Εθνικο Μετσοβιο Πολυτεχνειο
Σχολη Ηλεκτρολογων Μηχανικων και Μηχανικων Υπολογιστων
Τομεας Τεχνολογιας Πληροφορικης και Υπηρεσιων
Εργαστηριο Λογικης και Επιστημης Υπολογιστων

Τιμολόγηση Καθοδηγούμενη από Δεδομένα για Δρομολόγηση και Έλεγχο Αποδοχής σε Δίκτυα με Χωρητικότητες

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Βασίλειος Βαρσαμής

Επιβλέπων : Δημήτριος Φωτάκης
Καθηγητής Ε.Μ.Π.

Αθήνα, Ιούλιος 2024

Εθνικό Μετσόβιο Πολυτεχνείο
Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών
Τομέας Τεχνολογίας Πληροφορικής και Υπηρεσιών
Εργαστήριο Λογικής και Επιστήμης Υπολογιστών

Τιμολόγηση Καθοδηγούμενη από Δεδομένα για Δρομολόγηση και Έλεγχο Αποδοχής σε Δίκτυα με Χωρητικότητες

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Βασίλειος Βαρσαμής

Επιβλέπων : Δημήτριος Φωτάκης
Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 10η Ιουλίου 2024.

....	....	....
Δημήτριος Φωτάκης	Αριστείδης Παγουρτζής	Ευάγγελος Μαρκάκης
Καθηγητής Ε.Μ.Π.	Καθηγητής Ε.Μ.Π.	Αν.Καθηγητής Ο.Π.Α

Αθήνα, Ιούλιος 2024

....

Βασίλειος Βαρσαμής
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Περίληψη

Στην παρούσα διπλωματική εργασία, εξετάζουμε την απόδοση των prophet inequalities σε Συνδυαστικές Δημοπρασίες όταν οι αξιολογήσεις παρουσιάζουν συμπληρωματικότητες και τα αντικείμενα είναι διαθέσιμα σε πολλά αντίγραφα. Εστιάζουμε στο παραδειγματικό πρόβλημα της άμεσης δρομολόγησης και ελέγχου αποδοχής σε δίκτυα όταν οι ακμές έχουν χωρητικότητα που είναι λογαριθμικά μεγάλη σε σχέση με το πλήθος των ακμών και γνωρίζουμε την κατανομή των αιτημάτων. Η λογαριθμικά μεγάλη χωρητικότητα μας επιτρέπει να λύσουμε προσεγγιστικά την κλασματική χαλάρωση του προβλήματος που μας δίνει πιθανότητες σε μονοπάτια και προσεγγιστικό παράγοντα ίσο με $1 + \epsilon$. Υποθέτοντας ότι έχουμε γνώση των συνολικών αιτημάτων που έρχονται σε κάθε ζευγάρι κόμβων, έπειτα χρησιμοποιούμε την έννοια των balanced prices, όπως ορίζονται στο [32] οι οποίες υπολογίζονται με βάση τα μονοπάτια, και αποδομούμε το πρόβλημα σε στιγμιότυπα που πρακτικά συμπεριφέρονται ως ανεξάρτητα και λειτουργούν ως unit-demand.

Λέξεις Κλειδιά: Συνδυαστικές Δημοπρασίες, Σχεδιασμός Μηχανισμών, Έλεγχος Αποδοχής, Δρομολόγηση σε Δίκτυα, Prophet Inequalities, Balanced Prices, unit-demand

Abstract

In this thesis, we examine the performance of prophet inequalities in combinatorial auctions where valuations exhibit complementarities and items are available in multiple copies. We focus on the exemplary problem of online routing and admission control in networks where the edges have capacities that are logarithmically large relative to the number of edges, and we know the distribution of requests. The logarithmically large capacity allows us to approximately solve the fractional relaxation of the problem, which provides probabilities for paths and an approximation factor of $1 + \epsilon$. Assuming knowledge of the total requests for each node pair, we then use the concept of balanced prices, as defined in [32], which are calculated based on the paths, and decompose the problem into instances of node pairs that practically behave as independent and function as unit-demand.

Keywords: Combinatorial Auctions, Mechanism Design, Admission Control, Routing in Networks, Prophet Inequalities, Balanced Prices, unit-demand

Ευχαριστίες

Ολοκληρώνοντας τις σπουδές μου στο Ε.Μ.Π θα ήθελα να ευχαριστήσω ανθρώπους που υπήρξαν σημαντικό και αναπόσπαστο κομμάτι της διαδρομής μου έως τώρα. Πρώτα απ'όλα θέλω να ευχαριστήσω τον επιβλέποντά μου, κ. Δημήτρη Φωτάκη ο οποίος από τα πρώτα κιόλας έτη μέσα από την εμπνευμένη διδασκαλία του ενστάλαξε μέσα μου το ενδιαφέρον και την αγάπη για τη Θεωρητική Πληροφορική. Σε προσωπικό επίπεδο, τον ευχαριστώ για όλη τη σημαντική του καθοδήγηση και υποστήριξη καθ'όλη τη διάρκεια εκπόνησης της εργασίας καθώς και για την αμέριστη βοήθειά του με τις αιτήσεις για τα πανεπιστήμια του εξωτερικού.

Είμαι απίστευτα ευγνώμων για τους φίλους που απέκτησα στη σχολή: Αλέξανδρε, Βαγγέλη, Δήμητρα, Θοδωρή, Λευτέρη σας ευχαριστώ για όλες τις εμπειρίες που μοιραστήκαμε, εντός και εκτός σχολής και για τις όμορφες αναμνήσεις που δημιουργήσαμε παρέα. Ευχαριστώ επίσης όλα τα παιδιά του Corelab για το πολύ όμορφο και φιλικό κλίμα που δημιουργήσαμε στο εργαστήριο τον τελευταίο χρόνο και ιδιαίτερα τον Αλέξανδρο και τη Βάλια που ήταν συνοδοιπόροι σε αυτό το μακρύ ταξίδι των αιτήσεων και της διπλωματικής.

Το πιο σημαντικό ευχαριστώ το απευθύνω στην οικογένειά μου, τους γονείς μου Δημήτρη και Σέτη και τον αδερφό μου Παντελή. Μέσα από την ολοκληρωτική και ανιδιοτελή αγάπη τους προς εμένα με στήριξαν και με βοήθησαν να ξεπεράσω όποιες δυσκολίες εμφανίζονταν. Ήταν πάντα εκεί για εμένα, όποτε το χρειάστηκα. Τέλος, θα ήθελα να ευχαριστήσω την Κάλλη, η οποία ήταν δίπλα μου σε όλες τις φοιτητικές μου στιγμές, εύκολες ή δύσκολες, και χωρίς την οποία αυτό το ταξίδι θα ήταν πολύ πιο δύσκολο και λιγότερο όμορφο.

Βασίλης

Contents

Chapter 1

Εκτενής Ελληνική Περίληψη

1.1 Εισαγωγή

Μπορούμε να προβλέψουμε τη συμπεριφορά και την κατάσταση ενός συστήματος που αποτελείται από στρατηγικούς παίκτες που αλληλεπιδρουν επιδιώκοντας να επιτύχουν το πιο προτιμώμενο για αυτούς αποτέλεσμα; Ο John F. Nash στο βραβευμένο με Νόμπελ έργο του απέδειξε ότι, ανεξάρτητα από το πόσο περίπλοκο είναι το σύστημα, οι πράκτορες αλληλεπιδρούν με τρόπο που τελικά φτάνει σε μια κατάσταση όπου κανείς δεν μπορεί να επιτύχει καλύτερο αποτέλεσμα αλλάζοντας μονομερώς τη στρατηγική του. Η έννοια αυτή ονομάζεται ισορροπία Nash και βρίσκεται στην καρδιά του πεδίου των μαθηματικών που ονομάζεται Θεωρία Παιγνίων. Η έλευση του διαδικτύου έφερε πλήθος εφαρμογών της Θεωρίας Παιγνίων, επειδή πρόκειται για ένα σύστημα στο οποίο οι συμμετέχοντες αλληλεπιδρούν στρατηγικά για να το κατευθύνουν προς τα επιθυμητά αποτελέσματα. Ένα σημαντικό ερώτημα λοιπόν ήταν, αν μπορεί κανείς να υπολογίσει αποτελεσματικά την ισορροπία Nash (που είναι εγγυημένα υπαρκτή) ενός συστήματος που αποτελείται από στρατηγικούς παίκτες. Το ερώτημα αυτό, που τέθηκε κυρίως από επιστήμονες της πληροφορικής, έδωσε το έναυσμα για τη δημιουργία του πεδίου της Αλγοριθμικής Θεωρίας Παιγνίων. Η άποψη αυτή, αναλύει συστήματα που υπάρχουν "στη φύση" και αναζητά μεθόδους για την αποτελεσματική πρόβλεψη της έκβασής τους μετά από ένα χρονικό διάστημα. Ωστόσο, τι θα γινόταν αν αλλάζαμε οπτική γωνία και αναλαμβάναμε το ρόλο του σχεδιαστή του συστήματος; Τώρα το ερώτημα γίνεται, μπορούμε, εξοπλισμένοι με τη δύναμη να φτιάξουμε το δικό μας παιχνίδι (ή σύστημα), να επιβάλουμε μια επιθυμητή συμπεριφορά του παίκτη και ένα επιθυμητό αποτέλεσμα του συστήματος; Το κατεξοχήν παράδειγμα απάντησης σε αυτό το ερώτημα είναι οι δημοπρασίες. Από τη σκοπιά του πωλητή, είναι δυνατόν να σχεδιάσουμε και να υπολογίσουμε αποτελεσματικά τους κανόνες μιας δημοπρασίας κατά τρόπο ώστε η προτιμότερη ενέργεια των πλειοδοτών να είναι να αναφέρουν τις πραγματικές τους αποτιμήσεις; Το ερώτημα αυτό γέννησε τον τομέα του Αλγοριθμικού Σχεδιασμού Μηχανισμών, ο οποίος συχνά αποκαλείται και αντίστροφη Θεωρία Παιγνίων. Στο πλαίσιο των δημοπρασιών, για την εφαρμογή της φιλαλήθειας χρησιμοποιούνται χρηματικά ανταλλάγματα από τους πλειοδότες προς τον πωλητή. Με άλλα λόγια, ο πωλητής πρέπει να αποφασίσει τους νικητές και τις πληρωμές για να προκαλέσει την φιλαλήθεια. Στο Σχεδιασμό Μηχανισμών με Χρήματα, το θεμελιώδες έργο ενός άλλου νομπελίστα, του R. Myerson [53], έθεσε τα θεμέλια για το πεδίο αυτό, χαρακτηρίζοντας πλήρως τις φιλαλήθεις δημοπρασίες στο πλαίσιο μίας παραμέτρου, όπου οι αποτιμήσεις των προσφερόντων περιγράφονται από έναν μόνο αριθμό. Ο Myerson απέδειξε ότι αν ένας κανόνας κατανομής των αγαθών είναι αύξων σε σχέση με τοην αξιολόγηση, τότε υπάρχει ένας μοναδικός κανόνας πληρωμής, ο οποίος όταν συνδυάζεται με την εν λόγω κατανομή καθιστά τον μηχανισμό φιλαλήθη. Η εφαρμογή του χαρακτηρισμού του Myerson σε δημοπρασίες ενός στοιχείου δίνει τη διάσημη δημοπρασία Vickrey ή δημοπρασία δεύτερης υψηλότερης τιμής [60]. Οι Vickrey, Clarke και Groves [60, 20, 40] γενίκευσαν το αποτέλεσμα του Myerson σε περιβάλλοντα πολλών παραμέτρων, όπου οι αποτιμήσεις είναι συναρτήσεις. Ο μηχανισμός VCG, ο οποίος πήρε το όνομά του από τα αρχικά τους, χαρακτηρίζει τους φιλαλήθεις μηχανισμούς παρέχοντας πληρωμές ανάλογες με την περίπτωση ενός στοιχείου. Η μόνη επιφύλαξη είναι ότι για να λειτουργήσει ο μηχανισμός VCG, απαιτεί τη βέλτιστη λύση του υποκείμενου προβλήματος βελτιστοποίησης. Αυτό είναι ένα NP-δύσκολο έργο για τις περισσότερες κατηγορίες αποτίμησης. Ας σημειώσουμε ότι όλοι οι προαναφερθέντες μηχανισμοί μεγιστοποιούν επίσης την κοινωνική ευημερία, η οποία είναι το άθροισμα των αποτιμήσεων των νικητών. Στο κεφάλαιο 3 θα ορίσουμε επίσημα ορισμένες βασικές έννοιες στο Σχεδιασμό Μηχανισμών, οι οποίες είναι χρήσιμες σε όλη την εργασία.

Φανταστείτε ότι είστε ιδιοκτήτης μιας αίθουσας συναυλιών και θέλετε να πουλήσετε εισιτήρια (που αντιστοιχούν σε θέσεις μέσα στην αίθουσα) για την επερχόμενη συναυλία. Η συλλογική μας εμπειρία από την πραγματική ζωή υπαγορεύει ότι ο πωλητής δεν διοργανώνει δημοπρασία για να πουλήσει εισιτήρια, απλώς και μόνο επειδή υπάρχει ένας ευκολότερος και πιο διαισθητικός τρόπος πώλησης: δημοσιεύστε τιμές και αφήστε τους αγοραστές να διαλέξουν θέσεις και να αγοράσουν τα εισιτήρια. Πολλές άλλες ρεαλιστικές περιπτώσεις όπως η παραπάνω υποδηλώνουν την ανάγκη να βρεθούν απλούστεροι μηχανισμοί για την πώληση σε στρατηγικούς αγοραστές. Για το σκοπό αυτό, υπάρχει μια τεράστια σειρά ερευνών σε μια συγκεκριμένη κατηγορία μηχανισμών που ονομάζεται Posted-price. Μαζί με τις εγγυήσεις κινήτρων που παρέχουν, οι μηχανισμοί δημοσιευμένης τιμής (posted-price) παρέχουν πολυωνυ-μικό υπολογισμό των τιμών και απλότητα. Το μειονέκτημα είναι ότι προσεγγίζουν μόνο τη βέλτιστη κοινωνική ευημερία, ωστόσο τέτοιοι μηχανισμοί είναι πολύ χρήσιμοι στην πράξη και πυροδοτούν τη συζήτηση για τους απλούς έναντι των βέλτιστων μηχανισμών. Στο Κεφάλαιο 4 θα εμβαθύνουμε στις λεπτομέρειες τέτοιων μηχανισμών και στα αποτελέσματα αυτής της ερευνητικής γραμμής.

Τα εργαλεία που χρησιμοποιούνται για την ανάλυση των επιδόσεων των μηχανισμών δημοσιευμένων τιμών είναι το πλαίσιο των Prophet Inequalities, το οποίο εισήχθη για πρώτη φορά από τους μαθηματικούς Krengel, Sucheston και Garling [46, 47] τη δεκαετία του '70. Το πλαίσιο περιλαμβάνει έναν λήπτη διαδοχικών αποφάσεων που παρατηρεί τις τιμές με άμεσο τρόπο και πρέπει να αποφασίσει για μια βέλτιστη στρατηγική διακοπής που είναι προσεγγιστικά βέλτιστη σε σύγκριση με έναν προφήτη που μπορεί να δει το μέλλον και παίρνει πάντα το βέλτιστο αποτέλεσμα. Τα Prophet Inequalities έγιναν σημαντικά στον τομέα της επιστήμης των υπολογιστών όταν ανακαλύφθηκε η σύνδεσή τους με τις διαδοχικές δημοσιευμένες τιμές. Παρέχουν τις κατάλληλες τεχνικές για να επιχειρηματολογήσουμε για τον λόγο προσέγγισης που επιτυγχάνουν οι τιμές, υπό την υπόθεση ότι οι αποτιμήσεις των αγοραστών προέρχονται από γνωστές ανεξάρτητες κατανομές. Στο Κεφάλαιο 5 πρόκειται να διερευνήσουμε βασικά αποτελέσματα της βιβλιογραφίας για τα Prophet Inequalities τα οποία παρέχουν πληροφορίες για το σχεδιασμό μηχανισμών με βελτιωμένες εγγυήσεις προσέγγισης σε διάφορα πολυπαραμετρικά περιβάλλοντα.

Ένα σημαντικό και άκρως ενδιαφέρον έργο στα Prophet Inequalities έχει βελτιώσει τους συντελεστές προσέγγισης για τους μηχανισμούς δημοσιευμένων τιμών σε διάφορα πλαίσια. Ένα κοινό χαρακτηριστικό των περισσοτέρων από αυτά τα πλαίσια είναι ότι παρουσιάζουν μη συμπληρωματικές αποτιμήσεις, όπως στις Συνδυαστικές Δημοπρασίες με submodular αγοραστές. Αντίθετα, υπάρχουν πολυάριθμα πρακτικά σημαντικά παραδείγματα όπου οι αξίες των αγοραστών είναι συμπληρωματικές, δηλαδή ένα στοιχείο αποκτά αξία μόνο όταν αγοράζεται σε συνδυασμό με άλλα στοιχεία. Στο Κεφάλαιο 6 διερευνούμε ένα περιβάλλον που παρουσιάζει συμπληρωματικότητα. Οι αγοραστές έχουν μια αξία μίας παραμέτρου για να γίνουν δεκτοί και να δρομολογηθούν σε ένα δίκτυο και έχουν μη μηδενική αξία μόνο αν τους ανατεθεί μια διαδρομή που συνδέει την αφετηρία τους με τον κόμβο προορισμού τους. Σε αυτό το πλαίσιο, οποιαδήποτε δέσμη ακμών έχει μη μηδενική αξία μόνο αν σχηματίζει έγκυρο μονοπάτι που εξυπηρετεί έναν αγοραστή. Αξιοποιώντας το πλαίσιο των Prophet Inequalities, εξετάζουμε τα συστήματα τιμολόγησης και τις εγγυήσεις προσέγγισής τους σε όρους κοινωνικής ευημερίας.

1.2 Τα θεμέλια του Σχεδιασμού Μηχανισμών

Αν και τα Prophet Inequalities είναι ένα πλαίσιο ανεξάρτητου ενδιαφέροντος, η σύνδεσή της με το σχεδιασμό μηχανισμών επιβάλλει να ορίσουμε ορισμένες θεμελιώδεις έννοιες του πεδίου. Ξεκινάμε με την παροχή ενός γενικού ορισμού σχετικά με ένα πρόβλημα Σχεδιασμών Μηχανισμών. Στη συνέχεια, πραγματοποιούμε ορισμένες τυπικές απλοποιήσεις προκειμένου να καταλήξουμε σε ουσιαστικά αποτελέσματα. Ακολουθεί μια γενική διάταξη για ένα πρόβλημα σχεδιασμού μηχανισμού.

Definition 1.2.1 (Γενική Διάταξη)**.** Ας υποθέσουμε ότι έχουμε n πράκτορες και έναν χώρο αποτελεσμάτων Ω. Κάθε πράκτορας έχει έναν ιδιωτικό τύπο αποτίμησης $v_i : \Omega \to \mathbb{R}^+$ που αντιπροσωπεύει την αξία που ο πράκτορας αντλεί από ένα αποτέλεσμα, και ένα σύνολο ενεργειών $\mathcal{A}_i$. Ο μηχανισμός συλλέγει ένα διάνυσμα δράσεων: $\mathbf{a} = (a_1, a_2, \ldots, a_n)\,, a_i \in \mathcal{A}_i$ και αντιστοιχίζει αυτές τις δράσεις σε αποτελέσματα χρησιμοποιώντας μια συνάρτηση $f : (\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n) \to \Omega$. Οι μηχανισμοί αποφασίζουν επίσης για έναν κανόνα τιμολόγησης $\mathbf{p} = (p_1, p_2, \ldots, p_n)$ όπου για κάθε $i, p_i : (\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n) \to \mathbb{R}^+$. Κάθε πράκτορας i έχει μια χρησιμότητα u_i, την οποία θέλει να μεγιστοποιήσει εγωιστικά, που ποσοτικοποιεί το κέρδος του από την περιγραφόμενη διαδικασία. Εξαρτάται από το αποτέλεσμα, τον κανόνα τιμολόγησης και την ιδιωτική του αποτίμηση. Το ζεύγος $(f, \mathbf{p})$ ορίζει έναν μηχανισμό.

Προκειμένου να εργαστούμε με το παραπάνω πλαίσιο πρέπει να κάνουμε κάποιες απλοποιήσεις. Πρώτα απ' όλα, θα ορίσουμε και θα υποθέσουμε για το υπόλοιπο της διατριβής την ημι-γραμμική χρησιμότητα των αγοραστών. Άτυπα, η ημι-γραμμική χρησιμότητα ενός πλειοδότη για ένα αποτέλεσμα είναι απλά η αξία του μείον την τιμή που πληρώνει για το αποτέλεσμα αυτό.

Definition 1.2.2 (Quasi-linear utility)**.** Έστω $(f, \mathbf{p})$ ένας μηχανισμός, $v_i : \Omega \to \mathbb{R}^+$ μια ιδιωτική συνάρτηση αποτίμησης και $\mathbf{a}$ ένα διάνυσμα δράσης. Η χρησιμότητα u_i ενός πλειοδότη i είναι *quasi-linear*, αν

$$u_i = v_i(f(\mathbf{a})) - p_i(\mathbf{a})$$

1.2.1 Συναρτήσεις Αξιολόγησης

Σε αυτό το τμήμα της εργασίας ορίζουμε κάποιες από τις πιο σημαντικές κλάσεις αξιολογήσεων που υποθέτουμε για να έχουμε αποτελέσματα για τους αγοραστές. Οι παρακάτω είναι επίσης ορισμένες με βάση την πραγματική ζωή, γεγονός που προσπαθούμε να αιτιολογήσουμε με παραδείγματα μετά τον κάθε ορισμό.

Definition 1.2.3 (Submodular συνάρτηση)**.** Μια συνάρτηση αποτίμησης $v : 2^M \to \mathbb{R}^+$ είναι submodular αν για κάθε $S, T \subseteq M$, με $S \subseteq T$ και κάθε στοιχείο $j \notin S$

$$v(S \cup \{j\}) - v(S) \geq v(T \cup \{j\}) - v(T)$$

Αυτή η κατηγορία περιέχει τόσο προσθετικές όσο και μοναδιαίες αποτιμήσεις. Είναι το διακριτό ανάλογο των κοίλων συναρτήσεων και αποτυπώνει την έννοια της φθίνουσας απόδοσης: η πρόσθετη αξία για ένα νέο στοιχείο μειώνεται καθώς τα σύνολα μεγαλώνουν. Υποστηρίζεται ότι αυτή η κλάση είναι αρκετά εκφραστική για να μοντελοποιήσει τον πραγματικό κόσμο. Σκεφτείτε το ακόλουθο παράδειγμα: Η Alice έχει 0$, ο Bob έχει 1.000.000$ - και εκτιμούν τα χρήματα με τον ίδιο τρόπο. Τώρα δώστε 50$ στον καθένα τους. Ποιος παρατηρεί μεγαλύτερη αύξηση στην αξία του;

Definition 1.2.4 (Συνάρτηση XOS)**.** Μια συνάρτηση αποτίμησης $v : 2^M \to \mathbb{R}^+$ είναι XOS αν υπάρχει μια συλλογή από $a_1, a_2, \ldots, a_l$ προσθετικές συναρτήσεις αποτίμησης τέτοιες ώστε για κάθε $S \subseteq M$

$$v(S) = \max_{1 \leq i \leq l} a_i(S)$$

Είναι δύσκολο να βρούμε διαίσθηση για αυτή την κατηγορία αποτίμησης. Ωστόσο, αυτή η κλάση είναι ένα υπερσύνολο των υποδιαμετρικών αποτιμήσεων και είναι συχνά ευκολότερο να εργαστούμε με

Definition 1.2.5 (Subadditive συνάρτηση)**.** Μια συνάρτηση αποτίμησης $v : 2^M \to \mathbb{R}^+$ είναι subadditive αν για κάθε $S, T \subseteq M$

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

Αυτή η κατηγορία είναι η πιο γενική που χρησιμοποιείται στη βιβλιογραφία. Περιέχει αποτιμήσεις XOS και περιγράφει την έννοια των αποτιμήσεων χωρίς συμπλήρωμα. Δηλαδή, η ομαδοποίηση δύο συνόλων στοιχείων δεν μπορεί να αυξήσει την αξία κάθε μεμονωμένου συνόλου. Παρά το γεγονός ότι είναι η πιο γενική κλάση υπάρχουν πρακτικά αποκαλυπτικές περιπτώσεις συμπληρωματικότητας. Σκεφτείτε τα ακόλουθα: Ένας αγοραστής θέλει να αγοράσει ένα μονοπάτι που τον συνδέει από έναν αρχικό κόμβο s σε έναν τελικό κόμβο t. Ενώ δεν έχει καμία αξία για μεμονωμένες ακμές , όταν μια δέσμη ακμών σχηματίζει ένα εφικτό μονοπάτι από το s στο t, η αξία του ξαφνικά γίνεται μη μηδενική. Δηλαδή, οι ακμές αλληλοσυμπληρώνονται για να δώσουν αξία. Αυτό θα είναι το παράδειγμα κίνητρο για το μεγαλύτερο μέρος της διατριβής. Ο ακόλουθος ορισμός των *single-minded* πρακτόρων αποτυπώνει το παραπάνω παράδειγμα.

Definition 1.2.6 (Single-Minded συνάρτηση)**.** Έστω $v^* \in \mathbb{R}^+$ και S^* ένα σύνολο αντικειμένων. Τότε η συνάρτηση αποτίμησης $v : 2^M \to \mathbb{R}^+$ είναι single-minded αν για κάθε σύνολο S, αν $M \supseteq S \supseteq S^*$,τότε $v(S) = v^*$ και $v(S) = 0$, διαφορετικά. Δηλαδή, ένας αγοραστής με μονόπλευρη σκέψη έχει μη μηδενική αξία αν και μόνο αν του διατεθεί ένα σύνολο αντικειμένων που περιέχει την επιθυμητή δέσμη.

1.2.2 Μαντεία

Από τον ορισμό των συναρτήσεων αποτίμησης για τις Συνδυαστικές Δημοπρασίες, παρατηρούμε αμέσως ότι η περιγραφή τους είναι εκθετική στο μέγεθος του m.[1] Εξαιτίας

[1] Για την ακρίβεια, υπάρχουν κλάσεις αποτίμησης που είναι πολυωνυμικές στην περιγραφή τους (π.χ. για τις προσθετικές αποτιμήσεις χρειαζόμαστε μόνο m αριθμούς, έναν για κάθε μονό σύνολο), αλλά αυτές οι κλάσεις δεν είναι ούτε γενικές ούτε αρκετά εκφραστικές σε σχέση με τις subadditive ή submodular, για παράδειγμα.

αυτού, συνήθως υποτίθεται ότι υπάρχει πρόσβαση σε μαντεία. Μπορούμε να θεωρήσουμε τα μαντεία ως εργαλεία στα χέρια των αγοραστών που πρέπει να υπολογίσουν τη βέλτιστη γι' αυτούς απόφαση, δηλαδή ποια δέσμη διαθέσιμων στοιχείων να επιλέξουν. Για το σκοπό αυτό, παρακάτω θα ορίσουμε ορισμένα από τα μαντεία που χρησιμοποιούνται στη βιβλιογραφία.

Definition 1.2.7 (Μαντείο τιμών)**.** Ένα μαντείο τιμών δέχεται ως είσοδο ένα σύνολο S και εξάγει έναν αριθμό: την τιμή της συνάρτησης της αποτίμησης στο σύνολο S, δηλαδή $v_i(S)$.

