(S) = \max_{i \in [l]} \alpha_i(S)$. Για να είναι η $\bar{v}$ XOS, πρέπει να δείξουμε ότι υπάρχουν αθροιστικές συναρτήσεις $\alpha'_1, \ldots, \alpha'_k$ τ'τετοιες ώστε $\bar{v}(S) = \max_{i \in [k]} \alpha'_i(S)$. Για κάθε συνάρτηση α_i θα ορίσουμε $m!$ συναρτήσεις, μία για κάθε διαφορετική μετάθεση π των αντικειμένων. Έστω μια συγκεκριμένη μετάθεση π_t των αντικειμένων $\{1, 2, \ldots, m\}$ και έστω $\pi_t(j)$ η θέσει του αντικειμένου j στη μετάθεση π_t. Ορίζουμε $\beta_i^{\pi_t}$ ως:

$$\beta_i^{\pi_t}(\{j\}) = \begin{cases} \alpha_i(\{j\}), & \text{αν } \sum_{k:\pi_t(k) \leq \pi_t(j)} \alpha_i(\{k\}) \leq B \\ \max\left\{B - \sum_{k:\pi_t(k) < \pi_t(j)} \alpha_i(\{k\}), 0\right\}, & \text{αν } \sum_{k:\pi_t(k) \leq \pi_t(j)} \alpha_i(\{k\}) > B \end{cases}$$

Πρώτα, θα δείξουμε ότι για κάθε $S \subseteq U, \beta_i^{\pi_t}(S) \leq \min\{v(S), B\}, \forall i, \pi_t$.

Από τον ορισμό του $\beta_i^{\pi_t}$, είναι προφανές ότι $\beta_i^{\pi_t}(\{j\}) \leq \alpha_i(\{j\})$. Συνεπώς, αθροίζοντας σε όλα τα αντικείμενα στο S (αφού έχουμε αθροιστικές συναρτήσεις), παίρνουμε ότι:

$$\beta_i^{\pi_t}(S) \leq \alpha_i(S) \leq \max_k \alpha_k(S) = v(S)$$

Από τον ορισμό του $\beta_i^{\pi_t}$, έχουμε ακόμα ότι $\beta_i^{\pi_t}(S) \leq B$.

Στη συνέχεια, θα δείξουμε ότι για κάθε $S \subseteq U : \exists \beta_i^{\pi_t}$ τέτοια ώστε $\beta_i^{\pi_t}(S) = \min\{v(S), B\}$. Διακρίνουμε τις εξής περιπτώσεις:

1. $v(S) \leq B$. Έστω π_t μία μετάθεση, τέτοια ώστε όλα τα αντικείμενα στο S έρχονται πρώτα και έστω α_{i^*} η maximizing συνάρτηση για το σύνολο S, δηλ. $v(S) = \alpha_{i^*}(S)$. Τότε, επειδή $\sum_{j \in S} \alpha_{i^*}(\{j\}) \leq B$, έχουμε $\beta_{i^*}^{\pi_t}(S) = \sum_{j \in S} \beta_{i^*}^{\pi_t}(\{j\}) = \sum_{j \in S} \alpha_{i^*}(\{j\}) = v(S)$.

2. $v(S) > B$. Έστω π_t μία μετάθεση, τέτοια ώστε όλα τα αντικείμενα στο S έρχονται πρώτα και έστω α_{i^*} η maximizing συνάρτηση για το σύνολο S, δηλ. $v(S) = \alpha_{i^*}(S)$. Έστω j^* το τελευταίο αντικείμενο στη μετάθεση π_t για το οποίο $\sum_{r:\pi_t(r)\leq\pi_t(j^*)} \alpha_{i^*}(\{r\}) \leq B$. Τότε:

$$\sum_{r:\pi_t(r)\leq\pi_t(j^*)} \beta_{i^*}^{\pi_t}(\{r\}) = \sum_{r:\pi_t(r)\leq\pi_t(j^*)} \alpha_{i^*}(\{r\})$$

Για τα επόμενα αντικείμενα $z \in S$ στη μετάθεση π_t,έχουμε ότι $\beta_{i^*}^{\pi_t}(\{z\}) = \max\{B - \sum_{k:\pi_t(k)<\pi_t(z)} \alpha_i(\{k\}), 0\}$. Στην πραγματικότητα, το πρώτο αντικείμενο μετά το j^* θα συμπληρώσει την επιπλέον αξία, τέτοια ώστε να έχουμε: $\sum_{k:\pi_t(k)\leq\pi_t(j^*)+1} \beta_{i^*}^{\pi_t}(\{j\}) = B$, και για όλα τα επόμενα αντικείμενα q θα έχουμε $\beta_{i^*}^{\pi_t}(\{q\}) = 0$. Τότε, $\sum_{j\in S} \beta_{i^*}^{\pi_t}(\{j\}) = B$.

- (subadditive) Έστω v μία subadditive συνάρτηση. Τότε, από τον ορισμό του subadditivity, για τα σύνολα T, S έχουμε:

$$v\left(S \cup T\right) \leq v\left(T\right) + v\left(S\right)$$

Έχουμε τις εξής περιπτώσεις:

1. Αν $\bar{v}(S \cup T) = v(S \cup T) < B$. Τότε έχουμε σίγουρα ότι $\bar{v}(S) = v(S) < B$ και ότι $\bar{v}(T) = v(T) < B$. Τότε, $\bar{v}(S \cup T) = v(S \cup T) \leq v(S) + v(T) = \bar{v}(S) + \bar{v}(T)$, όπου η ανισότητα προκύπτει από το subadditivity της v.

2. Αν $\bar{v}(S \cup T) = B < v(S \cup T)$. Έπειτα, παίρνουμε τις εξής περιπτώσεις:
   (α) Αν $\bar{v}(S) = B < v(S), \bar{v}(T) = B < v(T)$. Τότε, $\bar{v}(S \cup T) = B \leq 2B = \bar{v}(S) + \bar{v}(T)$.

   (β) Αν $\bar{v}(S) = B < v(S), \bar{v}(T) = v(T) < B$. Τότε, $\bar{v}(S \cup T) = B \leq B + v(T) = \bar{v}(S) + \bar{v}(T)$, όπου η ανισότητα προκύπτει από το ότι το liquid valuation είναι μη αρνητικό.

   (ς) Αν $\bar{v}(S) = v(S) < B, \bar{v}(T) = B < v(T)$. Τότε, $\bar{v}(S \cup T) = B \leq v(S) + B = \bar{v}(S) + \bar{v}(T)$, όπου η ανισότητα προκύπτει από το ότι το liquid valuation είναι μη αρνητικό.

   (δ) $\bar{v}(S) = v(S) < B, \bar{v}(T) = v(T) < B$. Τότε, $\bar{v}(S \cup T) = B \leq v(S \cup T) \leq v(S) + v(T) = \bar{v}(S) + \bar{v}(T)$, όπου η τελευταία ανισότητα προκύπτει από το γεγονός ότι η v είναι subadditive.

Ωστόσο, έχουμε το εξής πρόβλημα, όπως φαίνεται στο παρακάτω παράδειγμα:

Παράδειγμα: Φανταστείτε ένα παίχτη με $B = 2$ και δύο αντικείμενα a και b διαθέσιμα στις τιμές $p_a = 2$ και $p_b = 1$. Αν υποθέσουμε ότι το valuation του είναι $v(\{a\}) = v(\{a, b\}) = 10$,

$v(\{b\}) = 2$ (και άρα το liquid valuation είναι $\bar{v}(\{a\}) = \bar{v}(\{b\}) = \bar{v}(\{a, b\}) = 2$), τότε θέλει το a στην τιμή 2, έχοντας ωφέλεια 8. Όμως, το demand query για το liquid valuation $\bar{v}$ επιλέγει το b, που του δίνει ωφέλεια 1. Συνεπώς, ο παίχτης έχει κίνητρο να πει ψέματα στο demand query oracle.

Το παράδειγμα αυτό αποτυπώνει ότι κάθε παίχτης θέλει να μεγιστοποιήσει την ωφέλειά του, ανεξάρτητα με το πόσα χρήματα έχει να διαθέσει. Συνεπώς, διαλέγει το σύνολο των αντικειμένων που θέλει μέσω του BCDQ$(v, U, \vec{p}, B)$ που του επιστρέφει, από ένα σύνολο αντικειμένων U, το σετ:

$$S_i = \arg\max_{S \in U} \left\{ v_i(S) - p\left(S\right) \middle| p(S) \leq B_i \right\}$$

Το πρόβλημα είναι ότι δε φαίνεται ξεκάθαρα να μπορούμε να πάρουμε μία αντίστοιχη σχέση για το liquid valuation σαν αυτή που δίνει το DQ$(\min\{v, B\}, U, \vec{p})$, δηλαδή:

$$\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$$

για κάθε $T \subseteq U$. Ωστόσο, έχουμε το παρακάτω λήμμα:

Λήμμα 2: Έστω $S \subseteq U$ το σύνολο που επιστρέφει το BCDQ για κάποιον παίχτη με valuation v και budget B. Τότε, για οποιοδήποτε σύνολο $T \subseteq U$, ισχύουν τα εξής:

1. $\bar{v}(S) \geq \bar{v}(T) - p(T)$
2. $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$.

Απόδειξη: Θα αποδείξουμε τους δύο ισχυρισμούς ξεχωριστά. Για την πρώτη σχέση, αν $p(T) > B$, τότε το δεξί μέλος την ανισότητας θα είναι αρνητικό και άρα η ανισότητα ισχύει τετριμμένα. Συνεπώς, θα εστιάσουμε στην περίπτωση που $p(T) \leq B$. Έχουμε τις εξής περιπτώσεις:

1. ($\bar{v}(S) = v(S)$ και $\bar{v}(T) = v(T)$.) Άρα, $B \geq v(T)$. Το σύνολο T εξετάστηκε από το query και παρ' όλα αυτά επέστρεψε το σύνολο S. Συνεπώς: $\bar{v}(S) \geq \bar{v}(S) - p(S) = v(S) - p(S) \geq v(T) - p(T) = \bar{v}(T) - p(T)$.
2. ($\bar{v}(S) = B$ και $\bar{v}(T) = B$) Τότε, η ανισότητα ισχύει τετριμμένα διότι: $B \geq B - p(T)$ και οι τιμές είναι μη αρνητικές.
3. ($\bar{v}(S) = B$ και $\bar{v}(T) = v(T)$) Η ανισότητα ισχύει αφού: $B \geq B - p(T) \geq v(T) - p(T) = \bar{v}(T) - p(T)$.
4. ($\bar{v}(S) = v(S)$ και $\bar{v}(T) = B$) Άρα, $B \leq v(T)$. Το σύνολο T εξετάστηκε από το query και παρ' όλα αυτά επέστρεψε το σύνολο S. Συνεπώς: $\bar{v}(S) \geq \bar{v}(S) - p(S) = v(S) - p(S) \geq v(T) - p(T) \geq B - p(T) = \bar{v}(T) - p(T)$.

Αυτό ολοκληρώνει την απόδειξη της πρώτης σχέσης.

Για τη σχέση 2, επειδή το S είναι το σετ που επέστρεψε το BCDQ, τότε μπορούσε να αγοραστεί: $\bar{v}(S) \geq p(S)$. Αθροίζοντας την ανισότητα αυτή στη σχέση 1, έχουμε ότι: $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$.

Έχοντας αυτές τις σχέσεις, μπορούμε να πάρουμε για το LW ίδιας τάξης προσεγγίσεις με

Algorithm 1 KV-Mechanism for Liquid Welfare

Fix an ordering π of bidders and set $U_1 = U$.
Set initial prices $p_1^{(1)} = \cdots = p_m^{(1)} = \frac{L}{4m}$.
for each bidder $i = 1, \ldots, n$ according to π **do**
    Let $S_i = \text{BCDQ}(v_i, U_i, \vec{p}^{\,(i)}, B_i)$
    With probability q, allocate $R_i = S_i$ to i and set $U_{i+1} = U_i \setminus S_i$. Otherwise, set $U_{i+1} = U_i, R_i = \emptyset$.
    Update prices $\forall j \in S_i$: $p_j^{(i+1)} = 2p_j^{(i)}$.
end for

αυτές που δίνουν posted-price μηχανισμοί για το SW. Παρακάτω φαίνεται πως μετασχηματίζεται ο αλγόριθμος των Krysta και Vöcking [41] με τη χρήση του BCDQ. Η αναλυτική απόδειξη φαίνεται στο section 6.4. Την ίδια μετατροπή μπορούμε να κάνουμε και στο [29] για το Bayesian setting, παίρνοντας πάλι ίδιας τάξης προσέγγιση.

Στη συνέχεια, δείχνουμε ότι η προσέγγιση των Lu και Xiao [46] για το large market assumption δεν εφαρμόζεται στην περίπτωση των συνδυαστικών δημοπρασιών με αδιαίρετα αντικείμενα και ορίζουμε με διαφορετικό τρόπο ένα competitive market assumption. Πιο συγκεκριμένα, οι Lu και Xiao [46] θεώρησαν την προϋπόθεση $B_i \leq \frac{OPT}{m \cdot c}$, όπου OPT η βέλτιστη τιμή του LW και c μια μεγάλη σταθερά για συνδυαστικές δημοπρασίες με διαιρέσιμα αντικείμενα. Ωστόσο, στην περίπτωση των αδιαίρετων αντικειμένων κάτι τέτοιο δεν είναι εφικτό, αφού το πολύ m παίχτες θα πάρουν κάποιο αντικείμενο και άρα $OPT \leq mB_{max}$. Τότε, η σχέση αυτή δίνει $B_{max} \leq \frac{B_{max}}{c}$ για $c > 1$, πράγμα άτοπο.

Η δική μας προσέγγιση στον ορισμό του competitive market βασίζεται στην ιδέα ότι αν αφαιρέσουμε ένα τυχαίο σύνολο από $n/2$ παίχτες, έστω $\mathbb{S}$ τότε με πιθανότητα τουλάχιστον $1 - \delta$, όπου $\delta < \frac{1}{2}$, η βέλτιστη λύση στο εναπομείναν σύνολο $\mathbb{T}$ θα είναι συγκρίσιμη με τη βέλτιστη λύση του προβλήματος. Πιο συγκεκριμένα, έχουμε:

Ορισμός: Έστω $0 \leq \varepsilon < 2$ και σταθερά $\delta \geq 0$. Μία αγορά καλείται (ε, δ) - Competitive, αν, αφαιρώντας ένα τυχαίο σύνολο $\mathbb{S}$, από $\frac{n}{2}$ παίχτες, για το εναπομείναν σύνολο $\mathbb{T}$, ισχύει:

$$\mathbb{P}\left[O\bar{P}T_{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) \cdot O\bar{P}T\right] \geq 1 - \delta \tag{1.3}$$

όπου με $O\bar{P}T_{\mathbb{T}}$ συμβολίζουμε το βέλτιστο LW στο σύνολο $\mathbb{T}$.

Λόγω συμμετρίας του συνόλου που αφαιρέσαμε και του εναπομείναντος, με πιθανότητα τουλάχιστον $1 - 2\delta$, και τα δύο σύνολα έχουν λύσεις συγκρίσιμες με τη βέλτιστη, και άρα μπορούμε να υπολογίσουμε τις τιμές από το $\mathbb{S}$ και να τις προσφέρουμε στους παίχτες του $\mathbb{T}$. Πιο αναλυτικά, η ιδέα αυτή παρουσιάζεται στο section 6.6.

Λήμμα 3: Έστω $\mathcal{C} = \left\{ j | q_j^{\mathbb{T}} > \frac{\bar{v}(\mathcal{A}_j^{\mathbb{S}})}{\beta} \right\}$ για σταθερά $\beta > 1$. Τότε, $\sum_{j \in \overline{\mathcal{C}}} q_j^{\mathbb{T}} \leq \frac{\varepsilon}{2(\beta - 1)} O\bar{P}T$ και $\sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} \geq \frac{\beta(2-\varepsilon)-2}{2(\beta-1)} O\bar{P}T$.

Απόδειξη: Από τον ορισμό 6.4, ισχύει με σταθερή πιθανότητα ότι:

$$O\bar{P}T \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \sum_{j \in \overline{\mathcal{C}}} q_j^{\mathbb{T}} = \sum_{j \in U} q_j^{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) \cdot O\bar{P}T$$

Algorithm 2 Competitive Market (CM) Αλγόριθμος

Divide the bidders into sets $\mathbb{S}, \mathbb{T}$ uniformly at random, s.t., $|\mathbb{S}| = \frac{n}{2} = |\mathbb{T}|$.
Run the greedy algorithm $\mathcal{A}$ for bidders in $\mathbb{S}$ and denote the solution obtained by $\mathcal{A}^{\mathbb{S}}$.
for $j \in U$ **do**
    Set $p_j = \frac{1}{2\beta}\bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right)$, where $\beta > 1$ is a constant
end for
Fix an internal ordering of bidders in $\mathbb{T}$, π, and set $U_1 = U$.
for each bidder $i \in \mathbb{T}$ arriving according to π **do**
    Let $S_i = \text{BCDQ}(v_i, U_i, \vec{p})$.
    Set $U_{i+1} = U_i \setminus S_i$.
end for

. Έστω $\mathbb{S}_{\overline{\mathcal{C}}} \subseteq \mathbb{S}$ το σύνολο των παιχτών που παίρνουν τα non-competitive αντικείμενα από τον άπληστο αλγόριθμο $\mathcal{A}$ όταν τρέχει στο σύνολο $\mathbb{S}$. Τότε, στο επαυξημένο σετ $\mathbb{T} \cup \mathbb{S}_{\overline{\mathcal{C}}}$, υπάρχει ανάθεση $\mathcal{Q}$ με liquid valuation

$$\bar{v}(\mathcal{Q}) \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \sum_{j \in \overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \tag{1.4}$$

και άρα:

$$\begin{aligned} O\bar{P}T \geq \bar{v}(\mathcal{Q}) \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \sum_{j \in \overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \geq \sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} + \beta \sum_{j \in \overline{\mathcal{C}}} q_j^{\mathbb{T}} \\ \geq \left(1 - \frac{\varepsilon}{2}\right) O\bar{P}T + (\beta - 1) \sum_{j \in \overline{\mathcal{C}}} q_j^{\mathbb{T}} \end{aligned}$$

Μετά από πράξεις, παίρνουμε:

$$\sum_{j \in \mathcal{C}} q_j + \frac{\varepsilon}{2(\beta - 1)} O\bar{P}T \geq \sum_{j \in U} q_j^{\mathbb{T}} \geq \left(1 - \frac{\varepsilon}{2}\right) O\bar{P}T$$

Συνεπώς, για τα αντικείμενα στο $\mathcal{C}$ ισχύει:

$$\sum_{j \in \mathcal{C}} q_j^{\mathbb{T}} \geq \frac{\beta(2 - \varepsilon) - 2}{2(\beta - 1)} O\bar{P}T$$

Στο επόμενο λήμμα, δίνουμε ένα κάτω φράγμα για τη συνεισφορά των competitive αντικειμένων στη λύση του άπληστου αλγορίθμου.

Λήμμα 4: $\sum_{j \in \mathcal{C}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \geq \frac{2(\beta-1) - \varepsilon \cdot (3\beta - 1)}{4(\beta - 1)} O\bar{P}T$.

Απόδειξη: Συνδυάζοντας την ανισότητα (1.4) και το Λήμμα 3 παίρνουμε ότι:

$$\sum_{j \in \overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \leq \frac{\beta\varepsilon}{2(\beta - 1)} O\bar{P}T \tag{1.5}$$

Ο αλγόριθμος $\mathcal{A}$ δίνει μία 2-προσέγγιση για το βέλτιστο LW για το σετ $\mathbb{S}$, συνεπώς έχουμε:

$$\sum_{j \in \mathcal{C}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) + \sum_{j \in \overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \geq \frac{1}{2} O\bar{P}T_{\mathbb{S}} \geq \frac{1-\frac{\varepsilon}{2}}{2} O\bar{P}T \tag{1.6}$$

Συνδυάζοντας τις 2 τελευταίες σχέσεις, παίρνουμε το αποτέλεσμα.

Θεώρημα: Ο CM Αλγόριθμος είναι φιλαλήθης και και πετυχαίνει $O(1)$ προσέγγιση του βέλτιστου Liquid Welfare. Πιο συγκεκριμένα:

$$\mathbb{E}\left[\bar{v}\left(\mathcal{S}\right)\right] \geq (1-2\delta) \cdot \frac{2(\beta-1) - \varepsilon \cdot (3\beta - 1)}{16\beta(\beta-1)} OPT$$

.

Chapter 2

Introduction

Imagine that you are a social planner wanting to auction-off the seats of a local stadium in an extremely wealthy neighborhood (i.e., people have *no* budget constraints for the seats) for a big concert. As a social planner, your goal is to allocate the seats in a way that maximizes (or at least approximates as closely as possible) the happiness of the people interested in these seats. However, different people have different seat preferences; some people are happy with two consecutive seats anywhere in the stadium, some might be happy with only one seat in front of the stage, and some might want a whole row. Phrased in mechanism design language, this is a *Combinatorial Auction*, where you seek to optimize the *Social Welfare* by a truthful mechanism. Combinatorial Auctions, like the one above, appear in many contexts (e.g., spectrum auctions, network routing auctions [38], airport time-slot auctions [53], etc.) and have been extensively studied by both Economists and Computer Scientists (see e.g.,[12] for a survey).

As if this problem was not hard enough to solve, imagine that you find out two unfortunate events; the stadium is in fact at a working-middle class neighborhood (i.e., people do have budget constraints) and your boss is concerned about the effect of budget constraints on potential revenue. Now, the objective function should balance between the willingness and the ability of the people to pay for their seats. Motivated by usual discrepancies between auction participants' ability and willingness to pay, Dobzinski and Leme [19] introduced the notion of *Liquid Welfare*, which is the minimum of an agent's budget and valuation for a bundle of goods. As such, maximizing the Liquid Welfare achieves a reasonable compromise between social efficiency and potential for revenue extraction (which is constrained by the budgets).

It is clear that in such complex and unprincipled systems the people's preferences may coincide, in a way that they cannot all be fulfilled at the same time, and, as an outgrowth, they behave in a selfish manner to get the maximum satisfaction, ignoring whether their actions disappoint other people or not. This context comprises a game, where many people, the so called *agents*, interact with each other in order to satisfy their own desire. These agents have their private strategy and their sole aim is to maximize their "happiness", the so called *utility*.

Game Theory is the field that tries to quantify in a principled way the interactions between the agents, calculating the benefits and the losses of every possible action and providing solutions that try to maximize every agent's payoff. However, these solutions may be inefficient for the whole welfare.*Mechanism design*- the science of rule-making-tries to fix this problem, providing guarantees that agents behave in such a manner, that the final solution is efficient for the whole system and as a result can be thought as *reverse game theory*. In this field, we design procedures that affirm that no agent has incentive to lie about their preference, and thus, maximize their ulitity by acting in a truthful manner. Consequently, we are able to predict the action of the players and, therefore, to reason about the outcome of the mechanism. In other words, we want to design systems with strategic participants that have good performance guarantees.

Informally, a mechanism is characterized by a set of feasible outcomes $\mathcal{O}$, an allocation rule f and a payment rule $\vec{p} = (p_1, \ldots, p_n)$. Each agent has a private *valuation* $v_i : \mathcal{O} \to \mathbb{R}_{\geq 0}$ and tries to maximize its *utility*, which is defined as the valuation for the outcome ω minus the payment, i.e. $u_i = v_i(\omega) - p_i$. This utility model is called *quasi-linear*. From now on, we will assume that the utilities of the agents are quasi-linear, unless mentioned otherwise. We call a mechanism truthful, if every agent has no incentive to lie about their private information, no matter how the rest of the agents play. This means that every agent maximizes its utility by reporting his true valuation, independently of what others' strategies are. Therefore, the payment rule is essential for succeeding the truthfulness of the mechanism.

