Εθνικο Μετσοβιο Πολυτεχνειο
Τμημα Ηλεκτρολογων Μηχανικων Και Μηχανικων Υπολογιστων

ΤΟΜΕΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ, ΤΕΧΝΟΛΟΓΙΑΣ & ΥΠΟΛΟΓΙΣΤΩΝ
ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ & ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΤΩΝ

# Βέλτιστη Επικοινωνία σε Φιλαλήθεις Μηχανισμούς

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

**Ιωάννη Γ. Αναγνωστίδη**

**Επιβλέπων**: Δημήτριος Φωτάκης
Αν. Καθηγητής Ε.Μ.Π.

Αθήνα, Οκτώμβριος 2020

**Εθνικο Μετσοβιο Πολυτεχνειο**
ΤΜΗΜΑ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
ΤΟΜΕΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ, ΤΕΧΝΟΛΟΓΙΑΣ & ΥΠΟΛΟΓΙΣΤΩΝ
ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ & ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΤΩΝ

# Βέλτιστη Επικοινωνία σε Φιλαλήθεις Μηχανισμούς

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

**Ιωάννη Γ. Αναγνωστίδη**

**Επιβλέπων**: Δημήτριος Φωτάκης
Αν. Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την $29^{\eta}$ Οκτωμβρίου 2020.

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

Αθήνα, Οκτώμβριος 2020.

...................................

**Ιωάννης Γ. Αναγνωστίδης**
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

## Περίληψη

Η παρούσα διπλωματική καταπιάνεται με το σχεδιασμό μηχανισμών με αποδοτική εξαγωγή πληροφορίας από τους παίκτες. Πιο συγκεκριμένα, ο βασικός μας στόχος είναι η ελαχιστοποίηση του συνολικού πλήθους από bits που λαμβάνουμε από τους χρήστες, χωρίς να θυσιάσουμε ταυτόχρονα άλλες επιθυμητές ιδιότητες του μηχανισμού όπως είναι η φιλαλήθεια καθώς και το social welfare που επιτυγχάνεται.

Αρχικά, δείχνουμε πως μπορούν να υλοποιηθούν ένα σύνολο από παραδειγματικές δημοπροσίες με ασυμπτωτικά βέλτιστη επικοινωνία. Ειδικότερα, ξεκινάμε τη μελέτη μας με single-parameter domains, περιλαμβάνοντας single item και multi-unit auctions. Για το πρότερο, δείχνουμε ότι το Vickrey's auction μπορεί να υλοποιηθεί με $1 + \epsilon$ αναμενόμενο communication complexity, για κάθε $\epsilon > 0$, θεωρώντας ότι κάθε valuation μπορεί να αναπαρασταθεί με ένα constant πλήθος από bits. Έτσι, προτείνουμε μία αποτελεσματική μέθοδο προσαρμογής της τιμής σε ένα English auction. Στη συνέχεια, σχεδιάζουμε αποδοτικά σχήματα κωδικοποίησης ώστε να πετύχουμε το ίδιο βέλτιστο φράγμα για multi-item auctions με additive valuations και constant πλήθος από αντικείμενα, αλλά και multi-unit auctions με unit demand bidders. Τα αποτελέσματά μας έπονται από απλές τεχνικές δειγματοληψίας και δεν απαιτούν επιπρόσθετες υποθέσεις για πρότερη γνώση ως προς τις παραμέτρους των παικτών.

Στο δεύτερο μέρος της δουλειάς μας εστιάζουμε στο γενικευμένο median μηχανισμό του Moulin σε μετρικούς χώρους εφοδιασμένους με την $L^1$ νόρμα. Ο συγκεκριμένος μηχανισμός έχει ιδιαίτερη σημασία στην περιοχή του Social Choice καθώς παρακάμπτει το Gibbard-Satterthwaite impossibility theorem για το πολύ φυσικό σενάριο των single-peaked preferences. Δείχνουμε ότι μία μέθοδος δειγματοληψίας μπορεί να προσεγγίσει το βέλτιστο κόστος με πολύ μικρό κομμάτι της συνολικής πληροφορίας. Το βασικό μας αποτέλεσμα βασίζεται στον χαρακτηρισμό μίας κατανομής, και πιστεύουμε ότι παρουσιάζει ανεξάρτητο ενδιαφέρον.

***Λέξεις Κλειδιά***— Σχεδιασμός Μηχανισμών, Facility Location Games, Communication Complexity, Αποδοτική Εξαγωγή Πληροφορίας, Vickrey's auction, English auction, Multi-Unit auctions, Multi-item auctions, Γενικευμένος Median, Δειγματοληψία

**Abstract**

This thesis is concerned with efficient preference elicitation in the field of Algorithmic Mechanism Design. More precisely, our goal is to minimize the elicited number of bits from the agents without sacrificing the other desired properties of the mechanism, namely the incentive compatibility guarantee and the social welfare. In this context, our main contribution is twofold.

First, we show how to implement a series of well-known mechanisms from Auction Theory with asymptotically optimal communication. Specifically, we initially turn our attention to single-parameter domains, namely single item and multi-unit auctions. For the former, we show that Vickrey's auction can be implemented with an expected communication complexity of at $1 + \epsilon$ bits – on average – per bidder, for any $\epsilon > 0$, assuming that the valuations can be represented with a constant number of bits. As a corollary, we provide a compelling method to increment the price in English auctions. Moreover, we design efficient encoding schemes in order to obtain the same asymptotic bound for multi-item auctions with additive bidders and a constant number of items, and for multi-unit auctions with unit demand bidders. Our results follow from simple sampling schemes and do not require any prior knowledge on the agents' parameters.