Definition 1.2.8 (Μαντείο ζήτησης)**.** Ένα μαντείο ζήτησης δέχεται ως είσοδο ένα διάνυσμα τιμών $\mathbf{p} = (p_1, p_2, \dots, p_m)$ και εξάγει μια δέσμη ζήτησης κάτω από αυτές τις τιμές, δηλαδή ένα σύνολο $S \subseteq M$ που μεγιστοποιεί το $v_i(S) - \sum_{j \in S} p_j$, το οποίο είναι η χρησιμότητα του αγοραστή κάτω από αυτές τις τιμές.

Ο παραπάνω ορισμός ισχύει για τιμές αντικειμένων, ωστόσο ο ορισμός μπορεί να επεκταθεί ώστε να συμπεριλάβει την τιμολόγηση δεσμίδων, δηλαδή όταν η τιμή μιας δέσμης δεν μπορεί απαραίτητα να εκφραστεί ως άθροισμα των τιμών των αντικειμένων που περιέχει.

Definition 1.2.9 (XOS Μαντείο)**.** Για μια XOS συνάρτηση v_i, το μαντείο δέχεται ως είσοδο ένα σύνολο T και επιστρέφει την αντίστοιχη προσθετική αντιπροσωπευτική συνάρτηση για το σύνολο T, δηλαδή μια προσθετική συνάρτηση $A_i(\cdot)$ τέτοια ώστε (i) $v_i(S) \geq A_i(S)$ για κάθε $S \subset M$, και (ii) $v_i(T) = A_i(T)$.

Τα XOS oracles δεν χρησιμοποιούνται ευρέως στη βιβλιογραφία, ωστόσο τα ορίζουμε επειδή ορισμένα από τα αποτελέσματα που θα παρουσιαστούν στη συνέχεια απαιτούν πρόσβαση σε αυτά. Ένα ενδιαφέρον γεγονός είναι ότι για υποπολλαπλάσιες αποτιμήσεις ένα XOS μαντείο μπορεί να υλοποιηθεί μέσω πολυωνυμικά πολλών ερωτημάτων σε ένα μαντείο τιμών.

1.2.3 Μεγιστοποίηση Κοινωνικής Ωφέλειας και Εσόδων

Από τη σκοπιά του δημοπράτη, ο απώτερος στόχος μιας δημοπρασίας είναι να παράγει μια ανάθεση αντικειμένων σε παίκτες και πληρωμές που μεγιστοποιούν έναν συνολικό στόχο, δηλαδή έναν στόχο που εξαρτάται από το συνολικό αποτέλεσμα του μηχανισμού. Οι δύο πιο συχνά χρησιμοποιούμενοι στόχοι είναι οι εξής: 1) *έσοδα*, που είναι οι συνολικές πληρωμές που πραγματοποιούν οι αγοραστές στον μηχανισμό και 2) *κοινωνική ωφέλεια* η συνολική αξία που εξάγεται από ένα συγκεκριμένο αποτέλεσμα. Παρακάτω ορίζουμε επίσημα τους παραπάνω στόχους

Definition 1.2.10 (Κοινωνική Ωφέλεια)**.** Έστω $(f, \mathbf{p})$ ένας μηχανισμός. Η Κοινωνική Ωφέλεια ενός αποτελέσματος $\omega \in \Omega$ είναι:

$$SW = \sum_i v_i(\omega)$$

Η μεγιστοποίηση της κοινωνικής ωφέλειας σημαίνει ότι αναζητούμε αποτελέσματα που βελτιστοποιούν τη συνολική αξία των αγοραστών. Κατά μία έννοια, σε ένα βέλτιστο αποτέλεσμα ευημερίας οι αγοραστές είναι συλλογικά ευχαριστημένοι με αυτό που παίρνουν.

Definition 1.2.11 (Έσοδα του πωλητή)**.** Έστω $(f, \mathbf{p})$ ένας μηχανισμός. Τα έσοδα του πωλητή για τον μηχανισμό $(f, \mathbf{a})$ για ένα δεδομένο διάνυσμα δράσης $\mathbf{b}$ είναι:

$$Rev = \sum_i p_i(\mathbf{a})$$

Η μεγιστοποίηση των εσόδων και η μεγιστοποίηση της κοινωνικής ωφέλειας είναι δύο εντελώς διαφορετικοί στόχοι. Οι τεχνικές που χρησιμοποιούνται για την ανάλυση του πρώτου είναι πολύ διαφορετικές σε σύγκριση με τον δεύτερο. Επειδή δεν είναι το κύριο θέμα της εργασίας θα συζητήσουμε εν συντομία για τέτοιους μηχανισμούς στα επόμενα κεφάλαια, αλλά πρώτα μας λείπει ένα πολύ σημαντικό συστατικό του σχεδιασμού μηχανισμών: τα κίνητρα.

1.2.4 Φιλαλήθεια και Ατομική Ορθολογικότητα

Ο τομέας του Σχεδιασμού Μηχανισμών επικεντρώνεται κυρίως στην ενσωμάτωση κινήτρων στο σχεδιασμό αλγορίθμων. Στην προηγούμενη ενότητα, είδαμε ορισμένα αποτελέσματα που μεγιστοποιούν την κοινωνική ευημερία όταν οι συναρτήσεις αποτίμησης όλων των παραγόντων είναι δημόσια γνωστές. Δηλαδή, τα προαναφερθέντα αποτελέσματα σχηματίζουν και επιλύουν ένα πρόβλημα βελτιστοποίησης για τον προσδιορισμό των νικητών της δημοπρασίας. Ωστόσο, οι συναρτήσεις αποτίμησης αποτελούν ιδιωτική πληροφορία για τους παίκτες. Δηλαδή, σε έναν μηχανισμό άμεσης αποκάλυψης, οι πλειοδότες αναφέρουν τις προσφορές τους (οι οποίες μπορεί να είναι διαφορετικές από τις πραγματικές τους αποτιμήσεις) και ο μηχανισμός τις συλλέγει, αλλά δεν έχει πρόσβαση στις πραγματικές τους προτιμήσεις. Επομένως, η κύρια πρόκληση του σχεδιασμού του μηχανισμού έγκειται στο πώς ο σχεδιαστής, εφοδιασμένος με την εξουσία να καθορίζει την κατανομή και τις πληρωμές, δίνει κίνητρα για την αποκάλυψη της αλήθειας.

Πρώτον, οι αγοραστές θα πρέπει να έχουν κίνητρο να συμμετέχουν στον μηχανισμό. Πιο συγκεκριμένα, κάθε πλειοδότης δεν θα πρέπει ποτέ να παρατηρεί αρνητική χρησιμότητα ανεξάρτητα από το αποτέλεσμα του μηχανισμού, δηλαδή δεν θα πρέπει να χάνει χρήματα απλώς και μόνο συμμετέχοντας. Η έννοια αυτή ονομάζεται ατομική ορθολογικότητα και ορίζεται επίσημα παρακάτω.

Definition 1.2.12 (Ατομική Ορθολογικότητα)**.** Ένας μηχανισμός $(f, \mathbf{p})$ είναι *individually rational* αν για κάθε πλειοδότη i, αποτίμηση v_i και διάνυσμα bit $\mathbf{b}_{-i}$ των άλλων πλειοδοτών

$$v_i(f(v_i, \mathbf{b}_{-i})) - p_i(f(v_i, \mathbf{b}_{-i}) \geq 0$$

Η φιλαλήθεια μπορεί να οριστεί με διάφορους τρόπους που εξαρτώνται από τις πληροφορίες που χρησιμοποιεί ένας πράκτορας κατά τη λήψη της απόφασής του. Η έννοια της συμβατότητας κινήτρων κυρίαρχης στρατηγικής είναι θεμελιώδης στον τομέα του σχεδιασμού μηχανισμών. Ανεπίσημα, συνεπάγεται ότι κάθε πράκτορας δεν μπορεί να είναι χειρότερος από το να αποκαλύψει τον πραγματικό του τύπο, ανεξάρτητα από το τι προσφέρουν οι άλλοι πράκτορες.

Definition 1.2.13 (Dominant Strategy Incentive Compatibility (DSIC))**.** Ένας μηχανισμός είναι *συμβατός με τα κίνητρα της κυρίαρχης στρατηγικής* αν η αποκάλυψη της αλήθειας είναι προς το συμφέρον ενός πράκτορα ακόμη και αφού ο πράκτορας παρατηρήσει τους τύπους των άλλων. Τυπικά, για όλα τα i, v_i, b_i, και όλες τις προσφορές $\mathbf{b}_{-i}$ των άλλων bidders :

$$v_i(f(v_i, \mathbf{b}_{-i})) - p_i(v_i, \mathbf{b}_{-i}) \geq v_i(f(b_i, \mathbf{b}_{-i})) - p_i(b_i, \mathbf{b}_{-i})$$

Συχνά θα αναφερόμαστε σε μηχανισμούς που είναι DSIC ως *φιλαλήθεις* ή *μη στρατηγικούς*. Επιπλέον, ένας μηχανισμός που είναι DSIC υπονοείται επίσης ότι είναι ατομικά ορθολογικός, εκτός αν αναφέρεται διαφορετικά.

1.3 Φιλαλήθεις Μηχανισμοί για Μεγιστοποίηση Κοινωνικής Ωφέλειεας

Αρχικά, θα εξετάσουμε την έννοια της Ουαλρασιανής Ισορροπίας. Ο καθορισμός των τιμών της Ουαλρασιανής Ισορροπίας για τα αντικείμενα και το να αφήσουμε απλά τους πλειοδότες να πάρουν ό,τι ζητούν, φαίνεται σαν ένας απλός και φιλαλήθης μηχανισμός. Δυστυχώς, μια Ουαλρασιανή Ισορροπία δεν υπάρχει πάντα, όπως φαίνεται στο ακόλουθο παράδειγμα.

Example 1.3.1 (Μη ύπαρξη Ουαλρασιανής Ισορροπίας). Έστω δύο αγοραστές, η Alice και ο Bob και δύο αντικείμενα a, b. Η Alice έχει αξία 2 για κάθε μη κενό σύνολο και ο Bob έχει αξία 3 μόνο για ολόκληρη τη δέσμη $\{a, b\}$. Διακρίνουμε δύο περιπτώσεις: Εάν η αντιστοιχία ζήτησης του Bob δεν είναι κενή,τότε η αντιστοιχία της Alice είναι κενή. Επομένως, το άθροισμα των τιμών στην Ισορροπία είναι το πολύ 3. Τότε υπάρχει ένα στοιχείο το οποίο έχει τιμή το πολύ 1,5. Αυτό σημαίνει ότι αυτό το στοιχείο βρίσκεται επίσης στην αντιστοιχία ζήτησης της Alice, πράγμα που αποτελεί αντίφαση. Αν η αντιστοιχία ζήτησης του Bob είναι κενή, αυτό σημαίνει ότι το άθροισμα των τιμών είναι μεγαλύτερο από 3. Επομένως, η ζήτηση της Alice δεν μπορεί να περιέχει $\{a, b\}$ αφού έχει τιμή 2. Η Alice μπορεί να ζητήσει μόνο ένα στοιχείο, ας πούμε b για τιμή μικρότερη από 2. Κατά συνέπεια, το στοιχείο a παραμένει απούλητο, αλλά έχει μη μηδενική τιμή, πράγμα που αποτελεί και πάλι αντίφαση.

Αν και η ύπαρξη τιμών εκκαθάρισης της αγοράς δεν είναι πάντα εγγυημένη, παρουσιάζουμε εδώ δύο θεωρήματα κοινωνικής ωφέλειας που παρέχουν αναγκαίες και ικανές συνθήκες για την ύπαρξη και τη βέλτιστη λειτουργία μιας τέτοιας ισορροπίας, χαρακτηρίζοντας έτσι πλήρως την έννοια. Θα αποδείξουμε τα δύο θεωρήματα καταφεύγοντας στη Θεωρία Γραμμικού Προγραμματισμού. Ας πάρουμε το δυϊκό του *LP διαμόρφωσης* που παρουσιάστηκε στη δεύτερη ενότητα. Οι περιορισμοί για κάθε αντικείμενο j και κάθε αγοραστή i μεταφράζονται σε δυϊκές μεταβλητές p_j και u_i αντίστοιχα. Όπως θα αποδείξουμε, οι μεταβλητές αυτές μπορούν να θεωρηθούν ως τιμές αντικειμένων και χρησιμότητα αγοραστών. Το δυϊκό πρόγραμμα είναι το ακόλουθο:

$$\min \sum_{i=1}^{n} u_i + \sum_{j \in M} p_j \tag{1.1}$$

$$s.t\ u_i + \sum_{j \in S} p_j \geq v_i(S) \qquad \forall i, S \tag{1.2}$$

$$p_j \geq 0 \qquad \forall j \in M \tag{1.3}$$

$$u_i \geq 0 \qquad \forall i \in [n] \tag{1.4}$$

Το πρώτο Θεώρημα Ωφέλειας ισχυρίζεται ότι μια Ουαλρασιανή Ισορροπία μεγιστοποιεί την κοινωνική ωφέλεια μεταξύ όλων των εφικτών κλασματικών λύσεων. Η τυπική δήλωση ακολουθεί.

Theorem 1.3.2 (Πρώτο θεώρημα ωφέλειας). *Έστω $(N, M, \mathbf{v})$ μια συνδυαστική δημοπρασία και $(\mathbf{S}^*, \mathbf{p}^*)$ μια Ουαλρασιανή ισορροπία για αυτή τη δημοπρασία. Τότε για κάθε εφικτή κλασματική λύση της διαμόρφωσης LP $\{x_{i,S}\}_{i,S}$ ισχύει ότι $\sum_i v_i(S_i^*) \geq \sum_{i,S} x_{i,S} v_i(S)$*

Proof. Σταθεροποιούμε τον αγοραστή i. Δεδομένου ότι $(\mathbf{S}^*, \mathbf{p}^*)$ είναι μια Ουαλρασιανή ισορροπία ισχύει εξ ορισμού ότι μεγιστοποιεί τη χρησιμότητά του:

$$v_i(S_i^*) - \sum_{j \in S_i^*} p_j^* \geq v_i(S) - \sum_{j \in S} p_j^*$$

για κάθε $S \subseteq M$. Πολλαπλασιάζοντας και τις δύο πλευρές με $x_{i,S}$ και αθροίζοντας σε όλα τα i, S έχουμε:

$$\sum_{i,S} x_{i,S} \cdot v_i(S_i^*) - \sum_{i,S} \sum_{j \in S_i^*} x_{i,S} \cdot p_j^* \geq \sum_{i,S} x_{i,S} \cdot v_i(S) - \sum_{i,S} \sum_{j \in S} x_{i,S} \cdot p_j^*$$

Αφού $\sum_S x_{i,S} \leq 1$ για όλα τα i, έχουμε:

$$\sum_{i,S} v_i(S_i^*) - \sum_{i} \sum_{j \in S_i^*} p_j^* \geq \sum_{i,S} x_{i,S} \cdot v_i(S) - \sum_{i,S} \sum_{j \in S} x_{i,S} \cdot p_j^*$$

Επομένως, αρκεί να δείξουμε ότι $\sum_i \sum_{j \in S_i^*} p_j^* \geq \sum_{i,S} \sum_{j \in S} x_{i,S} \cdot p_j^*$. Στην Ουαλρασιανή Ισορροπία, κάθε στοιχείο περιλαμβάνεται σε ένα το πολύ σύνολο, και για όλα τα στοιχεία j που δεν ανατίθενται σε κανέναν ισχύει ότι $p_j^* = 0$, επομένως το αριστερό μέρος της παραπάνω ανισότητας είναι ίσο με $\sum_{j \in M} p_j^*$. Η δεξιά πλευρά μπορεί να επαναδιατυπωθεί ως εξής: $\sum_{j \in M} p_j^* \sum_i \sum_{S|j \in S} x_{i,S}$ το οποίο είναι το πολύ $\sum_{j \in M} p_j^*$ από τον πρώτο περιορισμό της πρωταρχικής LP. □

Στο Πρώτο Θεώρημα Ωφέλειας δείξαμε ότι οι Ουαλρασιανές Ισορροπίες, όταν υπάρχουν, παρέχουν μια βέλτιστη ακέραια λύση στη διαμόρφωση LP. Στο Δεύτερο Θεώρημα Ωφέλειας θα δείξουμε ότι ισχύει και το αντίστροφο: αν υπάρχει μια βέλτιστη ακέραια λύση στο LP, τότε η λύση αυτή είναι μια Ουαλρασιανή Ισορροπία. Δηλαδή, οι Ουαλρασιανές Ισορροπίες υπάρχουν εάν και μόνο εάν η βέλτιστη λύση της LP διαμόρφωσης είναι ακέραια.

Theorem 1.3.3 (Δεύτερο θεώρημα ωφέλειας)**.** *Έστω* $(N, M, \mathbf{v})$ *μια Συνδυαστική Δημοπρασία, τότε αν υπάρχει μια βέλτιστη ακέραια λύση στη διαμόρφωση LP, τότε υπάρχει επίσης μια Ουαλρασιανή Ισορροπία.*

Proof. Έστω $\mathbf{S}^* = (S_1^*, S_2^*, \ldots, S_n^*)$ μια βέλτιστη ακέραια λύση που δίνεται από το πρωτέυον LP και $\mathbf{p}^* = (p_1^*, p_2^*, \ldots, p_n^*)$, $\mathbf{u}^* = (u_1^*, u_2^*, \ldots u_n^*)$ να είναι η αντίστοιχη λύση του δυϊκού LP. Θα δείξουμε ότι $(\mathbf{S}^*, \mathbf{p}^*)$ είναι μια Ουαλρασιανή Ισορροπία. Σταθεροποιούμε τον αγοραστή i. Σύμφωνα με τις comlementary slackness συνθήκες, $x_{i,S_i^*} = 1$ συνεπάγεται ότι $u_i^* = v_i(S_i^*) - \sum_{j \in S_i^*} p_j^*$. Επομένως, $v_i(S_i^*) - \sum_{j \in S_i^*} p_j^* \geq v_i(T) - \sum_{j \in T} p_j^*$ για κάθε $T \subseteq M$. Επιπλέον, αν ένα στοιχείο j δεν έχει ανατεθεί, ο πρώτος περιορισμός του πρωτεύοντος LP ισχύει αυστηρά με ανισότητα και από το complementary slackness αυτό συνεπάγεται ότι $p_j^* = 0$. Με το παραπάνω επιχείρημα, συμπεραίνουμε ότι $(\mathbf{S}^*, \mathbf{p}^*)$ είναι μια Ουαλρασιανή Ισορροπία. □

Τα δύο θεωρήματα ωφέλειας υποδηλώνουν ότι υπάρχει Ουαλρασιανή Ισορροπία αν και μόνο αν το integrality gap του LP είναι μηδέν. Επιπλέον, το δεύτερο θεώρημα μας παρέχει έναν τρόπο να βρούμε μια Ουαλρασιανή Ισορροπία: Λύστε την κλασματική χαλάρωση της LP διαμόρφωσης. Εάν η βέλτιστη λύση τυχαίνει να είναι ακέραια, τότε λύστε το δυϊκό πρόγραμμα για να βρείτε τις τιμές των στοιχείων που επιβάλλουν την Ισορροπία. Οι Ουαλρασιανές Ισορροπίες είναι εγγυημένο ότι υπάρχουν για μια κατηγορία συναρτήσεων αποτίμησης που ονομάζονται gross-substitutes (που περιέχει unit-demand και additive) [55]

Η μεγιστοποίηση της κοινωνικής ωφέλειας χωρίς να υποθέτουμε ότι έχουμε εκ των προτέρων πληροφορίες σχετικά με τις κατανομές πιθανότητας των αξιολογήσεων των αγοραστών έχει επίσης μελετηθεί εκτενώς στη βιβλιογραφία. Για τις αποτιμήσεις XOS, οι συγγραφείς στο [29] απέδειξαν μια προσέγγιση $O(\log^2 m)$ το 2006. Μεταγενέστερη εργασία το 2007 [24] βελτίωσε την προηγούμενη εγγύηση προσέγγισης σε $O(\log m \log \log \log m)$. Και οι δύο μηχανισμοί είναι τυχαιοποιημένοι, επομένως επιτυγχάνουν αυτή την αναμενόμενη αναλογία προσέγγισης. Το 2012 οι Krysta και Vöcking [48] έδειξαν έναν $O(\log m)$ -προσεγγιστικό τυχαιοποιημένο μηχανισμό ο οποίος βελτιώθηκε περαιτέρω σε $O(\sqrt{\log m})$ από τον Dobzinski το 2016 [26]. Πιο πρόσφατη εργασία στο [4] βελτίωσε εκθετικά το αποτέλεσμα του Dobzinski αποδεικνύοντας έναν $O((\log \log m)^3)$-προσεγγιστικό τυχαιοποιημένο μηχανισμό. Όλοι οι παραπάνω μηχανισμοί προϋποθέτουν πρόσβαση τόσο σε ερωτήματα αξίας όσο και σε ερωτήματα ζήτησης και είναι επίσης καθολικά φιλαληθείς. Το 2021, οι συγγραφείς στο [5] απέδειξαν έναν $O((\log \log m))^3$-προσεγγιστικό τυχαιοποιημένο φιλαλήθη μηχανισμό για subadditive αποτιμήσεις ο οποίος έσπασε το πρόβλημα του λογαριθμικού φράγματος το οποίο ήταν από το 2007 όταν ο Dobzinski [24] απέδειξε μια προσέγγιση $O(\log m \log \log m)$. Ας σημειωθεί ότι η εργασία αυτή βελτιώνει επίσης τον συντελεστή προσέγγισης για τις αποτιμήσεις XOS, από $O((\log \log m))^3$ σε $O((\log \log m))^2$. Τέλος, η [5] είναι επίσης η τρέχουσα εργασία αιχμής σε προσεγγιστικά βέλτιστους, αληθείς, τυχαιοποιημένους μηχανισμούς για υποαθροιστικές και XOS συνδυαστικές δημοπρασίες. Από τα παραπάνω αποτελέσματα μπορούμε να παρατηρήσουμε το εκφραστικό πλεονέκτημα της ζήτησης έναντι των χρηστικών μαντείων. Για να γίνουμε πιο συγκεκριμένοι, οι συγγραφείς στο [28] παρέχουν έναν αλγόριθμο που επιτυγχάνει ένα προσεγγιστικό κάτω φράγμα $O(\sqrt{m})$ για το πρόβλημα της φιλαλήθους μεγιστοποίησης της κοινωνικής ευημερίας χρησιμοποιώντας μόνο μαντεία αξίας, το οποίο στη συνέχεια αποδείχθηκε το καλύτερο δυνατό για ερωτήματα αξίας ακόμη και αν χρησιμοποιούνται τυχαιοποιημένοι μηχανισμοί [25].

1.4 Το πλαίσιο εργασίας των Prophet Inequalities

Τα Prophet Inequalities είναι ένα πρόβλημα από τη θεωρία βέλτιστης διακοπής που ανακαλύφθηκε τη δεκαετία του '70 από τους Krengel, Sucheston και Garling. [46, 47]. Περιλαμβάνει ένα *παίκτη* και ένα *προφήτη*. Από τη μία πλευρά, ο παίκτης βρίσκεται αντιμέτωπος με μια ακολουθία n τυχαίων μεταβλητών, X_i. Αυτές οι τυχαίες μεταβλητές είναι δείγματα που προέρχονται από ανεξάρτητες (αλλά όχι απαραίτητα ταυτόσημες) γνωστές κατανομές D_i. Κατά την άφιξη κάθε τυχαίας μεταβλητής, ο παίκτης πρέπει να αποφασίσει μεταξύ των δύο ακόλουθων εναλλακτικών λύσεων: είτε να πάρει την τιμή του δείγματος της τυχαίας μεταβλητής και να σταματήσει (χωρίς να παρατηρήσει μελλοντικές τυχαίες μεταβλητές) είτε να απορρίψει την τυχαία μεταβλητή και να χάσει την τιμή της για πάντα, συνεχίζοντας τις επόμενες. Από την άλλη πλευρά, ο προφήτης είναι παντογνώστης (δηλ. μπορεί να δει το μέλλον) και διεκδικεί πάντα την τυχαία μεταβλητή με την υψηλότερη τιμή. Η ανισότητα του προφήτη είναι το πρόβλημα της εύρεσης μιας *βέλτιστης στρατηγικής διακοπής* για τον τζογαδόρο που συλλέγει αξία που είναι συγκρίσιμη με τη βέλτιστη εκ των υστέρων (δηλ. με αυτή που παίρνει ο προφήτης). Οι Krengel και Sucheston έδειξαν μια στρατηγική που εγγυάται ότι η ανταμοιβή του τζογαδόρου είναι τουλάχιστον 50 % της ανταμοιβής του προφήτη. Η βέλτιστη στρατηγική που επιτυγχάνει την προηγούμενη προσέγγιση δίνεται με ανάδρομη επαγωγή. Εάν ο παίκτης φτάσει στην τυχαία μεταβλητή X_n είναι βέλτιστο να την αποδεχθεί, καθώς δεν υπάρχουν άλλες μεταβλητές που θα ακολουθ- ησουν. Τώρα ο επαγωγικός ορισμός της στρατηγικής έχει ως εξής: ο παίκτης αποδέχεται την X_i αν και μόνο αν η τιμή της είναι μεγαλύτερη από την αναμενόμενη αξία που εισπράττει ξεκινώντας από την X_{i+1} έως την X_n. Στην πραγματικότητα, με ένα απλό παράδειγμα μπορούμε να δείξουμε ότι ο παράγοντας 2 είναι ο καλύτερος δυνατός για το πρόβλημα. Θεωρήστε τις ακόλουθες δύο τυχαίες μεταβλητές οι οποίες φτάνουν με τη σειρά που είναι δεικτοδοτημένες: X_1 είναι ντετερμινιστικά 1 και X_2 είναι $\frac{1}{\epsilon}$ με πιθανότητα ϵ και 0 διαφορετικά. Οποιαδήποτε στρατηγική διακοπής αποδίδει στον παίκτη αναμενόμενα κέρδη της τιμής 1. Από την άλλη πλευρά, ο προφήτης, ο οποίος επιλέγει πάντοτε τη μέγιστη πραγματοποιηθείσα αξία, λαμβάνει αναμενόμενη αξία $2 - \epsilon$.

Στη δεκαετία του '80, ο Cahn [59] μελέτησε την απόδοση των στρατηγικών που βασίζονται σε ένα κατώφλι, δηλαδή των στρατηγικών που χρησιμοποιούν ένα μόνο κατώφλι για να αποφασίσουν αν θα αποδεχτούν ή θα απορρίψουν μια τυχαία μεταβλητή. Απέδειξε ότι θέτοντας ένα κατώφλι που είναι η διάμεσος της κατανομής $\max_i X_i$ και επιλέγοντας την πρώτη τυχαία μεταβλητή που η τιμή της είναι πάνω από αυτό το κατώφλι επιτυγχάνεται επίσης μια προσέγγιση 2. Ωστόσο, το πλεονέκτημα αυτής της στρατηγικής έναντι αυτής των Krengel και Sucheston, είναι ότι η προσέγγιση παραμένει αναλλοίωτη ακόμη και αν η σειρά των τυχαίων μεταβλητών είναι αυθαίρετη. Επομένως, μια τυπική υπόθεση στη βιβλιογραφία των Prophet Inequalities είναι ότι η σειρά επιλέγεται από έναν αντίπαλο, ο οποίος μπορεί να είναι προσαρμοστικός, δηλαδή η επιλογή της επόμενης τυχαίας μεταβλητής στην ακολουθία μπορεί να εξαρτάται από τις τιμές των προηγούμενων τυχαίων μεταβλητών και τις αποφάσεις του παίκτη.