The performance is measured, in the majority of the mechanisms in the literature, by the total utility or *Social Welfare*, which is defined as the sum of every agent's valuation on the outcome of the mechanism $\omega \in \mathcal{O}$, i.e. $\sum_{i=1}^{n} v_i(\omega)$. At this point, we have to mention that problems in Mechanism Design are either single-parameter, where the preference of each agent consists of only one number, or multi-parameter. In this diploma thesis, we focus on a certain type of mechanisms, the so called *Auctions*, which are mechanisms specifically for the exchange of goods and money, and has been studied by both Computer Scientists and Economists.

Single-Parameter Auctions

Imagine that you want to auction-off an item and there are n available buyers. Every buyer i has a private valuation v_i for the item and wants to maximize their utility, i.e. $u_i = v_i - p$, if buyer i gets the item, where p is the price he has to pay and 0 otherwise. Your goal is to give the item to the buyer who wants it more, i.e. the buyer with the highest valuation. Now, we are going to compare two different payment rules, that do not affirm truthfulness. First, we assume that the item is given for free, meaning that $p = 0$. Then, every buyer would misreport their true valuation, by announcing a much higher number, since they only have better chances of getting the item. Second, we assume that if buyer i gets the item, he has to pay the number he announced to the auctioneer. Then, it s clear, that every buyer would announce a lower number, as their utility stays 0 by reporting their true valuation, whether they acquire the item or not. This latter form is known as *first-price* auction and its performance has been studied extensively in

The solution to this problem is due to Vickrey [64], who introduced the so called *Second-Price Auction*. In this mechanism, the bidder with the highest bid gets the item and has to pay a price equal to the second highest bid. The truthfulness of the Second-Price auction is presented in Section 3.2.1.

In the case of multi-unit auctions, where there are multiple copies of the same item, Myerson [49] proved that a mechanism is truthful if and only if is monotone, meaning that the amount of staff an agent is allocated is monotone to his bid. At the same time, the payment rule is unique (see Section 3.2.2) and, intuitively, each agent pays for each unit of item the minimum report needed to win this unit. As the single-parameter environment is well understood, we will focus on the multi-parameter environment and *Combinatorial Auctions*.

Multi-Parameter Auctions

Now, assume that we have to face a more general problem, for example there are multiple items to be auctioned. This setting is known as a Combinatorial Auction. A Combinatorial Auction (CA) consists of a set U of m items to be allocated to n bidders. Each bidder i has a valuation function $v_i : 2^U \to \mathbb{R}_{\geq 0}$. Valuation functions, v, are assumed to be non-decreasing, i.e., $v(S) \leq v(T)$, for all $S \subseteq T \subseteq U$, and normalized $v(\emptyset) = 0$. For the objective of Social Welfare (SW), the goal is to compute a partitioning $\mathcal{S} = (S_1, \ldots, S_n)$ of the set of items, U, that maximizes $v(\mathcal{S}) = \sum_{i=1}^{n} v_i(S_i)$.

We focus on Combinatorial Auctions with submodular, XOS or subadditive bidders. A set function $v : 2^U \to \mathbb{R}_{\geq 0}$ is *submodular* if for every $S, T \subseteq U$, $v(S) + v(T) \geq v(S \cap T) + v(S \cup T)$ and *subadditive* if $v(S) + v(T) \geq v(S \cup T)$. A set function v is *XOS* (a.k.a. *fractionally subadditive*, see [28]) if there exist additive functions $w_k : 2^U \to \mathbb{R}_{\geq 0}$ such that for every $S \subseteq U$, $v(S) = \max_k\{w_k(S)\}$. The class of submodular functions is a proper subset of the class of XOS functions, which is a proper subset of the class of subadditive functions.

The question is how should an appropriate payment rule look like in a Combinatorial Auction in order to to retain truthfulness. In fact, there is a mechanism that terminates with the optimal solution for general valuation functions in a truthful way. This mechanism was introduced by Vickrey [64], Clarke [8] and Groves [35] and is known as VCG mechanism. VCG is the unique truthful welfare-maximizing mechanism and is applied in general mechanism design environments. The idea behind this mechanism is to associate the welfare maximization with the maximization of the utility of each bidder and charge him his externality, i.e. $p_i = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega) - \sum_{j \neq i} v_j(\omega^*)$, where ω^* is the optimal solution (see Section 3.3.3 for further details). However, VCG cannot be implemented in polynomial time for most valuation function classes and as a result, we have to look for approximation mechanisms. One would may ask why not compute an approximate VCG solution and calculate the prices according to it. The answer is that any price calculation based on approximation would result to non-truthful mechanisms. As far as the communication complexity of the mechanism is concerned, we assume that the auctioneer has oracle access to the players' utilities and that he can ask queries to them, since each bidder has to announce exponentially-sized information to the mechanis. The two main categories of queries are the value and the demand queries. In this thesis, we focus on Combinatorial Auctions with demand queries, as this kind of queries is a more powerful tool and allows the achievement of much better approximation guarantees.

Combinatorial Auctions via Value Queries

In mechanisms that operate with value queries, valuations are given via black boxes. In a value query, the mechanism defines a subset $S \subseteq U$ and the player i returns $v_i(S)$. As optimization problem, forgetting incentives, the computation of a near-optimal welfare maximizing allocation with value queries is still a challenging task; Khot et al. [40] proved that there is no approximation algorithm that uses *polynomially* many value queries and approximates the optimal welfare by a factor better than $(1 - 1/e)$, unless $P = NP$. Based on this result, Vondrak [65] developed an algorithm which achieves a $(1 - 1/e)$-approximation for the submodular welfare problem with polynomial number of value queries. Furthermore, Mirrokni et al. [48] proved that for any fixed $\varepsilon > 0$, achieving an approximation ratio of $(1 - \frac{1}{e} + \varepsilon)$ for welfare maximization with submodular bidders requires an exponential number of value queries. For the case of subadditive valuations, Dobzinski [20] introduced an $\frac{1}{\sqrt{m}}$-approximation algorithm, using a polynomial number of value queries. Later, Mirrokni et al. [48] proved that for any fixed $\varepsilon > 0$, achieving an approximation ratio of $\frac{1}{m^{\frac{1}{2}-\varepsilon}}$ for welfare maximization with subadditive bidders requires an exponential number of value queries. Dobzinski [15] proved that any truthful mechanism for submodular Combinarorial Auctions with approximation ratio better than $m^{\frac{1}{2}-\varepsilon}$ must use *exponentially* many value queries.

Combinatorial Auctions via Demand Queries

In mechanisms that operate with demand queries, valuations are given via black boxes. In a demand query $\mathrm{DQ}(v_i, U, \vec{p})$, the mechanism presents price vector $\vec{p}$ for the goods and each player i returns the set S_i of goods that maximizes his utility, i.e. $S_i = \arg\max_{T \in U} \{v_i(T) - p(T)\}$. In the *worst-case* setting, where we do not make any further assumptions on bidders' valuations, Dobzinski et al. [21] presented the first truthful mechanism with a non-trivial approximation guarantee of $O(\log^2 m)$. Dobzinski [17] improved the approximation ratio to $O(\log m \log\log m)$ for the more general class of subadditive valuations. Subsequently, Krysta and Vöcking [41] provided an elegant randomized online mechanism that achieves an approximation ratio of $O(\log m)$ for XOS valuations. Dobzinksi [16] broke the logarithmic barrier for XOS valuations, by providing an approximation guarantee of $O(\sqrt{\log m})$. We highlight that accessing valuations through demand queries is essential for these strong positive results. In the *Bayesian* setting, bidder valuations are drawn as independent samples from a known distribution. Feldman et al. [29] showed how to obtain item prices that provide a constant approximation ratio for XOS valuations. In Chapter 5 we analyze these mechanisms' components, presenting the way they operate.

In fact, these mechanisms are inspired from the notion of *clearing prices*, for which supply equals demand. We say that an allocation $S = (S_1, \ldots, S_n)$ and a price vector $\vec{p}$ are in Walrasian Equilibrium if (i) $S_i \in \arg\max_{T \subseteq U} \left\{v_i(T) - \sum_{j \in T} p_j\right\}$ and (ii) an item $j \in U$ is unsold, only if $p_j = 0$. According to the First Welfare Theorem, if $(S, \vec{p})$ is a Walrasian Equilibrium, then S is a welfare-maximizing allocation. The case of Walrasian Equilibrium in auction format has been first studied through *ascending price* auctions (see e.g. [47, 10, 1, 2]). Kelso and Crawford [10] introduced an ascending price auction with use of demand queries that terminates at a Walrasian Equilibrium for *gross substitutes* valuations. An agent is said to have a *gross substitutes* valuation if, whenever the prices of some items increase and the prices of other items remain constant, the agent's demand for the items whose price remain constant weakly increases. However, Gul and Staccheti [36] proved that *gross substitutes* is the largest valuation class, for which a Walrasian Equilibrium is guaranteed to exist (see Chapter 4 for further details). As a result, for more general valuation functions, such as submodular, XOS, or subadditive, where clearing prices are not guaranteed to exist, the above mechanisms tried to compute approximate clearing prices, i.e. prices that a posted-price mechanism would result with high welfare. In other words, these mechanisms try to find prices, neither too low, nor too high, such that there exist an allocation S with $\sum_{j \in S} p_j \geq a \cdot v(\text{OPT})$.

Value vs Demand Queries

It is apparent that demand queries are much more powerful than value queries. In fact, a value query can be simulated by polynomially many value queries, but exponentially many value queries may be required to simulate a single demand query, as Blumrosen and Nisan [5, Lemma 3] proved. It gets more clear with the following example: Assume there are 2 bidders and m items. The first bidder has valuations $2|S|$ for every set S, except for a set H of $\frac{m}{2}$ items with valuation $2|H| + 2$, and the second bidder has valuation $2|S| + 1$ for each set S. The optimal solution is to assign H to the first bidder and the rest $\frac{m}{2}$ items to the second bidder, with a total welfare of $2m + 3$. For price vector $\vec{p} = (p_1, \ldots, p_m)$ with $p_j = 2 + \varepsilon$ for all items j, the demand query of the first bidder returns set H. However, with the use of value queries, unless all the sets of $\frac{m}{2}$ items be queried, the optimal allocation cannot be determined. Therefore, $\Omega(2^m)$ value queries will be needed in the worst case. As this example highlights, a demand query can export and communicate exponentially larger amount of information than a value query, as it returns the utility maximizing set along all possible 2^m different subsets.

Budget Feasible Mechanisms

Until now, the mechanisms presented did not take into account a crucial parameter in real life application, the budget constraints. However, when bidders are budget constrained, the Social Welfare cannot be approximated by a factor better than n, where n is the number of agents. The classic VCG approach does not hold anymore, since the utilities of the bidders stop being quasi-linear. In fact, they are formed as:

$$u_i = \begin{cases} v_i - p, & \text{if } p \leq B_i \\ -\infty & \text{otherwise} \end{cases}$$

where B_i is bidder's i budget. Nonetheless, a reverse approach has been studied extensively, in the so called *Procurement Auctions*. In a Procurement Auction there are n sellers, competing each other to be preferred by a single buyer. Budget feasible mechanisms refer to the case where the buyer is budget constrained. This topic was introduced by Singer [62], where he introduced a randomized $O(1)$ factor budget feasible mechanism that is universally truthful for the class of submodular functions. Moreover, Dobzinski et al. [22] proved a $O(\log^2 n)$ approximation for the case of subadditive functions. This results were improved by Bei et al. [4], where they gave a $O(\frac{\log n}{\log \log n})$ approximation mechanism for subadditive functions. Budget feasible mechanism design focuses on payment optimization in reverse auctions, a setting almost orthogonal to the setting we consider in our work, but we refer to it for completeness.

Liquid Welfare

Liquid Welfare was introduced as an efficiency measure for auctions when bidders are budget constrained in [19], since it was known that getting any non-trivial approximation for the SW in these cases is impossible. Moreover, Dobzinski and Leme [19] proved a $O(\log n)$ (resp. $(\log^2 n)$)-approximation to the optimal LW for the case of a single divisible item and submodular (resp. subadditive) bidders. Dobzinski and Leme [19] and Lu and Xiao [45] proved that the optimal LW can be approximated truthfully within constant factor for a single divisible good, additive bidder valuations and public budgets. Closer to our setting, Lu and Xiao [46] provided a truthful mechanism that achieves a constant factor approximation to the LW for multi-item auctions with divisible auctions, under a large market assumption. Under similar large market assumptions, Eden et al. [26] obtained mechanisms that approximate the optimal revenue within a constant factor for multi-unit online auctions with divisible and indivisible items, and a mechanism that achieves a constant approximation to the optimal LW for general valuations over divisible items. However, prior to our work, there was no work on approximating the LW in CAs (in fact, that was one of the open problems in [19]).

Our Results: Intuition and Contribution

Our aim is to extend these results to the objective of Liquid Welfare. To this end, we exploit the fact that most of the mechanisms above (and the mechanisms of [41] and [29], in particular) follow a simple pattern: first, by exploring either part of the instance in [41] or the knowledge about the valuation distribution in [29], the mechanism computes appropriate (a.k.a. *supporting*[1]) prices for all items. Then, these prices are "posted" to the bidders, who arrive one-by-one and select their utility-maximizing bundle, through a demand query, from the set of available items,as in Algorithm 3.

Algorithm 3 Core Mechanism

Fix an ordering π of bidders and set $U_1 = U$.
Set initial prices for the items: $\vec{p}^{(1)} = (p_1^{(1)}, \dots, p_m^{(1)})$.
for each bidder $i = 1, \dots, n$ according to π **do**
    Let $S_i = \mathrm{DQ}(v_i, U_i, \vec{p}^{(i)},)$
    With probability q, allocate S_i to i and set $U_{i+1} = U_i \setminus S_i$. Otherwise, set $U_{i+1} = U_i$.
    Update item prices to $\vec{p}^{(i+1)} = (p_1^{(i+1)}, \dots, p_m^{(i+1)})$.
end for

The technical intuition behind the high level approach above is nicely explained in [16, Section 1.2]. Let $\mathcal{O} = (O_1, \dots, O_n)$ be an optimal solution for the SW (in fact, any con-

[1] A price vector $\vec{p} = (p_1, \dots, p_m)$ *supports* allocation $\mathcal{S} = (S_1, \dots, S_n)$, if $v(\mathcal{S}) \geq \sum_j p_j$.

stant factor approximation suffices). The supporting price of item i in $\mathcal{O}$ is $q_j = w_k(\{j\})$, where w_k is the additive valuation determining the value $v_i(O_i)$ (recall that valuation functions are XOS). Intuitively, q_j is how much item j contributes to the social welfare of $\mathcal{O}$. Then, a price of $p_j = q_j/2$ for each item j is appropriate in the sense that a constant approximation to $v(\mathcal{O})$ can be obtained by letting the bidders arrive one-by-one, in an arbitrary order, and allocating to each bidder i her utility maximizing bundle, chosen from the set of available items by a demand query (see [16, Lemma 4.2]).

Hence, approximating the SW by demand queries boils down to computing such prices p_j. In the Bayesian setting, prices p_j can be obtained by drawing n samples from the valuation disribution and computing the expected contribution of each item j to a constant factor approximation of the optimal allocation (see Section 3 and Lemma 3.4 in [29]). Similarly, the idea of estimating the contribution of the items would work under some market uniformity assumption, as the one introduced in Definition 6.4. In the worst-case setting, if we assume integral and polynomially-bounded valuations (i.e., that $\max_i\{v_i(U)\} \leq m^d$, for some constant d), a uniform price for all items selected at random from $1, 2, 4, 8, \dots, 2^{d \log m}$ results in an logarithmic approximation ratio. Krysta and Vöcking [41] show how to estimate supporting prices online, by combining binary search and randomized rounding. Importantly, as long as each bidder does not affect the prices offered to her, this general approach results in (randomized universally) truthful mechanisms.

Towards extending the above approach and results to the LW, our first observation (Lemma 6.2) is that if a valuation function v is submodular (resp. XOS), then the corresponding liquid valuation function $\bar{v} = \min\{v, B\}$ is also submodular (resp. XOS). Then, one can directly use the mechanisms of e.g., [41, 16, 29] with valuation functions $\bar{v} = \min\{v, B\}$ and demand queries of the form: $\mathrm{DQ}(\min\{v, B\}, U, \vec{p})$ (i.e., wrt. the liquid valuation of the bidders) and obtain the same approximation guarantees but now for the LW. However, the resulting mechanisms are no longer truthful; bidders still seek to maximize their *utility* (i.e., value minus price) from the bundle that they get, subject to their budget constraint, rather than their *liquid utility* (i.e., liquid value minus price). Specifically, given a set of items U available at prices p_j, $j \in U$, a budget-constrained bidder i wants to receive the bundle $S_i = \arg\max_{S \subseteq U}\{v_i(S) - p(S) \,|\, p(S) \leq B_i\}$, and might not be happy with the bundle $S'_i = \arg\max_{S \subseteq U}\{\bar{v}_i(S) - p(S)\}$ computed by the demand query for the liquid valuation.

To restore truthfulness, we replace demand queries with *budget-constrained demand queries*. A budget-constrained demand query, denoted by $\mathrm{BCDQ}(v, U, \vec{p}, B)$, specifies a valuation function v, a set of available items U, a price p_j for each $j \in U$ and a budget B, and receives the set $S \subseteq U$ maximizing $v(S) - p(S)$, subject to $p(S) \leq B$, i.e., the set of available items that maximizes bidder's utility subject to her budget constraint.

To establish the approximation ratio, we first observe that the fact that liquid valuations are XOS suffices for estimating supporting prices, as in previous work on the SW. Additionally, we show that the bundles allocated by $\mathrm{BCDQ}(v, U, \vec{p}, B)$ approximately satisfy the efficiency guarantees on the liquid welfare and the liquid utility of the allocated bundles (see Lemma 6.3). Specifically, we observe that the approximation guarantees of mechanisms for the SW mostly follow from the fact that a demand query $\mathrm{DQ}(v, U, \vec{p})$

guarantees that for the allocated bundle S and for any $T \subseteq U$:

1. $v(S) - p(S) \geq v(T) - p(T)$
2. $v(S) \geq v(T) - p(T)$

In Lemma 6.3, we show that a budget-constrained demand query, $\mathrm{BCDQ}(v, U, \vec{p}, B)$, guarantees that for the allocated bundle S and any $T \subseteq U$,

1. $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$
2. $\bar{v}(S) \geq \bar{v}(T) - p(T)$.

Using this property, we can prove the equivalent of [16, Lemma 4.2] in Lemma 6.5) and also the approximation guarantees of the mechanisms in [41, 29] but for the LW.

Formalizing the intuition above, we obtain (i) a randomized universally truthful mechanism that approximates the LW within a factor of $O(\log m)$ (Section 6.4), and (ii) a posted-price mechanism that approximates the LW within a constant factor when bidder valuations are drawn as independent samples from a known distribution (Section 6.5). Both mechanisms assume XOS bidder valuations; the former is based on the mechanism of [41] and the latter on the mechanism of [29].

Motivated by large market assumptions often used in Algorithmic Mechanism Design (see e.g., [6, 26, 46] and the references therein), we introduce a competitive market assumption in Section 6.6. The main idea is that when there is an abundance of bidders, even if we remove a random half of them, the optimal LW does not decrease by much. Then, computing supporting prices for all items based on a randomly chosen half of the bidders, and offering these prices through budget-constrained demand queries to the other half, yields a universally truthful mechanism that approximates LW within a constant factor (Theorem 6.4).

Organization of the Thesis

In Chapter 3 we will make a brief introduction to *Mechanism Design*, giving fundamental definitions and theorems, such as the notions of *dominant strategy, truthfulness, DSIC mechanisms* and *Revelation Principle*. For the single-parameter case, we will deal with the *Second-Price Auction* and *Myerson's Lemma*, while for the multi-parameter setting, we will state the *VCG mechanism*, its power and limitations, and the definition of a *Combinatorial Auction*. Afterwards, we will introduce a different class of mechanism, the *Posted-Price mechanisms* and finally, we will make a short notice on the *Linear Programming* machinery.

In Chapter 4 we will introduce a new notion of auctions, the *Ascending Price* auctions and define a new concept of truthful revelation, *sincere bidding* and *ex-post Nash equilibrium*. Furthermore, we are going to deal with the case of market equilibrium, the so called *Walrasian Equilibrium* and the prerequisites under which, ascending price auctions terminate at a *Walrasian Equilibrium*. Finally, we will express the *Combinatorial Auctions* via *Linear Programming* and examine under which conditions the VCG mechanism can be implemented in polynomial time.

In Chapter 5 we will introduce the *supporting prices* and the structure of *Posted-Price* mechanisms in *Combinatorial Auctions*. In the worst-case and the Bayesian setting, we will present the most important results, providing a deeper explanation and intuition to their technics. More specific, we are going to categorize mechanisms according to the way they post the prices, and connect it with the notions of *Walrasian Equilibrium* and *Ascending-Price* auctions.

In Chapter 6 we will present an alternative notion of efficiency in auctions with budget-constrained bidders, the so called *Liquid Welfare*, introduced by Dobzinski and Leme [19]. Thus far, there was no work approximating the *Liquid Welfare* in *Combinatorial Auctions*, except the case of multi-item divisible setting with additive bidders under a large market assumption. However, we will extend the approximation results of *Social Welfare* in *Combinatorial Auctions* with submodular (or XOS) bidders for the measure of *Liquid Welfare*. Finally, we will define a new notion of *Competitive Markets* and give a constant approximation mechanism for XOS bidders.

Chapter 3

Basics of Mechanism Design

First, we are going to give some fundamental definitions of the area of mechanism design and present some basic mechanisms for single and multi -parameter environment.

3.1 Preliminaries

The basic setup: Assume that there are n agents participating in the mechanism and let $\mathcal{O}$ be the set of the feasible outcomes. Each agent i has a private valuation $v_i : \mathcal{O} \to \mathbb{R}_{\geq 0}$ and reports its bid $b_i : \mathcal{O} \to \mathbb{R}_{\geq 0}$ to the mechanism. After collecting the bids, the mechanism uses a social choice function $f : \vec{b} \to \mathcal{O}$, which maps the bidding profile, $\vec{b} = (b_1, \ldots, b_n)$, to an allocation, and a payment scheme $\vec{p} = (p_1, \ldots, p_n)$ for this allocation. A deterministic mechanism is defined by the pair (f, p).
Note: A mechanism that operates as described above, is called *direct revelation* An *indirect mechanism* is a function $M : A^n \to \mathcal{O}$, which maps an action vector to a feasible outcome. We will focus later on indirect mechanisms.

Notation: By $\vec{v}_{-i}$ we express the valuation profile of all agents, except for i, i.e. $\vec{v}_{-i} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$

Definition 3.1 (Quasi-linear Utility). *In a mechanism (f, p) we say that the utility functions are* quasi-linear *if:*

$$u_i = v_i(f(\vec{b})) - p_i(\vec{b}) \tag{3.1}$$

where $\vec{b}$ is the bid profile of the agents.