Moreover, we consider Moulin's generalized median mechanism on metric spaces endowed with the $L^1$ norm. This mechanism is of fundamental importance in the realm of Social Choice as it circumvents the Gibbard-Satterthwaite impossibility theorem for the natural setting of single-peaked preferences. We show that a sampling approximation of the median achieves a $1 + \epsilon$ approximation of the optimal social cost, for any $\epsilon > 0$, with a constant sample $c = c(\epsilon)$. Thus, our sampling approximation incurs an arbitrarily small error with an arbitrarily small fraction of the total information. Our main result is established based on the asymptotic characterization of a distribution, and could be of independent interest.

***Keywords***— Mechanism Design; Facility Location Games; Communication Complexity; Preference Elicitation; Vickrey's auction; English auction; Multi-unit auctions; Multi-item auctions; Generalized Median mechanism; Sampling

# Ευχαριστίες

Με την ολοκλήρωση της παρούσας διπλωματικής αισθάνομαι την ανάγκη να ευχαριστήσω ένα σύνολο από άτομα που μου προσέφεραν τη βοήθεια και τη συμπαράστασή τους, τόσο στα πλαίσια εκπόνησης της διπλωματικής όσο και συνολικά στον κύκλο των προπτυχιακών μου σπουδών.

Πρώτα από όλα, θα ήθελα να ευχαριστήσω βαθύτατα τον καθηγητή Δημήτρη Φωτάκη, ο οποίος επίβλεψε την διπλωματική μου. Πιο συγκεκριμένα, θέλω να τον ευχαριστήσω για την καθοδήγηση στα πρώτα ερευνητικά μου βήματα, για τον πολύτιμο χρόνο που διέθσε, αλλά και την υποστήριξη και τα ενθαρρυντικά του λόγια ακόμα και όταν οι ιδέες που πρότεινα δεν ήταν ιδιαίτερα καλές. Τον ευχαριστήσω επίσης που με έκανε μέλος του εργαστηρίου (CoReLab), προσφέροντας μία μοναδική εμπειρία στα τελευταία έτη των προπτυχιακών μου σπουδών. Επιπλέον, θα ήθελα να ευχαριστήσω θερμά τον Παναγιώτη Πατσιλινάκο, ο οποίος πραγματοποιεί διδακτορικό με τον κ. Φωτάκη, για τη βοήθεια που μου προσέφερε κατά την διάρκεια εκπόνησης της διπλωματικής. Θέλω να ευχαριστήσω επίσης τους καθηγητές Άρη Παγουρτζή και Ευάγγελο Μαρκάκη που αφιέρωσαν πολύτιμο χρόνο ως μέλη της εξεταστικής επιτροπής της διπλωματικής μου. Ορισμένο από το υλικό που περιέχει η παρούσα δουλειά παρουσιάστηκε στο φετινό Symposium of Algorithmic Game Theory, και επομένως, θέλω να ευχαριστήσω τους ανώνυμους reviewers για τα πολύ βοηθητικά σχόλια και τις υποδείξεις που βοήθησαν να βελτιωθεί η παρουσίαση της δουλειάς μας.

Επιπλέον, θέλω να ευχαριστήσω θερμά όλους τους φίλους και τους συμφοιτητές μου για τις πολύ όμορφες στιγμές που περάσαμε κατά τη διάρκεια των προπτυχιακών μου σπουδών. Θα ήθελα επίσης να αναγνωρίσω την ανεκτίμητη επίδραση και επιρροή όλων των υπέροχων καθηγητών που είχα την τύχη να γνωρίσω όλα αυτά τα χρόνια. Τέλος, θέλω να ευχαριστήσω την οικογένεια μου για την κατανόηση και την υπομονή, και κυρίως τον αδερφό μου Σωτήρη για την αδιάκοπη συμπαράσταση, βοήθεια και πίστη όλα αυτά τα χρόνια.

# Contents

# Κεφάλαιο 1

# Εισαγωγή στα Ελληνικά

Το communication complexity αποτελεί ένα κεντρικό ζήτημα στην περιοχή του Σχεδιασμού Μηχανισμών. Μία πρώτη γραμμή έρευνας μελετάει την επικοινωνία που απαιτείται για την προσέγγιση μίας αντικειμενικής συνάρτησης, όπως είναι τα αναμενόμενα έσοδα ή το social welfare· το πεδίο των combinatorial auctions αποτελεί ένα χαρακτηριστικό παράδειγμα όπου ισχυρά αρνητικά αποτελέσματα έχουν καθιερωθεί. Πράγαμτι, το communication complexity framework χρησιμοποιείται τυπικά για την απόδειξη lower bounds σε δημοπρασίες. Μία δεύτερη περιοχή μελέτης προσπαθεί να σχεδιάσει τη διαδικασία αλληλεπίδρασης με τους παίκτες έτσι ώστε η αποτελεσματική επικοινωνία να είναι εγγενές χαρακτηριστικό του μηχανισμού. Ακολουθώντας αυτή τη γραμμή έρευνας, στην παρούσα διπλωματική προσπαθούμε να αναπτύξουμε ένα φυσικό πλαίσιο για την ανάπτυξη ασυμπτωτικά βέλτιστων μηχανισμών σε παραδειγματικά περιβάλλοντα από το Σχεδιασμό Μηχανισμών.