Αρκετές δεκαετίες μετά τα αποτελέσματα των Krengel, Sucheston και Cahn, οι συγγραφείς στο [41] διαπίστωσαν μια σύνδεση μεταξύ των Prophet Inequalities και των μηχανισμών των αναρτημένων τιμών. Παρατήρησαν ότι οι αλγόριθμοι που βασίζονται σε κατώτατα όρια για την ανισότητα του προφήτη μπορούν να θεωρηθούν ως τιμές για έναν (διαδοχικό) μηχανισμό αναρτημένων τιμών. Πάρτε, για παράδειγμα, το ακόλουθο σενάριο. Υπάρχει ένας πωλητής που έχει στην κατοχή του ένα αδιαίρετο αντικείμενο. Οι αγοραστές καταφθάνουν ένας-ένας με αυθαίρετη σειρά, έχοντας ανεξάρτητα (και όχι απαραίτητα πανομοιότυπα) κατανεμημένες τιμές για το αντικείμενο. Οι κατανομές είναι γνωστές στον πωλητή, ο οποίος υπολογίζει και δημοσιεύει μια τιμή. Κατά την άφιξη κάθε αγοραστή, το αντικείμενο πωλείται σε αυτόν εάν και μόνο εάν η αξία του είναι πάνω από την αναρτημένη τιμή (και δεν έχει πωληθεί σε κάποιον προηγούμενο αγοραστή). Στόχος του μηχανισμού είναι η μεγιστοποίηση της συνολικής αξίας που εξάγεται, σε σύγκριση με το καλύτερο εκ των υστέρων βέλτιστο, δηλαδή τη μέγιστη πραγματοποιηθείσα αξία κάθε αγοραστή. Παρατηρούμε ότι ο μηχανισμός που περιγράφεται παραπάνω δεν ζητά από τους αγοραστές να αποκαλύψουν τους τύπους τους, αφού η μόνη πληροφορία που χρησιμοποιεί για να παράγει μια κατανομή είναι αν η αξία είναι πάνω ή κάτω από το επιλεγμένο κατώφλι. Μπορούμε να συμπεράνουμε ότι ένας μηχανισμός με απλή δημοσίευση τιμών δίνει μια προσέγγιση 2 στη βέλτιστη ευημερία. Επιπλέον, αυτό το αποτέλεσμα είναι σφιχτό, με την έννοια ότι κανένας άλλος μηχανισμός δεν μπορεί να βελτιώσει αυτή την εγγύηση, κάτι που μπορεί εύκολα να γίνει αντιληπτό από την απόδειξη της σφιχτότητας της ανισότητας του προφήτη: θεωρήστε έναν αγοραστή που έχει μια ντετερμινιστική τιμή 1 για το αντικείμενο και έρχεται πρώτος. Οποιοσδήποτε μηχανισμός θα πρέπει να αποφασίσει αν θα δώσει ή όχι το αντικείμενο στον πρώτο αγοραστή, αποδίδοντας το χαμηλότερο όριο.

Λόγω της σύνδεσής τους με τους μηχανισμούς αναρτημένων τιμών, οι ανισότητες προφητών έγιναν ένα σχετικό και έγκυρο ερευνητικό μονοπάτι στον τομέα του αλγοριθμικού σχεδιασμού μηχανισμών. Το 2012, οι Kleinberg και Weinberg [45] έδωσαν ένα νέο κατώφλι για το αρχικό, μονοπαραμετρικό περιβάλλον, που επιτυγχάνει τον ίδιο παράγοντα. Ο καθορισμός του κατωφλίου ίσου με $\frac{1}{2}\mathbb{E}[\max_i X_i]$ είναι επίσης μια προσέγγιση 2.

Στα επόμενα χρόνια, οι μεταγενέστερες εργασίες γενίκευσαν τα παραπάνω αποτελέσματα σε μητροειδή, πολυματροειδή και διασταύρωση μητροειδών [45, 30], knapsack [32], k matroid [39] και προς τα κάτω κλειστά με περιορισμένο μέγιστο μέγεθος συνόλου [57] περιορισμούς feasibility.

Για Συνδυαστικές Δημοπρασίες με αποτιμήσεις XOS (οι οποίες περιλαμβάνουν επίσης τις submodular) οι συγγραφείς στο [38] παρέχουν μια $\frac{2e}{e-1}$-ανταγωνιστική Prophet Inequality που προϋποθέτει πρόσβαση black-box σε έναν αλγόριθμο του Vondrák [61] -ο οποίος επιλύει βέλτιστα το offline πρόβλημα- και ένα μαντείο που απαντά σε ερωτήματα XOS. Η μεταγενέστερη εργασία στο [32] βελτιώνει το προηγούμενο αποτέλεσμα παρέχοντας μια 2-ανταγωνιστική Prophet Inequality για αποτιμήσεις XOS, υποθέτοντας πρόσβαση σε μαντεία ζήτησης. Αυτό ταιριάζει με το κατώτερο όριο για αυτή την κατηγορία αποτιμήσεων που κληρονομείται από την περίπτωση ενός μόνο στοιχείου. Οι συγγραφείς στο [32] παρέχουν επίσης μια ενοποιητική προσέγγιση για την απόδειξη Prophet Inequalities εισάγοντας την έννοια των balanced-prices. Το επιχείρημά τους περιλαμβάνει μια αναγωγή από το πλαίσιο Bayes στην πλήρη πληροφόρηση και μια απόδειξη ότι όταν υπάρχουν "καλές" τιμές για το τελευταίο, κατάλληλα κλιμακωτές εκδοχές αυτών είναι επίσης "καλές" για το πρώτο. Τέλος, η προηγούμενη γραμμή εργασίας συνεπάγεται επίσης μια $O(\log m)$ -ανταγωνιστική Prophet Inequality για subadditive αποτιμήσεις προσεγγίζοντας τις subadditive με XOS αποτιμήσεις [9].

Για subadditive αποτιμήσεις, η [31] επιτυγχάνει εκθετική βελτίωση αποδεικνύοντας μια $o(\log \log m)$ Prophet Inequality. Το αποτέλεσμά τους είναι επίσης υπολογιστικό: εξετάζουν το δυϊκό πρόγραμμα της LP διαμόρφωσης και δείχνουν πώς να υπολογίζουν αποτελεσματικά τις τιμές εκτελώντας τον Ελλειψοειδή Αλγόριθμο με ένα separation oracle που μπορεί να υλοποιηθεί χρησιμοποιώντας ερωτήματα ζήτησης. Στην πραγματικότητα, σε αυτό το πλαίσιο, το μαντείο ζήτησης και το μαντείο διαχωρισμού συμπίπτουν. Παρατηρήστε ότι το δυϊκό πρόγραμμα (στην περίπτωση των συνδυαστικών δημοπρασιών) έχει εκθετικά πολλούς περιορισμούς - έναν για κάθε πιθανό υποσύνολο αντικειμένων και αγοραστή. Δεδομένης μιας ακολουθίας τιμών αντικειμένων $p_1, p_2, \ldots, p_m$, το μαντείο διαχωρισμού που χρησιμοποιείται από τον αλγόριθμο ελλειψοειδούς, πρέπει να υπολογίσει το σύνολο S_i που μεγιστοποιεί τη χρησιμότητα i του αγοραστή κάτω από αυτές τις τιμές. Αυτός ακριβώς είναι ο ορισμός του μαντείου ζήτησης στο οποίο οι συγγραφείς υποθέτουν πρόσβαση για τον αλγόριθμό τους. Πρόσφατη εργασία στο [21] αποδεικνύει την ύπαρξη μιας 6-προσεγγιστικής προφητικής ανισότητας για subadditve αξιολογήσεις, επιλύοντας έτσι ένα κεντρικό ανοιχτό πρόβλημα στην περιοχή. Η επακόλουθη εργασία στο [7] δείχνει πώς μπορεί να γίνει το προηγούμενο αποτέλεσμα υπολογιστικό, αποδεικνύοντας ότι οποιαδήποτε Prophet Inequality μπορεί να εφαρμοστεί ως μηχανισμός αναρτημένων τιμών με τουλάχιστον εξίσου καλή εγγύηση κοινωνικής ευημερίας.

1.5 Prophet Inequalities για δρομολόγηση και έλεγχο αποδοχής σε δίκτυα με χωρητικότητα.

Σκεφτείτε ένα τηλεπικοινωνιακό δίκτυο που εξυπηρετεί αιτήματα. Το δίκτυο αποτελείται από κόμβους, οι οποίοι για παράδειγμα μπορεί να είναι δρομολογητές, και διμερείς συνδέσεις μεταξύ των κόμβων που διαθέτουν ένα συγκεκριμένο εύρος ζώνης. Το εύρος ζώνης είναι αντιπροσωπευτικό της ικανότητας της σύνδεσης να εξυπηρετεί ταυτόχρονες αιτήματα. Τα αιτήματα φθάνουν διαδοχικά στους κόμβους εκκίνησης, απαιτώντας να δρομολογηθούν σε μια διαδρομή που θα τα οδηγήσει στους τερματικούς κόμβους τους. Μελετάμε το πρόβλημα από τη σκοπιά του σχεδιασμού μηχανισμών: κάθε αίτημα προέρχεται από έναν αγοραστή που έχει κάποια ιδιωτική αξία να λάβει ένα μονοπάτι που τον συνδέει με τον επιθυμητό προορισμό του. Ο μηχανισμός διεκδικεί μη μηδενική αξία από τον αγοραστή, εάν και μόνο εάν δρομολογήσει το αίτημα του αγοραστή στο δίκτυο. Υποθέτοντας ότι οι αξίες αντλούνται ανεξάρτητα από δημόσια γνωστές κατανομές, μπορούμε να σχεδιάσουμε έναν μηχανισμό δημοσιευμένης τιμής που να τιμολογεί τις συνδέσεις εύρους ζώνης ώστε να μεγιστοποιεί (κατά προσέγγιση) την κοινωνική ωφέλεια; Οι δύο κύριοι λόγοι για την τιμολόγηση των συνδέσεων (αντί των μονοπατιών) είναι οι εξής:1)υπάρχουν εκθετικά πολλές διαφορετικά μονοπάτια σε ένα δίκτυο, ενώ υπάρχουν μόνο πολυωνυμικά πολλές συνδέσεις, 2) διαφορετικοί κόμβοι μπορεί να ανήκουν σε διαφορετικούς παρόχους υπηρεσιών, επομένως για κάθε κόμβο που ανήκει σε έναν πάροχο, ο τελευταίος θα πρέπει να γνωρίζει μόνο τις τιμές των συνδέσεων προς τους άμεσους γείτονες κάθε κόμβου. Στα προηγούμενα κεφάλαια είδαμε ότι έχει γίνει τεράστια δουλειά στα Prophet Inequalities με complement-free αξιολογήσεις, με αποκορύφωμα ένα Prophet Inequality για subadditive με σταθερό competitive ratio. Αυτό το κεφάλαιο, από την άλλη πλευρά, εξετάζει θετικά και αρνητικά αποτελέσματα για Prophet Inequalities με αποτιμήσεις που παρουσιάζουν συμπληρωματικότητα. Παρατηρήστε ότι στο παραπάνω σενάριο, μια σύνδεση έχει αξία μόνο αν αγοραστεί μαζί με άλλες συνδέσεις με τρόπο που να σχηματίζουν έγκυρο μονοπάτι.

Εξετάζουμε το υποκείμενο πρόβλημα βελτιστοποίησης. Δηλαδή, υποθέτουμε ότι γνωρίζουμε το σύνολο της εισόδου των αιτήσεων r_i και θέλουμε να μεγιστοποιήσουμε την κοινωνική ωφέλεια με βάση τους περιορισμούς που επιβάλλονται από τις χωρητικότητες για κάθε ακμή. Ορίζουμε ως $f_i(v, w)$ τη ροή από το αίτημα r_i στην ακμή (v, w) E. Η διατύπωση του Γραμμικού Προγράμματος, είναι η ακόλουθη:

$$\max \quad \sum_i \sum_{w \in V} v_i f_i(s_i, w) \tag{1.5}$$

$$s.t \quad \sum_i f_i(v, w) \leq c(e) \qquad \forall (v, w) \in E \tag{1.6}$$

$$\sum_w f_i(v, w) - \sum_w f_i(w, v) = 0 \qquad \forall i \in [k], \forall v \neq s_i, t_i \tag{1.7}$$

$$\sum_w f_i(s_i, w) = \sum_w f_i(w, t_i) \qquad \forall i \in [k] \tag{1.8}$$

$$\sum_w f_i(s_i, w) \leq 1 \qquad \forall i \in [k] \tag{1.9}$$

$$f_i(v, w) \in [0, 1] \qquad \forall (v, w) \in E \tag{1.10}$$

Η παραπάνω διατύπωση μπορεί να θεωρηθεί ως ένα πρόβλημα multi-commodity flow, για το οποίο γνωρίζουμε ότι είναι NP-δύσκολο να βρεθεί μια βέλτιστη ακέραια λύση [34]. Στο κλασματικό καθεστώς υπάρχουν λύσεις που βασίζονται στην επίλυση του LP, καθώς και σχήματα προσέγγισης σε πλήρως πολυωνυμικό χρόνο [43].

Ωστόσο, ένα έγκυρο ερώτημα που εξετάζουμε στο υπόλοιπο της παρούσας ενότητας είναι το ακόλουθο: Τι συμβαίνει στη βέλτιστη λύση του παραπάνω LP καθώς οι χωρητικότητες αυξάνονται; Διαισθητικά, αν οι χωρητικότητες είναι αρκετά μεγάλες το πρόβλημα γίνεται ευκολότερο. Αυτό οφείλεται στο γεγονός ότι με μεγαλύτερες χωρητικότητες έχουμε περισσότερο "περιθώριο για σφάλματα": η κατανομή ενός μη βέλτιστου αιτήματος δεν βλάπτει τόσο πολύ τη λύση μας. Κατ' αρχάς, θα μετατρέψουμε το παραπάνω LP σε μια ισοδύναμη, αλλά πιο βολική μορφή.

Στόχος μας είναι να μετατρέψουμε μια βέλτιστη λύση που περιγράφεται από τις ροές ακμών $f_i(v, w)$ σε μια ισοδύναμη λύση που περιγράφεται από τις ροές μονοπατιών $f_{i,p}$, για όλες τις αιτήσεις r_i και για όλες τις διαδρομές $p \in \mathcal{P}_i$. Αυτό γίνεται με τον ακόλουθο αλγόριθμο.

Algorithm 1 Path Decomposition Algorithm

---

input f

for all requests r_i **do**
  while there is a $s_i - t_i$ path p using only edges with $f_i(e) > 0$ **do**
    $f_{i,p} \leftarrow \min_{e \in p} f_i(e)$
    for $e \in p$ **do**
      $f_i(e) \leftarrow f_i(e) - f_{i,p}$
    end for
  end while
end for

output $f_{i,p}$ for each request, for each path

Εκτελώντας τον αλγόριθμο μετατρέψαμε το πρώτο LP στο ακόλουθο:

$$\max \quad \sum_i \sum_{p \in \mathcal{P}_i} v_i f_{i,p} \tag{1.11}$$

$$s.t \quad \sum_i \sum_{p \in \mathcal{P}_i | e \in p} f_{i,p} \leq c(e) \qquad \forall e \in E \tag{1.12}$$

$$\sum_{p \in \mathcal{P}_i} f_{i,p} \leq 1 \qquad \forall i \in [k] \tag{1.13}$$

$$f_{i,p} \in [0,1] \qquad \forall i \in [k], \forall p \in \mathcal{P}_i \tag{1.14}$$

Από εδώ και στο εξής θα χρησιμοποιούμε αυτή την πιο βολική, αλλά παρόλα αυτά ισοδύναμη, μορφή του LP και μπορούμε να εκφράσουμε τη βέλτιστη κλασματική λύση σε όρους ροών μονοπατιών $f_{i,p}$. Ερμηνεύουμε τις ροές $f_{i,p}$ ως κατανομές πιθανοτήτων στα μονοπάτια. Δηλαδή, εκτελούμε μια *τυχαία στρογγυλοποίηση* της κλασματικής λύσης της LP ως εξής: το αίτημα r_i παίρνει το μονοπάτι $p \in \mathcal{P}_i$ με πιθανότητα $f_{i,p}$ και δεν παίρνει τίποτα με πιθανότητα $1 - \sum_{p \in \mathcal{P}_i} f_{i,p}$. Παρατηρούμε αμέσως ότι η αναμενόμενη τιμή αυτής της στρογγυλοποίησης είναι ίση με τη μέγιστη τιμή που επιτυγχάνεται από την LP (Εξίσωση 1.11). Ωστόσο, η 6.8 ικανοποιείται επίσης κατ'αναμενόμενη τιμή. Αυτό σημαίνει ότι υπάρχουν περιπτώσεις όπου η τυχαιοποιημένη στρογγυλοποίηση μας παραβιάζει τους περιορισμούς χωρητικότητας. Για να αντιμετωπίσουμε αυτό το πρόβλημα, μειώσουμε όλες τις χωρητικότητες κατά έναν παράγοντα $1 - \epsilon$, όπου $0 < \epsilon < 1$. Μια νέα λύση $f'_{i,p} = (1 - \epsilon) f_{i,p}$ ικανοποιεί τις 6.8 και 6.9 και είναι επίσης μια προσέγγιση $1 - \epsilon$ της βέλτιστης λύσης. Συνεχίζουμε δείχνοντας ότι το γεγονός της παραβίασης της **??** δεν συμβαίνει με μεγάλη πιθανότητα. Ορίζουμε $X_{i,p} \in \{0, 1\}$ την τυχαία μεταβλητή που δείχνει αν το αίτημα r_i δρομολογείται στο μονοπάτι $p \in \mathcal{P}_i$, για όλα τα i. Συνεπώς,

$$Pr[X_{i,p} = 1] = f'_{i,p}.$$

Επιπλέον, για κάθε ακμή e, ορίζουμε $L_e = \sum_i \sum_{p \in \mathcal{P}_i | e \in p} X_{i,p}$ να είναι το φορτίο της ακμής, δηλαδή το ποσό των ταυτόχρονων αιτήσεων που μεταφέρει. Παρατηρήστε ότι $\mathbb{E}[L_e] \leq (1 - \epsilon) c(e)$. Λόγω του γεγονότος ότι $\sum_{p \in \mathcal{P}_i} X_{i,p} \leq 1$, δηλαδή η τυχαιοποιημένη στρογγυλοποίηση επιστρέφει το πολύ ένα μονοπάτι για κάθε αίτημα, $\{X_{i,p}\}_{p \in \mathcal{P}_i}$ είναι ένα σύνολο αρνητικά εξαρτημένων τυχαίων μεταβλητών για όλα τα i. Επιπλέον, οι $X_{i,p}$ και $X_{j,p}$ είναι ανεξάρτητες για κάθε $i \neq j$. Επομένως, μπορούμε να εφαρμόσουμε το ακόλουθο όριο Chernoff από το [19].

Theorem 1.5.1. *Έστω X το άθροισμα n ανεξάρτητων (ή καλύτερα) τυχαίων μεταβλητών με μέση τιμή $\mathbb{E}[X] \leq \mu$. Τότε ισχύει ότι*

$$Pr[X \geq (1+\delta)\mu] \leq e^{\frac{-\delta^2 \mu}{2+\delta}}$$

Εφαρμόζοντας το θεώρημα με $X = L_e, \delta = \frac{\epsilon}{1-\epsilon}$ και $\mu = (1-\epsilon)c_e$ έχουμε ότι:

$$Pr[L_e \geq c_e] \leq e^{-\frac{\frac{\epsilon^2}{(1-\epsilon)^2}(1-\epsilon)c_e}{2+\frac{\epsilon}{1-\epsilon}}}$$

$$= e^{-\frac{\epsilon^2 c_e}{2-\epsilon}}$$

$$\leq e^{-\frac{\epsilon^2 c_e}{2}}$$

Ορίζουμε $\epsilon = \sqrt{\frac{4 \log m}{c_e}}$. Εφόσον $\epsilon < 1$ από την υπόθεση, απαιτούμε ότι $c_e \geq 4 \log m$. Με αυτή την επιλογή του ϵ έχουμε: $\Pr[L_e \geq c_e] \leq \frac{1}{m^2}$. Εφαρμόζοντας ένα union bound στο σύνολο των ακμών E, λαμβάνουμε ότι η πιθανότητα να παραβιάζεται η χωρητικότητα οποιασδήποτε ακμής είναι το πολύ $\frac{1}{m}$.

Σε αυτή την ενότητα, εξετάσαμε το πρόβλημα βελτιστοποίησης πλήρους πληροφόρησης, χωρίς σύνδεση, του αρχικού μας διαδοχικού στοχαστικού πλαισίου. Δείξαμε ότι αν η χωρητικότητα κάθε ακμής είναι $\Omega(\log m)$ τότε το πρόβλημα επιδέχεται μια $(1+\epsilon)$-προσέγγιση, με τη βοήθεια ενός τυχαιοποιημένου αλγορίθμου στρογγυλοποίησης που στρογγυλοποιεί προσεκτικά τη βέλτιστη κλασματική λύση που παράγεται από το LP. Παρά το γεγονός ότι είναι NP-hard στη γενική περίπτωση, αν οι πόροι είναι αρκετά μεγάλοι μπορούμε να φτάσουμε όσο πιο κοντά θέλουμε στη βέλτιστη λύση. Το ερώτημα που τίθεται τώρα είναι: **Είναι δυνατόν να χρησιμοποιήσουμε αυτό το θετικό αποτέλεσμα για να σχεδιάσουμε έναν άμεσο μηχανισμό δημοσιεύμενων τιμών για το στοχαστικό μας περιβάλλον, ο οποίος επιτυγχάνει την ίδια εγγύηση ωφέλειας, είναι order-oblivious και τιμολογεί αντικείμενα;**

Chapter 2

Introduction

Can we predict the behavior and state of a system that consists of selfish and strategic agents that seek to derive their most preferred outcome? John F. Nash in his nobel-winning work proved, that no matter how sophisticated the system, agents interact in a way that eventually reaches a state where no one can obtain a better outcome by unilaterally changing his strategy. This concept is called a Nash Equilibrium and lies at the heart of the field of mathematics called Game Theory. The advent of the internet has brought forth a multitude of applications of Game Theory because it is a system in which participants interact strategically to steer it toward their desired outcomes. An important question then was, if one can efficiently compute the Nash Equilibrium (that is guaranteed to exist) of a system comprising strategic players. This question, primarily posed by computer scientists, gave rise to the field of Algorithmic Game Theory. This point of view, analyzes systems that exist "in nature" and seeks methods to efficiently predict their outcome after a period of time. However, what if we changed perspective and assumed the role of the designer of the system? Now the question becomes, can we, equipped with the power of customizing our own game (or system), enforce a desired agent behavior and system outcome? The quintessential example of an answer to this question is auctions. From the seller's perspective, is it possible to design, and efficiently compute, the rules of an auction in a way that bidders' most preferable action is to report their true valuations? This question gave birth to the field of Algorithmic Mechanism Design, which is often also called inverse Game Theory. In the context of auctions, monetary transfers from the bidders to the seller are used in order to implement truthfulness. In other words, the seller must decide the winners and the payments to induce truth-telling. In Mechanism Design with money, the seminal work of another Nobel laureate, R.Myerson [53], laid the grounds for the field by fully characterizing truthful auctions in the single-parameter setting, where bidders' valuations are described by a single number. Myerson proved, that if an allocation rule is monotone, then there is a unique payment rule, that when it is paired with that allocation renders the mechanism truthful. The implementation of Myerson's characterization to single-item auctions yields the famous Vickrey or second price auction [60]. Vickrey, Clarke and Groves [60, 20, 40] generalized Myerson's result to multi-parameter settings, where valuations are functions. The VCG mechanism, which is named after their initials, characterizes truthful mechanisms by providing payments analogous to the single-item case. The only caveat is, that for the VCG mechanism to work, it requires the optimal solution to the underlying optimization problem. This is a NP-hard task for most valuation classes. Let us note that all of the aforementioned mechanisms also maximize social welfare which is the sum of the winning bidder's valuations. In Chapter 3 we are going to formally define some key concepts in mechanisms design, which are useful throughout the thesis.

Imagine that you own a concert hall and you want to sell tickets (that correspond to seats inside the hall) for the upcoming concert. Our collective real-life experience dictates that the seller does not host an auction to sell tickets, simply because there is an easier and more intuitive way to sell: post prices and let the buyers pick seats and buy the tickets. Many other realistic instances like the above suggest the need to find simpler mechanisms to sell to strategic buyers. To this end, there is a huge line of research in a specific class of mechanisms called Posted-price. Along with their incentive guarantees, posted-price mechanisms provide polynomial computation of prices and simplicity. The downside is that they only approximate the optimal social welfare, however such mechanisms are very useful in practice and spark the debate of simple vs optimal mechanisms. In Chapter 4 we are going to delve into the details of such mechanisms and the results of this line of research.

The machinery used to analyze the performance of posted-price mechanisms is the prophet inequality framework, first introduced by mathematicians Krengel, Sucheston and Garling [46, 47] in the 70's. The setting involves a sequential decision maker who observes values in an online manner and needs to decide on an optimal stopping strategy that is approximately optimal compared to a prophet who can see the future and always take the optimal outcome. Prophet inequalities became relevant in the field of computer science when their connection to sequential posted-pricing was discovered. They provide the appropriate techniques to argue for the approximation ratio that prices achieve, under the assumption that buyer's valuations come from known independent distributions. In Chapter 5 we are going to explore key results in the Prophet Inequality literature, that inform the design of mechanisms with improved approximation guarantees in various multi-parameter settings.

A significant and highly intriguing body of work in Prophet Inequalities has enhanced the approximation factors for posted-price mechanisms across various contexts. A common characteristic of most of these contexts is that they exhibit complement $-$ free valuations, such as in Combinatorial Auctions with submodular buyers. Conversely, there are numerous practically relevant examples where buyers' values are complementary, meaning an item gains value only when purchased in combination with other items. In Chapter 6 we explore a setting that exhibits complementarity. Buyers have a single-parameter value for being admitted and routed in a network and claim their value only if they are assigned a path that connects their start to their destination node. In this context, any bundle of edges has a non-zero value only if it forms a valid path that serves a buyer. Utilizing the Prophet Inequality framework, we examine pricing schemes and their approximation guarantees in terms of social welfare.

Chapter 3

Basics of Mechanism Design

3.1 Preliminaries

Although the prophet inequality is a framework of independent interest, its connection with mechanism design dictates that we define some fundamental concepts of the field. We begin by providing a general definition about a mechanism design problem. Then we perform some standard simplifications in order to arrive to meaningful results. The following is a generic setup for a mechanism design problem.

Definition 3.1.1 (Generic setup)**.** Suppose we have n agents and an outcome space Ω. Each agent has a private valuation type $v_i : \Omega \to \mathbb{R}^+$ that represents the value that the agent derives from an outcome, and a set of actions $\mathcal{A}_i$. The mechanism collects a vector of actions: $\mathbf{a} = (a_1, a_2, \ldots, a_n), a_i \in \mathcal{A}_i$ and maps these actions to outcomes using a function $f : (\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n) \to \Omega$. The mechanisms also decides on a pricing rule $\mathbf{p} = (p_1, p_2, \ldots, p_n)$ where for each $i, p_i : (\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n) \to \mathbb{R}^+$. Each agent i has a utility u_i, which he wants to selfishly maximize, that quantifies his gain from the described process. It depends on the outcome, the pricing rule and his private valuation. The pair $(f, \mathbf{p})$ defines a mechanism.