Note: For the rest of the analysis, we assume that all agents have quasi-linear utilities, until we introduce the notion of *Liquid Welfare.*

Definition 3.2 (Dominant Strategy). *A bidding strategy $b_i \in V_i$ is dominant, if it maximizes agent's i utility, regardless what others are doing. Formally:*

$$v_i\left(f(b_i, \vec{v}_{-i})\right) - p_i(b_i, \vec{v}_{-i}) \geq v_i\left(f(b_i', \vec{v}_{-i})\right) - p_i(b_i', \vec{v}_{-i}) \tag{3.2}$$

We call a mechanism truthfull, when truthtelling is the dominant strategy for every agent. Formally:

Definition 3.3 (Truthfulness, [51]). *Let (f, p) be a mechanism. Then, Mechanism (f, q) is* truthful *if for all $i \in [n]$ and for any v_i', it holds that:*

$$v_i\left(f(v_i, \vec{v}_{-i})\right) - p_i(v_i, \vec{v}_{-i}) \geq v_i\left(f(v_i', \vec{v}_{-i})\right) - p_i(v_i', \vec{v}_{-i}) \tag{3.3}$$

Our mechanisms in this work are going to be *randomized*, i.e., they are probability distributions over *deterministic* mechanisms. The incentives desiderata for randomized mechanisms are either *universal truthfulness* (when all the deterministic mechanisms are Dominant Strategy Incentive Compatible (DSIC)) or *truthfulness in expectation*(when bidders' *expected* utilities are maximized under truthful reporting of their private information). In this work, we are focusing on the former, stronger notion; the one of *universal truthfulness.*

Definition 3.4 (Universal Truthfulness). *Let $\left(\tilde{f}, \tilde{p}\right)$ be a randomized mechanism over a set of deterministic mechanisms $\left\{(f^1, p^1), \ldots, (f^k, p^k)\right\}$. Mechanism $\left(\tilde{f}, \tilde{p}\right)$ is* universally truthful *if for all $i \in [n], \kappa \in [k]$ and for any v_i', it holds that:*

$$v_i\left(f^{\kappa}(v_i, \vec{v}_{-i})\right) - p_i^{\kappa}(v_i, \vec{v}_{-i}) \geq v_i\left(f^{\kappa}(v_i', \vec{v}_{-i})\right) - p_i^{\kappa}(v_i', \vec{v}_{-i}) \tag{3.4}$$

When an agent enters the mechanism, it is important for him to know that whatever the outcome is, he can never have negative utility. The mechanisms that fulfill this property are called *individually rational.* Formally, we have:

Definition 3.5 (Individually Rational, [51]). *A mechanism (f, p) is* individually rational *if for all $i \in [n]$ and for all valuation profiles $\vec{v} = (v_1, \ldots, v_n)$, it holds:*

$$v_i\left(f(\vec{v})\right) - p_i(\vec{v}) \geq 0 \tag{3.5}$$

Definition 3.6 (DSIC). *A mechanism (f, p) is* DSIC (Dominant Strategy Incentive Compatible) *if it is truthful and individually rational.*

Assume that mechanism (f, p) is non-DSIC, in the sense that every agent has a dominant strategy, but it is not guaranteed that this strategy is truthtelling. The question that arrises is whether (f, q) can be simulated by another mechanism (f', p'), such that (f', p') is DSIC. The answer is positive and is called *Revelation Principle.* Formally, we have:

Theorem 3.1 (Revelation Principle). *For every mechanism (f, p) in which every agent has a dominant strategy, there is an equivalent direct-revelation DSIC mechanism (f', p').*

Proof. As defined above, every agent has a private valuation v_i and a dominant strategy $s_i(v_i)$. Hence, it means that each agent i would announce $s_i(v_i)$ to (f, p). Let us now construct an equivalent mechanism (f', p') that accepts each agent's bid b_i, operates the function s_i on the announced bid $b_i, \forall i \in [n]$ and then outputs the same allocation and payments as (f, p). Formally, we define $f'(\vec{b}) = f\left(s(\vec{b})\right)$ and $p'(\vec{b}) = p\left(s(\vec{b})\right)$, where $s(\vec{b}) = (s_1(b_1), \ldots, s_n(b_n))$.

Consequenlty, as agent i who has private valuation v_i and dominant strategy $s_i(v_i)$, would only diminish their utility by reporting a bid other than v_i and therefore, by possibly playing in (f, p) a strategy other than $s_i(v_i)$. Hence, mechanism (f', p') is DSIC.

□

By this time, we have not mention anything about the measure of efficiency of the mechanisms. It is reasonable that efficiency varies according to the mechanism's designer wishes. If a government is to decide about the realization or not of a public project, then the appropriate efficiency measure is the total welfare of the community. If a seller is to auction his car, he probably wants to get as much money as possible and wants to maximize their revenue. Formally, we have:

Definition 3.7 (Social Welfare)**.** *Let (f, p) a mechanism and $\mathcal{S} \in \mathcal{O}$ its output by the participation of n agents. Then,* Social Welfare *is defined as:*

$$\text{SW} = \sum_{i=1}^{n} v_i(\mathcal{S}) \tag{3.6}$$

On the other hand, it is logical that the revenue extracted from the agents is the sum of the prices each agent was charged. More formally, we define the *revenue* of the mechanism as follows:

Definition 3.8 (Revenue)**.** *Let (f, p) a mechanism and $\mathcal{S} \in \mathcal{O}$ its output by the participation of n agents. Then,* revenue *is defined as:*

$$\text{REV} = \sum_{i=1}^{n} p_i(\mathcal{S}) \tag{3.7}$$

In Chapter 6 we introduce another notion of effficiency, called *Liquid Welfare* which is defined when agents are restricted by their budget. Briefly, we have:

Definition 3.9 (Liquid Welfare)**.** *In a budgeted setting with n bidders, where each bidder i has budget B_i and valuation v_i, we define the Liquid Welfare of outcome $\omega \in \mathcal{O}$ by:*

$$LW = \sum_{i=1}^{n} \min\{v_i(\omega), B_i\} \tag{3.8}$$

Finally, we are going to define the concept of approximation and randomized algorithms, which are widely used in Mechanism Design.

Approximation in Mechanism Design. We say that a mechanism ρ-approximates the optimal solution if: $\text{ALG} \geq \rho \cdot v(\text{OPT})$, where $\rho \leq 1$, ALG is the solution of the mechanism and OPT the optimal solution.

Randomization in Mechanism Design. We say that a randomized mechanism ρ-approximates the optimal solution if: $\mathbb{E}[\text{ALG}] \geq \rho \cdot \text{OPT}$, where $\rho \leq 1$, ALG is the solution of the mechanism and OPT the optimal solution.

3.2 Single-Parameter Environments

3.2.1 Single-item Auctions

Suppose that you want to design an auction for an indivisible item, so as to maximize the *Social Welfare*, i.e. allocate the item to the agent that wants it more. As claimed before, we cannot give the item for free, nor charge the winning bidder his bid, as we incentivize the agents to misreport their true valuation and can lead to a really low welfare. Therefore, we have to find the appropriate payment rule, so as no agent can augment their utility by lying about their valuation. The solution to this problem is the so called *Second-Price Auction*, where the bidder with the highest bid gets the item and has to pay a price equal tothe second highest bid, and which was introduced by Vickrey.

Theorem 3.2. *The Second-Price Auction is DSIC.*

Proof. Let n be the number of bidders that participate in the auction. Fix an arbitrary player i an let v_i be its valuation. Our aim is to prove that i has dominant strategy to bid their true valuation, i.e. $b_i = v_i$. Let $\mathbf{b}_{-i}$ be the bidding profile of the rest of the agents and let L be the highest bid, among them, i.e. $L = \arg\max_{j \neq i} b_j$. Now, we are going to distinguish the following cases:

- $v_i < L$. If $b_i < v_i$ then the outcome of the auction remains the same and i does not get the item. If $b_i > v_i$, agent i compromises to exceed L and therefore, to acquire the item in the price of L, getting negative utility.

- $v_i \geq L$ If $b_i > v_i$ then the outcome of the auction remains the same and agent i gets the item in the price of L. If $b_i < v_i$, agent i compromises to fall behind L and therefore, to lose the item , getting zero utility. If they still win the item, then price will be still be L and their utility would remain the same.

Therefore, it is clear that agent i has no incentive to lie about their valuation and truthtelling is their dominant strategy. □

3.2.2 Multi-unit Auctions

Now, suppose that you have multiple or infinite copies of the same item. Second-Price Auction seems a powerless tool to be implemented in this setting. So, we have to think of a new feasible allocation and payment rule $f(\vec{b}), p(\vec{b}) \subseteq \mathbb{R}^n$. Agents have still to report their valuation v, which in this case is the valuation "per unit of staff". The solution to this setting was given by *Myerson*. Formally, we have:

Theorem 3.3 (Myerson's Lemma, [49]). *A mechanism (f, p) is DSIC if and only if, assuming for every bidder i with bid b_i and $\vec{b}_{-i}$, it holds:*

- *$f_i(b_i, \vec{b}_{-i})$ is non-decreasing in its bid b_i*

- *the unique payment rule is given by the formula:*

$$p_i(b_i, \vec{b}_{-i}) = \int_0^{b_i} z \frac{d}{dz} f_i(z, \vec{b}_{-i}) dz \tag{3.9}$$

3.3 Multi-Parameter Environments

In the previous sections, we discussed primarily for cases where agents should report to the mechanism a single valuation. What if we have to face a more general problem, for example there are multiple items to be auctioned? How should an appropriate payment rule look like in order to to retain truthfulness? One would reasonable think that, running separate *Second-Price Auctions* for each item, would result to DSIC mechanism. Hence, the answer is more complex than that. In fact, separate *Second-Price Auctions* is the solution if agents' valuations are additive, meaning that the valuation of a bundle of items equals to the sum of valuation of each item separately.

3.3.1 Combinatorial Auctions

At this point, we are ready to define the notion of a *Combinatorial Auction*. Informally, it is a special case of the multi-parameter environment defined previously, where valuation functions are defined on the 2^m subsets of the m different items, although there are $(n+1)^m$ different outcomes. The reason for this is that every agent is indifferent who gets an item, if it is not him.

Definition 3.10 (Combinatorial Auction). *A* Combinatorial Auction *(CA) consists of a set U of m items to be allocated to n bidders. Each bidder i has a valuation function $v_i : 2^U \to \mathbb{R}_{\geq 0}$. Valuation functions, v, are assumed to be non-decreasing, i.e., $v(S) \leq v(T)$, for all $S \subseteq T \subseteq U$, and normalized $v(\emptyset) = 0$. For the objective of Social Welfare (SW), the goal is to compute a partitioning $\mathcal{S} = (S_1, \ldots, S_n)$ of the set of items, U, that maximizes $v(\mathcal{S}) = \sum_{i=1}^{n} v_i(S_i)$.*

3.3.2 Valuation Classes

Now, we are going to give some basic definitions about the valuation classes, which determine the difficulty of the problem. Let v a valuation function and U the set of m items.

Definition 3.11 (Additive). *A set function $v : 2^U \to \mathbb{R}_{\geq 0}$ is* additive *if for every $S \subseteq U$:*

$$v(S) = \sum_{j \in S} v(\{j\}) \tag{3.10}$$

This is the least general class of valuation functions and entails that there are no dependencies between the items or the size of the set. This setting can be solved optimal through parallel second-price auctions.

Definition 3.12 (Gross Substitutes). *An agent is said to have a* gross substitutes *valuation if, whenever the prices of some items increase and the prices of other items remain constant, the agent's demand for the items whose price remain constant weakly increases.*

The above definition is informal, as it requires some more technical background, which is provided later in the thesis, and this class arises naturally as a necessity for the efficiency guarantees of an ascending price auction, presented later. Therefore, we give the formal definition in Section 4.5.

Definition 3.13 (Submodular). *A set function* $v : 2^U \to \mathbb{R}_{\geq 0}$ *is* submodular *if for every* $S, T \subseteq U$, *with* $T \subseteq S$ *and item* $j \notin S$*:*

$$v(S \cup \{j\}) - v(S) \leq v(T \cup \{j\}) - v(T) \tag{3.11}$$

Submodularity can be seen as the discrete analog of concavity and arises naturally in economic settings since it captures the property that marginal utilities are decreasing as we allocate more goods to a player.

Definition 3.14 (XOS). *A set function* v *is* XOS *(a.k.a.* fractionally subadditive*) if there exist additive functions* $w_k : 2^U \to \mathbb{R}_{\geq 0}$ *such that for every* $S \subseteq U$*:*

$$v(S) = \max_k \{w_k(S)\} \tag{3.12}$$

Definition 3.15 (Subadditive). *A set function* $v : 2^U \to \mathbb{R}_{\geq 0}$ *is* subadditive *if for every* $S, T \subseteq U,$*:*

$$v(S) + v(T) \geq v(T \cup S) \tag{3.13}$$

Subadditivity can be seen as complement-free market, as the combination of any two bundles of items does not increase their value.

The relation between the aformentioned classes is the following:

$$\textit{Additive} \subseteq \textit{Gross Substitutes} \subseteq \textit{Submodular} \subseteq \textit{XOS} \subseteq \textit{Subadditive} \tag{3.14}$$

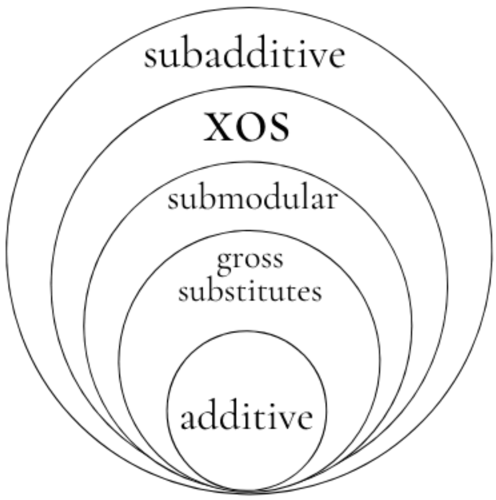

Figure 3.1: The relation between the valuation function classes.

3.3.3 VCG mechanism

Hence, there is a mechanism that solves much more complicated and general settings than the auctioning of multiple items, even for general valuation functions. Imagine the most abstract definition of a multi-parameter environment. Then, it should be defined like this:

- There are n agents.
- There is a finite set Ω of possible outcomes.
- Every agent has a private valuation $v_i(\omega)$, $\forall \omega \in \Omega$.

Now, we are ready to introduce a DSIC mechanism for the Welfare maximization of any multi-parameter environment, where the utilities are quasi-linear.

Theorem 3.4 (VCG mechanism, [64, 8, 35])**.** *In every general mechanism design environment, there is a DSIC welfare-maximizing mechanism (f, p), which is defined as follows:*

- $\omega^* = f(\vec{b}) = \arg\max_{\omega \in \Omega} \sum_{i=1}^{n} b_i(\omega)$
- $p_i(\vec{b}) = \max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) - \sum_{j \neq i} b_j(\omega^*)$.

(Intuitively, each agent is asked to pay its externality, since the first term in the RHS of the payment rule equals to the optimal solution, calculated on all the agents except for i, and the second term equals to the contribution of the rest of the agents in the optimal solution.)

Proof. Fix an arbitrary player i an let v_i be its valuation function. Our aim is to prove that i has dominant strategy to bid their true valuation function, i.e. $b_i = v_i$. Let $\vec{b}_{-i}$ be the bidding profile of the rest of the agents. Based on the allocation and payment rules, defined above, the utility of player i is defined as:

$$u_i = v_i(\omega^*) - p_i(\omega^*) = \left[v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) \right] - \max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) \tag{3.15}$$

It is clear, that agent i cannot influence the second term of the RHS of equation (1.9). However, the term $\left[v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) \right]$ is maximized when agent i reveals their true valuation profile, by the definition of the allocation rule, as misreporting can lead to another outcome with lower welfare. Therefore, it is dominant strategy for every agent i to bid their true valuation function. Moreover, it is easy to verify that truthtelling guarantees non-negative utilities for all agents, since the maximization space for the positive term of the RHS of (1.9) is only bigger than that of the negative term, and hence, the VCG mechanism is DSIC. □

However, as we will discuss now, VCG is not a panacea; imagine a *Combinatorial Auction* with n bidders and m items. As defined above, each bidder has to announce 2^m values, one for each bundle of items, assuming that valuations are not additive. So, for instance, for $m = 20$ each bidder has to communicate more than 1 million numbers to the auctioneer, and as a result, it is clear that it cannot be implemented in real life.
Even if we overlook the communication problem, another problem that appears is that VCG cannot be implemented in polynomial time, except for some very specific classes of valuation function, that we are going to define in the next chapter. As a result, the way to overcome this problem is to seek for near-optimal solutions, the so called *approximate solutions*.
Furthermore, although we proved that VCG is DSIC mechanism, it is vulnerable to collusion among the agents. In other words, it is not *group strategyproof.* The following example delineates this issue.

Example 3.1. *Think of a case with 3 agents and the mechanism wants to decide whether to build a road or not. The possible outcomes are "yes" and "no". Now, assuming* $v_1(yes) = v_2(yes) = 3$, $v_1(no) = v_2(no) = 0$ *and* $v_3(yes) = 0, v_3(no) = 5$, *then the outcome of the mechanism, implementing VCG, is "yes" and ends up with payments* $p_1 = p_2 = 2, p_3 = 0$. *Now, if agents* $1, 2$ *collude and report* $v_1'(yes) = v_2'(yes) = 5$, *then the outcome will remain the same, but the payments will decrease to 0, for both agents.*

The last example highlights one more problem, that VCG can have a really bad *revenue*. Finally, in many real-life situations, decisions about the allocation of resources must be made online, without information about the future. For example, think of a *Combinatorial Auction*, where bidders arrive sequentially; how can the seller be reassured that items will not end up early to agents with low valuations? The problem that arises is that VCG cannot be implemented in online settings, as it requires to compute an offline optimal solution.
Consequently, there have to be found another mechanisms to be implemented in *Combinatorial Auctions*. Until now, we have only talked about direct-revelation mechanisms, where the agents report all of their private information to the mechanism. Though, we have to introduce a new category of mechanisms, the so called *indirect* mechanism, where information is provided only on a "need-to-know" basis through queries.

3.4 Value and Demand Queries

As it has been clear, Combinatorial Auctions with m items have a strong disadvantage; the communication complexity. Each agent has 2^m different values, one for each bundle, except for the case of additive valuations. Of course, it is impossible for a mechanism to operate with exponential size of input, even in real life auctions with a small number of items. A solution to this problem is the implementation of queries, through which the mechanism can acquire the information needed. The two main categories are the *Value* and the *Demand* queries. More specifically, we have:

Definition 3.16 (Value Query (VQ))**.** *Let* S *be a set of items. Then, bidder* i*'s Value Query* (VQ(S)) *returns the value of the set:*

$$\mathrm{VQ}(S) = v_i(S) \tag{3.16}$$

The definition of a value query is simple enough; it just returns the value of a bundle. It is clear that this tool does not give great power to the mechanism, as it would require exponentially many value queries, in order to execute VCG. However, it allows the mechanism designer to define the way he learns the valuations. A more powerful tool is the demand query and is defined as follows:

Definition 3.17 (Demand Query (DQ))**.** *Let* U *be the set of items that are available and* $\vec{p}$ *their prices. Then, bidder* i*'s Demand Query* (DQ($v_i, U, \vec{p}$)) *returns set* $S_i \subseteq U$ *satisfying:*

$$S_i = \arg\max_{S \in U} \{v_i(S) - p(S)\} \tag{3.17}$$

It is apparent, that demand queries are much more powerful than value queries; they return the most profitable bundle between 2^m different bundles at specific prices. A naive way to simulate the function of demand queries through value queries would require to examine all 2^m possible bundles. However, it is necessary in many cases, as it is illustrated by the following example:

Example 3.2 ([5]). *Assume there are 2 bidders and m items. The first bidder has valuations $2|S|$ for every set S, except for a set H of $\frac{m}{2}$ items with valuation $2|H|+2$, and the second bidder has valuation $2|S|+1$ for each set S. The optimal solution is to assign H to the first bidder and the rest $\frac{m}{2}$ items to the second bidder, with a total welfare of $2m+3$. For price vector $\vec{p}=(p_1,\ldots,p_m)$ with $p_j = 2+\varepsilon$ for all items j, the demand query of the first bidder returns set H. However, with the use of value queries, unless all the sets of $\frac{m}{2}$ items be queried, the optimal allocation cannot be determined. Therefore, $\Omega(2^m)$ value queries will be needed in the worst case.*

More specifically, we have the following lemmas, according to Blumrosen and Nisan [5]:

Lemma 3.1. *A value query can be simulated by a polynomial number of demand queries.*

Lemma 3.2. *An exponential number of value queries may be required for simulating a single demand query.*

Mechanism design with the exclusive use of value queries is very restricted. In fact, as optimization problem, forgetting incentives, the computation of a near-optimal welfare maximizing allocation with value queries is still a challenging task; for example, Mirrokni et al. [48] proved that for any fixed $\varepsilon > 0$, achieving an approximation ratio of $(1-\frac{1}{e}+\varepsilon)$ for welfare maximization with submodular bidders requires an exponential number of value queries. Moreover, Dobzinski [15] proved that any truthful mechanism for submodular Combinarorial Auctions with approximation ratio better than $m^{\frac{1}{2}-\varepsilon}$ must use *exponentially* many value queries. Therefore, we have to look to obtain better approximation guarantees with the use of demand queries.

Informally, in a Combinatorial Auction with *Demand Queries*, items are posted with prices and agents arrive one-by-one, choosing their utility-maximizing bundle in the existing prices. Then, it is common to operate an update rule on the prices, which is necessary for the improvement of the approximation guarantees. It is important to mention that this whole procedure is executed in a way that retains truthfulness. Now, it is clear that the efficiency of such mechanisms depends on the way of posting and updating the prices.
In order to illustrate the notion of approximate randomized combinatorial auctions with demand queries, we give the following simple example:

Example 3.3. *Let a single-item auction with n agents. We divide the agents into sets A, B uniformly at random, such that $|A| = |B| = \frac{n}{2}$. The mechanism comprises of two phases; in the first phase, we ask the agents of set $|A|$ to reveal their valuation and let $L = \max_{i\in A}\{v_i\}$ be the maximum of these valuations. These agents will not be allocated the item, so they have no incentive to lie about their true valuation. In the second phase, we post price L to the item and fix an internal ordering π of agents in $|B|$. Then, agents*

from set $|B|$ arrive one-by-one and have to decide whether they buy the item or not. The moment the item gets sold (if any), the mechanism is terminated.
We claim that the above randomized mechanism $\frac{1}{4}$-approximates the optimal solution. Indeed, with probability $\frac{1}{4}$ the second largest bidder, i.e. the one whose valuation is the second largest among all the bidders, belongs to the set A and the largest bidder to set B. Conditioning on this event, the only bidder from set $|B|$ that can buy the item is the largest one, which results to the optimal solution. Therefore, $\mathbb{E}[ALG] \geq \frac{1}{4}OPT$.

Chapter 4

Ascending Implementations and Walrasian Equilibrium

When someone hears about auctions, the first thing that comes to their mind are big auction houses, where items of great value, such as art work and houses, are sold. In this context, the auctioneer raises constantly the price, until only one bidder remains in the auction and who finally gets the item. It is apparent that this type of mechanism is indirect, in the sense that bidders do not report their valuation to the auctioneer, but rather they answer "yes" or "no" queries, as the price increases. The question is why this kind of auctions are implemented in most cases in real life and whether we can reason about efficiency and incentive guarantees.

First of all, it is clear that it is easier for bidders to answer simple queries than to announce their valuation, especially in more tortuous settings. Secondly, it retains privacy, since the winner of the auction does not reveal their valuation, even not to the auctioneer; we just learn a lower bound on their valuation by the last price all other bidders quited the auction. Thirdly, it has more transparency, since all bidders are able to see the price augmentation. First, we are going to analyse the $k-$Vickrey auction in an ascending price format and define the equivalent notion of truthful revelation in ascending price auctions. This chapter is based on [54, 55, 60, 56, 57, 58].

4.1 Warm-up

Suppose an auction with n bidders and k identical items, the so called $k-$Vickrey Auction. Each bidder is unit demand, meaning that he only wants to acquire one copy of the item. Implementing VCG on this setting, the k highest bidders would be allocated a copy of the item, and they would be charged the $k+1$ highest bid. An ascending implementation of this setting is described on *Algorithm 1*.

Let us now reason about the incentives of this procedure. First, we are going to define an analog of truthful revelation in ascending price auctions.

Definition 4.1 (Sincere bidding). *We say that an agent bids sincere, if they answer honestly all queries.*

In this case, *sincere bidding* means that an agent aswers "yes" in *Step 3* of *Algorithm 1*, only if his valuation exceeds $p_t + \varepsilon$.

Algorithm 4 Ascending Implementation of $k-$Vickrey Auction

Initialize the price $p_0 = 0$ and the set of active bidders $S_0 = [n]$, $t = 0$.
while $|S_t| > k$ **do**
  Ask each bidder in S_t if they want the item in price $p_t + \varepsilon$ and let S_{t+1} be the set of the bidders they answered positive.
  $p_{t+1} = p_t + \varepsilon$
  $t \leftarrow t + 1$.
end while
Allocate an item to each bidder in $|S_t|$ at price p_{t-1}.
if $|S_t| < k$ **then**
  Choose at random $k - |S_t|$ bidders from $S_{t-1} - S_t$ and give them an item on price p_{t-1}.
end if

Proposition 4.1. *Sincere bidding is a dominant strategy in Algorithm 1 (up to ε).*

In this setting, it easy to reason about incentives, since if a bidder leaves earlier from the auction, then he loses the chance to get the item in a profitable price, while if he stays longer, he jeopardizes getting the item in a higher price than his valuation, and thus, leading to negative utility. Therefore, the welfare achieved by this auction is within $k\varepsilon$ of the maximum possible.