Αυτή η έμφαση στην επικοινωνία ενός μηχανισμού εξηγείται από πολλούς λόγους. Πρώτον, υπάρχει η ανάγκη σχεδιασμού μηχανισμών με ισχυρές εγγυήσεις απόδοσης σε περιπτώσεις με περιορισμένη επικοινωνία και πιθανώς περιορισμένου χώρου δράσης. Επιπλέον, η εξαγωγή δεδομένων σε κατανεμημένα περιβάλλοντα μπορεί να είναι πολύ ακριβής. Έχει επίσης αναγνωρισθεί ότι ο όγκος της επικοινωνίας ενθυλακώνει και το βαθμό πληροφορίας που μεταφέρεται από τους χρήστες. Στο πλαίσιο αυτό, η εξαγωγή πληροφορίας απαιτεί υψηλό cognitive cost, και οι παίκτες ενδέχεται να διστάσουν να αποκαλύψουν ολοκληρωτικά την ιδιωτική τους πληροφορία. Πράγματι, ο περιορισμός της πληροφορίας που εξάγεται από τους χρήστες προσφέρει πιο ισχυρές εγγυήσεις ιδιωτικότητας για τους χρήστες.

Είναι φανερό λοιπόν ότι δεν είναι πρακτικό να λάβει ο μηχανισμός όλες τις προτιμήσεις των χρηστών, και ο σχεδιαστής επιθυμεί να λάβει αποκλειστικά το κομμάτι της συνολικής πληροφορίας που απαιτείται για την υλοποίηση του μηχανισμού. Αυτός ακριβώς είναι ο στόχος του preference elicitation, μίας κεντρικής γραμμής έρευνας στο Social Choice.

Σαν παράδειγμα θεωρούμε ένα single item auction και συγκεκριμένα τις αδυναμίες που παρατηρούνται στα πιο καθιερωμένα formats, δηλαδή στο Vickrey's auction και στο English auction. Αρχικά, είναι σημαντικά να επισημάνουμε ότι αν και κάθε μηχανισμός μπορεί να υλοποιηθεί με direct revelation – όπως έπεται από το revelation principle – αυτή η ισοδυναμία έχει αμφισβητηθεί για πολλούς λόγους. Πράγματι, εκτός των άλλων η δουλειά μας θα δείξει ότι το communication complexity ενός sealed-bid auction είναι μακριά από το βέλτιστο. Επιπλέον, παρά τις πολλές επιθυμητές ιδιότητες που χαρακτηρίζουν ένα sealed-bid auction, το English auction χρησιμοποιείται εκτενώς σε πρακτικές εφαρμογές, όπου και εμφανίζει καλύτερη απόδοση. Παρ' όλα αυτά, η υλοποίηση ενός second price rule μέσα από ένα τυπικό ascending auction απαιτεί στη χειρότερη περίπτωση εκθετική επικοινωνία.

Στο δεύτερο μέρος της δουλειάς μας ασχολούμαστε με την τεχνική της δειγματοληψίας, μία κλασσική μέθοδος που χρησιμοποιείται στην περιοχή του Social Choice. Πράγματι, έχει αναγνωρισθεί σε πολλές πρακτικές εφαρμογές ότι είναι ανέφικτο να συλλέξει κανείς όλες τις προτιμήσεις των παικτών. Επιπλέον, σε άλλες περιπτώσεις επιθυμούμε να προβλέψουμε το αποτέλεσμα ενός εκλογικού σχήματος χωρίς να πραγματοποιήσουμε τις εκλογές. Σε αυτό το πλαίσιο, χρησιμοποιούμε δειγματοληψία με σκοπό να προσεγγίσουμε τον generalized median mechanism, δείχνοντας ότι ακόμα και με 'μικρά' μεγέθη δείγματος μπορούμε να πετύχουμε σχεδόν βέλτιστη προσέγγιση.

## 1.1 Η Συνεισφορά μας

Η δική μας δουλειά προσφέρει πολλές νέες συνεισφορές ως προς το communication complexity που απαιτείται για την υλοποίηση θεμελιωδών μηχανισμών από το Σχεδιασμό Μηχανισμό. Αν και η δουλειά μας αποτελεί φυσική συνέχεια έρευνας στην περιοχή του preference elicitation, ισχυροποιούμε και βελτιώνουμε προηγούμενα αποτελέσματα με διάφορους τρόπους, ενώ επίσης πιστεύουμε ότι ανοίγουμε και νέες κατευθύνσεις για μελλοντική έρευνα.

### 1.1.1 Δημοπρασίες

Αρχικά, μελετάμε ένα σύνολο από δημοπρασίες, και δείχνουμε ότι σε κάθε περίπτωση μπορούμε ασυμπτωτικά να προσεγγίσουμε το βέλτιστο lower bound ως προς το average communication complexity, δηλαδή το πλήθος των bits που μεταφέρει ένας τυχαίος χρήστης στο μηχανισμό.