In order to work with the above setting we need to perform some simplifications. First of all, we are going to define and assume for the rest of the thesis the quasi-linear utility of the bidders. Informally, a bidder's quasi-linear utility for an outcome is simply his value minus the price he pays for that outcome.

Definition 3.1.2 (Quasi-linear utility)**.** Let $(f, \mathbf{p})$ be a mechanism, $v_i : \Omega \to \mathbb{R}^+$ a private valuation function and $\mathbf{a}$ an action vector. The utility u_i of a bidder i is *quasi-linear*, if

$$u_i = v_i(f(\mathbf{a})) - p_i(\mathbf{a})$$

An intuitive example of a mechanism design problem is that of Combinatorial Auctions. In a Combinatorial Auction, the auctioneer possesses a set of items and agents have valuation functions for each one of the subsets of items they might get. The auctioneer (or equivalently the mechanism) should decide an allocation of items to agents and their payments. Below we define Combinatorial Auctions formally.

Definition 3.1.3 (Combinatorial Auction)**.** A Combinatorial Auction consists of a set of buyers, which we denote by N, a set of items, which we denote by M. We denote $|N| = n$ the cardinality of the buyers and $|M| = m$ the cardinality of the items. A vector

$\mathbf{v} = (v_1, v_2, \ldots, v_n)$ consists of n valuation functions, one for each buyer. Each valuation function $v_i : M \to \mathbb{R}^+$ maps every subset of items to a non-negative value, is normalized and non-decreasing, i.e for each buyer i $v_i(\emptyset) = 0$ and for all S, T, $S \subseteq T \subseteq M$ implies $v_i(S) \leq v_i(T)$.

The outcome of a mechanism is a feasible allocation of items to buyers. Below we define the feasible allocation of a Combinatorial Auction

Definition 3.1.4 (Feasible Allocation)**.** Let $(N, M, \mathbf{v})$ be a Combinatorial Auction. A Feasible Allocation $\mathbf{X} = (X_1, X_2 \ldots X_n)$ is a vector of n disjoint sets, $\forall i \neq j$, $X_i \cap X_j = \emptyset$, where each X_i denotes the (possibly empty) set of items each buyer gets from the mechanism.

In a Combinatorial Auction (and mechanism design, in general) buyers seek to maximize their utility. For a given price vector $\mathbf{p}$ the *demand correspondence* or simply *demand* of buyer is the set that contains every bundle that maximizes his utility. Below is a formal definition.

Definition 3.1.5 (Demand correspondence)**.** Given an item pricing vector $\mathbf{p} = (p_j)_{j \in M}$ and valuation function v_i of buyer i. The demand correspondence of buyer i is the set

$$D_i(\mathbf{p}) = arg \max_{S \subseteq M} \left\{ v_i(S) - \sum_{j \in S} p_j \right\}$$

Observe that by definition, the demand correspondence is allowed to contain more than one bundles. We will say that a buyer i is indifferent between bundles S and T under pricing $\mathbf{p}$ if $S \in D_i(\mathbf{p})$ and $T \in D_i(\mathbf{p})$.

A mathematical economist of the 19th century, Leon Walras considered an ideal equilibrium concept: select the right prices (and respective allocations) such that each buyer receives a bundle inside his demand correspondence. In this way, demand equals supply and the market clears. Observe that this equilibrium is Pareto Optimal, in that every buyer does not benefit by changing the bundle allocated to him as he cannot observe higher utility. This notion is called *Walrasian Equilibrium* (named after the person who invented it) and the prices of such an Equilibrium are often called *market clearing prices.* Below is a formal definition.

Definition 3.1.6 (Walrasian Equilibrium)**.** Let $(N, M, \mathbf{v})$ be a Combinatorial Auction, $p_1^* \ldots, p_m^*$ be item prices and $S_1^*, \ldots, S_n^*$ be bundles allocated to buyers in $[n]$, such that $S_i^* \cap S_j^* = \emptyset$ for all $i \neq j$. The tuple $(\mathbf{S}^*, \mathbf{p}^*)$ is a Walrasian Equilibrium iff for all buyers i, $S_i^* \in D_i(\mathbf{p})$ and for items that are not allocated, i.e $j \notin \bigcup_i S_i^*$, it holds that $p_j^* = 0$

A very well studied class of mechanisms is direct-revelation mechanisms. In these mechanisms agents' actions are simplified to be bids, which are are what they claim their private valuation function to be. Mechanisms under this assumption are called direct revelation because they ask from the bidders to directly reveal their types. A canonical example of a direct revelation mechanism is a sealed-bid auction where each bidder writes an offer that represents his valuation in a piece of paper. The mechanism collects every piece of paper and determines the winner and the price of the good to be sold.

Definition 3.1.7 (Direct Revelation Mechanism)**.** A mechanism $(f, \mathbf{p})$ is a *direct revelation mechanism*, if each bidder's action is to report $b_i : \Omega \to \mathbb{R}^+$.

There are mechanisms that are not direct-revelation, for example posted-price mechanisms do not collect bids. This class of mechanisms is going to be furthered explored in the following chapter.

3.2 Valuation Classes

Every mechanism should be able to compute a feasible allocation efficiently. However, performing this task when the valuations are general, can be computationally hard. Therefore, we define interesting valuation classes and reason about the extent they align with reality. For the following definitions M is the set of items.

Definition 3.2.1 (Additive Function)**.** A valuation function $v : 2^M \to \mathbb{R}^+$ is additive if for every $S \subseteq M$

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

That is, the value of any set is completely defined by the value of the singleton sets it contains. Nevertheless, this is not a realistic model of buyers' valuations. Consider the following example: Alice wants to buy 10 slices of cheese. She participates in a weird auction that sells cheese and she ends up going home with 50 slices. It cannot be the case that her value for 50 slices is 5 times her value for 10: she cannot consume all 50 before they go bad. Simply put, in the real world values do not scale linearly.

Definition 3.2.2 (Unit-Demand Function)**.** A valuation function $v : 2^M \to \mathbb{R}^+$ is unit-demand if for every $S \subseteq M$

$$v(S) = \max_{j \in S} v(\{j\})$$

This valuation class is the complete opposite of additive: it derives its name by the fact that for each subset, buyers only want the item for which they have maximum value. However, in the above scenario Alice is indifferent between 1 slice and 10 slices of cheese, which is also not very practical.

Definition 3.2.3 (Submodular Function)**.** A valuation function $v : 2^M \to \mathbb{R}^+$ is submodular if for every $S, T \subseteq M$, with $S \subseteq T$ and any item $j \notin S$

$$v(S \cup \{j\}) - v(S) \geq v(T \cup \{j\}) - v(T)$$

This class contains both additive and unit-demand valuations. It is the discrete analogue of concave functions and captures the notion of diminishing returns: the additional value for a new item decreases as sets grow larger. It is argued that this class is expressive enough to model real world. Consider the following example: Alice has 0\$, Bob has 1,000,000\$ - and they value money the same way. Now give 50\$ to each one of them. Which person observes a greater increase in their value?

Definition 3.2.4 (XOS function)**.** A valuation function $v : 2^M \to \mathbb{R}^+$ is XOS if there exists a collection of $a_1, a_2, \ldots, a_l$ additive valuation functions such that for every $S \subseteq M$

$$v(S) = \max_{1 \leq i \leq l} a_i(S)$$

It is hard to find intuition about this valuation class. However, this class is a superset of submodular valuations and it is often easier to work with

Definition 3.2.5 (Subadditive Function)**.** A valuation function $v : 2^M \to \mathbb{R}^+$ is subadditive if for every $S, T \subseteq M$

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

This class is the most general used in the bibliography. It contains XOS valuations and describes the notion of complement-free valuations. That is, bundling two item sets cannot increase the value of each individual set. Despite being the most general class there are practically revelant instances of complementarity. Consider the following: A buyer wants to buy a path connecting him from a start node s to a terminal node t. While he has no value for individual edges , when a bundle of edges forms a feasible path from s to t, his value suddenly becomes non-zero. That is, edges complement each other to give value. This is going to be the motivating example for the most part of the thesis.

The following definition of *single-minded* agents captures the above example.

Definition 3.2.6 (Single-Minded Valuations)**.** Let $v^* \in \mathbb{R}^+$ and S^* be a set of items. Then the valuation function $v : 2^M \to \mathbb{R}^+$ is single-minded if for every set S, if $M \supseteq S \supseteq S^*$,then $v(S) = v^*$ and $v(S) = 0$, otherwise. That is, a single-minded buyer incurs non-zero value if and only if he gets allocated a set of items that contains his desired bundle.

3.3 Oracles

By definition of the valuation functions for Combinatorial Auctions, we immediately observe that their description is exponential in the size of m.[1] Because of that, it is usually assumed access to oracles. We can view oracles as tools in the hands of the buyers who need to calculate the optimal decision for them, i.e which bundle of available items to choose. To this end, below we define some of the oracles utilized in the literature.

Definition 3.3.1 (Value Oracle)**.** A value oracle takes as input a set S and outputs a number: the valuation's function value on set S, i.e $v_i(S)$.

Definition 3.3.2 (Demand Oracle)**.** A demand oracle takes a input a price vector $\mathbf{p} = (p_1, p_2, \ldots, p_m)$ and outputs a demand bundle under these prices, i.e a set $S \subseteq M$ that maximizes $v_i(S) - \sum_{j \in S} p_j$, which is the buyer's utility under these prices.

The above definition is for item prices, however the definition can be extended to include bundle pricing, i.e when the price of a bundle cannot necessarily be expressed as a sum of the price of the items it contains.

Definition 3.3.3 (XOS Oracle)**.** For an XOS function v_i, the 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) $v_i(S) \geq A_i(S)$ for any $S \subset M$, and (ii) $v_i(T) = A_i(T)$

XOS oracles are not widely used in the literature, however we define them because some of the results presented later on require access to them. An interesting fact is that for submodular valuations an XOS oracle can be implemented via polynomially many queries to a value oracle.

[1] To be precise, there are valuation classes that are polynomial in their description (e.g for additive valuations we only need m numbers, one for each singleton set) but these classes are neither general nor expressive enough compared to subadditive or submodular, for instance.

3.4 Maximizing a global objective

From the auctioneer's perspective, the ultimate goal of an auction is to produce an allocation and payments that maximize a global objective, that is an objective that depends on the aggregate outcome of the mechanism. The two most commonly used objectives are: 1) *revenue*, which is the total payments made by the buyers to the mechanism and 2) *social* welfare the total value extracted by a specific outcome. Below we formally define the above objectives

Definition 3.4.1 (Social Welfare)**.** Let $(f, \mathbf{p})$ be a mechanism. The Social Welfare of an outcome $\omega \in \Omega$ is:

$$SW = \sum_i v_i(\omega)$$

Maximizing social welfare means that we search for outcomes that optimize the buyers' aggregate value. In a sense, in a welfare optimal outcome buyers are collectively happy with what they get.

By defining the above objectives we make a step towards meaningful problems and results: Consider, for instance, a Combinatorial Auction. Up to this point, we could set all payments to zero and produce a random feasible allocation. On the contrary, now we can formulate an optimization problem that requires that we produce a feasible allocation that maximizes social welfare, for example. The following linear program, which from now on will be referred as the *configuration LP* formalizes the underlying optimization problem.

$$\max \sum_{i=1}^{n} \sum_{S \subseteq M} v_i x_{i,S} \tag{3.1}$$

$$s.t \sum_{i=1}^{n} \sum_{S | j \in S} x_{i,S} \leq 1 \qquad \forall j \in M \tag{3.2}$$

$$\sum_{S \subseteq M} x_{i,S} \leq 1 \qquad \forall i \in [n] \tag{3.3}$$

$$x_{i,S} \in [0,1] \qquad \forall i \in [n], S \subseteq M \tag{3.4}$$

Observe that since the variables $x_{i,S} \in [0,1]$ the above is the fractional relaxation of the integer program. Equation (5.24) tells us that, in the optimal integer solution, for every i, j with $i \neq j$, $X_i \cap X_j = \emptyset$, where X_i is the (possibly empty) bundle allocated to buyer i in the optimal solution. Equation (5.25) is to ensure that every bundle is (fractionally) allocated at most once.

Welfare maximization for combinatorial auctions with subomodular valuation functions is NP-hard [50], while the optimal welfare cannot be approximated to a factor better than $1 - \frac{1}{e}$ unless $P = NP$ [44]. Moreover, the authors in [52] prove an information-theoretic lower bound of $1 - \frac{1}{e}$, regardless of whether $P = NP$, where a better approximation ratio would require exponentially many value queries.

Using only value oracles, Dobzinski and Schapira [27] show a $\frac{n}{2n-1}$-approximate algorithm and an $(1 - 1/e)$-approximate algorithm for the special case in which the agent's valuations are set coverage functions. Vondrák [61] proves an algorithm that is $\left(1 - \frac{1}{e}\right)$-approximate with high probability.

In the demand oracle model, Dobzinski and Schapira [27] also give a polytime $(1-1/e)$-approximation which was subsequently improved to $1-1/e+\epsilon$ by Feige and Vondrák [36]

For subadditive valuation functions [52] prove that a $\frac{1}{m^{1/2-\epsilon}}$-approximation would require exponentially many value queries. In the demand oracle model, Feige [35] presents a way of rounding any fractional solution to an LP relaxation to this problem to a feasible solution with welfare at least 1/2 the value of the fractional solution. This gives a 1/2-approximation for general subadditive agents, and (1-1/e)-approximation for the special case of XOS valuations.

For single-minded valuation functions it is NP-hard to maximize social welfare (reduction from set packing). It cannot be approximated within a constant factor and the best known algorithm approximates it within a factor of $O(\sqrt{m})$ [55].

The primary concern of this thesis is about mechanism that maximize social welfare. However, a practically relevant (and perhaps more realistic) objective of a mechanism is to maximize the seller's revenue. Revenue maximization is what the auctioneer strives for, especially when auctions are run online by large companies e.g E-bay, Amazon, Google. Observe that revenue does not depend on the outcome (or allocation) but rather on the payments. Below we formally define this objective.

Definition 3.4.2 (Seller's Revenue)**.** Let $(f, \mathbf{p})$ be a mechanism. The revenue of a mechanism $(f, \mathbf{a})$ for a given action vector $\mathbf{b}$ is:

$$Rev = \sum_i p_i(\mathbf{a})$$

Revenue and welfare maximization are two totally different objectives. The techniques employed to analyze the former are very different compared to the latter. Since it is not the main subject of the thesis we are going to briefly discuss about such mechanisms in the following chapters, but first we are missing a very important ingredient of mechanism design: incentives.

3.5 Incentive Compatibility and Individual Rationality

The field of Mechanism Design primarily focuses on incorporating incentives into Algorithm Design. In the previous section, we saw some results that maximize social welfare when the valuation functions of all agents are publicly known. That is, the aforementioned results form and solve an optimization offline problem to determine the winners of the auction. However, valuation functions are private information for the bidders. That is, in a direct revelation mechanism, bidders report their bids (which can be different from their true valuations) and the mechanism collects them, but it has no access to their true preferences. Therefore, the main challenge of mechanism design lies on how the designer, equipped with the power of determining the allocation and payments, incentivizes truth-telling.

Firstly, bidders should have an incentive to participate to the mechanism. To be more specific, every bidder should not ever observe negative utility regardless of the outcome of the mechanism, that is one should not lose money by simply participating. This concept is called Individual Rationality and it is defined formally below.

Definition 3.5.1 (Individual Rationality)**.** A mechanism $(f, \mathbf{p})$ is *individually rational* if for every bidder i, valuation v_i and bit vector $\mathbf{b}_{-i}$ of the other bidders

$$v_i(f(v_i, \mathbf{b}_{-i})) - p_i(f(v_i, \mathbf{b}_{-i}) \geq 0$$

Truth-telling can be defined in various ways that depend on the information an agent uses in making his decision. The concept of Dominant Strategy Incentive Compatibility is fundamental in the field of mechanism design. Informally, it implies that every agent cannot be worse of revealing his true type, regardless of what other agents bid.

Definition 3.5.2 (Dominant Strategy Incentive Compatibility (DSIC))**.** A mechanism is *dominant strategy incentive compatible* if truth-telling is in an agent's best interest even after the agent observes the types of others. Formally, for all i, v_i, b_i, and all bids $\mathbf{b}_{-i}$ of other bidders :

$$v_i(f(v_i, \mathbf{b}_{-i})) - p_i(v_i, \mathbf{b}_{-i}) \geq v_i(f(b_i, \mathbf{b}_{-i})) - p_i(b_i, \mathbf{b}_{-i})$$

We will often refer to mechanisms that are DSIC as *truthful* or *strategyproof.* Moreover, a mechanism that is DSIC is also implied to be Individual Rational, unless stated otherwise.

Mechanism design often needs to employ randomness, that is there are randomized mechanisms that flip coins to decide on the allocation and payments. Below, for completeness, we define truthfulness for such mechanisms.

Definition 3.5.3 (Truthful in Expectation)**.** A randomized mechanism is *truthful in expectation* if every bidder maximizes his *expected* utility by bidding truthfully.

Definition 3.5.4 (Universally truthful)**.** A randomized mechanism is *universally truthful* if every bidder maximizes his utility regardless of the instantiation of the mechanism's randomness. Randomized universally truthful mechanisms are probability distributions over deterministically truthful mechanisms.

3.6 Extension to the Bayesian Setting

In this section, we are going to assume that agent's types are independently drawn from publicly known distributions. Because of its statistical nature, this field is called *Bayesian Mechanism Design* and is formally defined as follows.

Definition 3.6.1 (Bayesian Mechanism Design)**.** A (direct-revelation) mechanism $(f, \mathbf{p})$ is *Bayesian* if agents' types $\mathbf{v} \sim \mathbf{F}$, where $\mathbf{F}$ is a (joint) product distribution, i.e $\mathbf{F} = F_1 \times F_2 \cdots \times F_n$, with $v_i \sim F_i$. We also assume quasi-linear utility as in the standard setting.

It is not totally unrealistic to assume that the mechanism has access to some distributional information about the agent's types. For example, large corporations that run auctions daily, are able to collect statistical data about the bidders who participate. Besides, revenue maximization depends critically on this distributional assumption. Now we need to redefine truthfulness to account for the stochasticity of valuations.

Definition 3.6.2 (Bayes-Nash Incentive Compatibility (BNIC))**.** A mechanism is *Bayes-Nash incentive compatible* if truth-telling is in an agent's best interest before observing the types of the others. Formally, for every bidder i, for every bid b_i the following holds:

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i}\sim\mathbf{F}_{-i}}[v_i(f(v_i,\mathbf{v}_{-i})-p_i(v_i,\mathbf{v}_{-i})] \geq \mathop{\mathbb{E}}_{\mathbf{v}_{-i}\sim\mathbf{F}_{-i}}[v_i(f(b_i,\mathbf{v}_{-i})-p_i(b_i,\mathbf{v}_{-i})]$$

where the expectation is taken over the types of all bidders except bidder i.

We observe that 3.5.2 defines a stronger notion of incentive compatibility than 3.6.2. The experienced reader might observe that these definitions refer to an equilibrium notion (i.e Bayes-Nash equilibrium and Dominant Strategy equilibrium, respectively). These two definitions differ in the amount of information they use. In the Bayesian setting, the designer designs the mechanism based on the information about agents' distributions. Then, agent's types are instantiated (a sample from buyers' product distribution is drawn). Finally, the mechanism is run on the sample from the distribution and outcome and payments are generated. When we refer to a quantity *ex ante*, we are referring to its expected value before the agents' types have been instantiated, where the expectation is taken over the distribution from which types are drawn. *Interim* refers to the time after the agents' types have been instantiated, but before the mechanism has been run; In particular, at this time, agents know their own types but not each others' instantiated types, and so the final outcome of the mechanism is as yet unknown to them. *Ex post* refers to the final realized value after the mechanism has been run. Similarly to incentive compatibility, individual rationality can be satisfied *ex post* or *interim*. From this point onward in this thesis, when referring to a mechanism as BNIC, it is implied to be individually rational, unless explicitly stated otherwise.

The global objectives defined in a previous section, can be redefined to be an expectation on bidders' types in order to hold in the Bayesian setting. That is, a mechanism designer seeks to maximize the *expected* social welfare (or revenue). Let us also remark that if a mechanism design problem is confined into a specific class of valuations (e.g submodular), then the distributions we consider are only on valuations from that class.

Chapter 4

Truthful Social Welfare Maximization

4.1 Mechanism Design meets Linear Programming

In this section we are going to revisit a concept defined in the previous chapter: the Walrasian Equilibrium. Setting Walrasian Equilibrium prices for the items and simply letting the bidders take what they demand, seems like a simple and truthful mechanism. Unfortunately, a Walrasian Equilibrium does not always exist, as it is shown in the following example.

Example 4.1.1 (Non-existence of Walrasian Equilibrium)**.** Consider two buyers, Alice and Bob and two items a, b. Alice has value of 2 for any non-empty set and Bob only has value of 3 only for the whole bundle $\{a, b\}$. We distinguish between two cases: If Bob's demand correspondence is not empty,then Alice's is. Therefore, the sum of prices at Equilibrium is at most 3. Then there is an item which has a price of at most 1.5. This implies that this item is also in Alice's demand correspondence which is a contradiction. If Bob's demand correspondence is empty, this means that the sums of prices is greater than 3. Hence, Alice's demand cannot contain $\{a, b\}$ since she has value of 2. Alice can only demand one item, say b for a price less than 2. Consequently, item a remains unsold, but it has a non-zero price which is again a contradiction.

Although the existence of market clearing prices is not always guaranteed, we present here two welfare theorems that provide necessary and sufficient conditions for the existence and optimality of such an equilibrium, thus fully characterizing the concept. We will prove the two theorems by resorting to Linear Programming Theory. Let us take the dual of the *configuration LP* presented in the second section. The constraints for each item j and each buyer i translate into dual variables p_j and u_i respectively. As we will prove, these variables can be seen as item prices and buyer utility. The dual program is the following:

$$\min \sum_{i=1}^{n} u_i + \sum_{j \in M} p_j \tag{4.1}$$

$$s.t\ u_i + \sum_{j \in S} p_j \geq v_i(S) \qquad \forall i, S \tag{4.2}$$

$$p_j \geq 0 \qquad \forall j \in M \tag{4.3}$$

$$u_i \geq 0 \qquad \forall i \in [n] \tag{4.4}$$

The first Welfare Theorem asserts that a Walrasian Equilibrium maximizes social welfare among all feasible fractional solutions. The formal statement follows.

Theorem 4.1.2 (First Welfare Theorem)**.** *Let $(N, M, \mathbf{v})$ be a Combinatorial Auction and $(\mathbf{S}^*, \mathbf{p}^*)$ be a Walrasian Equilibrium for this auction. Then for any feasible fractional solution of the configuration LP $\{x_{i,S}\}_{i,S}$ it holds that $\sum_i v_i(S_i^*) \geq \sum_{i,S} x_{i,S} v_i(S)$*

Proof. Fix buyer i. Since $(\mathbf{S}^*, \mathbf{p}^*)$ is a Walrasian Equilibrium it holds by definition that it maximizes his utility:

$$v_i(S_i^*) - \sum_{j \in S_i^*} p_j^* \geq v_i(S) - \sum_{j \in S} p_j^*$$

for every $S \subseteq M$. Multiplying both sides by $x_{i,S}$ and summing over all i, S we get:

$$\sum_{i,S} x_{i,S} \cdot v_i(S_i^*) - \sum_{i,S} \sum_{j \in S_i^*} x_{i,S} \cdot p_j^* \geq \sum_{i,S} x_{i,S} \cdot v_i(S) - \sum_{i,S} \sum_{j \in S} x_{i,S} \cdot p_j^*$$

Since $\sum_S x_{i,S} \leq 1$ for all i, we get:

$$\sum_{i,S} v_i(S_i^*) - \sum_{i} \sum_{j \in S_i^*} p_j^* \geq \sum_{i,S} x_{i,S} \cdot v_i(S) - \sum_{i,S} \sum_{j \in S} x_{i,S} \cdot p_j^*$$

Hence, it suffices to show that $\sum_i \sum_{j \in S_i^*} p_j^* \geq \sum_{i,S} \sum_{j \in S} x_{i,S} \cdot p_j^*$. At Walrasian Equilibrium, every item is included in at most one set, and for all items j that are not allocated it holds that $p_j^* = 0$, therefore the left hand side of the above inequality is equal to $\sum_{j \in M} p_j^*$. The right hand side can be rewritten as follows: $\sum_{j \in M} p_j^* \sum_i \sum_{S | j \in S} x_{i,S}$ which is at most $\sum_{j \in M} p_j^*$ by the first constraint of the primal LP. □

In the First Welfare Theorem we showed that Walrasian Equilibria, when they exist, provide an optimal integral solution to the configuration LP. In the Second Theorem we are going to show that the converse is also true: If there exists an optimal integral solution to the LP, then this solution is a Walrasian Equilibrium. That is, Walrasian Equilibria exist if and only if the optimal solution of the configuration LP is integral.

Theorem 4.1.3 (Second Welfare Theorem)**.** *Let $(N, M, \mathbf{v})$ be a Combinatorial Auction, then if there exists an optimal integral solution to the configuration LP, then a Walrasian Equilibrium also exists.*

Proof. Let $\mathbf{S}^* = (S_1^*, S_2^*, \ldots, S_n^*)$ be an optimal integral solution given by the primal LP and $\mathbf{p}^* = (p_1^*, p_2^*, \ldots, p_n^*)$, $\mathbf{u}^* = (u_1^*, u_2^*, \ldots u_n^*)$ be the respective solution to the dual LP. We will show that $(\mathbf{S}^*, \mathbf{p}^*)$ is a Walrasian Equilibrium. Fix buyer i. By the complementary slackness conditions, $x_{i,S_i^*} = 1$ implies that $u_i^* = v_i(S_i^*) - \sum_{j \in S_i^*} p_j^*$. Hence, $v_i(S_i^*) - \sum_{j\in S_i^*} p_j^* \geq v_i(T) - \sum_{j \in T} p_j^*$ for every $T \subseteq M$. Furthermore, if an item j is not allocated the first constraint of the primal LP holds strictly with inequality and by complementary slackness this implies that $p_j^* = 0$. By the above argument, we conclude that $(\mathbf{S}^*, \mathbf{p}^*)$ is a Walrasian Equilibrium. □

The two theorems imply that a Walrasian Equilibrium exists if and only if the integrality gap of the LP is zero. Furthermore, the second theorem provides us a way to find a Walrasian Equilibrium: Solve the fractional relaxation of the configuration LP. If the optimal solution happens to be integral, then solve the dual to find the item prices that impose the Equilibrium. Walrasian Equilibria are guaranteed to exist for a class of valuation functions called gross-substitutes (that contains unit-demand and additive) [55]

4.2 An introduction to Posted-Price Mechanisms

Imagine that you're in possession of an indivisible good, for instance a car or house, which you want to sell. The simplest and most natural mechanism is to decide on a price for the good and "post" it, that is to make a take-it-or-leave-it offer. Buyers arrive sequentially, they meet with you on a daily basis, having a valuation for the good which is of course not public knowledge. If the price is less than their valuation they buy, otherwise they do not. From the seller's point of view, the problem lies on how we decide (and efficiently compute) a price that (approximately) maximizes social welfare, which in this case is finding the buyer with the highest valuation.