Now, think of a case with n bidders, m non-identical items and each bidder has additive valuation function. It is easy to see that the VCG solution is the same as running m parallel *Second-Price* auctions, since there is no dependence between the items that an agent gets. Based on the previous setting, it is reasonable to think that *sincere bidding* is a dominant strategy. However, as we show on the next example, it is not a dominant strategy.

Example 4.1. *Suppose a Combinatorial Auction with items A, B and 2 bidders with additive valuations. Let $v_1(A) = v_1(B) = 2$ and $v_2(A) = v_2(B) = 1$. Asumming that both agents bid sincere, then agent 1 acquires both items in price 1 each. Now, think that agent 2 follows the strategy below: If agent 1 bids for item A, then keep bidding for both items until agent 1 leaves the auction, otherwise, bid sincere. Then, it is clear that agent 1 would prefer to bid only for item 2 since it guarantees him a utility of 1.*

The aformentioned example illustrates a fundamental problem of the ascending price auctions; the action space is much larger than that of the direct revelation mechanisms and is history-dependent. In other words, it means than an agent can act according to the information he receives by observing what other agents do, since it proceeds iteratively. Therefore, we have to define a new notion of incentive guarantees, implying that sincere bidding is dominant strategy, conditioning on the event that all other agents bid sincere. This guarantee does not sound implausible, since each agents has only to assume that all others bid according to their valuation profile, no matter what it looks like. In the next section we are going to define a class of mechanisms, according to this guarantee, the so called EPIC mechanisms.

4.2 EPIC Mechanisms

In this section, we are going to define a new equilibrium concept, the so called *ex post Nash equilibrium.* Assume there an n bidders, and each bidder i has a set of private valuations V_i. As *strategy* of bidder i, we define a function $s_i : V_i \to A_i$, where A_i is the set of actions. A_i is much more richer than that of direct revelation mechanisms, as ascending auctions proceed in iterations and can be history-dependent.

Definition 4.2 (Ex Post Nash Equilibrium). *A strategy profile $(s_1, \dots, s_n)$ is an ex post Nash equilibrium, if, for each bidder i, sincere bidding, i.e. $s_i(v_i)$, is dominant strategy to every action profile $s_{-i}(\vec{v}_{-i})$*

In other words, it means every bidder knows that sincere bidding is the best response, only knowing that all the other agents use strategies $s_{-i}(\vec{v}_{-i})$, without actually knowing their true valuation profile.

Note: *Ex post Nash equilibrium* is a weaker equilibrium concept than *dominant strategy equilibrium* as the latter does not need to assume anything about the action of the others.

Definition 4.3 (EPIC mechanism). *A mechanism (f, p) is ex post incentive compatible (EPIC), if sincere bidding is an ex post Nash equilibrium ,and is individually rational.*

Now, the question that needs to be answered is, how difficult is it to design an *EPIC* mechanism. Since *dominant strategy equilibrium* concept is subset of *ex post Nash equilibrium*, then it is obvious that every DSIC mechanism is EPIC. However, in direct-revelation mechanisms, the two concepts coincide, since every available action is consistent with the truthful revelation of a possible valuation profile, and as a result, truthtelling can only increase a bidder's utility. Therefore, designing an EPIC iterative mechanism is only harder than designing a DSIC direct-revelation mechanism; assume that (f, p) is an EPIC iterative mechanism. Then, implementing revelation principle, there is a mechanism (f', p') that is direct-revelation and has the same outcome as (f, p). Since the two equilibrium concepts are equivalent in direct revelation mechanisms, as described above, then (f', p') is DSIC.

As we claimed in the previous chapter, in a direct-revelation mechanism, VCG outcome is the unique welfare maximizing outcome, that guarantees truthfulness. Furthermore, as said before, for every EPIC mechanism, there is an equivalent direct-revelation DSIC mechanism, by implementation of the revelation principle. Therefore, we have the following proposition:

Proposition 4.2. *Every EPIC and welfare-maximizing mechanism must lead to the same allocation and payment rule as VCG.*

In the next section, we are going to define the notion of a new equilibrium concept, but this time an equilibrium concept by the side of the market.

4.3 Walrasian Equilibrium

In economics, markets are places for the exchange of money and goods, where buyers and sellers participate and act selfishly. If the price of an item is relative low, meaning that its demand exceeds its supply, then there is a great margin for the seller to augment the price, without losing revenue. On the other hand, if the price of the item is relative high, then supply surpasses the demand and the price diminution is inevitable. As we observe, there is a state, where supply equals demand, the so called *Market Equilibrium*. The price that succeeds this type of equilibrium is called *clearing price*; for instance, in the case of the k-Vickrey auction, every price between the k-highest and the $(k+1)$-highest bid is *clearing*, since the demand of the bidders equals to k for every price in that region. The question is how effective a *Market Equilibrium* is, in the sense of the welfare achieved. In this section we will define the notion of *competitive* or *Walrasian Equilibrium* and explore its efficiency guarantees.
Formally, in a *Combinatorial Auction* environment, the *Walrasian Equilibrium* is defined as follows:

Definition 4.4 (Walrasian Equilibrium)**.** *Suppose there are n agents and a set U of m non-identical items. A Walrasian Equilibrium (WE) is a non-negative price vector $\vec{p}$ and an allocation $S = (S_1, \ldots, S_n)$ such that:*

- *Each agent i is allocated its utility-maximizing bundle (or $\emptyset$), i.e.*

$$S_i \in \arg\max_{T \subseteq U} \left\{ v_i(T) - \sum_{j \in T} p_j \right\} \tag{4.1}$$

- *If $j \notin \cup_i S_i$, then $p_j = 0$.*

Then, we say that $(S, \vec{p})$ is a Walrasian Equilibrium.

It is clear, that some very interesting and fundamental properties emerge from the above definition. Suppose a *Walrasian price vector* $\vec{p}$, i.e. there exists an allocation S, such that $(S, \vec{p})$ is a *Walrasian Equilibrium.* Then, we have that every agent is allocated its favorite bundle and therefore, there are no collisions among the bundles that agents want to get. As a result, if we have a *Walrasian price vector* $\vec{p}$, there is no need for a central coordination of the mechanism, since there is a distributed solution, where each agent would choose their favorite bundle, without worrying that another agent would get some of the items included in it. The question that arises is how efficient is the allocation S in a *Walrasian Equilibrium* $(S, \vec{p})$. The answer is given by the *First Welfare Theorem.*

Theorem 4.1 (First Welfare Theorem)**.** *Let a mechanism with n agents, a set U of m non-identical items, an allocation $S = (S_1, \ldots, S_n)$ and a price vector $\vec{p}$. If $(S, \vec{p})$ is a Walrasian Equilibrium, then S is a welfare-maximizing allocation.*

Proof. Let $O = (O_1, \ldots, O_n)$ be a welfare-maximizing allocation rule and $P = \sum_{j \in U} p_j$. Since $(S, \vec{p})$ is a Walrasian Equilibrium, each agents is allocated its favorite bundle, meaning that for every bundle $T \subseteq U$, we have:

$$S_i \in \arg\max_{T \in U} \left\{ v_i(T) - \sum_{j \in T} p_j \right\}$$

Therefore, for each agent i it holds:

$$v_i(S_i) - \sum_{j \in S_i} p_j \geq v_i(O_i) - \sum_{j \in O_i} p_j$$

Summing the above equation for each agent, we have:

$$v(S) - \sum_{i \in [n]} \sum_{j \in S_i} p_j \geq v(O) - \sum_{i \in [n]} \sum_{j \in O_i} p_j \tag{4.2}$$

However, since in allocation S an item j is unsold only if $p_j = 0$, the negative term of the LHS of the above equation sums over all the items that have a non-zero price, and, thus, we have that:

$$\sum_{i \in [n]} \sum_{j \in S_i} p_j \geq \sum_{i \in [n]} \sum_{j \in O_i} p_j \tag{4.3}$$

Combining equations (2.2) and (2.3), we get that:

$$v(S) \geq v(O)$$

and, therefore, S is a welfare-mazimizing allocation. □

The question that emerges is in what cases there exists a *Walrasian Equilibrium* and how it can be reached.

4.4 Ascending-Price Combinatorial Auctions

Now, think of the most general valuation model you can think. The only assumptions that need to be made are the following:

- $v(\emptyset) = 0$
- if $T \subseteq S$, then $v(T) \leq v(S)$.

Next, we provide an algorithm, introduced by Kelso and Crawford, and then, we will examine what properties should the valuations meet, in order to terminate at a *Walrasian Equilibrium.*

Now, think the above algorithm, implemented on the following case of 2 bidders and a set U of 2 items, A and B, with valuations:

$$v_1(S) = \begin{cases} 2, & \text{if } S = U \\ 0 & \text{otherwise} \end{cases}$$

and:

$$v_2(S) = \begin{cases} 1.5, & \text{if } S \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

So, the first bidder, wants both items or non, meaning that he has zero valuation for each of the items seperately. On the other hand, the second bidder wants at least one item, being indifferent which of the two he gets. Now, if we run *Algorithm 2* on this setting,

Algorithm 5 Kelso-Crawford (KC) Auction

Initialize the price of each item $j, p_j = 0$ and the allocating set of each bidder $i, S_i = \emptyset$.
while true **do**
  Ask each bidder sequentially for their favorite bundle of items *not assigned to them, giving them the items they already have*, .i.e. for each bidder i:

$$T_i = \arg\max_{T \subseteq [m]-S_i} \{v_i(S_i \cup T) - \mathbf{p}_\varepsilon(S_i \cup T)\}$$

  where:

$$\mathbf{p}_\varepsilon(S_i \cup T) = \sum_{j \in S_i} p_j + \sum_{j \in T} (p_j + \varepsilon)$$

  if $T_i = \emptyset$ for each bidder i **then**
    Terminate the auction with allocation $(S_1, \ldots, S_n)$ and prices $\vec{p}$.
  else
    Choose uniformly at random a bidder i, with $T_i \neq \emptyset$
    $S_i \leftarrow S_i \cup T_i$
    $S_k \leftarrow S_k - T_i$, for all $k \neq i$
    $p_j \leftarrow p_j + \varepsilon$, for all $j \in T_i$
  end if
end while

the first bidder will bid on both items, at price zero, and the second bidder will bid for one of the two, say item A. Then, the first bidder will bid again on both the items, while this time, the second bidder will bid on item B, which has a lower price. This bidding procedure will continue until the price of each reaches 1. Then, the second bidder will choose one of the two items, say A, increasing its price to $1 + \varepsilon$. Now, the first bidder wants to drop item B at price 1, as he has zero valuation for it alone. However, this kind of action, i.e. drop an item you already have, is not allowed by the algorithm. As a result, it is clear, that this algorithm is suitable only for specific valuation classes.

Note: This situation, where a bidder wants to drop an item, unless he has not another item, is known as the *exposure problem.*

This example did not only show a fundamental problem of this algorithm, but also indicated what properties should a valuation class fulfill, in order to terminate at a *Walrasian Equilibrium.*

4.5 Gross Substitutes

As it is apparent from the implementation of *Algorithm 2* to the above setting, the problem that occured was a situtation, where a bidder preffered to drop an item he already had, by the price augmentation of another item. This is exactly the condition that our valuation functions need to meet and this class of functions are called *Gross Substitutes.*

Definition 4.5 (Gross Substitutes). *Let a set of m items U a valuation function v, a* *price vector $\vec{p}$ and let $D(\vec{p})$ the set, containing the most demanded set under these prices, i.e.:*

$$\arg\max_{T\subseteq U}\left\{v(T)-\sum_{j\in T}p_j\right\}$$

We say that valuation function v satisfies the gross substitutes condition, if the following condition hold:
For every set $S \in D(\vec{p})$ and price vector $\vec{q}$, with $q_j \geq p_j$ for each $j \in U$, there exists a set $T \subseteq U$, such that:

$$(S - A) \cup T \in D(\vec{q})$$

where $A = \{j|q_j > p_j\}$.

To simplify the notation above, the *gross substitutes* condition means that if a bidder has a "favorite" bundle under a specific price vector, the increase in the prices of some of the items of the bundle, does not make him want to drop the items, whose price remained stable, but rather there exist a set, containing those items that is his "favorite" under the new price vector.

As a result, it is clear that we have fixed the problem, presented in the example above. Formally, we have the following theorem:

Theorem 4.2. *The Kelso-Crawford mechanism terminates at a $m\varepsilon$-Walrasian Equilibrium, if all bidders have gross substitutes valuations and bid sincere.*

Proof. The first condition of the *Walrasian Equilibrium* is more challenging than the second. First, we are going to prove the claim that at each moment of KC mechanism, with price vector $\vec{p}$, S_i belongs to a set of $D_i(\mathbf{p}_\varepsilon)$, where $\mathbf{p}_\varepsilon$ equals p_j, if $j \in S_i$ and $p_j + \varepsilon$ if $j \notin S_i$. The proof is by induction:

- The base case where $S_i = \emptyset$ for each bidder i, holds trivial.
- Now, think of a bidder i.
  - If i is chosen to bid at this iteration, then, by inductive hypothesis, there is a set $T_i \subseteq U - S_i$, such that $S_i \cup T_i \in D_i(\mathbf{p}_\varepsilon)$. As a result, at the end of the iteration, $S_i \leftarrow S_i \cup T_i$ and, thus, $S_i \in D_i(\mathbf{p}_\varepsilon)$.
  - If i was not chosen to bid at this iteration, then think of the last time he was chosen and acquired bundle G_i. Then, $S_i = G_i - H_i$, where H_i is the bundle of items that were taken from him by the other agents in the next iterations until now. Because his valuation satisfies the gross substitutes condition, we have that $S_i \in D_i(\mathbf{p}_\varepsilon)$.

As a result, at the final iteration, where no one wants to acquire any items he does not possess, every agent is with their favorite bundle. The second condition of the *Walrasian Equilibrium*, that an item j has $p_j = 0$ only if it is unsold, it is trivial, since bidders relinquish an item only if it is taken by another bidder. As a result, if an item is unsold, it entails that no one ever bids on it, and therefore has a price of 0. Therefore, KC auction terminates at a *Walrasian Equilibrium.* □

It is clear that the $m\varepsilon$ term results from the descretization of the price increase. Therefore, we have:

Proposition 4.3. *Let a mechanism with n agents and m items. If $v_1, \ldots, v_n$ fulfill the gross substitutes condition, then there exists a Walrasian Equilibrium.*

Although we have proved that the Kelso-Crawford auction terminates at a $m\varepsilon$-Walrasian Equilibrium, if all bidders have gross substitutes valuations and bid sincere, we have not talked about incentives. So, is this auction EPIC?

As we claimed in *Proposition 2.2*, a welfare-mazimizing allocation must lead to the same allocation and payment rule as VCG, in order to be EPIC. For the *unit-demand* case, i.e. each bidder wants only one item, there is a variation of KC mechanism, the Crawford-Knoer Auction, which yields to a *Walrasian Equilibrium* and the VCG payment rule. However, for *gross substitutes* valuation function, there does not exist a similar guarantee. Formally, we have:

Theorem 4.3. *Let n bidders whose valuations fulfill the gross substitutes condition. Then, there is no ascending price auction, that results to VCG payment rule for every possible valuation profile $\vec{v} = (v_1, \ldots, v_n)$.*

The proof of the theorem is based on the idea that in an indirect mechanism, where bidders answer queries, there may not be enough information revealed, so as to compute VCG prices. For example, there may be two different valuation profiles, that lead to exact the same aswer queries by the bidders, and as a result, the mechanisms computes the same allocation and payment rule, as it depends only on the aswer of the queries.

Proposition 4.4. *If a valuation function v satisfies the gross substitutes condition, then v is submodular.*

Finally, as we will show in the next section, *gross substitutes* is largest valuation class, that guarantees the implementation of VCG in polynomial time.

4.6 Combinatorial Auctions via Linear Programming

In this section, we are going to show that *gross substitutes* is largest valuation class, that guarantees the implementation of VCG in polynomial time, and at the same time, enhance our intuition about the *Posted-Price mechanisms*, which are discussed in the next chapter. As we have already seen, VCG is reduced to the calculation of $n + 1$ welfare-maximizing allocations, as the payment of each bidder requires computing such an allocation. But the question is whether there is a way to calculate a welfare maximizing allocation in polynomial time. The KC auction is not an option, as it is pseudo-polynomial. The solution is provided by the *Linear Programming.* Formally, a *Combinatorial Auction* is expressed by the linear program, below:

$$\begin{aligned}
\max \quad & \sum_{i\in[n]} \sum_{S\subseteq[m]} v_i(S) x_{iS} \\
\text{s.t.} \quad & \sum_{i\in[n]} \sum_{S\subseteq[m]: j\in S} x_{iS} \leq 1 \qquad && \forall j \in [m] \\
& \sum_{S\subseteq[m]} x_{iS} \leq 1 && \forall i \in [n] \\
& x_{iS} \geq 0 && \forall i \in [n], \forall S \subseteq [m]
\end{aligned} \tag{4.4}$$

and its dual:

$$\begin{aligned}
\min \quad & \sum_{i\in[n]} u_i + \sum_{j\in[m]} p_j \\
\text{s.t.} \quad & u_i \geq v_i(S) - \sum_{j\in S} p_j \qquad \forall i \in [n], \forall S \subseteq [m] \\
& u_i, p_j \geq 0 \qquad\qquad\qquad \forall i \in [n], \forall j \in [m]
\end{aligned} \tag{4.5}$$

In fact, we should allow x_{iS} to take only integer values. However, since *integer programming* is NP-hard, we express its linear relaxation. Before we proceed to the main theorem, let us express the *Complementary Slackness* conditions for this linear program:

1. if $x_{iS} > 0$, then $u_i = v_i(S) - \sum_{j\in S} p_j$
2. if $u_i > 0$, then $\sum_{S\subseteq[m]} x_{iS} = 1$
3. if $p_j > 0$, then $\sum_{i\in[n]} \sum_{S\subseteq[m]:j\in S} x_{iS} = 1$

The first and third conditions resemble the *Walrasian Equilibrium* definition, as the first means that if i gets (even fractionally) bundle S, then it is i's "favorite" bundle and the third that since item j has a non-zero price, then it is sold. Now, we have the following theorem:

Theorem 4.4. *Let OPT_{IP} and OPT_{LP} be the optimal values of the integer program and its linear relaxation, respectively. Then:*

- *There exists a Walrasian Equilibrium if and only if $OPT_{IP} = OPT_{LP}$*
- *Conditioning on this, $\vec{p}$ is a Walrasian price vector if and only if there exists vector $\vec{u}$ such that $(\vec{u}, \vec{p})$ is an optimal solution to the dual problem (2.5).*

Proof. Let $\vec{p}$ be a price vector and A an allocation. We will prove the equivalent condition, that $(A, \vec{p})$ is a *Walrasian Equilibrium* if and only if there exists vector $\vec{u}$ such that $(\vec{u}, \vec{p})$ is an optimal solution to the dual and A is an optimal solution for the linear relaxation. First, assume that $(A, \vec{p})$ is a *Walrasian Equilibrium* and let $\vec{x}$ denote the integral solution induced by allocation A, meaning that $x_{iA_i} = 1$ and $x_{iS} = 0, \forall S \neq A_i$, for all bidders i. Since we have $\vec{x}$ and $\vec{p}$, the only thing that remains to be constructed is the vector $\vec{u}$. Now, setting:

$$u_i = \max\left\{0, \max_{S\subseteq[m]}\left\{v_i(S) - \sum_{j\in S} p_j\right\}\right\}$$

we have that if bidder i is allocated A_i, then

$$u_i = v_i(A_i) - \sum_{j\in A_i} p_j \geq v_i(S) - \sum_{j\in S} p_j, \forall S \subseteq [m]$$

since $(A, \vec{p})$ is a *Walrasian Equilibrium*. Furthermore, if $x_{iA_i} = \emptyset$, i.e. $\sum_{S \subseteq [m]} x_{iS} = 0$ and, by the above definition of u_i's, we have that $u_i = 0$. Finally, the third condition holds, since $(A, \vec{p})$ is a *Walrasian Equilibrium* and $\vec{x}$ is induced by A.
Now, assuming that $\text{OPT}_{IP} = \text{OPT}_{LP}$, let $\vec{x}$ be the optimal solution of the integer program, and $(\vec{u}, \vec{p})$ the dual solution. Then, because of optimality, the *Complementary Slackness* conditions are fulfilled and therefore the allocation A, induced by $\vec{x}$, and the price vector $\vec{p}$, form a *Walrasian Equilibrium*. □

Therefore, according to *Proposition 2.3*, we have the following corollary:

Corollary 4.1. *Let an auction with n bidders and m items. If $v_1, \ldots, v_n$ satisfy the Gross Substitutes condition, then $OPT_{IP} = OPT_{LP}$.*

Now, it is clear that if we are able to solve the LP in polynomial time, then we are also able to implement VCG in polynomial time. Thus, we can compute the LP in polynomial time via *ellipsoid algorithm*.

Chapter 5

Combinatorial Auctions with Demand Queries

5.1 Introduction

As we have already seen in the first chapter, there is a DSIC mechanism, VCG, that achieves a welfare-maximizing solution for *Combinatorial Auctions* with general valuation classes. The problem is that VCG cannot be implemented in polynomial time, except for some very limited cases. In the previous chapter we have claimed that *Gross Substitutes* is the "largest" valuation class that guarantees the existence of *Walrasian Equilibrium* and, therefore, the VCG implementation in polynomial time. The question that emerges is whether we can design a computational efficient mechanism that succeeds even a near-optimal solution in *Combinatorial Auctions* with *submodular, XOS* or *subadditive* valuation classes. The answer to this question is positive and the solution is achieved by the implementation of *Posted-Price mechanisms*, where bidders arrive in the auction according to a specific order, and buy their favorite bundle of items in the prices being posted. As a result, the efficiency of the mechanism depends on the way prices are calculated. At this point, we should mention again a fundamental property of *clearing prices*; once they have been calculated, running a posted-price auctions with agents arriving even in adversarial order, it terminates with an optimal solution. Based on the notion of *clearing prices* and *Walrasian Equilibrium*, it would be ideal if we were able to calcualte *approximate clearing prices*. However, approximate clearing prices are not enough; the following examples illustrate the problem may arise and highlights a fundamental property these prices should satisfy.

Example 5.1. *Suppose there are n bidders with additive valuations and n non-identical items. Since valuations are additive, there exist clearing prices and Walrasian Equilibrium. Now assume that agents have the following valuations:*

- $v_{ii} = n$*, for* $i = 1, \ldots, n$
- $v_{ij} = 1$*, for* $i \neq j$ *and* $i, j = 1, \ldots, n$

It is clear, that in OPT solution, agent i receives item i, for all $i = 1, \ldots, n$, and the welfare achieved is n^2. Every price vector $\vec{p} = (p_1, \ldots, p_n)$, where $p_i \in (1, n)$ is a Walrasian price vector, since only bidder i wants to buy item i in a price between 1 and n.

As a result, suppose a Walrasian price vector $\vec{p}^{\,*} = (p_1^*, \ldots, p_n^*)$, where $p_i^* = 1+\varepsilon$, for all $i = 1, \ldots, n$. Now, suppose that someone with full knowledge of the private information of the bidders, gives us an approximate price vector $\hat{\vec{p}} = (\hat{p}_1, \ldots, \hat{p}_n)$, where $\hat{p}_i \in (\frac{p_i^*}{2}, p_i^*)$. We expect that, running a posted-price auctions with the approximate Walrasian price vector $\hat{\vec{p}}$, then the welfare achieved will be near-optimal. However, assume a random order of the bidders. Then, the first bidder arriving at the auction, say bidder k, will buy not only item k, but also all the other items, since his valuation for them exceeds their price. As a result, the welfare achieved by the mechanism will be $2n-1$, which is a $O(n)$-approximation of the optimal solution.