#### Single Item Auction

Πρώτα, σχεδιάζουμε μία παραλλαγή του English auction με μία δυναμικά προσαρμοζόμενη τιμή. Πιο συγκεκριμένα, σε κάθε γύρο υλοποιούμε ένα sub-auction – υλοποιημένο μέσω ενός αλγορίθμου $\mathcal{A}$ – σε ένα δείγμα αποτελούμενο από $c$ χρήστες. Με αυτόν τον τρόπο, η τιμή που ανακοινώνεται στο γύρο αποτελεί το market clearing price στο sub-auction. Η βασική μας παρατήρηση είναι ότι καθώς αυξάνουμε το μέγεθος του δείγματος το ποσοστό των χρηστών που θα παραμείνει ενεργό σταδικά μειώνεται. Ειδικότερα, δείχνουμε ότι για κάθε $\epsilon > 0$, θα υπάρχει ένα αρκετά μεγάλο μέγεθος δείγματος $c = c(\epsilon)$ έτσι ώστε το average communication complexity της δημοπρασίας που προτείνουμε να είναι $1+\epsilon$. Για το συγκεκριμένο επιχείρημα υποθέτουμε ότι κάθε valuation μπορεί να αναπαρασταθεί με ένα constant πλήθος από $k$ bits. Επιπλέον, είναι φανερό ότι κάθε δημοπρασία που προσδιορίζει τον χρήστη με το ψηλότερο valuation – με πιθανότητα 1 – απαιτεί τουλάχιστον 1 bit από κάθε χρήστη. Έτσι, η δημοπρασία που εισάγουμε προσεγγίζει ασυμπτωτικά αυτό το κάτω φράγμα. Πιο αυστηρά, στη συνέχεια παρουσιάζουμε το ascending auction μέσω δειγματοληψίας, καθώς και τις βασικές ιδιότητες που αποδείξαμε.

**Algorithm 1:** Ascending Auction μέσω Δειγματοληψίας

```
Input: Σύνολο N, μέγεθος δείγματος c, αλγόριθμος A
while |N| > c do
    S := τυχαίο δείγμα από c χρήστες από το N
    w := νικητής στο A(S)
    Ανακοίνωσε p := τιμή στο A(S)
    Ανανέωσε τους χρήστες: N := {i ∈ N \ S | v_i > p} ∪ {w}
end
if |N| = 1 then
    επέστρεψε w, p
else
    επέστρεψε A(N)
end
```

**Θεώρημα 1.1.1.** *Υποθέτοντας ότι οι παίκτες είναι φιλαλήθεις, το ascending auction μέσω δειγματοληψίας υλοποιεί με πιθανότητα 1 το VCG allocation rule.*

**Θεώρημα 1.1.2.** *Αν ο αλγόριθμος $\mathcal{A}$ υλοποιείται μέσω ενός sealed-bid auction, το ascending auction μέσω δειγματοληψίας είναι strategy-proof.*

**Θεώρημα 1.1.3.** *Έστω ότι $t(n; c, k)$ είναι το αναμενόμενο communication complexity του ascending auction μέσω δειγματοληψίας, με το $k$ να θεωρείται constant. Τότε, για κάθε $\epsilon > 0, \exists c_0 = c_0(\epsilon)$ ώστε $\forall c \geq c_0$,*

$$t(n; c, k) \lesssim n(1 + \epsilon). \tag{1.1}$$

Εδώ θα πρέπει να αναφέρουμε ότι ο συμβολισμός $f(n) \lesssim g(n)$ είναι ισοδύναμος με το $\lim_{n\to\infty} f(n)/g(n) \leq 1$. Είναι σημαντικό να επισημάνουμε επίσης διάφορες επιπρόσθετες ιδιότητες της δημοπρασίας μας. Αρχικά, ο τρόπος που εξάγουμε πληροφορία είναι έντονα μη συμμετρικός, καθώς όσο πιο 'κοντά' βρίσκεται ένας χρήστης στο να κερδίσει το αντικείμενο τόσο πιο πολύ πληροφορία πρέπει να αποκαλύψει. Επιπλέον, παρατηρούμε ότι ο μηχανισμός μας θα είχε πρακτικό ενδιαφέρον σε περιπτώσεις όπου δεν διαθέτουμε πρότερη πληροφορία ως προς τις προτιμήσεις των παικτών.

### Multi-unit auction

Επιπλέον, θεωρούμε αρκετές επεκτάσεις του προηγούμενου σεναρίου. Πρώτα, σχειδάζουμε multi-unit auction με unit-demand bidders που ανακτά και πάλι το ίδιο βέλτιστο φράγμα του $1+\epsilon$ bits, για οποιοδήποτε $\epsilon > 0$. Πιο συγκεκριμένα, σε ένα τυπικό ascending auction ανακοινώνεται σε κάθε γύρο μία τιμή. Αντιθέτως, η ιδέα μας είναι να μεταδίδουμε δύο διαφορετικές τιμές. Οι παίκτες που είναι υψηλότερα από την υψηλή τιμή θα θεωρούνται αυτόματα νικητές, ενώ οι χρήστες που είναι χαμηλότερα από τη χαμηλή τιμή θα αποκλείονται από τη συνέχεια της δημοπρασίας. Επισημαίνουμε ότι η τιμή που θα πληρώσουν οι χρήστες είναι κοινή, και καθορίζεται στο τέλος της δημοπρασίας. Έτσι, στον επόμενο γύρο είναι αρκετό να εστιάσουμε την προσοχή μας στους χρήστες που βρίσκονται ενδιάμεσα στις δύο τιμές. Η βασική μας παρατήρηση είναι ότι αν οι δύο τιμές είναι κοντά, τότε το ποσοστό των χρηστών που βρίσκονται ενδιάμεσα θα είναι αντίστοιχα μικρό. Με αυτόν τον τρόπο θα μπορούσαμε να εισάγουμε ένα αποδοτικό σχήμα κωδικοποίησης.