The above is a simplified example if a posted-price mechanism. In the context of Combinatorial Auctions, posted-price mechanisms need to compute a price per item and extend these prices to bundles linearly. As it might be evident from the example, the challenge that the seller faces is to determine the right price: on the one hand, if the price is too high the good might be left unsold. On the other hand, if the price is too low, it might be the case that is is sold to a buyer with low valuation, simply because he arrived earlier than a higher-valued buyer.

In designing posted-price mechanisms, the seller needs to decide on the nature of the prices. Consider the example in the previous paragraph, does the price of the good change if it remains unsold after a period of time? Furthermore, do prices depend on the ordering of buyers, or do prices depend on the identity of buyers? As far as the last question is concerned, observe that the seller might benefit from increasing the price for certain types of buyers. For instance, a buyer who has an expensive car or a well-paid job is likely to have a high valuation. We proceed by giving a formal definition of posted-price mechanisms and the different types of prices.

Definition 4.2.1 (Posted-Price Mechanism)**.** Let $i = 1, \ldots, n$ be an arbitrary ordering of buyers, $\mathcal{F}$ an arbitrary downward-closed feasibility constraint, $\mathbf{x} = (x_1 \ldots, x_n)$ a (partial) allocation. For each buyer i we also define $\mathbf{x}_{[i-1]} = (x_1, \ldots, x_{i-1}, 0, \ldots, 0)$, that is the allocation confined to the first $i-1$ buyers, according to the ordering. A posted-price mechanism $\mathcal{M}(\mathbf{x}, \mathbf{p})$ prices allocations $\mathbf{x}$, and buyers arrive with respect to the ordering, having quasi-linear utility, and purchase their most preferred bundle. We define $p_i(x_i \mid \mathbf{y})$ as the price of outcome x_i offered to buyer i given partial allocation $\mathbf{y} \in \mathcal{F}$. We require also that $p_i(x_i \mid \mathbf{y}) = \infty$ for every outcome $(x_i, \mathbf{y}) \notin \mathcal{F}$.

Definition 4.2.2 (Anonymous vs Discriminatory Pricing)**.** Let $\mathcal{M}(\mathbf{x}, \mathbf{p})$ be a posted-price mechanism. The pricing scheme $\mathbf{p}$ is said to be *anonymous*, when the price of an allocation does not depend on the identity of the buyer. Formally, for each buyer i and each buyer j, $i \neq j$, $x_i = x_j$ implies $p_i(x_i) = p_j(x_j)$. Otherwise it said to be discriminatory.

For instance, item pricing for Combinatorial Auctions is an *anonymous* pricing scheme because the price of each bundle depends only on the items it contains. In the context of Combinatorial Auctions we also need to differentiate between item and bundle pricing.

Definition 4.2.3 (Item vs Bundle Pricing)**.** Consider a posted-price Combinatorial Auction on a universe of $|N| = n$ buyers, $|M| = m$ items and a pricing scheme $\mathbf{p} : 2^M \to \mathbb{R}_{\geq 0}$. The pricing scheme $\mathbf{p}$ is said to be an *item pricing scheme* if for every bundle $S \subseteq M$, $p(S) = \sum_{j \in S} p\{j\}$. That is, the price of every bundle is the sum of the prices of its singletons. Otherwise, the pricing scheme is said to be a *bundle pricing scheme*.

Definition 4.2.4 (Dynamic Pricing)**.** Let $\mathcal{M}(\mathbf{x}, \mathbf{p})$ be a posted-price mechanism. The pricing scheme $\mathbf{p}$ is said to be *dynamic*, when the price of an allocation to buyer i depends on the allocation to buyers $1, \ldots, i-1$, where buyers are indexed in the order they arrive. Formally, we denote the dynamic price of allocation x_i to buyer i by $p_i(x_i \mid \mathbf{x}_{[i-1]})$, where $\mathbf{x}_{[i-1]}$ denotes the allocation to previous buyers.

Definition 4.2.5 (Static Pricing)**.** Let $\mathcal{M}(\mathbf{x}, \mathbf{p})$ be a posted-price mechanism. The pricing scheme $\mathbf{p}$ is said to be *static* when the price of allocation x_i to buyer i does not depend on the allocation to previous buyers. Formally, price p_i is *static* if and only if $p_i\left(x_i \mid \mathbf{x}_{[i-1]}\right) = p_i(x_i)$.

4.3 Posted-Price mechanisms and incentive guarantees

The simplicity and practical relevance paired with their strong incentive guarantees renders posted-price a compelling class of mechanisms. In this section we define the notion of obviously strategy-proof mechanisms, first introduced by [51]. To this point, we have defined truthfulness with respect to direct revelation mechanisms, i.e mechanisms in which buyers report bids. In posted-price mechanisms, the seller does not rely on bids in fact, the only amount of information exchanged between the mechanism and the buyers is whether a buyer buys at the posted prices or not. Below is the formal definition of obviously strategy-proof mechanisms

Definition 4.3.1 (Obviously Strategyproof mechanism)**.** A mechanism is called obviously strategyproof if the optimality of truth-telling can be extracted without contingent reasoning.

The intuition behind this definition is the following: Suppose we want to explain to a non expert the optimality of his bidding his true valuation. This is harder when we design a second-price auction compared to a posted-price mechanism. Informally, contingent reasoning means to keep track of previous bids and results to make our current decisions. According to the previous observation, it is easy to see that posted-price mechanisms do not need contingent reasoning. All we need to decide is if we take or leave each item, depending on their price at the time of our arrival. Therefore, in contrary to a second-price auction, posted-price mechanisms are obviously strategyproof.

Truthfully maximizing social welfare without assuming any prior information about the probability distributions of bidder's types has also been extensively studied in the literature. For XOS valuations, the authors in [29] proved a $O(\log^2 m)$-approximation in 2006. Subsequent work in 2007 [24] improved the previous approximation guarantee to $O(\log m \log\log m)$. Both mechanisms are randomized, thus they achieve this approximation ratio in expectation. In 2012 Krysta and Vöcking [48] showed an $O(\log m)$ -approximate randomized mechanism which was further improved to $O(\sqrt{\log m})$ by Dobzinski in 2016 [26]. More recent work in [4] exponentially refined Dobzinski's result by proving an $O((\log\log m)^3)$-approximate randomized mechanism. All of the above mechanisms assume access to both value and demand queries and are also universally truthful. In 2021, the authors in [5] proved an $O((\log\log m))^3$-approximate randomized truthful mechanism for subadditive valuations which broke the logarithmic barrier problem which was from 2007 when Dobzinski [24] proved an $O(\log m \log\log m)$ approximation. Let us note that this work also improves the approximation factor for XOS valuations, from $O((\log\log m))^3$ to $O((\log\log m))^2$. Finally, [5] is also the current state-of-the-art work in approximately-optimal, truthful, randomized mechanisms for subadditive and XOS Combinatorial Auctions. From the above results we can observe the expressive advantage of demand over value oracles. To be more specific, the authors in [28] provide an algorithm that achieves an approximation bound of $O(\sqrt{m})$ for the problem of truthfully maximizing social welfare using only value oracles, which subsequently was proved tight for value queries even if randomized mechanisms are used [25].

4.3.1 A brief overview of an $O(logm)$-approximate mechanism for XOS Combinatorial Auctions

The purpose of this section is to provide a brief presentation of the result of Krysta and Vöcking [47]. The authors provide a simple algorithm that is $O(logm)$-approximate to the optimal social welfare of a Combinatorial Auction with XOS valuations. We choose this particular algorithm, of the list stated above, due to its simplicity and elegance.

A high level idea of the algorithm of Krysta and Vöcking is that it is a procedure that "learns" the correct prices. Suppose we have multiple copies of each item: every time an item gets sold we double its price. This guarantees that at some point we sell the item for the last time, the price doubles and no buyer wants to buy. Therefore, the last time we sold it, it got sold at a correct price. However, there is an obvious problem here: we oversell the item which in the original setting exists only in one copy. The authors solve this problem by employing randomization: every time a buyer wants an item, flip a coin with success probability q. If the coin flip is a success, sell the item (thus making it unavailable for future buyers) otherwise don't. This modified procedure, ensures that the allocation returned by the algorithm respects the supply constraint. By fine-tuning q, for XOS valuations they get the desired approximation ratio. The following is the overselling algorithm that produces an infeasible solution.

The following is the overselling algorithm with the oblivious randomized rounding we describe to produce a feasible allocation.

Setting $q^{-1} = O(\log m)$ we get the desired approximation ratio for XOS valuations.

Algorithm 2 Overselling MPU algorithm

For each good $e \in U$ do $p_e^1 = p_0$
For each bidder $i = 1, 2, \ldots, n$ do
    Set $S_i = D_i(U_i, p^i)$, for a suitable $U_i \subseteq U$
    Update for each good $e \in S_i$: $p_e^{i+1} = 2 \cdot p_e^i$

Algorithm 3 MPU algorithm with oblivious randomized rounding

For each good $e \in U$ do $p_e^1 = p_0$.
For each bidder $i = 1, 2, \ldots, n$ do
    Set $S_i = D_i(U_i, p^i)$, for $U_i = \{e \in U \mid b_e^i > 0\}$.
    Update for each good $e \in S_i$: $p_e^{i+1} = 2 \cdot p_e^i$.
    With probability q set $R_i = S_i$ else $R_i = \emptyset$.
    For each good $e \in R_i$: reduce its multiplicity to zero.

4.4 Truthful Revenue Maximization

Although the thesis is on social welfare maximization, in this section we briefly discuss some results in the area of truthful revenue maximization. Remember that in the previous chapter we commented on revenue and social welfare being two very different objectives as far as mechanisms and techniques are concerned. However, due to the fact that we do not live in an idealistic world, it is of great practical relevance to assume that sellers (which often are large corporations) seek to maximize their earnings. Below we state an example that showcases the difference between welfare and revenue.

Example 4.4.1 (Social Welfare vs Revenue Maximization)**.** Consider a trivial example of an auction: one buyer, one item. We need to decide on a price p for the item. If we seek to maximize social welfare, the answer is very simple: set $p = 0$ and give the item for free. However, how does one maximizes revenue in this setting? Since we claim revenue p, if and only if there is a successful purchase we need to set a non-zero value for p. However, any value could be bad because we do not know anything about the type of the buyer. To be more specific, if the buyer's value is greater than p we have good revenue, but if it is not, we have revenue of 0.

The above example is to show a crucial difference between welfare and revenue. In the former, we seek the maximum valuation, in the latter we also care about the magnitude of the maximum valuation. The above example is to justify that, in order to produce non-trivial results, revenue maximization must be studied under the Bayesian setting, defined in the previous chapter. The pioneering work of Myerson [53] also introduced Bayesian Mechanism Design. Myerson proved in the single-parameter setting a reduction from (truthful) revenue maximization to welfare maximization via the use of virtual valuations. To be more specific, he managed to show that revenue maximization is virtual welfare maximization. However, Myerson's result does not hold in multi-parameter settings. One reason for that is that in the single-parameter settings the optimal mechanism is deterministic, whilst this is not true for multi-parameter settings. Firstly, optimal multi-dimensional mechanisms are not always deterministic, and do not permit succinct representations. Secondly, a deterministic mechanism for a single agent prices each outcome at a certain price, and lets the buyer choose which the utility maximizing outcome. A randomized mechanism can set a random price for each possible outcome. But, interestingly, in addition it can also price random allocations or lotteries. A lottery is a distribution over outcomes. Selling lotteries can increase the seller's expected revenue [56]. In a series of papers by Cai, Daskalakis and Weinberg[11, 12, 13] give a reduction from revenue maximization to social welfare maximization that is black-box in the sense that the reduction is "generic" and does not need to understand the inner functioning of the algorithm for social welfare. The approach also has similarities with Myerson's approach for single-parameter revenue maximization in that it queries the social welfare algorithm at value vectors that are not the agents' real values, but are functions of the real values; these "fake" values can be thought of as virtual values. We comment that in revenue maximization settings, truthfulness is with respect to Bayes-Nash equilibria or in other words, they are (approximately) BNIC, as defined in the previous chapter.

Chapter 5

The Prophet Inequality Framework

5.1 Introduction

The prophet inequality is a problem from optimal stopping theory that was discovered in the 70's by Krengel, Sucheston and Garling. [46, 47]. It involves a *gambler* and a *prophet*. On the one hand, the gambler is faced with a sequence of n random variables, X_i. These random variables are samples drawn from independent (but not necessarily identical) known distributions D_i. Upon arrival of each random variable, the gambler must decide between the two following alternatives: either claim the value of the sampled random variable and stop (without observing future random variables) or discard the random variable and forfeit its value forever. On the other hand, the prophet is omniscient (i.e can see the future) and always claims the random variable with the highest value. The prophet inequality is the problem of finding an *optimal stopping* strategy for the gambler that collects value that is comparable to the optimum in hindsight (i.e to what the prophet gets). Krengel and Sucheston showed a strategy that guarantees that the gambler's reward is at least 50 % of the prophet's reward. The optimal strategy that achieves the previous approximation is given by backward induction. If the gambler reaches random variable X_n it is optimal that he accepts it, as there are no other variables to come. Now the inductive definition of the strategy is as follows: gambler accepts X_i if and only if its value is greater than the expected value he collects starting from X_{i+1} to X_n. In fact, with a simple example we can show that the factor of 2 is tight for the problem. Consider the following two random variables which arrive in the order they are indexed: X_1 is deterministically 1 and X_2 is $\frac{1}{\epsilon}$ with probability ϵ and 0 otherwise. Any stopping strategy yields to the gambler expected winnings of value 1. On the other hand, the prophet, who always chooses the maximum realized value, gets value of $2 - \epsilon$ in expectation.

In the 80's, Cahn [59] studied the performance of single-threshold based strategies, i.e strategies that use a single threshold to decide wheter to accept or reject a random variable. He proved that setting a threshold that is the median of the distribution of $\max_i X_i$ and picking the first random variable that its value is above this threshold also achieves a 2-approximation. However, the advantage of this strategy over that of Krengel and Sucheston, is that the approximation remains invariant even if the order of the random variables is arbitrary. Therefore, a standard assumption in the prophet inequality literature is that the order is chosen by an adversary, who can be adaptive, that is, his choice of the next random variable in the sequence can depend on the values of the previous random variables and the decisions of the gambler.

Several decades after the results of Krengel, Sucheston and Cahn, the authors in [41] established a connection between the prophet inequality and posted price mechanisms. They observed that threshold-based algorithms for the prophet inequality can be viewed as prices for a (sequential) posted price mechanisms. Take, for example, the following scenario. There is a seller who is in possession of an indivisible item. Buyers arrive one-by-one in an arbitrary order, having independently (and not necessarily identically) distributed values for the item. The distributions are known to the seller, who calculates and posts a take-it-or-leave-it price. Upon arrival of each buyer, the item is sold to them if and only if their value is above the posted price (and it was not sold to some previous buyer). The mechanism's objective is to maximize the total value extracted, in comparison to the best hindsight optimum, that is the maximum realized value of any buyer. We observe that the mechanism described above does not ask from the buyers to reveal their types, since the only information that it uses to produce an allocation is whether the value is above or below the chosen threshold. We can deduce that a simple-posted price mechanism gives a 2-approximation to the optimal welfare. Moreover, this result is tight, in the sense that no other mechanism can improve this guarantee, which can be easily seen by the tightness proof of the prophet inequality: consider a buyer who has a deterministic value of 1 for the item and comes first. Any mechanism should decide whether or not to give the item to the first buyer, yielding the lower bound.

Because of their connection to posted price mechanisms, prophet inequalities became a relevant and valid research path in the field of Algorithmic Mechanism Design. In 2012, Kleinberg and Weinberg [45] gave a new threshold for the original, single-item setting, that achieves the same ratio. Setting the threshold equal to $\frac{1}{2}\mathbb{E}[\max_i X_i]$ is also a 2-approximation.

In the following years, subsequent work generalized the above results to matroid, polymatroid and matroid intersection [45, 30], knapsack [32], k matroid [39] and downward closed with bounded maximum set size [57] feasibility constraints.

For Combinatorial Auctions with XOS valuations (which also include submodular) the authors in [38] provide a $\frac{2e}{e-1}$-competitive prophet inequality that assumes black box access to an algorithm by Vondrák [61]-which optimally solves the offline problem- and an oracle that answers XOS queries. Subsequent work in [32] improves the previous result by providing a 2-competitive prophet inequality for XOS valuations, assuming access to demand oracles. This matches the lower bound for this class of valuations inherited by the case of a single item. The authors in [32] also provide a unifying approach to proving prophet inequalities by introducing the notion of balanced prices. Their argument comprises a reduction from the Bayesian to the full information setting and a proof that when "good" prices for the latter exists, appropriately scaled versions of these are also "good" for the former. Finally, the previous line of work also implies an $O(\log m)$ -competitive prophet inequality for subadditive valuations by approximating subadditive with XOS valuations [9].

For subadditive valuations, [31] achieves an exponential improvement by proving an $o(\log \log m)$ prophet inequality. Their result is also computational: they consider the dual program of the configuration LP and show how to compute prices efficiently by running the Ellipsoid Algorithm with a separation oracle that can be implemented using demand queries. In fact, in this context, the demand oracle and the separation oracle coincide. Observe that the dual program (in the case of Combinatorial Auctions) has exponentially many constraints - one for each possible subset of items and buyer. Given a sequence of item prices $p_1, p_2, \ldots, p_m$, the separation oracle utilized by the Ellipsoid

Algorithm, needs to compute the set S_i that maximizes buyer's i utility under these prices. This is precisely the definition of the demand oracle to which the authors assume access for their algorithm. Recent work in [21] proves the existence of a 6-approximate prophet inequality for subadditive valuations, thus resolving a central open problem in the area. Subsequent work in [7] shows how to make the previous result computational, by proving that any prophet inequality can be implemented as a posted price mechanism with at least as a good welfare guarantee.

5.2 A proof of the singe-item case

In this section we provide a formal proof of the pricing scheme used in [45] for the single-item case. But first, we need to formally define an α-competitive prophet inequality.

Definition 5.2.1 (The Prophet Inequality)**.** Let $\mathbf{v} = (v_1, \dots, v_n)$ be a sequence of n valuation functions drawn independently from publicly known distributions $\mathcal{D}_i \in \Delta$, where Δ is the set of all distributions over $\mathbb{R}_{\geq 0}$. We also denote $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2, \dots, \times \mathcal{D}_{n-1} \times \mathcal{D}_n$ to be the product distribution. Since we have independently drawn samples from each individual distribution we can equivalently claim that the vector $\mathbf{v}$ is drawn from the product distribution $\mathcal{D}$. Consider an online algorithm (the *gambler*) over the input $\mathbf{v}$. We denote by $\mathbf{v}(ALG(\mathbf{v}))$ the social welfare that the algorithm achieves when presented with $v_1, v_2, \dots, v_n$ in the order they are indexed. Accordingly, we define $\mathbf{v}(OPT(\mathbf{v}))$ to be the optimal social welfare (the offline optimal on input $\mathbf{v}$). We say that ALG is α-competitive, $\alpha \geq 1$ when:

$$\sup_{\mathcal{D} \in \Delta} \frac{\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(OPT(\mathbf{v}))]}{\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\mathbf{v}(ALG(\mathbf{v}))]} \leq \alpha.$$

The online algorithm used in the definition is not necessarily a single-threshold strategy but it can be any rule that decides without having information about future events. However, we often analyze the performance of posted-price mechanisms for ALG, due to their strong incentive compatibility properties.

We will often refer to α as the (stochastic) competitive ratio of the algorithm. That is not to be confused with the competitive ratio defined in classic online algorithms. There are two subtleties: 1) the input in the above definition is assumed to come from known distributions, whereas competitive analysis assumes worst-case inputs, 2) the ratio is achieved only in expectation, that is, there might be sequences that produce a ratio greater than α. The prophet inequality, however, analyzes average perfmormace. On the other hand, competitive analysis gives guarantees that hold under any worst-case sequence. Let us also remark, that if the expected performance of ALG is zero, we will refer to the competitive ratio being *unbounded* or *infinite*.

Theorem 5.2.2 (Single-Item Prophet Inequality [45])**.** *Let $V^* = \max_i v_i$ be the maximum of the n realized values. Let $x^+ = max\{x, 0\}$ be the positive part of variable x. Then, the policy that sets price $p = \frac{1}{2}\mathbb{E}[V^*]$ and accepts first buyer with value above p is 2-competitive. Furthermore, this is independent of the decision to sell or not, when a value is equal to $\frac{1}{2}\mathbb{E}[V^*]$.*

Proof. We denote $\mathbb{E}[SW]$ as the expected social welfare of the policy. We break the expected social welfare into two parts: the expected revenue of the mechanism, i.e the price it charges when it allocates the item, and the expected surplus of the buyers, that quantifies how much additional valued we gain by allocating to a buyer i who has $(v_i - p)^+$ units of value above the posted price. It is easy to see that, by quasi-linear utility: $\mathbb{E}[SW] = \mathbb{E}[Rev] + \mathbb{E}[Surplus]$. We proceed by bounding the two quantities independently.

Revenue: The expected revenue of the mechanism is the price for the item multiplied by the probability it gets sold.

$$\mathbb{E}[Rev] = p \cdot \mathbb{P}[\textit{item is sold}] \tag{5.1}$$
$$= \frac{1}{2}\mathbb{E}[V^*] \cdot \mathbb{P}[\textit{item is sold}] \tag{5.2}$$

Surplus: The expected surplus of the mechanism is the sum of buyer's utilities, conditioned on the event that the item is still available when buyer i arrives.

$$\mathbb{E}[Surplus] = \sum_i \mathbb{E}[u_i] \geq \sum_i \mathbb{E}[(v_i - p)^+ \mathbb{1}\{i\, sees\, item\}] \tag{5.3}$$
$$= \sum_i \mathbb{E}[(v_i - p)^+] \cdot \mathbb{P}[i\, sees\, item] \tag{5.4}$$
$$\geq \sum_i \mathbb{E}[(v_i - p)^+] \cdot \mathbb{P}[item\, is\, not\, sold] \tag{5.5}$$
$$\geq \mathbb{E}[\max_i (v_i - p)] \cdot \mathbb{P}[item\, is\, not\, sold] \tag{5.6}$$
$$\geq \frac{1}{2}\mathbb{E}[V^*] \cdot \mathbb{P}[item\, is\, not\, sold] \tag{5.7}$$

Where (5.4) is because the value of buyer i is, by assumption, independent of the values of previous buyers and therefore independent from the event of the item being unsold when he arrives. (5.5) reduces the probability to the item remaining unsold at the end of the process, (5.6) reduces the summation to a maximum after taking linearity of expectation and (5.7) is by definition of p and V^*.

Summing (5.2) and (5.7) yields the theorem. Observe that since the probability terms in the above analysis cancel each other out, this pricing is indeed robust to any tie-breaking decision. This decision influences the probability that the item is sold or unsold. However, the competitive ratio is independent of this decision. □

The above proof implicitly assumes that, by having full knowledge of the underlying distributions of buyer's valuations, we can efficiently compute the mean value. Here we need to compute the expectation of the maximum value $\mathbb{E}[V^*]$. However, the realistic assumption to make is that our online algorithm (or gambler's strategy) only has sample access to each of the n independent distributions. What we want is that if we post a price $\hat{p}$ such that $|\hat{p} - p| < \epsilon$, for all $\epsilon \in (0, 1)$, we will incur an additive loss in the competitive ratio. Then $\hat{p}$ will be an estimation of p (i.e the empirical average) which by standard concentration bounds can be found by using $poly(n, 1/\epsilon)$ number of samples. Observe that in our analysis, if we substitute p with $p - \epsilon$ in the revenue part and p with $p + \epsilon$ in the utility part, we will get that the algorithm collects value at least $\frac{1}{2}\mathbb{E}[V^*] - \epsilon$.

5.3 Extension to k-uniform matroids

Our first extension is to consider the k-uniform matroid feasibility constraint. A k-uniform matroid is a matroid where any subset of at most k items is an independent set. To put it simply, the gambler can choose up to k out of n boxes, he claims a prize that is the sum of the prizes of the collected boxes and he compares with an omniscient prophet that always chooses the k-larger prizes. The mechanism design analogue here is the sequential k-unit auction where n buyers arrive sequentially in an arbitrary order and each one wants to buy a unit. A simple generalization of the proof in [45] gives a 2-approximation policy for the gambler. However, better approximation guarantees should exist as k increases. Intuitively, as k increases, the gambler has more choices of boxes to choose and consequently he has more "room for error", that is selecting sub-optimal boxes hurts increasingly less the performance of the gambler compared to that of the prophet. Hajiaghayi et.al [41] formalize that idea by proving an $(1 + O(\sqrt{\frac{logk}{k}}))$-competitive prophet inequality for k-uniform matroids, which is asymptotic. This line of work in k-uniform matroids prophet inequalities will be useful for the main part of the thesis, therefore we present the proofs of the above claims.

First we are going to extend the proof of Kleinberg and Weinberg [45] to derive the classic 2-approximation result.

Let $V^* = \max_{S:|S| \leq k} \sum_{i \in S} v_i$ the sum of the top-k of n realized values. It follows that the prophet collects expected value equal to $\mathbb{E}[V^*]$, where the expectation is taken over the randomness of the valuations.

Theorem 5.3.1 (k-Prophet Inequality)**.** *The policy that sets price $p = \frac{1}{2k}\mathbb{E}[V^*]$ and accepts buyers*

with value above this threshold, while supplies last, yields social welfare that is at least $\frac{1}{2}\mathbb{E}[V^*]$, which is half the optimum.

Furthermore, this is independent of the decision to sell or not, when a value is equal to $\frac{1}{2}\mathbb{E}[V^*]$.

Proof. The analysis is almost identical to the case of a single item. Once again we are going to invoke the "revenue-surplus" argument to account for the total value collected by the gambler in expectation. The only difference here is the slightly richer feasibility constraint. Let $A_i = \{j \mid j < i \text{ and } v_j > p\}$ be the set of buyers accepted upon arrival of buyer $i \in [n]$. We also define A_{n+1} to be the total buyers accepted at the end of the process. Now we are ready to proceed to bound the revenue.

Revenue: The expected revenue of the mechanism is the price for the items multiplied by the probability that all of them get sold.

$$\mathbb{E}[Rev] = kp \cdot \mathbb{P}[\mid A_{n+1} \mid = k] \tag{5.8}$$

$$= \frac{1}{2}\mathbb{E}[V^*] \cdot \mathbb{P}[\mid A_{n+1} \mid = k] \tag{5.9}$$

Surplus: The expected surplus of the mechanism is the sum of buyer's utilities, conditioned on the event that supplies last when buyer i arrives.

$$\begin{aligned}
\mathbb{E}[Surplus] = \sum_i \mathbb{E}[u_i] &\geq \sum_i \mathbb{E}[(v_i - p)^+ \mathbb{1}\{|A_i| < k\}] && (5.10)\\
&= \sum_i \mathbb{E}[(v_i - p)^+] \cdot \mathbb{P}[|A_i| < k] && (5.11)\\
&\geq \sum_i \mathbb{E}[(v_i - p)^+] \cdot \mathbb{P}[|A_{n+1}| < k] && (5.12)\\
&\geq \mathbb{E}[\max_{S:|S|\leq k} \sum_{i\in S}(v_i - p)] \cdot \mathbb{P}[|A_{n+1}| < k] && (5.13)\\
&\geq \frac{1}{2}\mathbb{E}[V^*] \cdot \mathbb{P}[|A_{n+1}| < k] && (5.14)
\end{aligned}$$

Where (5.11) is because the value of buyer i is, by assumption, independent of the values of previous buyers and therefore independent from the event that k items were sold before he arrived. (5.12) reduces the probability to some of the items remaining unsold at the end of the process, (5.13) reduces the summation to the maximum of the top-k values after taking linearity of expectation and (5.14) is by definition of p and V^*.