It is apparent, that it is not enough calculating approximate clearing prices; we have to calculate prices that achieve *revenue* comparable to the optimal welfare, meaning that if it terminates with allocation S, then:

$$\sum_{j \in S} p_j \geq a \cdot v(\text{OPT})$$

In this chapter, we are going to discuss and analyze the algorithmic technics and components some of the *Posted-Price Combinatorial Auctions*, that give the best approximation ratios, until now.

In the *worst-case* setting, where we do not make any further assumptions on bidders' valuations, Dobzinski et al. [21] presented the first truthful mechanism with a non-trivial approximation guarantee of $O(\log^2 m)$. Dobzinski [17] improved the approximation ratio to $O(\log m \log\log m)$ for the more general class of subadditive valuations. Subsequently, Krysta and Vöcking [41] provided an elegant randomized online mechanism that achieves an approximation ratio of $O(\log m)$ for XOS valuations. Dobzinksi [16] broke the logarithmic barrier for XOS valuations, by providing an approximation guarantee of $O(\sqrt{\log m})$. We highlight that accessing valuations through demand queries is essential for these strong positive results. Dobzinski [15] proved that any truthful mechanism for submodular CAs with approximation ratio better than $m^{\frac{1}{2}-\varepsilon}$ must use *exponentially* many value queries.

In the *Bayesian* setting, bidder valuations are drawn as independent samples from a known distribution. Feldman et al. [29] showed how to obtain item prices that provide a constant approximation ratio for XOS valuations.

In the next section, we are going to present some of the basic components that are found in *Posted-Price* mechanisms.

5.2 Structure of Posted-Price mechanisms

At this section, we are going to analyze the idea and the basic components of *Posted-Price* mechanisms.

5.2.1 Preliminaries

Before we proceed with the core components, we find it useful to provide a definition from [17], that will assist us in our analysis.

Definition 5.1 (γ-supporting prices)**.** *Prices $p_1, \ldots, p_m$ are called γ-supporting for a set of items S and valuation function v if the following conditions hold simultaneously:*

- S **is strongly profitable:** *for each* $T \subseteq S$ *it holds that:* $v(T) \geq \sum_{j \in T} p_j$.

- **the prices are high:** $\sum_{j \in S} p_j \geq \frac{v(S)}{\gamma}$

For XOS valutions, we have to give a slightly different definition of *supporting prices*. At this point we have to make clear that it is not the same as γ-supporting prices, defined above.
Let $O = (O_1, \ldots, O_n)$ be an allocation and w_i be the maximizing clause of O_i in the valuation v_i of bidder i. For each bidder i and $j \in O_i$, we say that j's *contribution* in the allocation O is $q_j = w_i(\{j\})$. Moreover, we say that allocation O is supported by prices $p_1, \ldots, p_m$, if $q_j \geq p_j$ for all items j. These prices are called *supporting prices* of an allocation. Once calculating supporting prices of an allocation, we have the following lemma of Dobzinski [16]:

Lemma 5.1. *Let* $T = (T_1, ..., T_n)$ *be an allocation that is supported by prices* $p'_1, ..., p'_m$. *A fixed price auction with prices* $p_j = p'_j/2$ *generates an allocation* $(S_1, ..., S_n)$, *with* $\sum_i v_i(S_i) \geq \frac{\sum_{j \in \cup_i T_i} p'_j}{2}$.

This lemma makes obvious, that if we to calculate *supporting prices* with high value, we achieve an allocation with high welfare.

5.2.2 Components

General Framework for aproximating Social Welfare in CAs. It was well understood in the literature that posted-price mechanisms for approximating the SW in CAs use some core components; the most crucial of them are the `Price-Exploration` and the `Price-Exploitation`. Both for the Bayesian and the worst-case setting, the core of the framework remains the same; it goes without saying that the more advanced the results become, the more advanced the techniques become, as well. As we realized, this core is much more general than it is originally stated.

1. [`Dominant Bidder`] With constant probability, run a second-price auction for the grand bundle U, in case there is a dominant bidder. If there is a dominant bidder, i.e. his valuation for all the items is a higher fraction of the optimal solution than the approximation ratio, we get in a truthful way the desired result by allocating all the items to the dominant bidder.

2. [`Sampling and Estimation`]: Agents are divided with some probability into some sets, most commonly into SAMPLE and TEST sets, in order to estimate a quantity from the SAMPLE set and apply it on the TEST set. Because sampling methods depend only on statistical properties and agents do not get any items, it does not break the truthfulness. In the absence of a dominant bidder, two randomly chosen groups of bidders have many common characteristics, such as a constant fraction of the total welfare and preferences.

3. [`Price-Exploration`] Find γ-supporting prices $p_j, j \in U$ for an allocation S:

$$\sum_{j \in U} p_j \geq \frac{v(S)}{\gamma}$$

that is a c_1-approximation to the optimal:

$$v(S) \geq \frac{v\left(\text{OPT}\right)}{c_1}$$

.

4. [`Price-Exploitation`] Post prices $p_j, j \in U$ and let bidders use *Demand Queries* (DQ) to choose their most preferred bundles.

The *key relation* for the analysis of the posted-price mechanism is the relaxed utility-maximizing property, obtained from the *Demand Query* (DQ), i.e. for each agent i, where S_i is the bundle he obtained and T_i every other bundle available at the moment he arrived:

$$v_i(S_i) \geq v_i(T_i) - p(T_i)$$

The above relation is used to relate the the value of the bundle an agent obtained with the value of the bundle he acquired in the optimal solution, restricted on the available items. The strong profitability of a set is used for ensuring that when an agent arrives at the auction, the bundle he acquired in the optimal solution, restricted on the available items, gives him non-negative utility, and therefore, he could have chosen it. Finally, we have to inspect under which conditions the *Demand Query* returns a *Strongly Profitable* set. In Dobzinski [17] lies the following lemma:

Lemma 5.2. *Let v be a subadditive valuation and let S be the set returned from the utility-maximizing Demand Query in the price vector $\vec{p}$. Then, S is strongly profitable in these prices.*

5.3 Worst-case Setting

In this section we are going to analyze the three main ways of implementing a *Posted-Price* mechanism in a worst-case setting - Uniform Prices, Distinguished Prices and a combination of the two - in order to enhance out intuition about the primal technics, and present the most important results. We also provide the mechanisms in an algorithmic form, explaining the ideas behind them and the key reasons they finally work.

5.3.1 Uniform Prices

Assume that you know somehow a good approximation of the optimal solution, APX, such that APX $\geq a \cdot$ OPT. Then, supposing that every item contributes the same to the optimal solution, i.e. $q_j = \frac{\text{OPT}}{m}$ for all $j \in U$, a plausible thought would be to run a posted price auction with fixed price $p = \frac{\text{APX}}{2m}$ for all items. According to Lemma 5.1, and since $\frac{\text{APX}}{m} \leq \frac{\text{OPT}}{m}$, a fixed price auction with price p would lead to an 2- approximation of the optimal welfare. In the absence of the assumption that all items contribute the same in the optimal solution, could we find a price that results to a fixed price auction with high welfare? Of course this price cannot be too low, neither too high and has to balance between the low and high valued items. The answer to this question is positive and a mechanism, operating in this way, was designed by Dobzinski et al. [21], which results to a $O(\log^2 m)$- approximation.

Algorithm 6 Dobzinski et al. [21] - $O(\log^2 m)$ mechanism for XOS valuations

1: Assign each bidder to exactly one of the following three sets: SEC-PRICE with probability $1-\varepsilon$, FIXED with probability $\varepsilon/2$, and STAT with probability $\varepsilon/2$.
2: Find an $O(1)$ approximation of the optimal solution with bidders from STAT, OPT_{STAT}.
3: From the above allocation, find a price p' and an allocation T, such that T is supported by p' and its supported value is $\Omega(\frac{\text{OPT}_{STAT}}{\log m})$.
4: Run a second price auction for the grand bundle U with the bidders from SEC-PRICE and reserve price $\frac{\varepsilon}{100}\cdot\frac{\text{OPT}_{STAT}}{\log^2 m}$ and if there is a winning bidder, terminate the auction.
5: Run a fixed-price auction with price $p'/2$ with bidders from STAT.

Theorem 5.1. *Algorithm 6 is universally truthful and achieves an approximation ratio of $O(\log^2 m)$ for the Social Welfare.*

First, we are going to give some definitions. We say that an allocation $T=(T_1,\ldots,T_n)$ is supported by a price p, if for each bidder i and possible bundle $S_i \subseteq T_i$, it holds that $v_i(S_i)\geq |S_i|\cdot p$. We call $\sum_i |T_i|\cdot p$ the supported value of T.

The key point of the mechanism lies in Step 2 and Step 3 of the algorithm. In the absence of a dominant bidder, `Sample and Estimation` component succeeds with high probability that $\text{OPT}_{STAT}\geq\frac{\varepsilon}{4}\cdot\text{OPT}$ and $\text{OPT}_{FIXED}\geq\frac{\varepsilon}{4}\cdot\text{OPT}$. The most crucial part of the algorithm is the `Price-Exploration` and the calculation of price p'. The success of the algorithm is owed to the following lemma:

Lemma 5.3. *There exists with high probability a price $p_k \in P=\left\{\frac{OPT}{2m\log m},\ldots,\frac{OPT}{\log m}\right\}$, where $|P|=\log m$, such that there exists an allocation T^k of the items to the bidders in FIXED only, such that T^k is supported by p_k and the supported value of T^k is $\Omega(\frac{OPT_{STAT}}{\log m})$.*

Therefore, according to Lemma 5.1, running a fixed price auction with price $p'/2$, gives the desired result.

5.3.2 Distinguished Prices

The above mechanism is based on the idea that, in the absence of a dominant bidder, we have to find a price, neither too low, nor too high, in order to sell the items at that price and achieve a revenue comparable to the optimal welfare. But, as it has been clear from the previous chapter and the definition of *Walrasian Equilibrium*, every item should be sold at a different price, according to its demand. Therefore, the matter is whether we can imitate the notion of *clearing prices* in order to achieve better approximations.

This idea was implemented from Krysta and Vöcking [41], achieving an approximation ratio of $O(\log m)$. Their algorithm is inspired from the *Ascending-Price* auctions, where the price of each item is increased, as long as it is preferred by the agents. Informally, the mechanism calculates supporting prices by combining binary search and randomized rounding and operates as follows:

Supposing there an n agents, arriving online, and m non-identical items, the mechanism assigns each item the same price at first. As long as the bidders arrive online and acquire their utility-maximizing bundle, through a demand query, the mechanism increases the price of the items being chosen, multiplicatively, allowing the overselling of copies. In this way, the mechanism learns the appropriate price of each item, terminating, however, with an infeasible solution. In order to fix the infeasibility of the allocation, the mechanism assigns each bidder "virtual copies", and decides the realization of the bundle via oblivious randomized rounding. Therefore, the mechanism succeeds in finding of "correct" prices and at the same time, keeping the allocation feasible. Below, we provide the randomized algorithm, terminating at a feasible solution.

Algorithm 7 KV [41] $O(\log m)$ -mechanism for submodular valuations

1: Ask the first half of the bidders their valuation for the whole set of the items U and let L be the highest value. Allocate them nothing.
2: Fix an ordering π of the rest bidders and set $U_1 = U$.
3: Set initial prices $p_1^{(1)} = \cdots = p_m^{(1)} = \frac{L}{4m}$.
4: **for** each bidder i according to π **do**
5: Let $S_i = \mathrm{DQ}(v_i, U_i, \vec{p}^{\,(i)})$
6: With probability q, allocate $R_i = S_i$ to i and set $U_{i+1} = U_i \setminus S_i$. Otherwise, set $U_{i+1} = U_i, R_i = \emptyset$.
7: Update prices $\forall j \in S_i$: $p_j^{(i+1)} = 2p_j^{(i)}$.
8: **end for**

Theorem 5.2. *Algorithm 10 is universally truthful and for $q = 1/\Theta(\log m)$, achieves an approximation ratio of $O(\log m)$ for the Social Welfare.*

In this mechanism, the `Sample and Estimation` step is used in order to calculate the parameter L for the prices. In fact, the mechanisms conditions on the event that the second-highest bidder belongs to the SAMPLE set and the highest bidder to TEST set. The truthfulness of the mechanism is retained, since the first half of the bidders do not have incentives to lie, as they will not receive any item, and the rest do not affect the prices offered to them and the realization of their chosen bundles. Unlike the previous mechanism, the `Price-Exploration` and `Price-Exploitation` components take place simultaneously, as prices are posted for each bidder and formed online for the next one. The key lemma that allows the above mechanism achieve a "good" approximation ratio is the following:

Lemma 5.4. *For Algorithm 10 with $q = \frac{1}{4(\log(4m)+1)}$, it holds that for each agent i and for all $A \subseteq U$, $\mathbb{E}[v_i(A \cap U_i)] \geq v_i(A)/2$.*

The above lemma ensures that, when an agent arrives at the auction, the expected value of every bundle, restricted on the set of the available items, is comparable to the value of the unrestricted bundle. Therefore, it allows us to compare for each agent the value of the chosen bundle with his utility in the optimal solution. The proof of the theorem is based on the combination of the above lemma and the property of the utility-maximizing bundle. Informally, for each agent i, where T_i is the bundle of items he receives in the optimal solution, the fundamental relation is the following:

$$\begin{aligned}\mathbb{E}[v_i(S_i)] &\geq \mathbb{E}[v_i(T_i \cap U_i)] - \mathbb{E}[p(T_i \cap U_i)] \\ &\geq \frac{1}{2}v_i(T_i) - \mathbb{E}[p(T_i)]\end{aligned}$$

Summing the above equation for all agents, we can relate the welfare obtained from the Algorithm 10 with the optimal solution, as follows:

$$\mathbb{E}[v(S)] \geq \frac{1}{2}v(OPT) - \mathbb{E}[p(OPT)]$$

Finally, by the starting price and the price update rule, we get that:

$$\mathbb{E}[v(S)] \geq \frac{1}{8}v(OPT)$$

and since $\mathbb{E}[v(R)] = q\,\mathbb{E}[v(S)]$, we get the result. The analysis is almost identical to the analysis in Section 6.4 for the case of *Liquid Welfare*.

5.3.3 Uniform and Distinguished Prices

The question that emerges is whether we can combine the above ideas, i.e. uniform and distinguished price for all the items, and achieve a better approximation ratio. Dobzinski [16] managed to do this in a much more involved way, achieving an approximation ratio of $O(\sqrt{\log m})$

First we are going to define some necessary notions. Let $A = (A_1, \ldots, A_n)$ be some allocation and $P = \{\frac{2^k v(A)}{m^2}\}_{k\in\mathbb{Z}}$. We say that bin k is associated with $p(k) = \frac{2^k v(A)}{m^2}$ in allocation A. Suppose that all valuations are XOS and let q_j^A denote the contribution of item j in A. Let q'_j be the maximal value in P such that $q_j^A \geq q'_j$. We say that an item j belongs in bin k in allocation A if $q'_j = p(k)$.

At this point, we are going to expand our intuition. If we assume that all valuations are constructed from integers in $\{1, 2, 4, 8, \ldots, m\}$, a uniform price for all items selected at random from $1, 2, 4, 8, \ldots, 2^{\log m}$ results in an logarithmic approximation ratio. To see this, let $\mathcal{O} = (O_1, \ldots, O_n)$ be the optimal allocation. For each bidder i and item $j \in O_i$, let $q_j = a_i(\{j\})$ be the contribution of j in the optimal solution. Now, we partition the optimal allocation into $\log m$ bins, and allocate item j to bin G_r if $q_j = r$. Let $c(G_r) = r|G_r|$ be the contribution of bin G_r to the optimal solution $\mathcal{O}$. It is clear that $v(\mathcal{O}) = \sum_r c(G_r)$. For *additive* valuations, it is straightforward that in a fixed price auction with price r, all the items in bins G_r or higher would be sold. For XOS valuations, we have a similar result; from Lemma 5.1, it holds that running a fixed price auction with price $\frac{r}{2}$ terminates to an allocation with welfare $\Omega(c(G_r))$. Therefore, choosing a price p from $1, 2, 4, 8, \ldots, 2^{\log m}$ uniformly at random, we obtain an allocation with approximation ratio $O(\log m)$.

Now, what if we use $O(1)$ numbers of bins? Intuitively, when an item j belongs to bin k, it holds that $p_k \leq q_j \leq p_{k+1}$. In the case of $\log m$ bins, we have that $p_k \leq q_j \leq 2p_k$ and as a result, $p_k \geq q_j/2$. Therefore, when we find the correct price of an item, we get the efficient guarantee for free, which is impossible in the case of $O(1)$ bins. An elegant way to diminish the number of bins is by arranging the $\log m$ them into $\sqrt{\log m}$ chunks, such that each chunk contains $\sqrt{\log m}$ bins. Let $val(C_k)$ be the value of chunk k, i.e. $val(C_k) = \sum_{r \in C_k} c(G_r)$. We call a chunk k *easily approximable* if, running a fixed price auction with price $\frac{r_k}{2}$, where r_k is the price of the first bin in chunk k, terminates to an allocation with welfare $\Omega(val(C_k))$. Then, if the value of *easily approximable* chunks is high, selecting uniformly at random a price from the first bins of all the $\sqrt{\log m}$ chunks, and running a fixed price auction with this price, the mechanism terminates to an allocation with approximation ratio $O(\sqrt{\log m})$. On the other hand, if the value of *easily approximable* chunks is low, we have think of something else. Intuitively, a fixed price auction with low price will end up with the items, having high contribution to the optimal solution, sold. The problem is that it does not entail that these items are bought by bidders with high valuation, and therefore, can lead to inefficient solution. However, the fact that the items are sold in fixed price auctions with price lower than their contribution, allows us to execute an ascending-price procedure, in order to find their correct price. The idea that leads to a $O(\sqrt{\log m})$ approximation ratio is running "imaginery" fixed price auctions with increasing price to find the chunk of each item in the optimal solution. Then, choosing uniformly at random one of the $\sqrt{\log m}$ prices in each chunk, we "guess" for each chunk C_k the correct price of items of expected value $(\sqrt{\log m})^{-1} val(C_k)$. Formally, the mechanism is described in Algorithm 8.

Theorem 5.3. *Algorithm 8 is universally truthful and achieves an approximation ratio of $O(\sqrt{\log m})$ for the Social Welfare.*

The most innovative and interesting fact at this mechanism is that there are two `Price-Exploration` and two `Price-Exploitation` phases, intersecting with each other. To be more specific, in Step 7 of the algorithm, the `Price-Exploitation` from bidders in UNIFORM leads to `Price-Exploration` for bidders in FINAL. The key claim of the mechanism is that if `Price-Exploration` of Step 7(b) gives a bad approximation ratio, it results to an effective `Price-Exploration` for Step., i.e. calculates "good" supporting prices, and is given by the following lemma:

Lemma 5.5. *Let $A = (A_1, ..., A_n)$ be an allocation that is supported by $p'_1, ..., p'_m$. For every item $j \in A_i$, let q_j be its supporting price. Let $o = \sum_i \sum_{j \in A_i} q_j$. Let N' be a random set of bidders where each bidder is in N' independently at random with probability r. Let $T = (T_1, ..., T_n)$ be the random variable that denotes the allocation of the fixed price auction with price $p_j = p'_j/2$ for every item j when N' is constructed at random as above and the order over bidders in N' is chosen uniformly at random.*

Let $c = 1 - \frac{\sum_i \sum_{j \in A_i \cap (\cup_k T_k)} q_j}{o}$. Then $\mathbb{E}\left[\sum_i v_i(T_i)\right] \geq \frac{o \cdot \mathbb{E}[c] \cdot r}{4}$, where expectations are taken over the random choices of N' and its internal order.

Let us make the above lemma more clear. We have to interpret the role of variable c. From its definition, c expresses the fraction of the welfare the items, that were unsold in allocation T, give in allocation A. Therefore, when c is small, it means that the majority of the items were bought in allocation A, while when c is large, it ensures that allocation

Algorithm 8 Dobzinski [16] - $O(\sqrt{\log m})$ mechanism for XOS valuations

1: With probability $\frac{1}{2}$ run a second-price auction for the grand bundle U and terminate the auction.
2: Divide the agents uniformly at random at sets STAT, UNIFORM, FINAL.
3: Run a 2-approximation greedy algorithm for bidders in STAT and let APX be the allocation.
4: Partition bins $1, \ldots, \log m$ of the allocation APX into $\sqrt{\log m}$ disjoint chunks, such that chunk C_k contains bins $\{(k-1) \cdot \frac{4 \log m}{\sqrt{\log m}}, \ldots, k \cdot \frac{4 \log m}{\sqrt{\log m}}\}$
5: Select uniformly at random an integer from $\{1, 2, 3, \cdots, \frac{4 \log m}{\sqrt{\log m}}\}$
6: Select uniformly at random an ordering π of the bidders in UNIFORM and set $p_j = 0$ for all items $j \in U$.
7: Consider each chunk C_k in an ascending order:
   a) Let p' be the smallest price associated with any bin in C_k. Run a fixed price auction restricted to bidders in UNIFORM with price $p'/2$ for every item, where the order of the bidders in UNIFORM is π (Step 6). Denote the allocation that this fixed-price auction outputs by $T^k = (T_1^k, \ldots, T_n^k)$ and by $p^{i,k}$ the payment of bidder i.
   b) With probability $\frac{1}{\sqrt{\log m}}$ the mechanism ends with the allocation T^k. Each bidder i pays $p^{i,k}$.
   c) Let p'' be the price of the bin with the r'th smallest value in C_k. Update the price p_j of every item $j \in \cup_i T_i^k$ to $p_j = p''/2$.
8: Run a fixed price auction with prices $p_1, \ldots, p_m$ with the participation of bidders in FINAL.

T has high welfare.
Now, assume that A is the optimal allocation from bidders in UNIFORM $\cup$ FINAL, restricted on the items of chunk C_k. Then, A is supported by every price p below the price of the first bin of C_k. Let T be the allocation from the fixed price auction from bidders in UNIFORM with price p_k, where p_k is the value of the first bin of a chunk C_k. When chunk C_k is *not easily approximable*, it holds that $\mathbb{E}[c] \leq \frac{1}{4}$ and as a result, according to the above explanation, most of the items of chunk C_k were bought in allocation T. Therefore, when running "fictitious" fixed price auctions, most items are sold until reaching the chunk they belong and finally, we are able to guess their correct price. This step embodies the ascending-price procedure.Finally, the fact that the first bins of consecutive chunks are at a distance of $\sqrt{\log m}$ bins, ensures that the value of items that become repriced, even after they reach their chunk, is small. Consequently, after having calculated $\sqrt{\log m}$*-supporting prices* for bidders in FINAL, according to Lemma 5.1, the posted price auction in Step 8 gives the result.

5.4 Bayesian Setting

As it has been made clear through the previous section, a *Posted-Price* mechanism aims to calculate *approximate clearing prices* either through sampling methods, or dynamically. However, if we were able to see the full valuation profile of the bidders, what would we do? Perhaps, we would calculate a greedy solution, (for XOS valuations there is a 2-approximation greedy algorithm), and post prices according to every item's contribution in the solution. More specifically, for XOS bidders, we would calculate a 2-approximation allocation $\mathcal{A}$ and set for every item j price $p_j = q_j^{\mathcal{A}}/2$, where $q_j^{\mathcal{A}}$ is the contribution of j in $\mathcal{A}$. Then, according to Lemma 5.1, we would get an allocation S with:

$$v(S) \geq \frac{1}{2} \cdot \sum_{j=1}^{m} q_j^{\mathcal{A}} \geq \frac{1}{4} \cdot v(OPT)$$

In the *Bayesian* setting, where agents' valuations are drawn independently from $\mathcal{V}_1, \ldots, \mathcal{V}_n$, we denote by $\mathcal{V} = \mathcal{V}_1 \times \cdots \times \mathcal{V}_n$ so that $\vec{v}$ is drawn from $\mathcal{V}$. We suppose that $\mathcal{V}$ is public knowledge and we have sample access, but the realization of v_i is known only to bidder i. In the *Bayesian* setting, we say that an allocation $\mathcal{A}$ is an a - approximation of optimal solution if:

$$\mathop{\mathbb{E}}_{\vec{v} \sim \mathcal{V}}[v(\mathcal{A})] \geq \frac{1}{a} \cdot \mathop{\mathbb{E}}_{\vec{v} \sim \mathcal{V}}[v(OPT)]$$

Feldman et al. [29] implemented this idea in a slightly different, more probabilistic analysis. The aim is to assign for each item j price $p_j = \frac{1}{2}\mathbb{E}_{\vec{v} \sim \mathcal{V}}[q_j^{\mathcal{A}}]$. However, we have only sample access to the distribution $\mathcal{V}$.Therefore, we have to estimate p_j.