**Algorithm 2:** $\mathcal{M}(N,m)$: Multi-unit Auction μέσω Δειγματοληψίας

**Input**: Σύνολο παικτών $N$, πλήθος αντικειμένων $m := \gamma n$
Αρχικοποίησε τους νικητές $W := \emptyset$ και τους ηττημένους $L := \emptyset$
$p_h :=$ εκτιμώμενο άνω φράγμα στην τιμή
$p_\ell :=$ εκτιμώμενο κάτω φράγμα στην τιμή
Ανακοίνωσε $p_\ell$ και $p_h$
Ανανέωσε τους νικητές: $W := W \cup \{i \in N \mid v_i > p_h\}$
Ανανέωσε τους ηττημένους: $L := L \cup \{i \in N \mid v_i < p_\ell\}$
**if** $p_h = p_\ell$ **then**
    επέστρεψε $W, p_h$
**else**
    $N := N \setminus (W \cup L)$
    Ανανέωσε το πλήθος των αντικειμένων $m$
    επέστρεψε $\mathcal{M}(N,m)$
**end**

Η βασική δυσκολία της συγκεκριμένης μεθόδου έγκειται στον προσδιορισμό των δύο τιμών έτσι ώστε να είναι tight φράγματα ως προς το market clearing price. Για το σκοπό αυτό, προτείνουμε έναν πολύ φυσικό αλγόριθμο. Πιο συγκεκριμένα, θεωρούμε ένα δυαδικό δέντρο που αναπαριστά το discretized valuation space. Ξεκινάμε από τη ρίζα του δέντρου και σε κάθε επίπεδο εκτιμούμε ένα ακόμα bit μέσα από δειγματοληψία. Ειδικότερα, ρωτάμε κάθε χρήστη που ανήκει στο δείγμα αν το valuation του ξεπερνάει μία κατάλληλα επιλεγμένη τιμή. Ένα σημαντικό σημείο εδώ είναι ότι το δείγμα ενδέχεται να μην οδηγεί σε ξεκάθαρο branch, με την έννοια ότι περίπου μισοί χρήστες προτιμούν την κάθε επιλογή. Σε αυτην την περίπτωση είναι αρκετό η χαμηλή τιμή να ακολουθήσει το χαμηλό μονοπάτι και η ψηλή τιμή το υψηλό μονοπάτι. Με χρήση Chernoff bounds, δείχνουμε ότι αυτός ο αλγόριθμος θα τερματίσει με μεγάλη πιθανότητα σε δύο διαφορετικές στάθμες, έτσι ώστε το ποσοστό των χρηστών που βρίσκονται ενδιάμεσα να είναι μικρό. Είναι σημαντικό να αναφέρουμε εδώ ότι είναι πιθανό για έναν παίκτη που συμμέτεχει στο δείγμα η φιλαλήθεια να μην είναι dominant strategy. Έτσι, αποδεικνύουμε ένα weaker guarantee της μορφής του ex-post incentive compatibility. Πιο συγκεκριμένα, η δημοπρασία μας παρουσιάζει τις ακόλουθες ιδιότητες:

**Θεώρημα 1.1.4.** *Το multi-unit auction μέσω δειγματοληψίας είναι ex-post incentive compatible.*

**Θεώρημα 1.1.5.** *Έστω ότι $t(n; c, k)$ είναι το αναμενόμενο communication complexity του multi-unit auction μέσω δειγματοληψίας, με $k \in \mathcal{O}(n^{1-\ell})$, για κάποιο $\ell > 0$. Τότε, για κάθε $\epsilon > 0, \exists c_0 = c_0(\epsilon, k)$ ώστε $\forall c \geq c_0$,*

$$t(n; c, k) \lesssim n(1 + \epsilon). \tag{1.2}$$

### Multi-item auction

Στη συνέχεια, μελετάμε multi-item auctions με additive valuations. Αν υποθέσουμε ότι $m$ είναι το πλήθος των αντικειμένων, είναι φανερό ότι μπορούμε να πετύχουμε φράγμα $m(1 + \epsilon)$ bits για έναν τυχαίο χρήστη με χρήση $m$ δημοπρασιών σύμφωνα με το format που ορίσαμε προηγούνως. Παρ΄ όλα αυτά, η βασική μας παρατήρηση είναι ότι μπορούμε να μειώσουμε σημαντικά την επικοινωνία αν υλοποιήσουμε τις δημοπρασίες παράλληλα, σχεδιάζοντας ένα αποδοτικό σχήμα κωδικοποίησης (όπως συμβάινει για παράδειγμα στο Huffman coding). Έτσι, δείχνουμε ότι όταν το πλήθος των αντικειμένων είναι constant μπορούμε να πετύχουμε ξανά το φράγμα των $1 + \epsilon$ bits. Πιο συγκεκριμένα, δείχνουμε το ακόλουθο θεώρημα:

**Θεώρημα 1.1.6.** *Έστω $t(n; m, c, k)$ το αναμενό communication complexity που απαιτείται για την υλοποίηση $m$ παράλληλων ascending auctions μέσω δειγματοληψίας, με $m$ και $k$ να θεωρούνται constant. Τότε, υπάρχει κατάλληλο σχήμα κωδικοποίησης ώστε για κάθε $\epsilon > 0, \exists c_0 = c_0(\epsilon)$ ώστε $\forall c \geq c_0$,*

$$t(n; m, c, k) \lesssim n(1 + \epsilon). \tag{1.3}$$

Συνεπώς, παρουσιάζουμε δημοπρασίες με ασυμπτωτικά βέλτιστη επικοινωνία για τις ακόλουθες περιπτώσεις:

- single item auction
- multi-unit auction με unit-demand bidders
- mutli-item auction με additive valuations

Είναι σημαντικό να αναφέρουμε ότι με δεδομένα ορισμένα αρνητικά αποτελέσματα σε combinatorial auctions, υπάρχουν πολύ φυσικά εμπόδια για την επίτευξη αποδοτικής επικοινωνίας σε πιο γενικά domains.

### 1.1.2 Προσεγγίζοντας τον Median

Προχωρώντας στα αποτελέσματά μας στο Social Choice, θεωρούμε τον generalized median mechanism στο πλαίσιο facility location games. Ο συγκεκριμένος μηχανισμός έχει ιδιαίτερη σημασία καθώς παρακάμπτει το Gibbard-Satterthwaite impossibility theorem σε single-peaked domains. Η βασική μας συνεισφορά είναι να δείξουμε ότι ο median επιδέχεται σχεδόν βέλτιστη προσέγγιση μέσω δειγματοληψίας.

Σε πλαίσιο facility location games, ο median μηχανισμός τοποθετεί μία εγκατάσταση στη διάμεσο των θέσεων των παικτών. Έτσι, αν $x_i$ είναι η τοποθεσία του χρήστη $i$, τότε το social cost μίας εγκατάστασης στη θέση $x$ ορίζεται ως $\sum_{i=1}^{n} \mathrm{d}(x, x_i)$. Φυσικά, ο στόχος του μηχανισμού είναι να ελαχιστοποιήσει το social cost, δηλαδή να τοποθετήσει την εγκατάσταση όσο γίνεται κοντά στους χρήστες.

Αρχικά αναλύουμε το μονοδιάστατο median. Η πρώτη μας παρατήρηση είναι ότι ο median είναι sensitive ως προς το social cost. Πράγματι, δείχνουμε ότι ένας παίκτης μπορεί να επηρεάσει το social cost του median κατά $\mathcal{O}(1/n)$, όπου $n$ είναι το πλήθος των παικτών. Έτσι, καθώς το πλήθος των χρηστών αυξάνεται η επίδροση οποιοδήποτε παίκτη θα είναι αμελητέα. Αν και τέτοιες ιδιότητες έχουν παρατηρηθεί σε απλά σχήματα εκλογών, είμαστε οι πρώτοι που κάνουμε τέτοιου είδους σύνδεση σε facility location games. Πιο συγκεκριμένα, αποδείξαμε το ακόλουθο θεώρημα:

**Θεώρημα 1.1.7.** *Έστω ότι $x_{opt} \in \mathbb{R}$ είναι ο median της εισόδου και $x \in \mathbb{R}$ είναι μία τοποθεσία τέτοια ώστε το πολύ $\epsilon \cdot n$ χρήστες να βρίσκονται ανάμεσα στο $x$ και στο $x_{opt}$. Τότε, αν $D_{opt}$ είναι το ελάχιστο social cost, τοποθετώντας την τοποθεσία στο $x$ δίνει social cost $D$ ώστε*

$$D \leq D_{opt} \left(1 + \frac{4\epsilon}{1 - 2\epsilon}\right). \tag{1.4}$$

Στη συνέχεια, θεωρούμε την κατανομή που προκύπτει εκτιμώντας το rank $X_r$ – ως προς τον συνολικό πληθυσμό – του median του δείγματος, με πιθανότητα μάζας πιθανότητας ορισμένη ως

$$\Pr\left(X_r = \frac{i}{\kappa}\right) = \frac{\binom{\kappa - i}{\rho}\binom{\kappa + i}{\rho}}{\binom{2\kappa + 1}{2\rho + 1}}, \tag{1.5}$$

όπου $c = 2\rho + 1$ και $n = 2\kappa + 1$. Το προηγούμενο sensitivity επιχείρημα ανάγει την ανάκτηση μίας σχεδόν βέλτιστης προσέγγισης στην υψηλή συγκέντρωση της προηγούμενης κατανομής. Για το σκοπό αυτό, δείχνουμε ότι όταν ο πληθυσμός του μηχανισμού είναι υψηλός, η προηγούμενη κατανομή συμπεριφέρεται σαν μία μετασχηματισμένη κατανομή βήτα. Έτσι, δίνουμε ένα πολύ ακριβή χαρακτηρισμό της συγκέντρωσής της. Πιο συγκεκριμένα, αποδείξαμε τα ακόλουθα θεωρήματα:

**Θεώρημα 1.1.8.** *Για $\kappa \to \infty$, η κατανομή της τυχαίας μεταβλητής $X_r$ συγκλίνει σε μία μετασχηματισμένη βήτα κατανομή με την ακόλουθη συνάρτηση κατανομής* *πιθανότητας:*

$$f(t) = \frac{(2\rho+1)!}{(\rho!)^2 2^{2\rho+1}}(1-t^2)^\rho. \tag{1.6}$$

**Θεώρημα 1.1.9.** *Για κάθε $\epsilon > 0$ και $\delta > 0$, υπάρχει μία σταθερά $\rho = \rho(\epsilon, \delta)$ ώστε για κάθε $\rho \geq \rho_0$,*

$$\Pr(|X| \geq \epsilon) \leq \delta, \tag{1.7}$$

*όπου η τυχαία μεταβλητή $X$ ακολουθεί τη μετασχηματισμένη βήτα κατανομή που ορίσαμε στο προηγούμενο θεώρημα.*

**Θεώρημα 1.1.10.** *Ο προσεγγιστικός μονοδιάστατος median μηχανισμός έχει κατά μέση τιμή $1+\epsilon$ λόγο προσέγγισης ως προς το social cost του μηχανισμού που χρησιμοποιεί ολόκληρη την πληροφορία, για κάθε $\epsilon > 0$ και για σταθερό μέγεθος δείγματος $c = c(\epsilon)$, ενώ $n \to \infty$.*

Επισημάνουμε ότι δεν χρησιμοποιούμε πρότερη γνώση ως προς τις προτιμήσεις των παικτών. Επίσης, γενικεύουμε το αποτέλεσμα σε κάθε μετρικό χώρο εφοδιασμένο με την $L^1$ νόρμα. Πιο συγκεκριμένα, δείχνουμε το ακόλουθο θεώρημα:

**Θεώρημα 1.1.11.** *Ο προσεγγιστικός γενικευμένος median μηχανισμός έχει κατά μέση τιμή $1+\epsilon$ λόγο προσέγγισης ως προς το social cost του μηχανισμού που χρησιμοποιεί ολόκληρη την πληροφορία, για κάθε $\epsilon > 0$ και για σταθερό μέγεθος δείγματος $c = c(\epsilon)$, ενώ $n \to \infty$.*

Ένα ενδιαφέρον ερώτημα είναι κατά πόσο τα αποτελέσματά μας μπορούμε να γενικευτούν σε πιο γενικούς μετρικούς χώρους όπου ο median μπορεί να μην είναι καν ορισμένος.

# Chapter 2

# Introduction

Communication complexity has been a primary concern from the inception of Mechanism Design. The first consideration relates to the tractability of the communication exchange required to approximate an underlying objective function, such as the social welfare or the expected revenue; the domain of combinatorial auctions provides such an example where strong negative results have been established [DV13]. Indeed, the communication complexity framework is commonly used to establish lower bounds in Auction Theory. A second active area of research endeavors to design the interaction process so that efficient communication is an inherent feature of the mechanism [CS02]. Following this line of work, we aim to establish a natural framework for developing asymptotically optimal mechanisms in well-studied environments from Mechanism Design.

This emphasis is strongly motivated for a number of reasons. First, there is a need to design mechanisms with strong performance guarantees in settings with communication restrictions and possibly truncated action spaces, due to technical, behavioral or regulatory purposes [BNS07; AT06]. Moreover, extracting data from distributed parties can be burdensome, an impediment magnified in environments with vast participation. It has been also understood that the amount of communication captures the extent of information leakage from the participants. In this context, behavioral economists have recognized that soliciting information requires a high cognitive cost (e.g. [PUF99; Li17]) and bidders may be even reluctant to completely reveal their private valuation. Finally, truncating the information disclosure would provide stronger information privacy guarantees [SDX11] for the agents.

Indeed, it should be intuitively clear that having every agent communicate all of her preferences is impractical, and the designer aims to elicit only the relevant parts of the information. This endeavor lies at the heart of *preference elicitation*, a central theme in Social Choice and Mechanism Design that has engendered a vast and diverse literature; we refer to the pivotal works of Conitzer [Con09] and Oren et al. [OFB13], as well as references therein.

As a motivating example, we consider the single item auction, and in particular, the shortcomings of the most well-established formats, namely the *sealed-bid* and the *English* auction. First, it is important to point out that although every mechanism can be simulated with direct revelation – as implied by the revelation principle, this equivalence has been heavily criticized in the literature of Economics, not least due to the communication cost of revealing the entire valuation space. Indeed, our work will show that the communication complexity of Vickrey's sealed bid auction [Vic61] is suboptimal. Moreover, despite the theoretical appeal of Vickrey's auction, the ascending or English auction exhibits superior performance in practice [Aus04; KHL87; KL93; AM06], for reasons that mostly relate to the simplicity, the transparency and the privacy guarantees of the latter format. However, a faithful implementation of Vickrey's rule through a standard English auction requires – in the worst case – exponential communication and indeed, time complexity since the auctioneer has to increment the price by a single bit. In principle, the lack of prior knowledge on the agents' valuations would dramatically impede its performance.