The events of accepting at most $k-1$ buyers and exactly k buyers are complementary, since no algorithm can allocate more than k units and each buyer receives at most one unit. Therefore, summing (5.9) and (5.14) cancels out the probability terms and yields the theorem □

Now we are going to present a theorem that asymptotically improves the previous result, thus validating our intuition described in the introduction.

First, we are going to need the following Chernoff bound from [19]:

Theorem 5.3.2. *Let $X_1, \dots, X_n$ be independent $\{0,1\}$ random variables. Let $X = X_1 + X_2 + \dots + X_{n-1} + X_n$ and $\mu = \mathbb{E}[X]$. Then for any $0 < \epsilon < 1$:*

$$Pr[|X - \mu| > \epsilon\mu] \leq 2e^{-\epsilon^2\mu/3}$$

Now we are ready to state and prove our main theorem.

Theorem 5.3.3. *For the k-prophet inequality problem there is a price p such that accepting the first k values above p gives with probability $1 - \frac{2}{k}$ value V that satisfies the following:*

$$V \geq (1 - \sqrt{\frac{12 \ln k}{k}}) \cdot V^*$$

Proof. Fix a price p. Define the set $S_p = \{i \in [n] \mid v_i \geq p\}$. The idea of the proof is to reduce the supplies by a fixed amount δ and find a price p such that $\mathbb{E}[|S_p|] = k - \delta$. Then by applying a Chernoff bound we argue that the buyers that belong to the set S_p are concentrated around the mean value with high probability. Define the random variable $X_i = \mathbb{1}\{i \in S_p\}$, that is a 0-1 random variable that is equal to 1 when buyer i is in S_p. The defined $\{X_i\}_{i\in[n]}$ are independent 0-1 random variables. Set $X = |S_p| = X_1 + X_2 + \dots X_n$ and $\mu = \mathbb{E}[X] = k - \delta$, we have the following:

$$\mathbb{P}[|X - \mu| > \epsilon\mu] \leq 2e^{-\epsilon^2\mu/3}$$

By setting $\epsilon = \frac{\delta}{k-\delta}$ we have that $X \in [k-2\delta, k]$ w.p $1-2e^{-\frac{\delta^2}{3(k-\delta)}} \geq 1-2e^{-\frac{\delta^2}{3k}}$. Therefore, we set $\delta = \sqrt{3k \ln k}$ and the previous probability becomes at least $1-\frac{2}{k}$.

Now we know $k - 2\delta \leq |S_p| \leq k$ with high probability.

For the analysis below we are going to adopt the "revenue-surplus" approach.

Revenue: Since $S_p \geq k-2\delta$ we have that $Rev \geq (k-2\delta)p$.

Surplus: Since $|S_p| \leq k$ we accept every buyer above p. Also since $|S_p| \geq k-2\delta$ we accept at least $k-2\delta$ buyers. These buyers get also accepted by the optimal algorithm of the prophet that we denote by V^*. This is easy to see by an exchange argument since the prophet uses a greedy algorithm, he simply selects the top-k realized values. That is, the algorithm collects value which is at least a $\frac{k-2\delta}{k}$ fraction of the optimal. Hence for our price choice we achieve surplus:

$$Surplus = \sum_{i \in S_p} (v_i - p) \geq (\frac{k-2\delta}{k}) \cdot V^* - (k-2\delta)p = (1 - \frac{2\delta}{k})V^* - (k-2\delta)p$$

Adding revenue and surplus gives us value

$$V \geq (1 - \frac{2\delta}{k})V^*$$

For $\delta = \sqrt{3k \ln k}$ we get:

$$V \geq (1 - \sqrt{\frac{12lnk}{k}})V^*$$

□

Hajiaghayi et.al [41] also provided a lower bound for this setting. Subsequent work by Alaei [1], showed an $(1 - \frac{1}{\sqrt{k+3}})$-approximate prophet inequality for the setting which asymptotically matches the lower bound given in [41]. Observe that for $k=1$ Alaei's result is optimal. His algorithm is adaptive, i.e he selects a price for buyer i that depends on the realizations of the values of the previous buyers. Let us remark that Jiang, Ma, Zhang [54], solved the k-uniform case completely by proving tight guarantees for all k, instead of asymptotically optimal ones.

5.4 Introducing Balanced Prices: a unified approach

So far we have seen a standard pattern in our proofs of prophet inequalities: we break the value collected into two parts, that is revenue and surplus and we lower bound each one of these quantities separately. The probability terms cancel out and we end up with a lower bound for the initial aggregate value. However, it is not clear how this technique generalizes to richer feasibility constraints and more expressive valuation classes (eg. matroid constraints and submodular functions respectively). In 2017 Dütting et.al [32] provided a unifying approach to proving prophet inequalities, that generalizes the "surplus-revenue" argument to more general valuation classes and combinatorial feasibility constraints. In this section we are going to present some key parts of their work and the results when their technique is applied to specific problems.The following definition of *balanced prices* is key to their contribution.

Let $\mathcal{F}$ be an arbitrary downward-closed feasibility constraint. We also define $OPT(\mathbf{v} \mid \mathbf{x})$ to be the optimal residual allocation, i.e the allocation that maximizes $\sum_i v_i(x'_i)$ over $\mathbf{x}' \in \mathcal{F}$ with $\mathbf{x}, \mathbf{x}'$ disjoint and $\mathbf{x} \cup \mathbf{x}' \in \mathcal{F}$

Definition 5.4.1 (((α, β)-balanced prices)**.** Let $\alpha, \beta > 0$. A pricing rule $\mathbf{p}^{\mathbf{v}} = (p_1^{\mathbf{v}}, \ldots, p_n^{\mathbf{v}})$, where $p_i^{\mathbf{v}} : 2^M \to \mathbb{R}_{\geq 0}$, is (α, β)-balanced with respect to a valuation profile $\mathbf{v}$ if for all $\mathbf{x} \in \mathcal{F}$ and all $\mathbf{x}' \in \mathcal{F}$ with $\mathbf{x}, \mathbf{x}'$ disjoint and $\mathbf{x} \cup \mathbf{x}' \in \mathcal{F}$ the following hold:

- $\sum_i p_i^{\mathbf{v}}(x_i) \geq \frac{1}{\alpha}\left(\mathbf{v}(OPT(\mathbf{v})) - \mathbf{v}(OPT\left(\mathbf{v} \mid \mathbf{x}\right)\right)$
- $\sum_i p_i^{\mathbf{v}}(x_i') \leq \beta \mathbf{v}\left(OPT\left(\mathbf{v} \mid \mathbf{x}\right)\right)$

The intuitive meaning of the optimal residual allocation denoted by $OPT\left(\mathbf{v} \mid \mathbf{x}\right)$ is the following: Consider an auction where buyers arrive having a value vector $\mathbf{v}$, but for some reason a subset $S \subseteq M$ of the items is not available to them. The quantity $OPT\left(\mathbf{v} \mid \mathbf{S}\right)$ then denotes the optimal allocation of the remaining $S' = M \setminus S$ set of items, while $\mathbf{v}(OPT\left(\mathbf{v} \mid \mathbf{S}\right))$ denotes the value.

The first condition of Definition 5.4.1 generalizes the notion of high revenue, while the second condition generalizes the notion of high surplus. To see the previous fact more clearly, consider the full-information case of a sequential combinatorial auction. Each buyer has a valuation v_i known to us, but buyers arrive in a an arbitrary order. Let OPT_i denote the items that buyer i gets at the optimal allocation. Because of the sequential nature of the auction, there might be some items $j \in OPT_i$ that end up being purchased by buyers that arrived previous to buyer i. Hence, we want the aggregate price of such items to be high enough so as to cover the welfare lost due to allocating them suboptimally. The right-hand side of the first condition (without the multiplicative $1/\alpha$ term) is precisely the welfare lost due to the aforementioned fact. Furthermore, the second condition formalizes the fact that prices should be low enough so that buyers can afford their optimal allocations. In line with the following example, consider buyer i arriving, when only a subset of OPT_i being available. Nevertheless, he should be able to purchase the remaining items inside OPT_i in a price that guarantees a high utility. To be more accurate, the second condition of Definition 5.4.1 requires that the aggregate utility of buyers (i.e the surplus) be high, that is, buyers collectively can pay for their optimal outcomes.

The novel contribution of this approach is that it completely decouples the Bayesian, from the full information setting. That is, as it is proved in Theorem 5.4.2, in order to prove a prophet inequality for the stochastic setting, it suffices to find balanced prices for the full information, non-stochastic setting. Furthermore, this definition as it generalizes the "revenue-surplus" approach, is natural and intuitive. Previous work from Kleinberg and Weinberg [45], introduced a similar definition of balanced thresholds, which, however, applied directly to the stochastic setting, thus making the argument probabilistic. As we mentioned previously, using the above definition gives us the power to argue about the simpler, full-information setting. Below we state the theorem which reduces the stochastic to the full-information setting.

Theorem 5.4.2. *Consider the setting where valuations are drawn from product distributions. If there exists a pricing rule $\mathbf{p}^{\mathbf{v}}$ that is (α, β)-balanced with respect to a valuation profile $\mathbf{v}$. Then posting prices:*

$$p_i(x_i) = \frac{\alpha}{1 + \alpha\beta} \mathbb{E}_{\tilde{\mathbf{v}} \sim D}[p_i^{\tilde{\mathbf{v}}}(x_i)]$$

yields welfare at least $\frac{1}{1+\alpha\beta} \mathbb{E}[\mathbf{v}(OPT(\mathbf{v}))]$

Proof. Once again, we are going to break the value into revenue and surplus and lower bound them separately. The novelty in the argument is that we are going to use the pointwise inequalities that Definition 5.4.1 requires for (α, β)-balanced prices.

Revenue:

$$\sum_{i \in N} p_i(x_i(\mathbf{v})) = \frac{\alpha}{1+\alpha\beta} \sum_i \mathbb{E}_{\tilde{\mathbf{v}}}[p_i^{\tilde{\mathbf{v}}}(x_i)] \tag{5.15}$$

$$\geq \frac{1}{1+\alpha\beta} \mathbb{E}_{\tilde{\mathbf{v}}}[\tilde{\mathbf{v}}(OPT(\tilde{\mathbf{v}})) - \tilde{\mathbf{v}}(OPT(\tilde{\mathbf{v}} \mid x(\mathbf{v})))] \tag{5.16}$$

Where the inequality is by the first condition of balanced prices. Taking expectations on both sides of Equation 5.16 we get:

$$\mathbb{E}_{\mathbf{v}}\left[\sum_{i \in N} p_i(x_i(\mathbf{v}))\right] \geq \frac{1}{1+\alpha\beta} \mathbb{E}_{\tilde{\mathbf{v}}}[\tilde{\mathbf{v}}(OPT(\tilde{\mathbf{v}}))] - \frac{1}{1+\alpha\beta} \mathbb{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}(OPT(\tilde{\mathbf{v}} \mid x(\mathbf{v})))\right] \tag{5.17}$$

Surplus: To lower bound the expected surplus we consider the bundle a buyer i could have bought upon his arrival. First, we consider an independent sample $\mathbf{v}' \sim D$ from the distribution. Now when i arrives, he can choose to buy nothing, thus obtaining non-negative utility, or he can choose to buy the bundle: $OPT_i((v_i, \mathbf{v}'_{-i}) \mid x(v'_i, \mathbf{v}_{-i}))$. Therefore, his expected utility can be lower bounded as follows:

$$\mathbb{E}_{\mathbf{v}}[u_i(\mathbf{v})] \geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[v_i(OPT_i((v_i, \mathbf{v}'_{-i}) \mid x(v'_i, \mathbf{v}_{-i}))) - \frac{\alpha}{1+\alpha\beta}\, p_i(OPT_i((v_i, \mathbf{v}'_{-i}) \mid x(v'_i, \mathbf{v}_{-i})))\right] \tag{5.18}$$

Now summing over all buyers and by observing that $\mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_i v_i(OPT_i((v_i, \mathbf{v}'_{-i}) \mid x(v'_i, \mathbf{v}_{-i})))\right] = \mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(OPT\left(\mathbf{v}'\right) \mid x(\mathbf{v})\right)\right]$:

$$\mathbb{E}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right] \geq \mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(OPT\left(\mathbf{v}'\right) \mid x(\mathbf{v})\right) - \frac{\alpha}{1+\alpha\beta}\, \sum_i p_i(OPT_i((v_i, \mathbf{v}'_{-i}) \mid x(v'_i, \mathbf{v}_{-i})))\right] \tag{5.19}$$

By invoking the second condition of balanced prices we can upper bound the second term in the right-hand side of Equation 5.19:

$$\sum_i p_i(OPT_i((v_i, \mathbf{v}'_{-i}) \mid x(v'_i, \mathbf{v}_{-i}))) \leq \beta\, \mathbb{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}(OPT(\tilde{\mathbf{v}} \mid x(\mathbf{v})))\right] \tag{5.20}$$

Taking expectations on both sides an replacing $\tilde{\mathbf{v}}$ with $\mathbf{v}'$ we get:

$$\mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_i p_i(OPT_i((v_i, \mathbf{v}'_{-i}) \mid x(v'_i, \mathbf{v}_{-i})))\right] \leq \beta\, \mathbb{E}_{\mathbf{v}'}\left[\mathbf{v}'(OPT(\mathbf{v}' \mid x(\mathbf{v})))\right] \tag{5.21}$$

Combining 5.19 and 5.21 we finally get:

$$\mathbb{E}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right] \geq \frac{1}{1+\alpha\beta}\, \mathbb{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'(OPT(\mathbf{v}' \mid x(\mathbf{v})))\right] \tag{5.22}$$

Putting it all together, we replace $\tilde{\mathbf{v}}$ with $\mathbf{v}'$ in 5.16 (and we are allowed to, because they are independent samples from the same distribution) and we sum 5.16 and 5.22 to conclude the proof.

□

In general, prices can be dynamic, that is, the price that buyer i sees can depend on the partial allocation to previous buyers. Moreover, balanced prices can also be defined with respect to an algorithm, which we denote by ALG, that computes an allocation $\mathbf{x}(\mathbf{v}) = ALG(\mathbf{v})$ which respects the feasibility constraints but is not necessarily social welfare maximizing. For the sake of completeness, we state refined versions of Definition 5.4.1 and Theorem 5.4.2 which incorporate the aforementioned observations.

First, we need some more notation. Let $\mathbf{x} = (x_1 \ldots, x_n)$ be a (partial) allocation. Then for each buyer i we define $\mathbf{x}_{[i-1]} = (x_1, \ldots, x_{i-1}, 0, \ldots, 0)$, that is the allocation confined to the first $i-1$ buyers. We define $p_i(x_i \mid \mathbf{y})$ as the price of outcome x_i offered to buyer i given partial allocation $\mathbf{y} \in \mathcal{F}$. We require also that $p_i(x_i \mid \mathbf{y}) = \infty$ for every outcome $(x_i, \mathbf{y}) \notin \mathcal{F}$. That is, we cannot price outcomes that violate the feasibility constraints. Finally, we use an exchange-compatible set $\mathcal{F}_{\mathbf{x}}$ to denote all the outcomes $\mathbf{y}$ that remain feasible after partially allocating $\mathbf{x}$. In general, the following definition takes into account that prices can be dynamic and discriminatory.

Definition 5.4.3 ((α, β)-balanced prices **(general case)**)**.** Let $\alpha > 0$, $\beta \geq 0$. Given a set of arbitrary downward-closed feasibility constraints $\mathcal{F}$ and a valuation profile $\mathbf{v}$, a pricing rule $\mathbf{p}$ is (α, β)-balanced with respect to an allocation rule ALG, an exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x}\in X}$, and an indexing of the players $i = 1, \ldots, n$ if for all $x \in \mathcal{F}$:

- $\sum_i p_i(x_i \mid \mathbf{x}_{[i-1]}) \geq \frac{1}{\alpha}\left(\mathbf{v}(ALG(\mathbf{v})) - \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))\right)$
- $\forall \mathbf{x}' \in \mathcal{F}_x$: $\sum_i p_i(x_i' \mid \mathbf{x}_{[i-1]}) \leq \beta \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$

Theorem 5.4.4. *Suppose that the collection of pricing rules $(\mathbf{p}^{\mathbf{v}})_{\mathbf{v}\in V}$ for feasible outcomes $\mathcal{F}$ and valuation profiles $\mathbf{v} \in V$ is (α, β)-balanced with respect to allocation rule ALG and indexing of the players $i = 1, \ldots, n$. Then the posted-price mechanism with pricing rule $\frac{\alpha}{1+\alpha\beta} \cdot \mathbf{p}$, where $p_i(x_i|y) = \mathbb{E}_{\tilde{\mathbf{v}}}[p_i^{\tilde{\mathbf{v}}}(x_i|y)]$, generates welfare at least $\frac{1}{1+\alpha\beta}\mathbb{E}_{\mathbf{v}}[\mathbf{v}(ALG(\mathbf{v}))]$ when approaching players in the order they are indexed.*

We omit the proof of the theorem above, as it is almost identical to that of Theorem 5.4.2. Observe that $\mathbf{x}_{[i-1]}$ does not depend on v_i. For completeness we also state the definition of weakly balanced prices and the respective theorem.

Definition 5.4.5 (weakly $(\alpha, \beta_1, \beta_2)$-balanced prices **(general case)**)**.** Let $\alpha > 0$, $\beta_1, \beta_2 \geq 0$. Given a set of arbitrary downward-closed feasibility constraints $\mathcal{F}$ and a valuation profile $\mathbf{v}$, a pricing rule $\mathbf{p}$ is *weakly* $(\alpha, \beta_1, \beta_2)$*-balanced* with respect to an allocation rule ALG, an exchange-compatible family of sets $(\mathcal{F}_{\mathbf{x}})_{\mathbf{x}\in X}$, and an indexing of the players $i = 1, \ldots, n$ if for all $x \in \mathcal{F}$:

- $\sum_i p_i(x_i \mid \mathbf{x}_{[i-1]}) \geq \frac{1}{\alpha}\left(\mathbf{v}(ALG(\mathbf{v})) - \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))\right)$
- $\forall \mathbf{x}' \in \mathcal{F}_x$: $\sum_i p_i(x'_i \mid \mathbf{x}_{[i-1]}) \leq \beta_1 \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) + \beta_2 \mathbf{v}(ALG(\mathbf{v}))$

Theorem 5.4.6. *Suppose that the collection of pricing rules $(\mathbf{p}^{\mathbf{v}})_{\mathbf{v}\in V}$ for feasible outcomes $\mathcal{F}$ and valuation profiles $\mathbf{v} \in V$ is weakly $(\alpha, \beta_1, \beta_2)$-balanced with respect to allocation rule ALG and indexing of the players $i = 1 \ldots, n$ with $\beta_1 + \beta_2 \geq \frac{1}{\alpha}$. Then for $\delta = \frac{1}{\beta_1 + \max\{2\beta_2, 1/\alpha\}}$ the posted-price mechanism with pricing rule $\delta \cdot \mathbf{p}$, where $p_i(x_i|y) = \mathbb{E}_{\tilde{\mathbf{v}}}[p_i^{\tilde{\mathbf{v}}}(x_i|y)]$, generates welfare at least $\frac{1}{\alpha(2\beta_1+4\beta_2)}\mathbb{E}_{\mathbf{v}}[\mathbf{v}(ALG(\mathbf{v}))]$ when approaching players in the order they are indexed.*

The proof of the above theorem is similar to that of 5.4.2 and is thus omitted.

Again in this case, if a price vector $\mathbf{p}$ is balanced with respect to an allocation rule, we can compute a price vector $\hat{\mathbf{p}}$, for which it holds that $||\hat{\mathbf{p}} - \mathbf{p}||_\infty < \epsilon$ using $poly(n, m, 1/\epsilon)$ samples. Then by standard concentration bounds we achieve an additive loss of $O(n\epsilon)$ in the competitive ratio.

5.4.1 Applying balanced prices to XOS Combinatorial Auctions

In this part we use the balanced prices machinery to obtain a posted price mechanism that is a 2-approximation to the optimal social welfare of a Combinatorial Auction with XOS buyers. Observe that this approxiamation ratio is the best possible as the setting inherits the lower bound of the single-item prophet inequality. The authors in [32] prove the following theorem

Theorem 5.4.7 (Combinatorial Auctions with XOS valuations)**.** *For Combinatorial Auctions with XOS valuations a $(2+\epsilon)$-approximate posted-price mechanism, with static item prices, can be computed in $poly(m, n, 1/\epsilon)$ demand and XOS queries.*

The authors make use of a fractional solution to the *configuration LP* of Combinatorial Auctions which is the following:

$$\max \sum_{i=1}^{n} \sum_{S \subseteq M} v_i(S) x_{i,S} \tag{5.23}$$

$$s.t \sum_{i=1}^{n} \sum_{S|j \in S} x_{i,S} \leq 1 \qquad \forall j \in M \tag{5.24}$$

$$\sum_{S \subseteq M} x_{i,S} \leq 1 \qquad \forall i \in [n] \tag{5.25}$$

$$x_{i,S} \in [0, 1] \qquad \forall i \in [n], S \subseteq M \tag{5.26}$$

Let's assume for simplicity that we have the optimal *integer* solution which we denote by $S^* = OPT(\mathbf{v})$. Let w_i^S be the representative additive function of buyer i on set S according to the definition of XOS valuations. Firstly, we prove the foolowing lemma:

Lemma 5.4.8. *Fix a valuation profile $\mathbf{v}$ and an exchange-compatible family of sets defined as $\mathcal{F}_{\mathbf{x}} = \{\mathbf{y} \in \mathcal{F} \mid (\bigcup_i x_i) \cap (\bigcup_i y_i) = \emptyset\}$. The prices $p_j = w_i^{S_i^*}(\{j\})$ if $j \in S_i^*$ are $(1, 1)$-balanced with respect to OPT and $\mathcal{F}_{\mathbf{x}}$.*

Proof. Observe that we can lower bound $\mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$ by considering the set $S_i^* \setminus \bigcup_i x_i$ for buyer i. Hence,

$$\mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \geq \sum_{k \in N} v_k(S_k^* \setminus (\bigcup_i x_i)) \geq \sum_{k \in N} \sum_{j \in S_k^* \setminus (\bigcup_i x_i)} w_k^{S_k^*}(\{j\})$$

Where the second inequality is by definition of XOS functions. For each set of goods x, the prices extend linearly, i.e $p(x) = \sum_{j \in S} p(\{j\})$. For each agent i and allocation x_i, we will have $p(x_i \mid \mathbf{z}) = p(x_i)$ whenever x_i is disjoint from $\mathbf{z}$ and ∞ otherwise. For the first condition of balanced prices, take an allocation $\mathbf{y} \in \mathcal{F}_{\mathbf{x}}$, then

$$\sum_i p(y_i \mid \mathbf{x}) = \sum_i \sum_{j \in y_i} p(\{j\}) \tag{5.27}$$

$$= \sum_i \sum_k \sum_{j \in S_k^* \cap y_i} w_k^{S_k^*}(\{j\}) \tag{5.28}$$

$$\geq \sum_{k \in N} \sum_{j \in S_k^*} w_k^{S_k^*}(\{j\}) - \sum_{k \in N} \sum_{j \in S_k^* \setminus (\bigcup_i x_i)} w_k^{S_k^*}(\{j\}) \tag{5.29}$$

$$\geq \mathbf{v}(OPT(\mathbf{v})) - \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \tag{5.30}$$

For the second condition of balanced prices for all $\mathbf{x}$ and all $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$ we have the following:

$$\sum_i p(x_i' \mid \mathbf{x}) = \sum_i \sum_{j \in x_i'} p(\{j\}) \tag{5.31}$$

$$= \sum_i \sum_k \sum_{j \in S_k^* \cap y_i} w_k^{S_k^*}(\{j\}) \tag{5.32}$$

$$\geq \sum_{k \in N} \sum_{j \in S_k^*} w_k^{S_k^*}(\{j\}) - \sum_{k \in N} \sum_{j \in S_k^* \setminus (\bigcup_i x_i)} w_k^{S_k^*}(\{j\}) \tag{5.33}$$

$$\geq \mathbf{v}(OPT(\mathbf{v})) - \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \tag{5.34}$$

5.5 Walrasian Equilibrium vs Prophet Inequality

A Walrasian Equilibrium is a set of allocations and item prices such that the market clears, or in other words demand equals supply. Due to this fact, it is instructive we compare Walrasian Equilibrium to the Prophet Inequality framework. Let us consider the single-item setting, where buyers arrive having a value for getting the item. We also study the full-information version of the problem where buyer's valuations are a priori known. Now any price between the largest value and the second largest (assuming buyers buy in case of indifference) value yields the optimal outcome: the item is allocated to the bidder with highest value. In comparison, prophet inequalities only guarantee an approximation to the optimal social welfare, but they hold even when buyers' values are stochastic.

A second important difference between Walrasian Equilibrium and Prophet Inequalities lies in tie-breaking. On the one hand, Prophet Inequalities provide results that do not depend on the nature of tie-breaking [1] (e.g the decision when a realized value is equal to the price). On the other hand, Walrasian Equilibrium rely on tie-breaking. In [42] it is argued that prices of a Walrasian Equilibrium cannot on their own coordinate a market. That is because, a buyer's demand correspondence might contain more that one bundles. If we allow buyers to arbitrarily choose one of them, that might lead to an over-demand of items and loss in welfare due to lack of coordination. Furthermore, consider the example of the previous paragraph: one could claim that posting a price that is strictly between the largest and the second largest value will coordinate the market and will yield the optimal outcome. However, when buyers arrive sequentially in markets, it is impossible to know the largest and second largest value without relying to an external coordinator that knows the values of all bidders. The authors also prove that for every tie-breaking rule there is an instance where the over-demand is $\Omega(n)$, i.e every item is demanded by every buyer. Finally, they develop sufficient conditions for the valuations that the over-demand is bounded by one.

5.6 Extensions and Variants

In this section we are going to briefly explore some extensions and variants to the classic prophet inequality framework. The variants tweak the assumptions on the ordering of buyers and the distributions. Furthermore, we explore a variant where we can use only a limited number of samples of each underlying distribution, an assumption that is more practically relevant than knowing the exact distribution for every buyer. Last but not least, we extend the prophet inequality framework for the revenue maximization objective.

5.6.1 Prophet Secretary

The first variant of the prophet inequality lies on the relaxation of the assumption of worst-case (adversarial) ordering of buyers. While this assumption gives the seller a machinery that is robust to market manipulation, it is too constricting. In practice, buyers arrive in a uniformly random order because the choice of when to arrive is influenced by random real-life noise. It turns out that if we modify the original setting of prophet inequality assuming that the order is random we can do better than a $1/2$-competitive algorithm for the single-item setting. The new setting is called *prophet secretary* and takes its name due to its resemblance to the folklore secretary problem. Esfandiari et.al [33] prove a $1 - 1/e$ competitive ratio, which is not tight. The current state-of-the-art is 0.669, due to Correa et. al [22], via a multiple-threshold strategy (note that $1 - 1/e \approx 0.632$). In the same paper, they also show an upper bound; no algorithm can achieve a competitive ratio better than $\sqrt{3} - 1 \approx 0.732$. For the cardinality case of selecting at most k values (or selling at most k items) Arnosti and Ma [3] showed that if one sets a single threshold T such that, in expectation, we have $k \cdot \gamma_k$ realizations above T, where $\gamma_k = 1 - e^{-k} \frac{k^k}{k!}$, then one obtains a γ_k-competitive ratio and this is tight for every k among single-threshold strategies. Observe that for $k = 1$, we retrieve the known $1 - 1/e$ ratio.