Lemma 5.6. *Let p'_j be the average of $t = (\log m + \log n - \log \varepsilon)4m^2/\varepsilon^2$ identical samples from distribution $\mathcal{V}$, with expected value p_j. Then, with probability at least $1 - \varepsilon/n$, we have*

$$|p_j - p'_j| < \frac{\varepsilon}{2m}$$

for all items j.

Now, we are ready to proceed to the main lemma:

Lemma 5.7. *Given a distribution $\mathcal{V}$ over XOS valuations, let $\vec{p}$ be the price vector defined as $p_j = \frac{1}{2}\mathbb{E}_{\vec{v} \sim \mathcal{V}}[q_j^{\mathcal{A}}]$. Let $\vec{p}\,'$ be any price vector such that $|p_j - p'_j| < \frac{\varepsilon}{2m}$ for all items j. Then for any arrival order π, a posted price mechanism with prices $\vec{p}\,'$ results in an expected welfare at least $\frac{1}{2}\mathbb{E}_{\vec{v} \sim \mathcal{V}}[v(\mathcal{A})] - \frac{\varepsilon}{2}$.*

Algorithm 9 Feldman et al. [29] - $O(1)$ mechanism for XOS valuations in Bayesian setting

for each item $j \in U$ **do**
  $Q_j \leftarrow 0$
  Repeat t times:
    Draw $\vec{v} \sim \mathcal{V}$ and let $T = \mathcal{A}(\vec{v})$.
    Let q_j be the contribution of item j in $\mathcal{A}$.
    $Q_j \leftarrow Q_j + q_j$
  $p'_j \leftarrow \frac{1}{2t} Q_j$.
end for
Run a posted price auction with prices $p'_1, \ldots, p'_m$.

Formally, the mechanism is described in Algorithm 9. It is clear, that in this mechanism, the `Sample and Estimation` component is omitted, since we can learn all the information needed through samples from the distribution. The analysis is almost identical to the analysis in Section 6.5 for the case of *Liquid Welfare*.

5.5 Incentives: Dominant Strategy vs Ex-Post Nash Equilibrium

In Chapter 4 we defined a new notion of incentives' guarantee, the *Ex-Post Nash Equilibrium*. This definition was necessary for reasoning about the strategy profile of the bidders in iterative mechanisms, such as ascending price auctions. The main reason is that in an ascending price procedure, the actions' space can be history dependent, since it proceeds in rounds and every bidder can act according to the actions of the others. The question that may arises is how the mechanisms, presented in the previous sections, differ from the framework of chapter 4. The main difference is that when bidders arrive at the auction and choose their utility-maximizing bundle, they do not affect the price increase, as long as they are to participate again in the auction. Even in case of Algorithm 8, where bidders in UNIFORM participate in an ascending price auction in Step 7, the prices offered to them does not get influenced by their decisions in previous rounds, as they are pre-defined by the values of the bins; they only affect the prices that are posted to the bidders in FINAL. Therefore, it is *Dominant Strategy* for the agents to choose their utility-maximizing bundle and these mechanisms are truthful.

Chapter 6

Liquid Welfare

6.1 Introduction

In the previous chapters, we studied *Combinatorial Auctions* and *Posted-Price* mechanism, trying to maximize the *Social Welfare*, i.e. $\sum_{i=1}^{n} v_i(S_i)$, where S is the allocation of the mechanism. However, we overlooked a crucial parameter for real life applications; the *Budget* constraints. The budget constraints become more important as the magnitude of the financial transactions augments, as in spectrum auctions, where hundreds of millions of dollars are paid. Even in Google Adwords, the bidders are asked their available budget, even before bids or keywords. Until now, we designed truthful mechanism with good approximating guarantees, without taking into account the potential of the bidders to pay the asked amount. But what should we do in a second price auction, if the winning bidder has budget below the second highest bid? One thought would be to give him the item and get all of his budget. Nonetheless, such an action would violate the most fundamental property of mechanism design, the truthfulness of the mechanism. In order to make it more clear, imagine a single item auction, where bidder i has $v_i = 2$ and budget $B_i = 1$. Then, knowing that we can announce a much higher bid, for example $b_i = 100$ and avoid paying more than 1, why would he report his true preference? However, we said that VCG is a DSIC mechanism for much more general environments. Why does it fail now? The reason is, that the *utility* of the bidders stops being *quasi-linear* and is modified as follows:

$$u_i = \begin{cases} v_i - p, & \text{if } p \leq B_i \\ -\infty & \text{otherwise} \end{cases}$$

The aformentioned example illustrates an essential problem for the efficiency measure under budget constraints. In fact, *Social Welfare* cannot be approximated by a factor better than n, where n is the number of the agents, even when bidders' budgets are known to the auctioneer. The n-approximation is succeeded by allocating the item uniformly at random and charge nothing. Dobzinski and Leme [19] introduced a new notion of efficiency for budgeted settings, a measure that balances between the willingness and ability to pay, the so called *Liquid Welfare*. Informally, *Liquid Welfare* is defined as the minimum between each agent's budget and his valuation for the acquired bundle of items. In other words, this measure of efficiency respects the available funds of each bidder and as a result, denotes a compromise between the social welfare and the maximum revenue extraction, which is upper bounded by the budgets.

In this chapter, we are going to extend the results of *Posted-Price* mechanisms in *Social Welfare* to *Liquid Welfare*, but first we will present the some results in the multi-unit case.

6.2 Definitions and Previous Results

First, we are going to give the notion of *Liquid Valuation* and *Liquid Welfare*, as defined in [19]. Formally, we have:

Definition 6.1 (Liquid Valuation). *In a budgeted setting , we define the liquid valuation of a bidder in allocation $S = (S_1, \dots, S_n)$ by:*

$$\bar{v}_i(S_i) = \min\{v_i(S_i), B_i\} \tag{6.1}$$

Definition 6.2 (Liquid Welfare). *In a budgeted setting with n bidders, we define the Liquid Welfare of allocation $S = (S_1, \dots, S_n)$ by:*

$$LW = \sum_{i=1}^{n} \bar{v}_i(S_i) \tag{6.2}$$

At this point, the question that arises is whether we can implement VCG on liquid valuations. In fact, for single-item case, the aswer is positive. Formally we have, according to [19]:

Theorem 6.1. *For single-item case with indivisible item, VCG on modified values $\bar{v}_i = \min\{v_i, B_i\}$ is DSIC and exactly optimizes the liquid welfare objective.*

The proof is similar to that of Theorem 3.2, since if an agents overbids, he may has to pay more than his budget. Nevertheless, for multi-unit auctions with divisible item, it holds:

Lemma 6.1. *For multi-unit auctions with divisible item, the allocation that maximizes the Liquid Welfare occurs for $x_i^* = \min\left(\frac{B_i}{v_i}, \left[1 - \sum_{j<i} x_j^*\right]^+\right)$, where agents are sorted in non-increasing order of value, i.e. $v_1 \geq \cdots \geq v_n$.*

However, it is clear that this allocation rule is not monotone. To see this, the amount the highest bidder is allocated, is decreasing by the increase of his bid. Therefore, we had to design approximation mechanisms that guarantee truthfulness and the main technic was *Posted-Price* mechanisms.

Dobzinski and Leme [19] proved a 2-approximation to the optimal *Liquid Welfare* for the case of a single divisible item and additive bidders with public budgets and a $O(\log n)$ (resp. $(\log^2 n)$)-approximation to the optimal *Liquid Welfare* for submodular (resp. subadditive) bidders with private budgets. Lu and Xiao [45] proved that the optimal *Liquid*

Welfare can be approximated truthfully within constant factor for a single divisible good, general valuations and private budgets. Closer to our setting, Lu and Xiao [46] provided a truthful mechanism that achieves a constant factor approximation to the *Liquid Welfare* for multi-item auctions with divisible items and additive valuations, under a large market assumption. Under similar large market assumptions, [26] obtained mechanisms that approximate the optimal revenue within a constant factor for multi-unit online auctions with divisible and indivisible items, and a mechanism that achieves a constant approximation to the optimal *Liquid Welfare* for additive valuations over divisible items. However, prior to our work, there was no work on approximating the *Liquid Welfare* in *Combinatorial Auctions* (in fact, that was one of the open problems in [19]). In the next sections, we are going to extend the results of *Posted-Price* mechanisms for *Social Welfare* to *Liquid Welfare*.

6.3 Liquid Welfare in Combinatorial Auctions

As we said in the previous section, prior to our work, there was no work on approximating the *Liquid Welfare* in *Combinatorial Auctions*. The main difference between the multi-unit and multi-item setting is the heterogeneity of the items; in multi-item setting, agents have to "divide" their budget into many items, whereas in the multi-unit setting they have to buy the most their budget allows them.

6.3.1 Approach

If we are going to extend the results of chapter 3, we have to ensure that liquid valuation functions belong to the same class as the valuation function. Formally, we have:

Lemma 6.2. *Let v be a non-negative monotone submodular (resp. XOS, subadditive) function. Then, for any $B \in \mathbb{R}_{\geq 0}$, $\bar{v} = \min\{v, B\}$ is also monotone submodular (resp. XOS, subadditive).*

Proof. Clearly, capping valuation with budget does not affect monotonicity. We provide the proof for each case (i.e., submodular, XOS, subbaditive) separately.

- (submodular) Let v be a monotone submodular set function. Then, by the definition of submodularity, for sets $T \subseteq S$ and $j \notin S$ we have:

  $$v\left(S \cup \{j\}\right) - v\left(S\right) \leq v\left(T \cup \{j\}\right) - v\left(T\right)$$

  Further, since v is monotone: $v(T) \leq v(S)$, which implies that $\bar{v}(T) \leq \bar{v}(S)$. We distinguish the following cases:

  1. If $B \leq v\left(T \cup \{j\}\right) \leq v\left(S \cup \{j\}\right)$. Then, for the liquid valuations we have: $\bar{v}\left(S \cup \{j\}\right) - \bar{v}\left(S\right) = B - \bar{v}\left(S\right) \leq B - \bar{v}(T) \leq \bar{v}\left(T \cup \{j\}\right) - \bar{v}(T)$, where the first inequality is due to monotonicity.

  2. If $\bar{v}(T \cup \{j\}) \leq \bar{v}(S \cup \{j\}) \leq B$. Then, $\bar{v}(S \cup \{j\}) - \bar{v}(S) = v(S \cup \{j\}) - v(S) \leq v(T \cup \{j\}) - v(T) = \bar{v}(T \cup \{j\}) - \bar{v}(T)$.

3. If $v(T \cup \{j\}) \leq B \leq v(S \cup \{j\})$. This breaks down to the following two cases; on the one hand, if $v(S) \geq B$ then, $\bar{v}(S \cup \{j\}) - \bar{v}(S) = 0 \leq v(T \cup \{j\}) - v(T) = \bar{v}(T \cup \{j\}) - \bar{v}(T)$. On the other hand, if $v(S) < B$, then $\bar{v}(S \cup \{j\}) - \bar{v}(S) = B - v(S) \leq v(S \cup \{j\}) - v(S) \leq v(T \cup \{j\}) - v(T) = \bar{v}(T \cup \{j\}) - \bar{v}(T)$. Finally, we remark that due to monotonicity, these cases are the only possible ones.

- (XOS) Let v be an XOS set function; there exist additive functions $\alpha_1, \ldots, \alpha_l$ s.t. $v(S) = \max_{i \in [l]} \alpha_i(S)$. In order for $\bar{v}$ to XOS, we need to prove that there exist additive functions $\alpha'_1, \ldots, \alpha'_k$ s.t. $\bar{v}(S) = \max_{i \in [k]} \alpha'_i(S)$. For each function α_i we are going to define $m!$ functions, one for each permutation π of the items. Suppose a specific ordering π_t of the items $\{1, 2, \ldots, m\}$ and let $\pi_t(j)$ be the position of item j in ordering π_t. We define $\beta_i^{\pi_t}$ as:

$$\beta_i^{\pi_t}(\{j\}) = \begin{cases} \alpha_i(\{j\}), & \text{if } \sum_{k:\pi_t(k) \leq \pi_t(j)} \alpha_i(\{k\}) \leq B \\ \max\left\{B - \sum_{k:\pi_t(k) < \pi_t(j)} \alpha_i(\{k\}), 0\right\}, & \text{if } \sum_{k:\pi_t(k) \leq \pi_t(j)} \alpha_i(\{k\}) > B \end{cases}$$

First, we are going to prove that for each $S \subseteq U, \beta_i^{\pi_t}(S) \leq \min\{v(S), B\}, \forall i, \pi_t$.

By the definition of $\beta_i^{\pi_t}$, it is clear to see that $\beta_i^{\pi_t}(\{j\}) \leq \alpha_i(\{j\})$. Therefore, summing upon all items in S (since we have additive functions), we get that:

$$\beta_i^{\pi_t}(S) \leq \alpha_i(S) \leq \max_k \alpha_k(S) = v(S)$$

By the definition of $\beta_i^{\pi_t}$, we also have that $\beta_i^{\pi_t}(S) \leq B$.

Next, we are going to prove that for each $S \subseteq U : \exists \beta_i^{\pi_t}$ s.t. $\beta_i^{\pi_t}(S) = \min\{v(S), B\}$. We distinguish the following cases:

1. $v(S) \leq B$. Let π_t be a permutation, s.t. all items in S come first and let α_{i^*} be the maximizing function for set S, i.e. $v(S) = \alpha_{i^*}(S)$. Then, because $\sum_{j \in S} \alpha_{i^*}(\{j\}) \leq B$, we have $\beta_{i^*}^{\pi_t}(S) = \sum_{j \in S} \beta_{i^*}^{\pi_t}(\{j\}) = \sum_{j \in S} \alpha_{i^*}(\{j\}) = v(S)$.

2. $v(S) > B$. Let π_t be a permutation, s.t. all items in S come first and let α_{i^*} be the maximizing function for set S, i.e. $v(S) = \alpha_{i^*}(S)$. Let j^* be the last item in the permutation π_t s.t. $\sum_{r:\pi_t(r) \leq \pi_t(j^*)} \alpha_{i^*}(\{r\}) \leq B$. Then, $\sum_{r:\pi_t(r) \leq \pi_t(j^*)} \beta_{i^*}^{\pi_t}(\{r\}) = \sum_{r:\pi_t(r) \leq \pi_t(j^*)} \alpha_{i^*}(\{r\})$. For the next items $z \in S$ in permutation π_t, we have $\beta_{i^*}^{\pi_t}(\{z\}) = \max\{B - \sum_{k:\pi_t(k) < \pi_t(z)} \alpha_i(\{k\}), 0\}$. In fact, the first item after j^* will complete the missing value, in order to have: $\sum_{k:\pi_t(k) \leq \pi_t(j^*)+1} \beta_{i^*}^{\pi_t}(\{j\}) = B$, and all subsequent items, q will have $\beta_{i^*}^{\pi_t}(\{q\}) = 0$. Therefore, $\sum_{j \in S} \beta_{i^*}^{\pi_t}(\{j\}) = B$.

- (subadditive) Let v be a monotone subadditive set function. Then, by the definition of subadditivity, for sets T, S we have:

$$v\left(S \cup T\right) \leq v\left(T\right) + v\left(S\right)$$

  We distinguish the following cases:

  1. If $\bar{v}(S \cup T) = v(S \cup T) < B$. Then, we know for a fact that $\bar{v}(S) = v(S) < B$ and that $\bar{v}(T) = v(T) < B$. Then, $\bar{v}(S \cup T) = v(S \cup T) \leq v(S) + v(T) = \bar{v}(S) + \bar{v}(T)$, where the inequality comes from the subadditivity of v.

  2. If $\bar{v}(S \cup T) = B < v(S \cup T)$. We have to further distinguish the following cases:

     (a) $\bar{v}(S) = B < v(S), \bar{v}(T) = B < v(T)$. Then, $\bar{v}(S \cup T) = B \leq 2B = \bar{v}(S) + \bar{v}(T)$.

     (b) $\bar{v}(S) = B < v(S), \bar{v}(T) = v(T) < B$. Then, $\bar{v}(S \cup T) = B \leq B + v(T) = \bar{v}(S) + \bar{v}(T)$, where the inequality comes from the non-negativity of the liquid valuation.

     (c) $\bar{v}(S) = v(S) < B, \bar{v}(T) = B < v(T)$. Then, $\bar{v}(S \cup T) = B \leq v(S) + B = \bar{v}(S) + \bar{v}(T)$, where the inequality again comes from the non-negativity of the liquid valuation.

     (d) $\bar{v}(S) = v(S) < B, \bar{v}(T) = v(T) < B$. Then, $\bar{v}(S \cup T) = B \leq v(S \cup T) \leq v(S) + v(T) = \bar{v}(S) + \bar{v}(T)$, where the last inequality comes from the fact that v is subadditive.

□

Then, one might think that he can directly use the mechanisms of e.g., [41, 16, 29] with valuation functions $\bar{v} = \min\{v, B\}$ and demand queries of the form: $\mathrm{DQ}(\min\{v, B\}, U, \vec{p})$ (i.e., wrt. the liquid valuation of the bidders) and obtain the same approximation guarantees but now for the LW. However, the resulting mechanisms are no longer truthful; bidders still seek to maximize their *utility* (i.e., value minus price) from the bundle that they get, subject to their budget constraint, rather than their *liquid utility* (i.e., liquid value minus price). Specifically, given a set of items U available at prices p_j, $j \in U$, a budget-constrained bidder i wants to receive the bundle $S_i = \arg\max_{S \subseteq U}\{v_i(S) - p(S) \,|\, p(S) \leq B_i\}$, and might not be happy with the bundle $S_i' = \arg\max_{S \subseteq U}\{\bar{v}_i(S) - p(S)\}$ computed by the demand query for the liquid valuation, as the following example highlights:

Example 6.1. *Consider a bidder with budget $B = 2$ and two items a and b available at prices $p_a = 2$ and $p_b = 1$. Assume that the bidder's valuation function is $v(\{a\}) = v(\{a, b\}) = 10$, $v(\{b\}) = 2$ (and therefore, her liquid valuation is $\bar{v}(\{a\}) = \bar{v}(\{b\}) = \bar{v}(\{a, b\}) = 2$). The bidder wants to get item a at price 2, which gives her utility 8. However, the demand query for her liquid valuation function $\bar{v}$ allocates item b, which gives her a utility of 1. Clearly, in this example, the bidder would have incentive* *to misreport her preferences to the demand query oracle.*

To restore truthfulness, we replace demand queries with *budget-constrained demand queries*. A budget-constrained demand query, denoted by BCDQ$(v, U, \vec{p}, B)$, specifies a valuation function v, a set of available items U, a price p_j for each $j \in U$ and a budget B, and receives the set $S \subseteq U$ maximizing $v(S) - p(S)$, subject to $p(S) \leq B$, i.e., the set of available items that maximizes bidder's utility subject to her budget constraint. Formally, we have:

Definition 6.3 (BCDQ)**.** *Let U be the set of items that are available. Then, bidder i's BCDQ returns set $S_i \subseteq U$ satisfying:*

$$S_i = \arg\max_{S \in U} \left\{ v_i(S) - p\left(S\right) \middle| p(S) \leq B_i \right\} \tag{6.3}$$

To establish the approximation ratio, we first observe that the fact that liquid valuations are XOS suffices for estimating supporting prices, as in previous work on the SW. Additionally, we have to show that the bundles allocated by BCDQ$(v, U, \vec{p}, B)$ approximately satisfy the efficiency guarantees of DQ on the liquid welfare and the liquid utility of the allocated bundles. In fact, it holds:

Lemma 6.3. *Let $S \subseteq U$ be the set allocated by the BCDQ for a bidder with valuation v and budget B. Then, for every other $T \subseteq U$, the following hold:*

1. $\bar{v}(S) \geq \bar{v}(T) - p(T)$
2. $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$.

Proof. We will prove each claim of the lemma separately. For claim 1, notice that if $p(T) > B$, then the Right Hand Side (RHS) of the inequality will be negative and thus, the inequality trivially holds. So, we will focus on the case where $p(T) \leq B$. We distinguish the following cases:

1. ($\bar{v}(S) = v(S)$ and $\bar{v}(T) = v(T)$.) Hence, $B \geq v(T)$. Bundle T was *considered* at the time of the query and yet, the query returned set S. Thus: $\bar{v}(S) \geq \bar{v}(S) - p(S) = v(S) - p(S) \geq v(T) - p(T) = \bar{v}(T) - p(T)$.

2. ($\bar{v}(S) = B$ and $\bar{v}(T) = B$) Then, the inequality trivially holds since: $B \geq B - p(T)$ and prices are non-negative.

3. ($\bar{v}(S) = B$ and $\bar{v}(T) = v(T)$) The inequality holds since: $B \geq B - p(T) \geq v(T) - p(T) = \bar{v}(T) - p(T)$.

4. ($\bar{v}(S) = v(S)$ and $\bar{v}(T) = B$) Hence, $B \leq v(T)$. Bundle T was *considered* at the time of the query and yet, the query returned set S. Thus, $\bar{v}(S) \geq \bar{v}(S) - p(S) = v(S) - p(S) \geq v(T) - p(T) \geq B - p(T) = \bar{v}(T) - p(T)$.

This concludes our proof for claim 1.

For claim 2, notice that since S is the set received from the BCDQ, then it is *affordable*: $\bar{v}(S) \geq p(S)$. Adding this inequality to the inequality of claim 1, we have that: $2\bar{v}(S) - p(S) \geq \bar{v}(T) - p(T)$.

□

Now that we have ensured the truthfulness and the efficiency guarantees of the BCDQs, we can examine whether we can implement the same components as in Section 5.2.2 for the objective of *Liquid Welfare*.

6.3.2 Components

First of all, we have to check whether the $\gamma - supportingprices$ definitions can hold with *liquid valuations*. The condition we have to examine is the *strong profitability* of a set chosen by the BCDQ. Therefore, we have the following lemma:

Lemma 6.4. *Let* $S = \arg\max_{S' \subseteq U} \{v(S') - p(S') | p(S') \leq B\}$ *a* strongly profitable *set under item prices* $p_1, \ldots, p_m$ *for valuation* v*. Then,* S *is also a strongly profitable set for the liquid valuation* $\bar{v}$*.*

Proof. Let $T \subseteq S$. We want to show that $\bar{v}(T) \geq p(T)$. If $\bar{v}(T) = v(T)$, then the property holds, since S is strongly profitable for valuation v. If $\bar{v}(T) = B$, then, due to monotonicity of $\bar{v}$, $\bar{v}(T) = \bar{v}(S) \geq p(S) \geq p(T)$, where the first inequality comes from individual rationality.

□

Now, if we want to find *supporting prices* for *Liquid Welfare*, we need a variant of Lemma 5.1 for *liquid valuations*, in order to guarantee high welfare. Therefore, we have:

Lemma 6.5 (Extension of Lemma 4.2 in [16])**.** *Let* $\alpha = (\alpha_1, \ldots, \alpha_n)$ *be an allocation that is supported by prices* $p_1, \ldots, p_m$*. A fixed price auction where budget constrained bidders make BCDQs and the items have prices* $p'_j = \frac{p_j}{2}$ *generates an allocation* $\hat{A} = (\hat{\alpha}_1, \ldots, \hat{\alpha}_n)$ *with LW:* $\sum_{i \in [n]} \bar{v}_i(\hat{\alpha}_i) \geq \frac{\sum_{j \in \cup_i \alpha_i} p_j}{4}$.

Proof. We will follow closely the proof presented by Dobzinski [16], changing the analysis only slightly when it is required to reason about the behavior of the BFDO.

For every bidder i, let $W_i = \cup_{i' < i} \hat{\alpha}_{i'}$ denote the set of competitive items that were allocated before bidder i arrives to the auction. Let $\text{OPT}_i = \sum_{j \in (\cup_{i' \geq i} \alpha_{i'}) \setminus W_i} p_j$. Then, $\text{OPT}_1 = \sum_{j \in \cup_i \alpha_i} p_j$ and $\text{OPT}_{n+1} = 0$. For every bidder $i \in [n]$ it holds that $W_{i+1} = W_i + \hat{\alpha}_i$ and that the allocation $(\emptyset, \ldots, \emptyset, \alpha_i \setminus W_i, \alpha_{i+1} \setminus W_i, \ldots, \alpha_n \setminus W_i)$ is still supported by $p_1, \ldots, p_m$. Thus,

$$\text{OPT}_i - \text{OPT}_{i+1} = \sum_{j \in (\alpha_i \setminus W_i)} p_j + \sum_{j \in \hat{\alpha}_i} p_j \tag{6.4}$$

Now notice that bidder i could buy set $\alpha_i - W_i$ which implies that the liquid valuation that he got from the set that was ultimately received by the BCDQ was lower bounded by:

$$\begin{aligned} \bar{v}_i(\alpha_i \setminus W_i) - \sum_{j \in (\alpha_i \setminus W_i)} q_j &= \bar{v}_i(\alpha_i \setminus W_i) - \sum_{j \in (\alpha_i \setminus W_i)} \frac{p_j}{2} \\ &\geq \sum_{j \in (\alpha_i \setminus W_i)} \frac{p_j}{2} \end{aligned} \tag{6.5}$$

Since the bidder had enough budget to buy set $\hat{\alpha}_i$ (otherwise, it would not have been received as the answer of the BCDQ) we have that:

$$\bar{v}_i(\hat{\alpha}_i) \geq \sum_{j \in \hat{\alpha}_i} \frac{p_j}{2} \tag{6.6}$$

Summing up Equations(6.5) and (6.6) and using Equation (6.4) we get:

$$2\bar{v}_i(\hat{\alpha}_i) \geq \sum_{j \in (\alpha_i \setminus W_i)} \frac{p_j}{2} + \sum_{j \in \hat{\alpha}_i} \frac{p_j}{2} = \frac{\mathrm{OPT}_i - \mathrm{OPT}_{i+1}}{2}$$

which concludes our proof.

□

At this point, we need to test if the components Section 5.2.2 are violated because of the *liquid valuations*. The `Dominant Bidder` phase remains the same, since as we showed in Theorem 6.1, the VCG mechanism for single-item setting is DSIC and maximizes the *Liquid Welfare*. The same holds for the `Sample and Estimation` component, since *liquid valuation* functions belong to the same class as *valuation* functions (for submodular, XOS, subadditive functions) and therefore, we can implement the same algorithms, as agents have no incentive to lie. Finally, the `Price-Exploration` and `Price-Exploitation` phases are executed with the usage of BCDQs instead of *Demand Queries* and therefore, we have the efficiency guarantees of BCDQs, according to Lemma 6.3.

Conceptually, we present a general approach through which known truthful approximations to the SW, that access valuations through demand queries, can be adapted so that they retain truthfulness and achieve similar approximation guarantees for the LW. The important properties required are that liquid valuation functions $\bar{v}$ belong to the same class as valuation functions v (proven for submodular, XOS and subadditive valuations), and that the efficiency guarantees of budget-constrained demand queries on liquid welfare and liquid utility are similar to the corresponding efficiency guarantees of standard demand queries for liquid valuations (proven for all classes of valuations functions). Indeed, applying this approach to the mechanism of [16], we obtain a universally truthful mechanism that approximates the LW for CAs with XOS bidders within a factor of $O(\sqrt{\log m})$. In the next sections, we will prove the approximation guarantees of the mechanisms in Krysta and Vöcking [41], Feldman et al. [29] but for the LW, using BCDQs, since our technics become more clear in these mechanisms than in [16] .

6.4 Worst-Case setting

In this section, we present Algorithm 10, which is based on the mechanism of [41] and achieves the same approximation guarantee for the objective of *Liquid Welfare*. The only difference is that budget-constrained bidders in Algorithm 10 are restricted to using BCDQs, instead of DQs, thus making the mechanism universally truthful. Resembling the analysis of [41], we show that for $1/q = \Theta(\log m)$, Algorithm 10 achieves an approximation ratio of $O(\log m)$ for the LW. First, we note that parameter[1] L is selected so that there exists only one bidder whose liquid valuation for U (weakly) exceeds it.

[1] L can be computed with standard techniques, as explained in [41].

Algorithm 10 KV-Mechanism for Liquid Welfare

1: Fix an ordering π of bidders and set $U_1 = U$.
2: Set initial prices $p_1^{(1)} = \cdots = p_m^{(1)} = \frac{L}{4m}$.
3: **for** each bidder $i = 1, \ldots, n$ according to π **do**
4: Let $S_i = \text{BCDQ}(v_i, U_i, \vec{p}^{\,(i)}, B_i)$
5: With probability q, allocate $R_i = S_i$ to i and set $U_{i+1} = U_i \setminus S_i$. Otherwise, set $U_{i+1} = U_i, R_i = \emptyset$.
6: Update prices $\forall j \in S_i$: $p_j^{(i+1)} = 2p_j^{(i)}$.
7: **end for**

Theorem 6.2. *Algorithm 10 is universally truthful and for $q = 1/\Theta(\log m)$, achieves an approximation ratio of $O(\log m)$ for the LW.*

We present a series of Lemmas that will lead us naturally to the proof of the Theorem. Let $\mathcal{S} = (S_1, \ldots, S_n)$ and $\mathcal{R} = (R_1, \ldots, R_n)$ the provisional and the final allocation of Algorithm 10 respectively. First, we provide two useful bounds on $\bar{v}(\mathcal{S})$. We find it important to also discuss an overselling variant of Algorithm 10. In the *Overselling* variant, allow us to assume that for Step 5 of Algorithm 10, $q = 1$ (i.e., S_i is allocated to bidder i with certainty) and $U_{i+1} = U_i = U$ (thus the name of the variant). The *Overselling* variant allocates at most $k = \log(4m) + 2$ copies of each item and collects a liquid welfare within a constant factor of the optimal LW. To see that, observe that for $q = 1$, after allocating $k - 1$ copies of some item j, j's price becomes $\frac{L}{4m} 2^{\log(4m)+1} = 2L$. Then, there is only one agent with liquid valuation larger than L who can get a copy of j.

Lemma 6.6. *Let p_j denote the final price of each item j. Then, for* any *sets $U_1, \ldots, U_n \subseteq U$ of items available when the bidders arrive, Algorithm 10 with $q = 1$ satisfies $\bar{v}(\mathcal{S}) \geq \sum_{j \in U} p_j - L/4$.*

Proof. Since bidders are individually rational and do not violate their budget constraints, for every bidder i it holds that $B_i \geq \sum_{j \in S_i} p_j^{(i)}$ and $v_i(S_i) \geq \sum_{j \in S_i} p_j^{(i)}$. The rest of the proof is identical to the proof of [41, Lemma 2] for $b = 1$. Specifically, let $\ell_j^{(i)}$ be the number of copies of item j allocated just before bidder i arrives, and let ℓ_j be the total number of copies of item j allocated by Algorithm 10 with $q = 1$. Then,

$$\begin{aligned}\bar{v}(\mathcal{S}) &\geq \sum_{i=1}^{n} \sum_{j \in S_i} p_j^{(i)} = \frac{L}{4m} \sum_{i=1}^{n} \sum_{j \in S_i} 2^{\ell_j^{(i)}} \\ &= \frac{L}{4m} \sum_{j \in U} (2^{\ell_j} - 1) = \sum_{j \in U} p_j - L/4\,,\end{aligned}$$

where we have changed the order of summation and we have used the fact that $p_j = \frac{L}{4m} 2^{\ell_j}$. $\square$

Lemma 6.7. *For sets $U_1 = \cdots = U_n \subseteq U$, the* Overselling *variant of Algorithm 10 with $q = 1$ satisfies $\bar{v}(\mathcal{S}) \geq \overline{OPT} - \sum_{j \in U} p_j$.*

Proof. Let $\mathcal{O} = (O_1, \ldots, O_n)$ be the optimal allocation. From Lemma 6.3, we get that for each bidder i:

$$\bar{v}(S_i) \geq \bar{v}(O_i) - \sum_{j \in O_i} p_j^{(i)} \geq \bar{v}(O_i) - \sum_{j \in O_i} p_j$$

where we use that the final price of each item is the largest one. Summing over all bidders, we have that:

$$\bar{v}(\mathcal{S}) \geq \bar{v}(\mathcal{O}) - \sum_{i=1}^{n} \sum_{j \in O_i} p_j \geq \overline{\mathrm{OPT}} - \sum_{j \in U} p_j$$

where the last inequality uses the fact that the optimal solution is feasible and thus, each item is allocated at most once in $\mathcal{O}$.

□

Lemma 6.8. *The* Overselling *variant of Algorithm 10 with $q = 1$ allocates at most $\log(4m) + 2$ copies of each item and computes an allocation $\mathcal{S}$ with liquid welfare $\bar{v}(\mathcal{S}) \geq \frac{3}{8}\overline{OPT}$.*

Proof. Summing the equations from Lemma 6.6, Lemma 6.7 we have:

$$2\bar{v}(\mathcal{S}) \geq \overline{\mathrm{OPT}} - \frac{L}{4}$$

Since $\overline{\mathrm{OPT}} \geq L$, we have:

$$2\bar{v}(\mathcal{S}) \geq \overline{\mathrm{OPT}} - \frac{1}{4}\overline{\mathrm{OPT}} \Rightarrow \bar{v}(\mathcal{S}) \geq \frac{3}{8}\overline{\mathrm{OPT}}$$

□

Of course, the allocation $\mathcal{S}$ in Lemma 6.8 is highly infeasible, since it allocates a logarithmic number of copies of each item. The remedy is to use an allocation probability $q = 1/\Theta(\log m)$. For such values of q, we can plugin the proof of [41, Lemma 6] (which just uses that the valuation functions are fractionally subadditive) and show that for each agent i and for all $A \subseteq U$, $\mathbb{E}[\bar{v}_i(A \cap U_i)] \geq \bar{v}_i(A)/2$. We are now ready to conclude the proof of Theorem 6.2.

Lemma 6.9. *For Algorithm 10 with $q = \frac{1}{4(\log(4m)+1)}$, it holds that $\mathbb{E}[\bar{v}(\mathcal{S})] \geq \overline{OPT}/8$ and $\mathbb{E}[\bar{v}(\mathcal{R})] \geq q\,\overline{OPT}/8$.*

Proof. Let $\mathcal{O} = (O_1, \ldots, O_n)$ be the optimal allocation. For each bidder i, Lemma 6.3 implies that the response S_i of BCDQ satisfies:

$$\bar{v}_i(S_i) \geq \bar{v}_i(O_i \cap U_i) - \sum_{j \in O_i \cap U_i} p_j^{(i)}$$

for any U_i resulted from the outcome of the random coin flips. Therefore:

$$\mathbb{E}[\bar{v}_i(S_i)] \geq \mathbb{E}[\bar{v}_i(O_i \cap U_i)] - \mathbb{E}[\sum_{j \in O_i \cap U_i} p_j^{(i)}]$$

By the choice of q, for any bidder i, $\mathbb{E}[\bar{v}_i(O_i \cap U_i)] \geq \bar{v}_i(O_i)/2$. Then, working with the expectations as in the proofs of Lemma 6.6, Lemma 6.7, we have:

$$\mathbb{E}[\bar{v}(\mathcal{S})] \geq \mathbb{E}\left[\sum_{j \in U} p_j\right] - L/4$$

and:

$$\mathbb{E}[\bar{v}(\mathcal{S})] \geq \frac{1}{2}\overline{\text{OPT}} - \mathbb{E}\left[\sum_{j \in U} p_j\right]$$

Summing the above equations and using the fact that $\overline{\text{OPT}} \geq L$ we have:

$$\mathbb{E}[\bar{v}(\mathcal{S})] \geq \overline{\text{OPT}}/8$$

Finally, one can use linearity of expectation and show that $\mathbb{E}[\bar{v}(\mathcal{R})] = q\,\mathbb{E}[\bar{v}(\mathcal{S})]$. □

6.5 Bayesian setting

In this setting, let $\vec{v} = (v_1, \ldots, v_n)$ be a profile of bidder valuations and $\vec{B} = (B_1, \ldots, B_n)$ a profile of bidder budgets. Assume that the bidders' valuations are drawn independently from distributions $\mathcal{V}_1, \ldots, \mathcal{V}_n$ and the budgets from distributions $\mathcal{B}_1, \ldots, \mathcal{B}_n$. For simplicity, let us assume that their liquid valuations are drawn independently from distributions $\mathcal{D}_1, \ldots, \mathcal{D}_n$. We will denote by $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ the product distribution where liquid valuations profiles, $\bar{\mathbf{v}} = (\bar{v}_1, \ldots, \bar{v}_n)$, are independently drawn from.

We are going to show that the results presented in [29] can be extended for budget-constrained bidders. Specifically, we are going to show that, if liquid valuations are fractionally subadditive, then we can create appropriate prices such that, when presented to the bidders in a posted-price mechanism and bidders are making BCDQs, then we can obtain universally truthful constant-factor approximation mechanisms for the LW in Bayesian CAs. Our Lemma 6.10 establishes the existence of such appropriately scaled prices. The key component activating our results is that instead of reasoning about the *utility* achieved from the bundle purchased by bidder i (as received by the BCDQ), we instead have to use Lemma 6.3.

Theorem 6.3. *Let distribution $\mathcal{D}$ over XOS liquid valuation profiles be given via a sample access to $\mathcal{D}$. Suppose that for every $\bar{\mathbf{v}} \sim \mathcal{D}$, we have:*

1. *black-box access to a LW maximization algorithm, ALG*[2] *for CAs.*
2. *an XOS value query oracle (for liquid valuations sampled from $\mathcal{D}$)*[3]*.*

Then, for any $\varepsilon > 0$, we can compute item prices in $POLY(m, n, 1/\varepsilon)$ time such that, for any bidder arrival order, the expected liquid welfare of the posted price mechanism is at least $\frac{1}{4}\mathbb{E}_{\bar{\mathbf{v}} \sim \mathcal{D}}[\bar{\mathbf{v}}(ALG(\bar{\mathbf{v}}))] - \varepsilon$, where by $ALG(\bar{\mathbf{v}})$ we denote the solution produced by algorithm ALG.

[2]ALG can be any algorithm that provides a $O(1)$-approximation to the optimal LW, since we do not care about incentives (access to ALG will only happen for ghost samples). For example, it could be the greedy algorithm by [42].

[3]An XOS value oracle takes as input a set T and returns the corresponding additive representative function for the set T, i.e., an additive function $A_i(\cdot)$, such that (i) $\bar{v}_i(S) \geq A_i(\hat{S})$ for any $\hat{S} \subset [m]$ and (ii) $\bar{v}_i(T) = A_i(T)$.

Lemma 6.10. *Given a distribution $\mathcal{D}$ over XOS liquid valuations, let $\vec{p}$ be the price vector s.t. $p_j = \frac{1}{2}\mathbb{E}_{\bar{\mathbf{v}}\sim\mathcal{D}}[LW_j(\bar{\mathbf{v}})]$, where $LW_j(\bar{\mathbf{v}})$ is the contribution of item j to the liquid welfare under valuation profile $\bar{\mathbf{v}}$. Let $\vec{p}\,'$ be any price vector such that $|p'_j - p_j| < \delta$ for all $j \in [m]$. Then, for any arrival order, π, bidders buying bundles by making BCDQs under prices $\vec{p}\,'$ results in expected liquid welfare at least $\frac{1}{4}\mathbb{E}_{\bar{\mathbf{v}}\sim\mathcal{D}}[\bar{\mathbf{v}}(ALG(\bar{\mathbf{v}}))] - \frac{m\delta}{2}$.*

Proof. We are going to follow the proof presented by [29]. For each $j \in S_i$, we denote by $\mathrm{LW}_j(\bar{\mathbf{v}}) := A_i(\{j\})$ (i.e., $\mathrm{LW}_j(\bar{\mathbf{v}})$ corresponds to the contribution of item j to the liquid welfare, under liquid valuation profile $\bar{\mathbf{v}}$), where $A_i(\cdot)$ is the corresponding additive representative function for the set S_i. From the definition of p_j:

$$\begin{aligned} p'_j &= \mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\mathrm{LW}_j(\bar{\mathbf{v}}) - p'_j\right] + 2(p'_j - p_j) \\ &= \sum_{i\in[n]} \mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\left(\mathrm{LW}_j(\bar{\mathbf{v}}) - p'_j\right)\mathbb{1}\left\{j \in S_i(\bar{\mathbf{v}})\right\}\right] + 2(p'_j - p_j) \end{aligned} \tag{6.7}$$

Let $\mathrm{SOLD}_i(\bar{\mathbf{v}}, \pi)$ be the set of items that have been sold prior to the arrival of bidder i. Bidder i's BCDQ receives set S_i as the answer, from the items in $U \setminus \mathrm{SOLD}_i(\bar{\mathbf{v}}, \pi)$ that maximizes $v(S_i) - p(S_i)$ subject to the fact that $p(S_i) \leq B_i$. Consider another random liquid valuation profile $\bar{\mathbf{v}}'_{-i} \sim \mathcal{D}_{-i}$, independent of $\bar{\mathbf{v}}$. Let $S_i(\bar{v}_i, \bar{\mathbf{v}}'_{-i})$ be the allocation returned by ALG on input $(\bar{v}_i, \bar{\mathbf{v}}'_{-i})$. For the additive representative function A_i for the set $S_i(\bar{v}_i, \bar{\mathbf{v}}'_{-i})$ it holds that $A_i(\{j\}) = \mathrm{LW}_j(\bar{v}_i, \bar{\mathbf{v}}'_{-i})$ for each $j \in S_i(\bar{v}_i, \bar{\mathbf{v}}'_{-i})$. Let $S_i(\bar{v}_i, \bar{\mathbf{v}}_{-i}, \bar{\mathbf{v}}'_{-i}) := S_i(\bar{v}_i, \bar{\mathbf{v}}'_{-i}) \setminus \mathrm{SOLD}_i(\bar{\mathbf{v}}, \pi)$ be the subset of items in $S_i(\bar{v}_i, \bar{\mathbf{v}}_{-i})$ that are available for purchase when bidder i arrives. Since bidder i could have bought set $S_i(\bar{v}_i, \bar{\mathbf{v}}_{-i}, \bar{\mathbf{v}}'_{-i})$ but instead did not, using Lemma 6.3 we get that:

$$2\bar{v}_i(S_i(\bar{\mathbf{v}})) - p(S_i(\bar{\mathbf{v}})) \geq \mathop{\mathbb{E}}_{\bar{\mathbf{v}}'_{-i}}\left[\max\left\{\mathrm{LW}_j(\bar{v}_i, \bar{\mathbf{v}}'_{-i}) - p'_j, 0\right\}\right]$$

Summing up for all the bidders and taking the expectation over all $\bar{\mathbf{v}} \sim \mathcal{D}$ we have:

$$\begin{aligned} &2\mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} \bar{\mathbf{v}}_i(S_i(\bar{\mathbf{v}}))\right] - \mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} p(S_i(\bar{\mathbf{v}}))\right] \\ &\geq \sum_{j\in\mathcal{M}}\sum_{i\in[n]} \mathop{\mathbb{E}}_{\bar{v}_i, \bar{\mathbf{v}}_{-i}, \bar{\mathbf{v}}'_{-i}}\left[\mathbb{1}\left\{j \in S_i(\bar{v}_i, \bar{\mathbf{v}}'_{-i})\right\} \cdot \max\left\{\mathrm{LW}_j(\bar{v}_i, \bar{\mathbf{v}}'_{-i}) - p'_j, 0\right\} \cdot \mathbb{1}\left\{j \neq \mathrm{SOLD}_i(\bar{\mathbf{v}}, \pi)\right\}\right] \end{aligned} \tag{6.8}$$

Following exactly the same steps as in [29] we can rewrite the above as:

$$\begin{aligned} &2\mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} \bar{\mathbf{v}}_i(S_i(\bar{\mathbf{v}}))\right] - \mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} p(S_i(\bar{\mathbf{v}}))\right] \\ &\geq \sum_{j\in U} \mathop{\mathbb{P}}_{\bar{\mathbf{v}}}\left[j \neq \mathrm{SOLD}(\bar{\mathbf{v}}, \pi)\right] \cdot (p_j + (p_j - p'_j)) \end{aligned} \tag{6.9}$$

For the expected revenue, due to individual rationality of the bidders it holds that:

$$\mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\mathrm{Rev}(\bar{\mathbf{v}}, \pi)\right] = \sum_{j\in U} \mathop{\mathbb{P}}_{\bar{\mathbf{v}}}\left[j \in \mathrm{SOLD}(\bar{\mathbf{v}}, \pi)\right] \cdot (p_j - (p_j - p'_j)) \tag{6.10}$$

Adding Equations (6.10) and (6.9) we get:

$$
\begin{aligned}
&2 \mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} \bar{\mathbf{v}}_i(S_i(\bar{\mathbf{v}}))\right] - \mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} p(S_i(\bar{\mathbf{v}}))\right] + \mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\mathrm{Rev}(\bar{\mathbf{v}},\pi)\right] \\
&\geq \sum_{j\in U} + \sum_{j\in U}(p_j - p'_j)\left(1 - 2\mathop{\mathbb{P}}_{\bar{\mathbf{v}}}\left[j \in \mathrm{SOLD}\left(\bar{\mathbf{v}},\pi\right)\right]\right) \\
&\geq \frac{1}{2}\mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} \bar{\mathbf{v}}_i(S_i)\right] - \sum_{j\in U}\left|p_j - p'_j\right| \\
&\geq \frac{1}{2}\mathop{\mathbb{E}}_{\bar{\mathbf{v}}\sim\mathcal{D}}\left[\sum_{i\in[n]} \bar{\mathbf{v}}_i(S_i)\right] - m\delta
\end{aligned}
$$

□

6.6 Large and Competitive Market

In this section, we are going to extend the notion of *Large Market Assumption* for multi-unit settings, as presented in [6, 26, 46], and introduce a new notion of a *competitive market* in the multi-item setting. Finally, we give a $O(1)$-approximation algorithm for *Liquid Welfare* with XOS bidders under the *competitive market* assumption.

6.6.1 Introduction

Intuitively, the large market assumption says that the contribution of a single bidder or a small group of agents to the whole market is negligent. This assumption is well founded in cases, such as the internet economy. In budgeted settings, this assumption can be expressed through the ratio of the budget of each agent to the optimal solution for multi-unit cases. Borgs et al. [6] were the first ones to define a *budget dominance parameter* that corresponded to the ratio of the maximum budget of all the bidders to the value of the optimum SW in the context of multi-unit auctions with budget-constrained bidders. More recently, Eden et al. [26] and Lu and Xiao [46] used similar notions of budget dominance[4] (termed *large market assumptions*) as a means to achieve constant factor approximation to the LW in multi-unit auctions and auctions with divisible items respectively, when bidders have additive valuation functions. However, for the case of *divisible* items, it is clear that the definition of a large market used in the previous cases, becomes almost void.

Example 6.2. *Imagine, a large market with m indivisible items and n bidders, s.t. $B_i \leq \frac{\overline{OPT}}{m\cdot c}$ for some large constant $c > 1$. The number of bidders who receive* at least *one item is* at most *m and therefore, $\overline{OPT} \leq m \cdot B_{\max}$, which leads to $B_{\max} \leq B_{\max}/c$, which is a contradiction.*

In reality, the previous settings possessed another crucial property, that made it possible for the large market assumption to activate the results about the constant factor

[4] Namely, that $\forall i \in [n] : B_i \leq \frac{\overline{\mathrm{OPT}}}{m\cdot c}$, where c is a large constant.

approximation of the optimal LW. This property was the *homogeneity* of the goods being auctioned; every bidder wanted exactly the same item or at least some portion of every item. The homogeneity of the goods, coupled with the large market assumption, essentially established *competitive markets*.