[1] That is not completely true, since the threshold set by Samuel-Cahn depends on tie-breaking: they consider two policies that either accept or reject respectively and show that one of them yields a 2-approximation but neither works all of the time. However, the proof of Kleinberg and Weinberg and balanced prices do not rely on tie-breaking.

5.6.2 The I.I.D case

In this variant, we assume that each one of the n random variables is a sample of the same distribution, i.e $D_1 = D_2 = \dots D_n = D$. Here the concept of order arrival is irrelevant because we have n samples from the same distribution. We simply draw n samples from D, randomly permute them and give them as input to the online algorithm. Correea et.al [23] obtain a tight 0.745-approximation via an adaptive multi-threshold algorithm (which translates into a dynamic posted pricing scheme). Their proof formalizes the intuition that characterizes an optimal multi-threshold strategy. If ones arrive at the last buyer without having sold the item, give it to the last buyer for a price of 0. Otherwise, the threshold (or equivalently price) for the $i-th$ buyer should be the expected value achieved when running the optimal online algorithm on the following $n-i$ buyers.

5.6.3 Prophet Inequalities with limited information

Pioneered by the work of Azar, Kleinberg and Weinberg [6] this variant examines the prophet inequality setting where the gambler has access to a limited number of samples from each distribution. Azar et al. [6] showed that there is a connection between this model and the secretary problem, as many algorithms for the secretary problem can be adapted to obtain constant-factor sample-based prophet inequalities. The authors prove prophet inequalities with constant, but not optimal, competitive ratio for classes of matroids using only one sample. Surprisingly, Rubinstein, Wang and Weinberg [58] manage to retrieve the optimal 2-competitive ratio for the single-item case using only one sample. Caramanis et al. [15] consider sample-based greedy algorithms, which are, in a sense, a refinement of the framework of Azar et al [6]. With this framework, they obtained improved factors for various classes of matroids.

5.6.4 Prophet Inequalities for Revenue Maximization

The mechanism design implications of prophet inequalities to social welfare have been the exclusive topic of study in this thesis. However, we are going to briefly discuss work on prophet inequalities that inform the design of posted price mechanisms for (approximate) revenue maximization. For single-parameter settings, the connection to revenue maximization makes use of Myerson's theorem that equates expected revenue with virtual value [53]. For each buyer i we can define a virtual value function that depends on D_i, the distribution over agent i's value, and maps each value to a (possibly negative) virtual value. A standard result in Bayesian mechanism design equates the revenue of a mechanism with its expected virtual welfare. One can therefore approximate the revenue of the optimal mechanism by applying the prophet inequality policy to the virtual values, rather than the original values. E.g., for a single item, one could accept the first prize whose virtual value is greater than half the expected maximum virtual value. This yields an order-oblivious posted-price mechanism, albeit one with potentially personalized prices. We refer the interested reader to [17] for more details on this approach. An interesting question is how well one can approximate the optimal revenue using an anonymous price, rather than personalized prices; Alaei et al. [2] show that an e-approximation is possible using a single posted price, under a standard regularity assumption on the value distributions.

Myerson's characterization does not generally apply in multi-dimensional settings, so prophet inequalities cannot be directly used for revenue maximization in these cases.

However, there have been significant advances in applying prophet inequalities to specific revenue-maximization problems. One of the pioneering innovations was the development of an approximately revenue-optimal sequential posted-price mechanism for matching markets, assuming that each agent's values for the items are independent of each other [16, 17]. This approach directly bounds the optimal revenue in the unit-demand scenario by relating it to a single-parameter problem. Chawla et al. [17] also demonstrate the use of prophet inequalities, utilizing virtual values in the single-parameter problem, to achieve an approximation result for revenue. These techniques have been extended beyond the unit-demand case to situations involving, for instance, matroid constraints, still assuming independence of values across items.

Further research in algorithmic mechanism design has strengthened the link between virtual welfare and revenue for multi-dimensional problems. This connection interprets virtual values in terms of marginal revenue and dual solutions in a related allocation program [13, 14]. This development has led to better upper bounds on optimal revenue in multi-item mechanism design problems, paving the way for approximation results and the use of multi-dimensional prophet inequalities for revenue maximization. This line of inquiry has produced constant approximations to optimal revenue using posted prices in broader classes of multi-item problems. For instance, Cai and Zhao [10] demonstrate that sequential posted price mechanisms can O(1)-approximate the optimal revenue in submodular Combinatorial Auctions and other contexts, assuming independence across items.

Chapter 6

Prophet inequalities for routing and admission control in capacitated networks.

6.1 Introduction

Consider a telecommunication network that serves requests. The network comprises nodes, which for instance could be routers or hubs, and bilateral connections between nodes that possess a certain bandwidth. The bandwidth is representative of the capacity of the connection to serve concurrent requests. Requests arrive sequentially to their start nodes, demanding to be routed in a path that gets them to their terminal nodes. We study the mechanism design version of the problem : each request comes from a buyer who has some private value of getting a path that connects him to his desired destination. The mechanism claims non-zero value from the buyer, if and only if it routes the buyer's request in the network. Assuming that values are drawn independently by publicly known distributions, can we design a posted-price mechanism that prices bandwidth connections in order to (approximately) maximize social welfare? The two main reasons behind pricing connections (instead of paths) are the following: 1) there are exponentially many different paths in a network, while there are only polynomially many connections, 2) different nodes might be owned by different service providers, hence for every node owned by a provider, the latter should only know the prices of the connections to the immediate neighbors of each node. In previous chapters we saw that there has been tremendous work in prophet inequalities with complement-free valuations, culminating in a prophet inequality for subadditive bidders with constant (stochastic) competitive ratio. This chapter, on the other hand, examines positive and negative results for prophet inequalities with valuations that exhibit complementarity. Observe that in the above scenario, a connection has value only if it is bought together with other connections in a way that they form a valid path.

6.2 Model

In this section we define a model which formalizes the problem we described in the introduction. Let $\mathcal{G} = (V, E)$ be an undireccted graph defined on a set of $|V| = n$ vertices and $|E| = m$ edges. Furthermore, let $c : E \to \mathbb{R}_{\geq 0}$ be an *edge capacity* function, which quantifies the bandwidth of a connection. Let $r_i = (s_i, t_i, v_i), i \in [k]$ be a request form start node s_i, to a terminal (or destination) node t_i, which carries value v_i and $\mathcal{P}_i$ be the set of all $s_i - t_i$ paths. We can think of requests as buyers who want to be routed in the network and have a value for doing so. We will use the terms requests and buyers interchangeably. A feasible allocation assigns bundles of edges to each buyer in a way that capacities are not violated. For each buyer i we define his valuation function $w_i : 2^E \to \mathbb{R}_{\geq 0}$ as follows: for any bundle S, if there exists a subset of edges $S' \subseteq S$ such that $S' \in \mathcal{P}_i$ then $w_i(S) = v_i$. For any other bundle T, $w_i(T) = 0$. That is, a bundle yields value to buyer i if and only if it contains a valid $s_i - t_i$ path. Since it is a single-parameter valuation function, from now on we will describe buyers' valuations with a single number: v_i. We study the problem under the prophet inequality paradigm by evaluating pricing schemes on their (stochastic) competitive ratio α as defined in Definition 5.2.1. Unless stated otherwise, we assume that the order in which the requests are processed by the online algorithm is chosen by an *adaptive adversary*.

6.3 LP formulation

In this section we address the underlying optimization problem. That is, we assume that we know the whole input of requests r_i and we want to maximize social welfare subject to the constraints imposed by the capacities for each edge. We define $f_i(v, w)$ to be the flow by request r_i on edge $(v, w) \in E$. The Linear Programming formulation, is the following:

$$\max \quad \sum_i \sum_{w \in V} v_i f_i(s_i, w) \tag{6.1}$$

$$s.t \quad \sum_i f_i(v, w) \leq c(e) \qquad \forall (v, w) \in E \tag{6.2}$$

$$\sum_w f_i(v, w) - \sum_w f_i(w, v) = 0 \qquad \forall i \in [k], \forall v \neq s_i, t_i \tag{6.3}$$

$$\sum_w f_i(s_i, w) = \sum_w f_i(w, t_i) \qquad \forall i \in [k] \tag{6.4}$$

$$\sum_w f_i(s_i, w) \leq 1 \qquad \forall i \in [k] \tag{6.5}$$

$$f_i(v, w) \in [0, 1] \qquad \forall (v, w) \in E \tag{6.6}$$

Where Equation 6.1 maximizes the value of request r_i multiplied by the total amount of flow leaving s_i, equation 6.2 formalizes that concurrent flows in an edge should not violate its capacity. Equations 6.3 and 6.4 are the flow conversation constraints of any intermediate node that the flow passes and of the start and the terminal node respectively. Finally, equation 6.5 restricts the total flow exiting start node s_i to at most 1 and equation 6.6 is because we study the fractional relaxation of the integer program.

The above formulation can be seen as a multicommodity flow problem, which we know that it is NP-hard to find an optimal integer solution [34]. In the fractional regime there exist solutions based on solving the LP, as well as Fully Polynomial Time Approximation Schemes [43].

However, a valid question we examine in the rest of this section is the following: What happens to the optimal solution of the above LP as capacities grow larger? Intuitively, if capacities are large enough the problem becomes easier. That is because with larger capacities we get more "room for error": allocating a suboptimal request does not hurt our solution that much. Firstly, we are going to transform the above LP to an equivalent, but more convenient, form.

Our objective is to transform an optimal solution that is described by edge flows $f_i(v, w)$ to an equivalent one that is described by path flows $f_{i,p}$, for all requests r_i, and for all paths $p \in \mathcal{P}_i$. This is done by the following algorithm.

Algorithm 4 Path Decomposition Algorithm

input f
for all requests r_i **do**
  while there is a $s_i - t_i$ path p using only edges with $f_i(e) > 0$ **do**
    $f_{i,p} \leftarrow \min_{e \in p} f_i(e)$
    for $e \in p$ **do**
      $f_i(e) \leftarrow f_i(e) - f_{i,p}$
    end for
  end while
end for
output $f_{i,p}$ for each request, for each path

Running the algorithm we have transformed the first LP into the following:

$$\max \quad \sum_{i} \sum_{p \in \mathcal{P}_i} v_i f_{i,p} \tag{6.7}$$

$$s.t \quad \sum_{i} \sum_{p \in \mathcal{P}_i | e \in p} f_{i,p} \leq c(e) \qquad \forall e \in E \tag{6.8}$$

$$\sum_{p \in \mathcal{P}_i} f_{i,p} \leq 1 \qquad \forall i \in [k] \tag{6.9}$$

$$f_{i,p} \in [0, 1] \qquad \forall i \in [k], \forall p \in \mathcal{P}_i \tag{6.10}$$

From now on we will use this more convenient, but nonetheless equivalent, form of LP and we can express the optimal fractional solution in terms of path flows $f_{i,p}$. We interpret the flows $f_{i,p}$ as probability distributions on paths. That is, we perform a *randomized rounding* of the fractional solution of the LP as follows: request r_i gets path $p \in \mathcal{P}_i$ with probability $f_{i,p}$ and gets nothing with probability $1 - \sum_{p \in \mathcal{P}_i} f_{i,p}$. We immediately observe that the expected value of this rounding is equal to the maximum value attained by the LP (Equation 6.7). However, 6.8 is also satisfied in expectation. This means that there exist instances where our randomized rounding violates the capacity constraints. To address this issue we scale down all the capacities by a factor of $1 - \epsilon$, where $0 < \epsilon < 1$. A new solution $f'_{i,p} = (1 - \epsilon) f_{i,p}$ satisfies 6.8 and 6.9 and is also an $1 - \epsilon$ approximation of the optimal solution. We proceed by showing that the event of violating 6.8 does *not* happen with high probability. We define $X_{i,p} \in \{0, 1\}$ the random variable that indicates if the request r_i gets routed path $p \in \mathcal{P}_i$, for all i. Hence,

$$Pr[X_{i,p} = 1] = f'_{i,p}.$$

Furthermore, for each edge e, define $L_e = \sum_i \sum_{p \in \mathcal{P}_i | e \in p} X_{i,p}$ to be the load of the edge, i.e the amount of concurrent requests it carries. Observe that $\mathbb{E}[L_e] \leq (1-\epsilon)c(e)$. Due to the fact that $\sum_{p \in \mathcal{P}_i} X_{i,p} \leq 1$, i.e the randomized rounding returns at most one path for each request, $\{X_{i,p}\}_{p \in \mathcal{P}_i}$ is a set of negatively dependent random variables for all i. Additionally, $X_{i,p}$ and $X_{j,p}$ are independent for every $i \neq j$. Therefore, we can apply the following Chernoff bound from [19].

Theorem 6.3.1. *Let X be the sum of n independent (or better) random variables with mean value $\mathbb{E}[X] \leq \mu$. Then it holds that*

$$\Pr[X \geq (1+\delta)\mu] \leq e^{\frac{-\delta^2 \mu}{2+\delta}}$$

Applying the theorem with $X = L_e, \delta = \frac{\epsilon}{1-\epsilon}$ and $\mu = (1-\epsilon)c_e$ we have that:

$$\begin{aligned} Pr[L_e \geq c_e] &\leq e^{-\frac{\frac{\epsilon^2}{(1-\epsilon)^2}(1-\epsilon)c_e}{2+\frac{\epsilon}{1-\epsilon}}} \\ &= e^{-\frac{\epsilon^2 c_e}{2-\epsilon}} \\ &\leq e^{-\frac{\epsilon^2 c_e}{2}} \end{aligned}$$

Set $\epsilon = \sqrt{\frac{4\log m}{c_e}}$. Since $\epsilon < 1$ by assumption, we require that $c_e \geq 4\log m$. By this choice of ϵ we obtain: $\Pr[L_e \geq c_e] \leq \frac{1}{m^2}$. Applying a union bound on the set of edges E, we obtain that the probability a capacity of any edge is violated is at most $\frac{1}{m}$.

In this section, we examined the complete information, offline optimization problem of our original sequential stochastic setting. We showed that if the capacity of every edge is $\Omega(\log m)$ then the problem admits an (1-ϵ)-approximation, with the means of a randomized rounding algorithm that rounds carefully the optimal fractional solution produced by the LP. Despite being NP-hard in the general case, if the resources are large enough we can get as close as we want to the optimal solution. The question now becomes: **Is it possible to use this positive result to design a sequential posted-price mechanism for our stochastic setting that achieves the same welfare guarantee, is order-oblivious and prices items**? In the following sections, we try to answer this question.

6.4 An $O(m)$-approximate item pricing

The problem formulation that we introduced in the previous section falls into the class of Packing Integer Programs (PIPs). The *column sparsity* d of a PIP is defined to be an upper bound to the number of constraints a variable (here $f_{i,p}$) can participate. Since, in general, we can have a graph with only long paths, and 6.8 sums on all edges that belong to a path, the sparsity of our packing program is $d = \Theta(m)$, where m is the number of edges in the graph. The authors in [8] give an (non-truthful) algorithm for the optimization problem that has an approximation factor of $O(d^{1/b})$ where b is the minimum capacity. Furthermore, they prove that this approximation ratio is tight (up to constant factors) by proving an integrality gap of $\Omega(d^{1/b})$. Observe that for $b = \Omega(\log m)$ and $d = \Theta(m)$, the approximation factor becomes constant. A similar problem that involves multiplicity is Combinatorial Auctions when each item has b copies. The complexity theoretic lower bound for the problem, proved in [55] is $\Omega(m^{1/b})$. The framework of Lavi and Swamy [49] gives mechanisms that match this lower bound, but are only truthful in expectation.

Dütting et.al [32] prove a $8d - competitive$ prophet inequality for d-sparse packing problems using balanced prices as defined in Section 5.4. Below we present their argument, tailored to our model. We assume that the sparsity $d = \Theta(m)$ and that for all edges $e \in E$, $c_e = \Omega(\log m)$.

We assume that we have access to an allocation $\mathbf{f}^* = ALG(\mathbf{v})$ given by an offline algorithm that is given as input the valuation vector $\mathbf{v}$. For now, we can think of ALG as the optimal fractional solution that can be computed by solving the LP.Then the following theorem holds

Theorem 6.4.1. *For the d-spare linear packing program defined in the previous section and $\mathcal{F}$ being all fractional solutions, there exist $(1, 0, d)$-balanced prices with respect to ALG. The prices can be computed by running ALG once.*

Proof. Let $\mathcal{F}_{\mathbf{x}} = \left\{\mathbf{z} \mid \sum_i \sum_{p \in \mathcal{P}_i | e \in p}(x_{i,p} + z_{i,p}) \leq c(e)\ \forall e\right\}$ Let $\rho_e = \frac{1}{c(e)} \sum_i \sum_{p \in P_i | e \in p} v_i f^*_{i,p}$ be a per-edge pricing scheme. Observe that prices tend to grow bigger when the corresponding edge participates in many paths. Then $p_i(\mathbf{f}_i) = \sum_{p \in \mathcal{P}_i} \sum_{e \in p} f_{i,p} \rho_e$, where $\mathbf{f}_i = (f_{i,p})_{p \in \mathcal{P}_i}$. For the first condition of balanced prices we define a residual allocation $\mathbf{z}$ as follows:

$$z_{i,p} = f^*_{i,p}\left(1 - \max_{e' \in p} \frac{\sum\limits_{i'} \sum\limits_{p' \in P_{i'} | e' \in p'} f_{i',p'}}{c_{e'}}\right)$$

. Then $\mathbf{z} \in \mathcal{F}_x$ because for every edge $e \in E$ we have:

$$\begin{aligned}
\sum_i \sum_{p \in P_i | e \in p} f_{i,p} + \sum_i \sum_{p \in P_i | e \in p} z_{i,p} &= \sum_i \sum_{p \in P_i | e \in p} f_{i,p} + \sum_i \sum_{p \in P_i | e \in p} f^*_{i,p}\left(1 - \max_{e' \in p} \frac{\sum\limits_{i'} \sum\limits_{p' \in P_{i'} | e' \in p'} f_{i',p'}}{c_{e'}}\right) \\
&\leq \sum_i \sum_{p \in P_i | e \in p} f_{i,p} + \sum_i \sum_{p \in P_i | e \in p} f^*_{i,p}\left(1 - \frac{\sum\limits_{i'} \sum\limits_{p' \in P_{i'} | e \in p'} f_{i',p'}}{c_e}\right) \\
&\leq c_e
\end{aligned}$$

Where in the fist inequality we reduced the maximum, to edge e that is in path p. Now it follows that:

$$\sum_i p_i(\mathbf{f}_i) = \sum_i \sum_{p \in P_i} \sum_{e \in p} f_{i,p} \rho_e \tag{6.11}$$

$$= \sum_e \rho_e \left(\sum_i \sum_{p \in P_i | e \in p} f_{i,p} \right) \tag{6.12}$$

$$= \sum_e \left(\frac{1}{c_e} \sum_{i'} \sum_{p' \in P_{i'} | e \in p'} v_{i'} f^*_{i',p'} \right) \left(\sum_i \sum_{p \in P_i | e \in p} f_{i,p} \right) \tag{6.13}$$

$$= \sum_{i'} \sum_{p' \in P_{i'}} \sum_{e \in p'} \frac{v_{i'} f^*_{i',p'}}{c_e} \left(\sum_i \sum_{p \in P_i | e \in p} f_{i,p} \right) \tag{6.14}$$

$$\geq \sum_{i'} \sum_{p' \in P_{i'}} v_{i'} (f^*_{i',p'} - z_{i',p'}) \tag{6.15}$$

$$\geq \mathbf{v}(ALG(\mathbf{v})) - \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \tag{6.16}$$

The first and third equality are by definition of $p_i(\mathbf{f}_i)$ and ρ_e respectively. The second and fourth equality are rearranging terms. The last inequality follows from the fact that $\mathbf{z} \in \mathcal{F}_x$ and the welfare from allocating $\mathbf{z}$ is upper bounded by the optimal $\mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$. The first inequality holds because:

$$\sum_{e \in p'} \sum_i \sum_{p \in P_i | e \in p} \frac{f_{i,p}}{c_e} \geq \max_{e' \in p'} \frac{\sum_i \sum_{p \in P_i | e' \in p} f_{i,p}}{c_{e'}}$$

and

$$f^*_{i',p'} - z_{i',p'} = f^*_{i',p'} \left(\max_{e' \in p'} \frac{\sum_i \sum_{p \in P_{i'} | e' \in p} f_{i,p}}{c_{e'}} \right)$$

Now for the second condition of balanced prices:

$$\sum_i p_i(\mathbf{f}'_i) = \sum_i \sum_{p \in P_i} \sum_{e \in p} f'_{i,p} \rho_e \tag{6.17}$$

$$= \sum_e \rho_e \sum_i \sum_{p \in P_i | e \in p} f'_{i,p} \tag{6.18}$$

$$\leq \sum_e \rho_e c_e \tag{6.19}$$

$$= \sum_e \sum_i \sum_{p \in P_i | e \in p} v_i f^*_{i,p} \tag{6.20}$$

$$= \sum_i \sum_{p \in P_i} \sum_{e \in p} v_i f^*_{i,p} \tag{6.21}$$

$$\leq d \cdot \mathbf{v}(ALG(\mathbf{v})) \tag{6.22}$$

Where the first inequality is the capacity constraint and the second is by the sparsity definition. The first,second and fourth equalities are rearranging terms, while the third is by definition of ρ_e □

Combining Theorem 6.4.1 with Theorem 5.4.6 we get the following corollary.

Corollary 6.4.1.1. *Let ALG be the (optimal) fractional allocation rule. For $\delta = \frac{1}{2d}$ the posted price-mechanism with pricing rule $\delta \mathbf{p}$ where $p_i(\mathbf{f}_i) = \mathbb{E}_{\tilde{\mathbf{v}}}[p_i^{\tilde{\mathbf{v}}}(\mathbf{f}_i)]$ generates welfare at least $\frac{1}{4d}\mathbb{E}_{\mathbf{v}}[\mathbf{v}(ALG(\mathbf{v}))]$ when approaching the players in the order they are indexed.*

Observe that since ALG is the optimal fractional solution (which can be computed in polynomial time) the corollary implies an $4d$-competitive prophet inequality for the fractional d-sparse packing problem.

Now let us prove the theorem for integral solutions. We consider ALG to be an approximation algorithm to the k-sparse packing problem (eg. the algorithm in [8]) which, obviously, produces an integral solution.

Theorem 6.4.2. *For the d-spare linear packing program defined in the previous section and $\mathcal{F}$ being all integral solutions, there exist $(2, 0, d)$-balanced prices with respect to ALG. The prices can be computed by running ALG once.*

Proof. To define $\mathcal{F}_{\mathbf{x}}$, let for all edges $e \in E$: $c'(e) = c(e)$ if $\sum_i \sum_{p \in \mathcal{P}_i | e \in p} f_{i,p} \leq \frac{1}{2}$ and $c(e) = 0$ otherwise. That is, an edge keeps its capacity if and only if it is at most half full after adding $\mathbf{x}$. Then, $\mathcal{F}_x = \{\mathbf{z} \mid \sum_i \sum_{p \in \mathcal{P}_i | e \in p} (x_{i,p} + z_{i,p}) \leq c'(e)\}$. Define $z_{i,p} = f^*_{i,p}$ if $\sum_{i'} \sum_{p' \in \mathcal{P}'_i | e \in p'} f_{i',p'} \leq \frac{1}{2}$ $\forall e \in p$ and $z_{i,p} = 0$ otherwise. Observe that $z_{i,p} \geq f^*_{i,p}\left(1 - 2\max_{e \in p} \sum_{i'} \sum_{p' \in \mathcal{P}'_i | e \in p'} f_{i',p'}\right)$

Now the calculations are identical with the fractional case.

For the first condition of balanced prices we have:

$$\sum_i p_i(\mathbf{f}_i) = \sum_i \sum_{p \in P_i} \sum_{e \in p} f_{i,p} \rho_e \tag{6.23}$$

$$= \sum_e \rho_e \left(\sum_i \sum_{p \in P_i | e \in p} f_{i,p} \right) \tag{6.24}$$

$$= \sum_e \left(\frac{1}{c_e} \sum_{i'} \sum_{p' \in P_{i'} | e \in p'} v_{i'} f^*_{i',p'} \right) \left(\sum_i \sum_{p \in P_i | e \in p} f_{i,p} \right) \tag{6.25}$$

$$= \sum_{i'} \sum_{p' \in P_{i'}} \sum_{e \in p'} \frac{v_{i'} f^*_{i',p'}}{c_e} \left(\sum_i \sum_{p \in P_i | e \in p} f_{i,p} \right) \tag{6.26}$$

$$\geq \frac{1}{2} \sum_{i'} \sum_{p' \in P_{i'}} v_{i'} (f^*_{i',p'} - z_{i',p'}) \tag{6.27}$$

$$\geq \frac{1}{2} \left(\mathbf{v}(ALG(\mathbf{v})) - \mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}})) \right) \tag{6.28}$$

The first and third equality are by definition of $p_i(\mathbf{f}_i)$ and ρ_e respectively. The second and fourth equality are rearranging terms. The last inequality follows from the fact that $\mathbf{z} \in \mathcal{F}_x$ and the fact that the welfare from allocating $\mathbf{z}$ is upper bounded by the optimal $\mathbf{v}(OPT(\mathbf{v}, \mathcal{F}_{\mathbf{x}}))$. The first inequality holds because:

$$\sum_{e \in p'} \sum_i \sum_{p \in P_i | e \in p} \frac{f_{i,p}}{c_e} \geq \max_{e' \in p'} \frac{\sum_i \sum_{p \in P_i | e' \in p} f_{i,p}}{c_{e'}}$$

and

$$f^*_{i',p'} - z_{i',p'} \leq 2 f^*_{i',p'} \left(\max_{e' \in p'} \frac{\sum_i \sum_{p \in P_{i'} | e' \in p} f_{i,p}}{c_{e'}} \right)$$

For the second condition of balanced prices we have:

$$\sum_i p_i(\mathbf{f}'_i) = \sum_i \sum_{p \in P_i} \sum_{e \in p} f'_{i,p} \rho_e \tag{6.29}$$

$$= \sum_e \rho_e \sum_i \sum_{p \in P_i | e \in p} f'_{i,p} \tag{6.30}$$

$$\leq \sum_e \rho_e c_e \tag{6.31}$$

$$= \sum_e \sum_i \sum_{p \in P_i | e \in p} v_i f^*_{i,p} \tag{6.32}$$

$$= \sum_i \sum_{p \in P_i} \sum_{e \in p} v_i f^*_{i,p} \tag{6.33}$$

$$\leq d \cdot \mathbf{v}(ALG(\mathbf{v})) \tag{6.34}$$

Where the first inequality is the capacity constraint and the second is by the sparsity definition. The first,second and fourth equalities are rearranging terms, while the third is by definition of ρ_e □

Combining Theorem 6.4.2 with Theorem 5.4.6 we get the following corollary.

Corollary 6.4.2.1. *Let ALG be the algorithm in [8] for d-sparse packing that produces an integral allocation rule. For* $\delta = \frac{1}{2d}$ *the posted price-mechanism with pricing rule* $\delta \mathbf{p}$ *where* $p_i(\mathbf{f}_i) = \mathbb{E}_{\tilde{\mathbf{v}}}[p_i^{\tilde{\mathbf{v}}}(\mathbf{f}_i)]$ *generates welfare at least* $\frac{1}{8d}\mathbb{E}_{\mathbf{v}}[\mathbf{v}(ALG(\mathbf{v}))]$ *when approaching the players in the order they are indexed.*

Observe that since ALG is an $O(d^{1/b})$ - approximate algorithm, which is also the optimal approximation ratio, the corollary implies an $O(d^{1/b+1})$-competitive prophet inequality for the integral d-sparse packing problem.