6.6.2 Preliminaries

Below, we first introduce our definition of Competitive Markets for indivisible goods and then, show how one can obtain a constant factor approximation of the optimal LW, when bidders have XOS liquid valuations. To enhance our intuition, by *competitive market* we mean that if we remove a randomly selected group of agents from the auction, the value optimal solution does not get affected so much. This is a reasonable assumption, especially in budget-restricted agents, since it denotes the existence of many agents with similar budgets and preferences under the budget constraints.

Definition 6.4 ((ε, δ) - Competitive Market)**.** *Let $0 \leq \varepsilon < 2$ and a constant $\delta \geq 0$. A market is called (ε, δ) - Competitive Market, if for any randomly removed set of bidders, $\mathbb{S}$, with cardinality $\frac{n}{2}$, then for the remaining set of bidders, $\mathbb{T}$, it holds that:*

$$\mathbb{P}\left[\overline{OPT}_{\mathbb{T}} \geq \left(1-\frac{\varepsilon}{2}\right)\cdot\overline{OPT}\right] \geq 1-\delta \tag{6.11}$$

where by $\overline{OPT}_{\mathbb{T}}$ we denote the optimal LW achieved by bidders in set $\mathbb{T}$.

Proposition 6.1. *In an (ε, δ) - Competitive Market, let $\mathbb{S} \subseteq [n]$ be randomly chosen s.t. $|\mathbb{S}| = \frac{n}{2}$ and let $\mathbb{T} = [n] \setminus \mathbb{S}$. Then:*

$$\mathbb{P}\left[\left\{\overline{OPT}_{\mathbb{T}} \geq \left(1-\frac{\varepsilon}{2}\right)\overline{OPT}\right\} \cap \left\{\overline{OPT}_{\mathbb{S}} \geq \left(1-\frac{\varepsilon}{2}\right)\overline{OPT}\right\}\right] \geq 1-2\delta$$

Proof. Let $X_{\mathbb{S}}$ the event that $\overline{\mathrm{OPT}}_{\mathbb{S}} \geq \left(1-\frac{\varepsilon}{2}\right)\overline{\mathrm{OPT}}$ and $X_{\mathbb{T}}$ the event that $\overline{\mathrm{OPT}}_{\mathbb{T}} \geq \left(1-\frac{\varepsilon}{2}\right)\overline{\mathrm{OPT}}$. Then, we have:

$$\mathbb{P}\left[X_{\mathbb{S}} \cap X_{\mathbb{T}}\right] = 1 - \mathbb{P}\left[\overline{X_{\mathbb{S}}} \cup \overline{X_{\mathbb{T}}}\right] \geq 1-2\delta$$

where the inequality follows from the Union Bound. □

6.6.3 CM mechanism

Our mechanism operates as follows. First we divide the agents uniformly at random into two equal sets $\mathbb{S}$ and $\mathbb{T}$. Then, we run a 2-approximation greedy algorithm on the sample set $\mathbb{S}$. Our claim is that, since both $\mathrm{OPT}_{\mathbb{S}}$, $\mathrm{OPT}_{\mathbb{T}}$ are close to the optimal solution, the items whose contribution in $\mathrm{OPT}_{\mathbb{S}}$ is much greater than in $\mathrm{OPT}_{\mathbb{T}}$ cannot be large, since we could transfer their owners from $\mathbb{S}$ to $\mathbb{T}$ and take a solution whose value is greater than OPT. Now, we are ready to state our Competitive Market mechanism that will be used for approximating the optimal LW. We note here that the greedy algorithm $\mathcal{A}$ is due to [42].

As usual, we denote $\mathcal{S} = (S_1, \ldots, S_n)$ the final allocation from mechanism presented in Algorithm 11. Valuations of bidders are XOS (and so are the liquid valuations (Lemma

Algorithm 11 Competitive Market (CM) Algorithm

Divide the bidders into sets $\mathbb{S}, \mathbb{T}$ uniformly at random, s.t., $|\mathbb{S}| = \frac{n}{2} = |\mathbb{T}|$.
Run the greedy algorithm $\mathcal{A}$ for bidders in $\mathbb{S}$ and denote the solution obtained by $\mathcal{A}^{\mathbb{S}}$.
for $j \in U$ **do**
    Set $p_j = \frac{1}{2\beta}\bar{v}\left(\mathcal{A}^{\mathbb{S}}_j\right)$, where $\beta > 1$ is a constant
end for
Fix an internal ordering of bidders in $\mathbb{T}$, π, and set $U_1 = U$.
for each bidder $i \in \mathbb{T}$ arriving according to π **do**
    Let $S_i = \text{BCDQ}(v_i, U_i, \vec{p})$.
    Set $U_{i+1} = U_i \setminus S_i$.
end for

6.2)); let a_i be the maximizing clause of S_i in the liquid valuation $\bar{v}_i$ of bidder i. Since a_i's are additive, for each bidder i and $j \in S_i$ let $q_j = a_i(\{j\})$. Notice that $\sum_{i\in[n]} \bar{v}(S_i) = \sum_{j\in\cup_{i\in[n]}S_i} q_j$. We denote by $\overline{\text{OPT}}_{\mathbb{T}} = \sum_{j\in U} q^{\mathbb{T}}_j$, where $q^{\mathbb{T}}_j$ is the contribution of item j in $\overline{\text{OPT}}_{\mathbb{T}}$. We divide the set of all items U into two sets; the set of *competitive* items, denoted by $\mathcal{C}$ and the set of *non-competitive* items, denoted by $\overline{\mathcal{C}} = \mathcal{M}\setminus\mathcal{C}$. The following lemma upper bounds the contribution of non-competitive items in the optimal solution.

Lemma 6.11. *Let* $\mathcal{C} = \left\{ j \middle| q^{\mathbb{T}}_j > \frac{\bar{v}\left(\mathcal{A}^{\mathbb{S}}_j\right)}{\beta} \right\}$ *for constant* $\beta > 1$*. Then,* $\sum_{j\in\overline{\mathcal{C}}} q^{\mathbb{T}}_j \leq \frac{\varepsilon}{2(\beta-1)}\overline{OPT}$ *and* $\sum_{j\in\mathcal{C}} q^{\mathbb{T}}_j \geq \frac{\beta(2-\varepsilon)-2}{2(\beta-1)}\overline{OPT}$.

Proof. From Definition 6.4, it holds with constant probability (w.c.p) that: $\overline{\text{OPT}} \geq \sum_{j\in\mathcal{C}} q^{\mathbb{T}}_j + \sum_{j\in\overline{\mathcal{C}}} q^{\mathbb{T}}_j = \sum_{j\in U} q^{\mathbb{T}}_j \geq \left(1 - \frac{\varepsilon}{2}\right)\cdot\overline{\text{OPT}}$. Let $\mathbb{S}_{\overline{\mathcal{C}}} \subseteq \mathbb{S}$ be the set of the bidders that are allocated the non-competitive items from the greedy algorithm $\mathcal{A}$ when running on set $\mathbb{S}$. Then, in the augmented set $\mathbb{T}\cup\mathbb{S}_{\overline{\mathcal{C}}}$, there exists an allocation $\mathcal{Q}$[5] with liquid valuation,

$$\bar{v}(\mathcal{Q}) \geq \sum_{j\in\mathcal{C}} q^{\mathbb{T}}_j + \sum_{j\in\overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}^{\mathbb{S}}_j\right) \tag{6.12}$$

and therefore we have w.c.p:

$$\begin{aligned}\overline{\text{OPT}} \geq \bar{v}(\mathcal{Q}) &\geq \sum_{j\in\mathcal{C}} q^{\mathbb{T}}_j + \sum_{j\in\overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}^{\mathbb{S}}_j\right) \geq \sum_{j\in\mathcal{C}} q^{\mathbb{T}}_j + \beta\sum_{j\in\overline{\mathcal{C}}} q^{\mathbb{T}}_j \\ &\geq \left(1 - \frac{\varepsilon}{2}\right)\overline{\text{OPT}} + (\beta-1)\sum_{j\in\overline{\mathcal{C}}} q^{\mathbb{T}}_j\end{aligned}$$

Re-arranging the latter and using the fact that:

$$\sum_{j\in\mathcal{C}} q_j + \frac{\varepsilon}{2(\beta-1)}\overline{\text{OPT}} \geq \sum_{j\in U} q^{\mathbb{T}}_j \geq \left(1 - \frac{\varepsilon}{2}\right)\overline{\text{OPT}}$$

As a result, for the items in $\mathcal{C}$ it holds w.c.p that: $\sum_{j\in\mathcal{C}} q^{\mathbb{T}}_j \geq \frac{\beta(2-\varepsilon)-2}{2(\beta-1)}\overline{\text{OPT}}$. □

[5] Allocation $\mathcal{Q}$ is realized by allocating all items in $\mathcal{C}$ to bidders in $\mathbb{T}$ that also had them in the $\overline{\text{OPT}}_{\mathbb{T}}$ allocation and all items in $\overline{\mathcal{C}}$ to the bidders in $\mathbb{S}_{\overline{\mathcal{C}}}$ that had them in the allocation of the greedy $\mathcal{A}$. The claim is completed by submodularity.

In the next Lemma, we prove a lower bound on the contribution of competitive items to the solution obtained by the greedy algorithm, with respect to $\overline{\mathrm{OPT}}$.

Lemma 6.12. $\sum_{j\in\mathcal{C}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \geq \frac{2(\beta-1)-\varepsilon\cdot(3\beta-1)}{4(\beta-1)}\overline{OPT}.$

Proof. Combining Inequality (6.12) and Lemma 6.11 we get that:

$$\sum_{j\in\overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \leq \frac{\beta\varepsilon}{2(\beta-1)}\overline{\mathrm{OPT}} \tag{6.13}$$

Algorithm $\mathcal{A}$ provides a 2-approximation to the optimal LW of set $\mathbb{S}$ [42], so w.c.p we have:

$$\sum_{j\in\mathcal{C}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) + \sum_{j\in\overline{\mathcal{C}}} \bar{v}\left(\mathcal{A}_j^{\mathbb{S}}\right) \geq \frac{1}{2}\overline{\mathrm{OPT}}_{\mathbb{S}} \geq \frac{1-\frac{\varepsilon}{2}}{2}\overline{\mathrm{OPT}} \tag{6.14}$$

Combining the last two equations, we get the result. □

Theorem 6.4. *The CM Algorithm is* universally truthful *and achieves, on expectation, a constant approximation to the optimal LW, i.e.,* $\mathbb{E}\left[\bar{v}\left(\mathcal{S}\right)\right] \geq (1-2\delta)\cdot\frac{2(\beta-1)-\varepsilon\cdot(3\beta-1)}{16\beta(\beta-1)}\overline{OPT}.$

Proof. Since the bidders that control the prices being posted belong to set $\mathbb{S}$ and they never get any item, it is their (weakly) dominant strategy to report their valuations and their budgets truthfully. Furthermore, the bidders that are buying under the said posted prices belong to set $\mathbb{T}$ and they make BCDQs, which we shown to be truthful. Finally, the bidders are *uniformly at random* split at sets $\mathbb{S}$ and $\mathbb{T}$.

For each item $j \in \mathcal{C}$ we have $q_j^{\mathbb{T}} > \bar{v}(\mathcal{A}_j^{\mathbb{S}})/\beta$. Therefore, there exists an allocation for bidders in $\mathbb{T}$ and items in $\mathcal{C}$ that is supported by prices $p_1,\ldots,p_m$, where $p_j = \frac{\bar{v}(\mathcal{A}_j^{\mathbb{S}})}{\beta}$. Thus, from a modification of [16, Lemma 4.2] (formally presented in Lemma 6.5), setting $p'_j = p_j/2$, for each $j \in \mathcal{C}$, and running a fixed price auction in $\mathbb{T}$ with prices $p'_1,\ldots,p'_m$, we get that: $\bar{v}(\mathcal{S}) \geq \sum_{j\in\mathcal{C}} p_j/4$. Using the latter, along with the prices of the items, we have:

$$\bar{v}(\mathcal{S}) = \frac{1}{4\beta}\sum_{j\in\mathcal{C}} \bar{v}(\mathcal{A}_j^{\mathbb{S}}) \geq \frac{2(\beta-1)-\varepsilon(3\beta-1)}{16\beta(\beta-1)}\overline{\mathrm{OPT}}$$

where the last inequality is due to Lemma 6.12. Thus, we conclude that:

$$\mathbb{E}\left[\bar{v}\left(\mathcal{S}\right)\right] \geq (1-2\delta)\,\frac{2(\beta-1)-\varepsilon\cdot(3\beta-1)}{16\beta(\beta-1)}\overline{\mathrm{OPT}}$$

□

6.7 Conclusion and Future Work

In this work, we showed how some of the best known truthful mechanisms that approximate the SW, can be adapted in order to yield the same order approximations for the LW, when bidders are budget-constrained in the worst-case and Bayesian instances. Additionally, we introduced a notion of market competitiveness, for markets with indivisible goods and provided a constant factor approximation to the LW in this case. The most meaningful question that arises from our work (apart, of course, from the ever existent one of lowering the approximation guarantee in worst-case instances) is related to the competitive markets. We conjecture that the condition that we provide can be made even weaker, and leave it to future research. We hope that the results and the techniques presented in this work, will serve future researchers in obtaining improved *same* order approximations for both the SW and the LW.

Bibliography

[1] Lawrence M. Ausubel. “An Efficient Dynamic Auction for Heterogeneous Commodities”. In: *American Economic Review* 96.3 (June 2006), pp. 602–629.

[2] Lawrence M Ausubel and Paul R Milgrom. “Ascending auctions with package bidding”. In: *Advances in Theoretical Economics* 1.1 (2002).

[3] Eric Balkanski and Jason D Hartline. “Bayesian budget feasibility with posted pricing”. In: *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee. 2016, pp. 189–203.

[4] Xiaohui Bei et al. “Budget feasible mechanism design: from prior-free to bayesian”. In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM. 2012, pp. 449–458.

[5] Liad Blumrosen and Noam Nisan. “On the computational power of demand queries”. In: *SIAM Journal on Computing* 39.4 (2009), pp. 1372–1391.

[6] Christian Borgs et al. “Multi-unit auctions with budget-constrained bidders”. In: *Proceedings of the 6th ACM conference on Electronic commerce*. ACM. 2005, pp. 44–51.

[7] Ning Chen, Nick Gravin, and Pinyan Lu. “On the approximability of budget feasible mechanisms”. In: *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics. 2011, pp. 685–699.

[8] Edward Clarke. “Multipart pricing of public goods”. In: *Public Choice* 11.1 (1971), pp. 17–33.

[9] William J. Cook et al. *Combinatorial Optimization*. John Wiley & Sons, Inc., 1998.

[10] Vincent Crawford. “Job Matching, Coalition Formation, and Gross Substitutes”. In: *Econometrica* 50.6 (1982), pp. 1483–1504.

[11] Vincent P. Crawford and Elsie Marie Knoer. “Job Matching with Heterogeneous Firms and Workers”. In: *Econometrica* 49.2 (1981), pp. 437–450.

[12] Sven De Vries and Rakesh V Vohra. “Combinatorial auctions: A survey”. In: *INFORMS Journal on computing* 15.3 (2003), pp. 284–309.

[13] Gabrielle Demange and David Gale. “The Strategy Structure of Two-Sided Matching Markets”. In: *Econometrica* 53.4 (1985), pp. 873–888.

[14] Gabrielle Demange, David Gale, and Marilda Sotomayor. “Multi-Item Auctions”. In: *Journal of Political Economy* 94.4 (1986), pp. 863–872.

[15] Shahar Dobzinski. "An impossibility result for truthful combinatorial auctions with submodular valuations". In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. ACM. 2011, pp. 139–148.

[16] Shahar Dobzinski. "Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders". In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM. 2016, pp. 940–948.

[17] Shahar Dobzinski. "Two randomized mechanisms for combinatorial auctions". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 2007, pp. 89–103.

[18] Shahar Dobzinski, Hu Fu, and Robert Kleinberg. "Truthfulness via proxies". In: *arXiv preprint arXiv:1011.3232* (2010).

[19] Shahar Dobzinski and Renato Paes Leme. "Efficiency guarantees in auctions with budgets". In: *International Colloquium on Automata, Languages, and Programming*. Springer. 2014, pp. 392–404.

[20] Shahar Dobzinski, Noam Nisan, and Michael Schapira. "Approximation algorithms for combinatorial auctions with complement-free bidders". In: *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM. 2005, pp. 610–618.

[21] Shahar Dobzinski, Noam Nisan, and Michael Schapira. "Truthful randomized mechanisms for combinatorial auctions". In: *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. ACM. 2006, pp. 644–652.

[22] Shahar Dobzinski, Christos H Papadimitriou, and Yaron Singer. "Mechanisms for complement-free procurement". In: *Proceedings of the 12th ACM conference on Electronic commerce*. ACM. 2011, pp. 273–282.

[23] Shahar Dobzinski and Jan Vondrák. "Communication Complexity of Combinatorial Auctions with Submodular Valuations". In: *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '13. Society for Industrial and Applied Mathematics, 2013, pp. 1205–1215.

[24] Paul Düetting et al. "Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs". In: *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*. IEEE. 2017, pp. 540–551.

[25] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. "From convex optimization to randomized mechanisms: toward optimal combinatorial auctions". In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. ACM. 2011, pp. 149–158.

[26] Alon Eden, Michal Feldman, and Adi Vardi. "Online Random Sampling for Budgeted Settings". In: *International Symposium on Algorithmic Game Theory*. Springer. 2017, pp. 29–40.

[27] FCC. *Webpage for different spectrum auctions*. http://wireless.fcc.gov/auctions/. [Online; accessed 25-July-2018].

[28] Uriel Feige. "On maximizing welfare when utility functions are subadditive". In: *SIAM Journal on Computing* 39.1 (2009), pp. 122–142.

[29] Michal Feldman, Nick Gravin, and Brendan Lucier. “Combinatorial auctions via posted prices”. In: *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM. 2014, pp. 123–135.

[30] Dimitris Fotakis, Kyriakos Lotidis, and Chara Podimata. “A Bridge between Liquid and Social Welfare in Combinatorial Auctions with Submodular Bidders”. In: *arXiv preprint arXiv:1809.01803* (2018).

[31] Hu Fu, Robert Kleinberg, and Ron Lavi. “Conditional Equilibrium Outcomes via Ascending Price Processes with Applications to Combinatorial Auctions with Item Bidding”. In: *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC ’12. ACM, 2012, pp. 586–586.

[32] D. Gale and L. S. Shapley. “College Admissions and the Stability of Marriage”. In: *The American Mathematical Monthly* 69.1 (1962), pp. 9–15.

[33] Andrew V Goldberg et al. “Competitive auctions”. In: *Games and Economic Behavior* 55.2 (2006), pp. 242–269.

[34] Martin Grötschel, Lászlo Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Vol. 2. Springer, 1988.

[35] Theodore Groves. “Incentives in Teams”. In: *Econometrica* 41.4 (1973), pp. 617–31.

[36] Faruk Gul and Ennio Stacchetti. “Walrasian Equilibrium with Gross Substitutes”. In: *Journal of Economic Theory* 87.1 (1999), pp. 95–124.

[37] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C Parkes. “Adaptive limited-supply online auctions”. In: *Proceedings of the 5th ACM conference on Electronic commerce*. ACM. 2004, pp. 71–80.

[38] John Hershberger and Subhash Suri. “Vickrey prices and shortest paths: What is an edge worth?” In: *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*. IEEE. 2001, pp. 252–259.

[39] Howard Karloff. *Linear Programming*. Birkhauser Boston Inc., 1991.

[40] Subhash Khot et al. “Inapproximability Results for Combinatorial Auctions with Submodular Utility Functions”. In: *Internet and Network Economics*. 2005, pp. 92–101.

[41] Piotr Krysta and Berthold Vöcking. “Online mechanism design (randomized rounding on the fly)”. In: *International Colloquium on Automata, Languages, and Programming*. Springer. 2012, pp. 636–647.

[42] Benny Lehmann, Daniel Lehmann, and Noam Nisan. “Combinatorial auctions with decreasing marginal utilities”. In: *Games and Economic Behavior* 55.2 (2006), pp. 270–296.

[43] Renato Paes Leme. “Gross substitutability: An algorithmic survey”. In: *Games and Economic Behavior* 106 (2017), pp. 294–316.

[44] Herman B Leonard. “Elicitation of Honest Preferences for the Assignment of Individuals to Positions”. In: *Journal of Political Economy* 91.3 (1983), pp. 461–79.

[45] Pinyan Lu and Tao Xiao. “Improved efficiency guarantees in auctions with budgets”. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM. 2015, pp. 397–413.

[46] Pinyan Lu and Tao Xiao. "Liquid welfare maximization in auctions with multiple items". In: *International Symposium on Algorithmic Game Theory*. Springer. 2017, pp. 41–52.

[47] Paul R. Milgrom and Robert J. Weber. "A Theory of Auctions and Competitive Bidding". In: *Econometrica* 50.5 (1982), pp. 1089–1122.

[48] Vahab Mirrokni, Michael Schapira, and Jan Vondrak. "Tight Information-theoretic Lower Bounds for Welfare Maximization in Combinatorial Auctions". In: *Proceedings of the 9th ACM Conference on Electronic Commerce*. EC '08. Chicago, Il, USA: ACM, 2008, pp. 70–77. ISBN: 978-1-60558-169-9. DOI: 10.1145/1386790.1386805. URL: http://doi.acm.org/10.1145/1386790.1386805.

[49] Roger B. Myerson. "Optimal Auction Design". In: *Math. Oper. Res.* 6.1 (1981), pp. 58–73.

[50] Noam Nisan and Ilya Segal. "The communication requirements of efficient allocations and supporting prices". In: *Journal of Economic Theory* 129.1 (2006), pp. 192–224.

[51] Noam Nisan et al. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[52] David C. Parkes and Lyle H. Ungar. "Iterative Combinatorial Auctions: Theory and Practice". In: *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*. AAAI Press, 2000, pp. 74–81.

[53] Stephen J Rassenti, Vernon L Smith, and Robert L Bulfin. "A combinatorial auction mechanism for airport time slot allocation". In: *The Bell Journal of Economics* (1982), pp. 402–417.

[54] Tim Roughgarden. "CS364B: Frontiers in Mechanism Design Lecture# 1: Ascending and Ex Post Incentive Compatible Mechanisms". In: (2014).

[55] Tim Roughgarden. "CS364B: Frontiers in Mechanism Design Lecture# 2: Unit-Demand Bidders and Walrasian Equilibria". In: (2014).

[56] Tim Roughgarden. "CS364B: Frontiers in Mechanism Design Lecture# 5: The Gross Substitutes Condition". In: (2014).

[57] Tim Roughgarden. "CS364B: Frontiers in Mechanism Design Lecture# 6: Gross Substitutes: Welfare Maximization in Polynomial Time". In: (2014).

[58] Tim Roughgarden. "CS364B: Frontiers in Mechanism Design Lecture# 7: Submodular Valuations". In: (2014).

[59] Tim Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016.

[60] Tim Roughgarden et al. "CS364B: Frontiers in Mechanism Design Lecture# 3: The Crawford-Knoer Auction". In: (2014).

[61] L. S. Shapley and M. Shubik. "The assignment game I: The core". In: *International Journal of Game Theory* 1.1 (Dec. 1971), pp. 111–130.

[62] Yaron Singer. "Budget feasible mechanisms". In: *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*. IEEE. 2010, pp. 765–774.

[63] Ning Sun and Zaifu Yang. "Equilibria and Indivisibilities: Gross Substitutes and Complements". In: *Econometrica* 74.5 (2006), pp. 1385–1402.

[64] William Vickrey. "Counterspeculation, auctions, and competitive sealed tenders". In: *Journal of Finance* 16.1 (1961), pp. 8–37.

[65] Jan Vondrak. "Optimal Approximation for the Submodular Welfare Problem in the Value Oracle Model". In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. STOC '08. 2008, pp. 67–74.

[66] Weiwei Wu, Xiang Liu, and Minming Li. "Budget-feasible Procurement Mechanisms in Two-sided Markets." In: *IJCAI*. 2018, pp. 548–554.