However, instead of relying to an offline approximation algorithm, we can derive pricing schemes for the integral version of the problem based on the (optimal) fractional solution. The technique we use is essentially randomized rounding: we define $z_{i,p} = 1^*$ with probability $f^*_{i,p}$ if $\sum_{i'} \sum_{p' \in \mathcal{P}'_i | e \in p'} f_{i',p'} \leq \frac{1}{2}$ $\forall e \in p$ and $z_{i,p} = 0$ otherwise. Hence, $\mathcal{F}_{\mathbf{x}}$ contains distributions over outcomes. Performing the same calculations as before we can derive a $8d$-competitive prophet inequality, which this time is with respect to the fractional optimum.

Another example, from the same paper, in which assuming black-box access to an offline algorithm yields worse results is the case of Combinatorial Auctions with XOS (or submodular) bidders. Using the algorithm of Vondrák [61] gives an $\frac{2e}{e-1}$-competitive prophet inequality, while rounding the fractionally optimal solution of the configuration LP, gives an improved competitive ratio of 2. The latter is also the best possible, as the setting inherits the lower bound from the single-item case.

6.5 Going beyond $O(d)$

To improve the competitive ratio, we need to modify our approach. The problem with our current method is that we argue for item pricing, where the price of a bundle is the sum of the prices of the items it contains. However, the complementarity of value that our model exhibits calls for bundle pricing. Intuitively, if items have increased value when they are bundled together (consider, for example, that single-edges have no value unless they are bundled in a way that results in feasible paths) bundling should help in terms of social welfare. Firstly, we are going to prove that the $O(d)$-competitive anonymous, static item pricing from the previous section is tight (up to constant factors) when we are using item pricing. Let $\mathcal{P}$ be a single path of $\Theta(d)$ edges. Consider buyers arriving that belong to exactly one of the following categories: 1) buyers that demand only the whole path $\mathcal{P}$ and each one of them has valuation $\Theta(d)$ - for example, consider buyers that arrive on the leftmost node of path $\mathcal{P}$, wanting to go to the rightmost node, 2) buyers that only demand a single edge $e \in E$, all of them the same, and each one of them has value $1 + \epsilon$. If for all edges $e \in \mathcal{P}$ $p_e > 1$ then no buyer gets routed and we obtain total value of zero. Therefore, any item pricing that produces a finite competitive ratio should price each item below 1, i.e $p_e \leq 1$ for all $e \in E$. Now consider $\Theta(logm)$ buyers of category 2 arriving first. Because we have assumed that $c_e = \Omega(logm)$ and we set $p_e \leq 1$ all of them get allocated edge $e \in E$. Now e is at full capacity. This means that no one of the $\Theta(logm)$ subsequent buyers of category 1 get allocated the path, due to the fact that an edge inside the path is no longer available. Hence, the social welfare achieved in this scenario is $\Theta(logm)$. However, the optimal social welfare is $\Theta(d \log m)$, achieved by allocating the $\Theta(logm)$ buyers of category 1. We conclude that, item pricing cannot achieve a better than linear, competitive ratio.

Chawla et. al [18] consider a simpler model where the graph is a tree. They define $H = \frac{\max_i v_i}{\min_i v_i}$ and the prove an $O(\log H)$ competitive ratio for unit capacities and that this ratio decreases linearly with capacity, that is $O(\frac{1}{B} \log H)$, where $B = \min_e c_e$. Observe that if $B = \Omega(\log H)$ the competitive ratio is constant. For the unit-capacity setting they also prove a $\Omega\left(\sqrt{\frac{d}{\log d}}\right)$ lower bound on the competitive ratio. This work is the first one that resembles our model and validates our intuition that augmented resources (e.g if we have at least logarithmic capacity) help the competitive ratio of our online algorithm. Feldman et. al [37] also use bundling to tackle complementarity. They extend the definition of a Walrasian Equilibrium in a way that it is on bundles and not individual items. They prove that, despite the fact that the market may not clear, there is a polytime algorithm that achieves a 2-approximation to the optimal social welfare for any valuation class. The difference in their setting is that they consider the full information case, hence their algorithm is offline.

6.6 Warmup: Deterministic-like distributions

We study a simplified version of our problem. Firstly, we assume that all edges have the same capacity, which we denote by B. We also assume that buyers arriving at nodes can only have k discrete values. Let $\{(a_{ji}, v_{ji})\}_{j\in[k]}$ for path p_i. This means that the amount of buyers demanding path p_i and have value v_{ji} is a_{ji}. The underlying optimization problem must decide a fraction y_{ji} of the a_{ji} demands to accept, for each path and each of the k discrete values, in order to maximize the aggregate value, subject to not violating the capacity constraints. The linear program capturing these constraints is the following:

$$\max \sum_i \sum_{j=1}^{k} a_{ji}\ v_{ji}\ y_{ji} \tag{6.35}$$

$$s.t \sum_{p_i|e\in p_i} \sum_{j=1}^{k} a_{ji} y_{ji} \leq B \qquad \forall e \tag{6.36}$$

$$y_{ji} \in [0,1] \qquad \forall i,j \tag{6.37}$$

We study the admission control problem: Buyers arrive at nodes sequentially and in an adversarial order, demanding to be routed. Our objective is to design a strategy that irrevocably decides whether to accept or reject a buyer, without knowing the future, and compare its total value with the offline optimum, i.e an omniscient prophet that can select the best combination of buyers to maximize total value. We design a threshold-based algorithm that achieves at least half the value of the offline optimum.

Let $\mathbf{y}^*$ be the optimal fractional allocation given by the LP above. W.l.o.g, for each path p_i we relabel in terms of the index j, so that the vector $\{y^*_{ji}\}_{j\in[k]}$ has its components in increasing order of value. By a simple greedy-exchange argument we observe that these vectors, indexed in the way mentioned before, have some zeros in the first components, then a single fraction and a suffix of all ones. Let $m_i = \arg\min_j \{y^*_{ji} \neq 0\}$ be the index of the first non-zero entry of the vector. Now we are ready to state our theorem.

Theorem 6.6.1. *The following prices:*

$$p_i = v_{j_i,i}$$

where $j_i = \arg\max_{j=m_i,m_i+1}\{v_{j,i} \cdot \sum_{l=j}^{k} a_{ji} y^*_{ji}\}$ *yield a 2-approximation with respect to the optimal fractional allocation.*

Proof. Let $OPT_i = \sum_{j=1}^{k} a_{ji}\ v_{ji}\ y^*_{ji}$. be the welfare of the optimal fractional allocation for requests $s_i - t_i$ and ALG_i be the welfare that the algorithm with the above threshold achieves . By definition of m_i OPT_i can be equivalently written as $OPT_i = \sum_{j=m_i}^{k} a_{ji}\ v_{ji}\ y^*_{ji}$. We distinguish between two cases:

- If $j_i = m_i{+}1$ because any algorithm allocates at most $\sum_{j=m_i}^{k} a_{ji} y^*_{ji} \geq \sum_{j=m_i+1}^{k} a_{ji} y^*_{ji}$ it produces welfare at least $\sum_{j=m_i+1}^{k} a_{ji} v_{ji}$ at price $p_{j_i,i}$
- If $j_i = m_i$ any algorithm produces welfare at least $v_{m_i,i} \cdot \sum_{j=m_i}^{k} a_{ji} y^*_{ji}$

Adding the two cases gives:

$$\sum_{j=m_i+1}^{k} a_{ji} v_{ji} + v_{m_i,i} \cdot \sum_{j=m_i}^{k} a_{ji} y^*_{ji} \tag{6.38}$$

$$= \sum_{j=m_i+1}^{k} a_{ji}(v_{ji} + v_{m_i,i}) + v_{m_i,i} a_{m_i,i} y^*_{m_i,i} \tag{6.39}$$

$$\leq 2 \sum_{j=m_i+1}^{k} a_{ji} v_{ji} + v_{m_i,i} a_{m_i,i} y^*_{m_i,i} \tag{6.40}$$

$$\leq 2OPT_i \tag{6.41}$$

Where inequality 6.40 follows from the fact that we have relabeled j in increasing order of value, hence $v_{m_i,i} \leq v_{j,i}$ for all $j > m_i$ and 6.41 is by definition of OPT_i. We also observe that 6.39 is lower bounded by OPT_i. By taking the maximum of the two above cases we have that $ALG_i \geq \frac{1}{2} OPT_i$. Summing for all $s_i - t_i$ requests, yields the theorem. □

6.7 Reduction to the k-unit Prophet Inequality

In this section we are going to reduce our setting to the k-unit prophet inequality and then prove $(1,1) - balanced$ prices for this setting. We look into pairs of nodes, namely $s_i{-}t_i$, where we know a priori that n_i are expected to come. From these requests we solve the fractional relaxation with the *reduced* capacities to decide a fraction $a^*_i = \sum_p f^*_{i,p}$ of the n_i requests to accept. We define for each node pair $k_i = a^*_i n_i$. We price each pair independently with prices defined as:

$$p_i = \frac{1}{k_i} \max_{|S| \leq k_i} \sum_{i \in S} v_i$$

We proceed by showing that this pricing scheme is $(1,1) - balanced$

Theorem 6.7.1. *The above pricing scheme is* $(1, 1)-$*balanced with respect to a valuation profile* $\mathbf{v}$ *and the algorithm that rounds the fractional relaxation of the LP formulation.*

Proof. Fix a node pair i. Without loss of generality we sort the values in decreasing order: $v_1, v_2, \ldots, v_{k_i}, \ldots, v_{n_i}$. The allocation $\mathbf{x}$ might allocate to buyers that are not in the top-k_i of valuations. Hence if $\mathbf{x}$ allocates to $j \leq k_i$ suboptimal buyers we have that $\mathbf{v}(OPT(\mathbf{v})) - \mathbf{v}(OPT(\mathbf{v} \mid \mathbf{x})) = v_{k_i-j+1} + \cdots + v_{k_i}$, because after allocating $\mathbf{x}$ we can still select the top-$(k_i - j)$ remaining buyers that the optimal also selects. The total amount of payments for this allocation is $\sum_i p_i = \frac{j}{k_i}(v_1 + \cdots + v_{k_i})$. We can break the sum into chunks of $\frac{k_i}{j}$ valuations and because of the decreasing order of values we have, for example, $\frac{k_i}{j} \cdot v_{k_i-j+1} \leq v_1 + \cdots + v_{k_i/j}$ and $\frac{k_i}{j} \cdot v_{k_i-j+2} \leq v_{k_i/j+1} + \cdots + v_{2k_i/j}$ and so forth. Hence the first condition of balanced prices is fulfilled. For the second condition we take an allocation $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$ that can be feasibly added to a partial allocation $\mathbf{x}$. If there are j buyers remaining to be allocated after allocating $\mathbf{x}$ then $\mathbf{v}(OPT(\mathbf{v} \mid \mathbf{x})) = v_1 + \cdots + v_j$ and $\sum_i p_i(x'_i) = \frac{j}{k_i}(v_1 + \cdots + v_{k_i})$, we can prove by induction that the $\sum_i p_i(x'_i) \leq \mathbf{v}(OPT(\mathbf{v} \mid \mathbf{x}))$. The idea comes from the fact that since we have the values in decreasing order, for example $\frac{1}{2}(v_1+v_2) \leq v_1$, $\frac{2}{3}(v_1+v_2+v_3) \leq v_1+v_2$ and so forth. □

We use this pricing scheme to provide admission control inside the network: when a request arrives we route it in any $s_i - t_i$ path available provided that it is above the price. In the previous sections we showed, that by doing so for each node pair independently we violate the reduced capacity constraints. By the same Hoeffding bounds we show that we violate with high probability by at most ϵc_e which is feasible in terms of the original capacities. Since we prove $(1, 1) -$ *balanced* prices for each pair of nodes we get a factor of 2 that is due to the theorem of [32]. We know that in the high-k regime we have results that are asymptotically optimal with respect to the prophet, however we analyse for the worst case of k_i for *any* $s_i - t_i$ pair, which can be in the low-k regime, for instance it might be $k_i = 1$, where the best possible is a 2-approximation. That is the reason, our approach yields necessarily an approximation factor of 2 in the worst case.

6.8 Conclusions and Open Problems

We comment on our pricing scheme not being truthful in terms of paths. That is, a buyer might have an incentive to buy a superset of his desired path at a lower price. The prices are not monotonically increasing on paths as it is showcased by the following example:

Example 6.8.1 (Non-monotonic pricing scheme)**.** Consider a path graph with m edges where 3 buyers with value 9 and 3 buyers with value 12 want the whole path, for simplicity we represent them as $(3, 9), (3, 12)$ and 4 buyers with value 12, that is $(4, 12)$, who want only the edge in the middle of the path. This edge has capacity 10 and all the other edges have capacity 6. The optimal solution prices the big path at a price of 9 or lower,i.e accepts both $(3, 9)$ and $(3, 12)$, and prices the edge in the middle at a price lower than 12, i.e accepts all $(4, 12)$ buyers. Observe that we respect the capacity constraints. However, if $(4, 12)$ buyers arrive first they see a path that contains their desired path, namely the big path, at a lower price (it has a price of at most 9, compared to a price of 12). As a result they have an incentive to buy the big path, thus excluding 4 of the 6 in total buyers originally wanting that path.

This thesis identifies the non-monotonicity of pricing paths as an unresolved issue and thus an open problem for future research. In conclusion, despite the amazing work that has been done in the Prophet Inequality regime for complement-free valuations, very few work tries to address complementarity. Specifically, we mention that [18] is the only work that starts the discussion of pricing intervals and paths. However, as it is showcased by our paradigmatic problem of online network routing and admission control, there are relevant instances in real life which we need to price for complementary items. We show that, if we have a logarithmically large supply of items (in our problem this is the capacity) this allows us to solve a fractional releaxation that is a good approximation to the integral due to the small integrality gap. Then by pricing, node pairs independently we reduce the problem to the unit-demand setting (which by the way is the high level technique used in [18]) and invoke the k-unit Prophet Inequality for each separate pair. We achieve a $2(1+\epsilon)$-approximation to the optimal social welfare, where the factor 2 comes from the k-unit prophet inequality and the $1+\epsilon$ comes from the approximation of the underlying multi-commodity flow problem. Are there any other settings that exhibit complementarity, that have a small integrality gap (perhaps by assuming a large enough supply like in our setting) that can translate into a constant-approximate prophet inequality. Is there a way, by means of a different technique and analysis, to improve the factor 2 in the approximation ratio of our setting?

Bibliography

[1] Alaei, S. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014. doi: 10.1137/120878422. URL https://doi.org/10.1137/120878422.

[2] Alaei, S., Hartline, J. D., Niazadeh, R., Pountourakis, E., and Yuan, Y. Optimal auctions vs. anonymous pricing. *CoRR*, abs/1507.02615, 2015. URL http://arxiv.org/abs/1507.02615.

[3] Arnosti, N. and Ma, W. Tight guarantees for static threshold policies in the prophet secretary problem. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, pp. 242, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391504. doi: 10.1145/3490486.3538377. URL https://doi.org/10.1145/3490486.3538377.

[4] Assadi, S. and Singla, S. Improved truthful mechanisms for combinatorial auctions with submodular bidders. volume 0, pp. FOCS19–253–FOCS19–275, 0. doi: 10.1137/20M1316068. URL https://doi.org/10.1137/20M1316068.

[5] Assadi, S., Kesselheim, T., and Singla, S. Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 653–661, 2021. doi: 10.1137/1.9781611976465.40. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611976465.40.

[6] Azar, P. D., Kleinberg, R., and Weinberg, S. M. Prophet inequalities with limited information. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pp. 1358–1377, USA, 2014. Society for Industrial and Applied Mathematics. ISBN 9781611973389.

[7] Banihashem, K., Hajiaghayi, M., Kowalski, D. R., Krysta, P., and Olkowski, J. Power of posted-price mechanisms for prophet inequalities, 2024. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611977912.163.

[8] Bansal, N., Korula, N., Nagarajan, V., and Srinivasan, A. On k-column sparse packing programs. In *Conference on Integer Programming and Combinatorial Optimization*, 2009. URL https://api.semanticscholar.org/CorpusID:15478280.

[9] Bhawalkar, K. and Roughgarden, T. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 2011 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 700–709, 2011. doi: 10.1137/1.9781611973082.55. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611973082.55.

[10] Cai, Y. and Zhao, M. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pp. 170–183, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450345286. doi: 10.1145/3055399.3055465. URL https://doi.org/10.1145/3055399.3055465.

[11] Cai, Y., Daskalakis, C., and Weinberg, S. M. *Reducing Revenue to Welfare Maximization: Approximation Algorithms and other Generalizations*, pp. 578–595. doi: 10.1137/1.9781611973105.42. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611973105.42.

[12] Cai, Y., Daskalakis, C., and Weinberg, S. M. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. *CoRR*, abs/1207.5518, 2012. URL http://arxiv.org/abs/1207.5518.

[13] Cai, Y., Daskalakis, C., and Weinberg, S. M. Understanding incentives: Mechanism design becomes algorithm design. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 618–627, 2013. doi: 10.1109/FOCS.2013.72.

[14] Cai, Y., Devanur, N. R., and Weinberg, S. M. A duality based unified approach to bayesian mechanism design. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pp. 926–939, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341325. doi: 10.1145/2897518.2897645. URL https://doi.org/10.1145/2897518.2897645.

[15] Caramanis, C., Dütting, P., Faw, M., Fusco, F., Lazos, P., Leonardi, S., Papadigenopoulos, O., Pountourakis, E., and Reiffenhäuser, R. Single-sample prophet inequalities via greedy-ordered selection. In Naor, J. S. and Buchbinder, N. (eds.), *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pp. 1298–1325. SIAM, 2022. doi: 10.1137/1.9781611977073.54. URL https://doi.org/10.1137/1.9781611977073.54.

[16] Chawla, S., Hartline, J. D., and Kleinberg, R. Algorithmic pricing via virtual valuations. *CoRR*, abs/0808.1671, 2008. URL http://arxiv.org/abs/0808.1671.

[17] Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10, pp. 311–320, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781450300506. doi: 10.1145/1806689.1806733. URL https://doi.org/10.1145/1806689.1806733.

[18] Chawla, S., Miller, J. B., and Teng, Y. Pricing for online resource allocation: Intervals and paths. In Chan, T. M. (ed.), *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pp. 1962–1981. SIAM, 2019. doi: 10.1137/1.9781611975482.119. URL https://doi.org/10.1137/1.9781611975482.119.

[19] Chernoff, H. A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the sum of Observations. *The Annals of Mathematical Statistics*, 23(4):493 – 507, 1952. doi: 10.1214/aoms/1177729330. URL https://doi.org/10.1214/aoms/1177729330.

[20] Clarke, E. H. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971. doi: 10.1007/BF01726210. URL https://doi.org/10.1007/BF01726210.

[21] Correa, J. and Cristi, A. A constant factor prophet inequality for online combinatorial auctions. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pp. 686–697, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399135. doi: 10.1145/3564246.3585151. URL https://doi.org/10.1145/3564246.3585151.

[22] Correa, J., Saona, R., and Ziliotto, B. Prophet secretary through blind strategies. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pp. 1946–1961, USA, 2019. Society for Industrial and Applied Mathematics.

[23] Correa, J., Foncea, P., Hoeksma, R., Oosterwijk, T., and Vredeveld, T. Posted price mechanisms and optimal threshold strategies for random arrivals. *Mathematics of Operations Research*, 46(4):1452–1478, 2021. doi: 10.1287/moor.2020.1105. URL https://doi.org/10.1287/moor.2020.1105.

[24] Dobzinski, S. Two randomized mechanisms for combinatorial auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, APPROX '07/RANDOM '07, pp. 89–103, Berlin, Heidelberg, 2007. Springer-Verlag. ISBN 9783540742074. doi: 10.1007/978-3-540-74208-1_7. URL https://doi.org/10.1007/978-3-540-74208-1_7.

[25] Dobzinski, S. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, pp. 139–148, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450306911. doi: 10.1145/1993636.1993656. URL https://doi.org/10.1145/1993636.1993656.

[26] Dobzinski, S. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pp. 940–948, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341325. doi: 10.1145/2897518.2897569. URL https://doi.org/10.1145/2897518.2897569.

[27] Dobzinski, S. and Schapira, M. An improved approximation algorithm for combinatorial auctions with submodular bidders. SODA '06, pp. 1064–1073, USA, 2006. Society for Industrial and Applied Mathematics. ISBN 0898716055.

[28] Dobzinski, S., Nisan, N., and Schapira, M. Approximation algorithms for combinatorial auctions with complement-free bidders. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pp. 610–618, New York, NY, USA, 2005. Association for Computing Machinery. ISBN 1581139608. doi: 10.1145/1060590.1060681. URL https://doi.org/10.1145/1060590.1060681.

[29] Dobzinski, S., Nisan, N., and Schapira, M. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the Thirty-Eighth Annual ACM*

Symposium on Theory of Computing, STOC '06, pp. 644–652, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595931341. doi: 10.1145/1132516.1132607. URL https://doi.org/10.1145/1132516.1132607.

[30] Duetting, P. and Kleinberg, R. Polymatroid prophet inequalities. *CoRR*, abs/1307.5299, 2013. URL http://arxiv.org/abs/1307.5299.

[31] Dütting, P., Kesselheim, T., and Lucier, B. An o(log log m) prophet inequality for subadditive combinatorial auctions, 2020. URL https://doi.org/10.1109/FOCS46700.2020.00037.

[32] Dütting, P., Feldman, M., Kesselheim, T., and Lucier, B. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs, 2017.

[33] Esfandiari, H., Hajiaghayi, M., Liaghat, V., and Monemizadeh, M. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017. doi: 10.1137/15M1029394. URL https://doi.org/10.1137/15M1029394.

[34] Even, S., Itai, A., and Shamir, A. On the complexity of time table and multi-commodity flow problems. In *16th Annual Symposium on Foundations of Computer Science (sfcs 1975)*, pp. 184–193, 1975. doi: 10.1109/SFCS.1975.21.

[35] Feige, U. On maximizing welfare when utility functions are subadditive. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, pp. 41–50, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595931341. doi: 10.1145/1132516.1132523. URL https://doi.org/10.1145/1132516.1132523.

[36] Feige, U. and Vondrák, J. The submodular welfare problem with demand queries. *Theory of Computing*, 6(11):247–290, 2010. doi: 10.4086/toc.2010.v006a011. URL https://theoryofcomputing.org/articles/v006a011.

[37] Feldman, M., Gravin, N., and Lucier, B. Combinatorial walrasian equilibrium. STOC '13, pp. 61–70, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450320290. doi: 10.1145/2488608.2488617. URL https://doi.org/10.1145/2488608.2488617.

[38] Feldman, M., Gravin, N., and Lucier, B. Combinatorial auctions via posted prices, 2015. URL https://doi.org/10.1137/1.9781611973730.10.

[39] Feldman, M., Svensson, O., and Zenklusen, R. Online contention resolution schemes. SODA '16, pp. 1014–1033, USA, 2016. Society for Industrial and Applied Mathematics. ISBN 9781611974331.

[40] Groves, T. Incentives in teams. *Econometrica*, 41(4):617–631, 1973. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1914085.

[41] Hajiaghayi, M. T., Kleinberg, R. D., and Sandholm, T. Automated online mechanism design and prophet inequalities, 2007. URL http://www.aaai.org/Library/AAAI/2007/aaai07-009.php.

[42] Hsu, J., Morgenstern, J., Rogers, R., Roth, A., and Vohra, R. Do prices coordinate markets? 15(1):84–88, sep 2016. doi: 10.1145/2994501.2994508. URL https://doi.org/10.1145/2994501.2994508.

[43] Karakostas, G. Faster approximation schemes for fractional multicommodity flow problems. *ACM Trans. Algorithms*, 4(1), mar 2008. ISSN 1549-6325. doi: 10.1145/1328911.1328924. URL https://doi.org/10.1145/1328911.1328924.

[44] Khot, S., Lipton, R. J., Markakis, E., and Mehta, A. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica*, 52(1): 3–18, 2008. doi: 10.1007/s00453-007-9105-7. URL https://doi.org/10.1007/s00453-007-9105-7.

[45] Kleinberg, R. and Weinberg, S. M. Matroid prophet inequalities, 2012. URL https://doi.org/10.1145/2213977.2213991.

[46] Krengel, U. and Sucheston, L. Semiamarts and finite values. *Bulletin of The American Mathematical Society - BULL AMER MATH SOC*, 83, 10 1977. doi: 10.1090/S0002-9904-1977-14378-4.

[47] Krengel, U. and Sucheston, L. On semiamarts, amarts, and processes with finite value, 1978.

[48] Krysta, P. and Vöcking, B. Online mechanism design (randomized rounding on the fly). In Czumaj, A., Mehlhorn, K., Pitts, A., and Wattenhofer, R. (eds.), *Automata, Languages, and Programming*, pp. 636–647, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-31585-5.

[49] Lavi, R. and Swamy, C. Truthful and near-optimal mechanism design via linear programming. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pp. 595–604, 2005. doi: 10.1109/SFCS.2005.76.

[50] Lehmann, B., Lehmann, D., and Nisan, N. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, EC '01, pp. 18–28, New York, NY, USA, 2001. Association for Computing Machinery. ISBN 1581133871. doi: 10.1145/501158.501161. URL https://doi.org/10.1145/501158.501161.

[51] Li, S. Obviously strategy-proof mechanisms, November 2017. URL https://www.aeaweb.org/articles?id=10.1257/aer.20160425.

[52] Mirrokni, V., Schapira, M., and Vondrak, J. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, EC '08, pp. 70–77, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605581699. doi: 10.1145/1386790.1386805. URL https://doi.org/10.1145/1386790.1386805.

[53] Myerson, R. B. Optimal auction design. *Mathematics of Operations Research*, 6 (1):58–73, 1981. ISSN 0364765X, 15265471. URL http://www.jstor.org/stable/3689266.

[54] Naor, J. S. and Buchbinder, N. Proceedings of the 2022 annual acm-siam symposium on discrete algorithms (soda). Philadelphia, PA, 2022. Society for Industrial and Applied Mathematics. doi: 10.1137/1.9781611977073. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611977073.

[55] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. *Algorithmic Game Theory*. Cambridge University Press, USA, 2007. ISBN 0521872820.

[56] Pavlov, G. Optimal mechanism for selling two goods. *University of Western Ontario, Department of Economics, UWO Department of Economics Working Papers*, 11, 01 2010. doi: 10.2202/1935-1704.1664.

[57] Rubinstein, A. and Singla, S. Combinatorial prophet inequalities. In Klein, P. N. (ed.), *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pp. 1671–1687. SIAM, 2017. doi: 10.1137/1.9781611974782.110. URL https://doi.org/10.1137/1.9781611974782.110.

[58] Rubinstein, A., Wang, J. Z., and Weinberg, S. M. Optimal single-choice prophet inequalities from samples. In Vidick, T. (ed.), *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pp. 60:1–60:10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi: 10.4230/LIPICS.ITCS.2020.60. URL https://doi.org/10.4230/LIPIcs.ITCS.2020.60.

[59] Samuel-Cahn, E. Comparison of threshold stop rules and maximum for independent nonnegative random variables, 1984.

[60] Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961. ISSN 00221082, 15406261. URL http://www.jstor.org/stable/2977633.

[61] Vondrak, J. 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, pp. 67–74, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605580470. doi: 10.1145/1374376.1374389. URL https://doi.org/10.1145/1374376.1374389.