One of the issues we address in this thesis is how to increment the price in an ascending auction, without any prior knowledge, so that the communication cost is minimized and the desirable properties of each format are retained. More broadly, we apply sampling techniques in order to establish mechanisms with asymptotically optimal communication complexity guarantees, without sacrificing the social welfare and the incentive properties of the interaction process. In particular, we employ random samples of agents and we either request the full information, or we query on whether their valuations exceed a particular threshold. In this way, our mechanism elicits – asymptotically – only the necessary information in order to implement the optimal allocation rule.

In the second part of our thesis we are concerned with a particular branch of preference elicitation; namely sampling approximations [CEG95; DB15], a standard technique employed in Social Choice. Indeed, it has been recognized that in many real-world scenarios it may be infeasible to gather the preferences from all of the agents; online surveys serve as such an example. Moreover, in many applications we want to predict the outcome of a voting rule without actually holding the election for the entire population of voters; for instance, we are quite familiar with polls and exit polls in political elections and beyond. In this context, we employ a sampling framework in order to approximate the celebrated median mechanism, which circumvents the Gibbard-Satterthwaite impossibility theorem in the natural domain of *single-peaked* preferences. This setting is motivated – among others – in political spectrum theory and facility location games. Surprisingly, we show that even a sample of constant size can yield an allocation with near-optimal social cost.

## 2.1 Related Work

Communication efficiency has been a cardinal desideratum in the literature of Algorithmic Mechanism Design. The first consideration relates to the interplay between communication constraints and incentive compatibility; in particular, Van Zandt [Van07] articulated conditions under which they can be studied separately, while the authors in [Rei84; FS09] investigated the communication overhead induced in truthful implementations, i.e. the communication cost of truthfulness. In a closely related direction, Blumrosen et al. [BNS07] (see also [MT14]) considered the design of optimal single-item auctions under severely bounded communication: every bidder can only transmit a limited number of bits. One of their key results was a 0.648 social welfare approximation for 1-bit auctions and uniformly distributed valuations. In addition, the design of optimal – with respect to the obtained revenue – bid levels in English auctions was addressed in [Dav+05], where the authors had to posit known prior distributions.

Moreover, the solution concept of efficient preference elicitation has induced a significant amount of research in the field of Social Choice. In particular, Segal [Seg07] provided bounds on the communication required to realize a social choice rule through the notion of *budget sets*, with applications in resource allocation tasks and stable matching. Furthermore, the boundaries of computational tractability and the strategic issues that arise were investigated by Conitzer and Sandholm in [CS02], while the same authors established in [CS05] the worst-case number of bits required to execute common voting rules. Efficient aggregate preference in social networks was considered in [DN13], where they elicited preferences from a small subset of critical nodes in the network. The trade-off between accuracy and information leakage in facility location games was tackled by Feldman et. al [FFG16], where they investigated the behavior of truthful mechanisms with truncated input space – *ordinal* and *voting* information models – and constitutes the main focus of our work as well. Finally, our approximation scheme is founded on Moulin's generalized median rule [Mou80] (see also [Bla48]).

## 2.2 Overview of Contributions

Our thesis provide several new insights with regards to the communication complexity required to implement fundamental mechanisms from Mechanism Design. Although our work constitutes a natural continuation of research in preference elicitation, we strengthen and improve prior results along several lines, while we believe that our work also opens several interesting avenues for future research. Specifically, our contribution lies in the following.

### 2.2.1 Communication in Auctions

We first consider a series of environments from Auction Theory, and we show that for every instance we can asymptotically match the lower bound with respect to the average communication complexity, i.e. the number of bits a bidder transmitted on average during the interaction process. We first design a variant of the English auction with an adaptive ascending price. Specifically, in every round of the ascending auction we simulate a sub-auction on a sample of size $c$ from the active agents. Then, the announced price of the round in the ascending auction will simply be the market clearing price in the simulated sub-auction. The main intuition behind our algorithm is that the fraction of the agents that will remain active in the following round scales as $\mathcal{O}(1/c)$; thus, as we augment the size of the sample most of the agents will withdraw from a given round of the auction. In other words, in our auction most of the agents will transmit in total a single bit. In fact, we show that for any $\epsilon > 0$, there will be a sufficiently large $c = c(\epsilon)$ such that the average communication complexity will be $1 + \epsilon$. For this argument we make the hypothesis that every valuation can be represented with a constant number of bits (e.g. 32 bits), although it can be relaxed. Naturally, this assumption is well motivated in single-parameter domains. Moreover, a trivial lower bound implies that every bidder has to transmit a single bit to the mechanism, so that the agent with the highest valuation can be recovered with probability 1. As such, our simple mechanism asymptotically matches this lower bound.

It is important to point out several features of our proposed single item auction. First, it introduces an interaction process that very naturally couples different formats. For instance, it would be natural to implement the sub-auction through a sealed-bid format. Moreover, the information elicitation is highly asymmetrical, as the closer a bidder is to winning the item, the more information she has to reveal. We argue that this property is actually desirable, and has been motivated for example in deferred acceptance auctions. We also remark that our auction would be of practical relevance in settings where the auctioneer does not possess adequate prior information with regards to the agents' valuations. Indeed, a strong feature of our mechanism is that no prior information is required to recover the optimal bound, unlike some prior work [Dav+05]. Therefore, we establish a compelling method to increment the price of an English auction, strongly supplementing the findings of Blumrosen, Nisan, and Segal [BNS07] in a symmetric model of communication.

Furthermore, we consider several extensions of the previous setting. First, we design a multi-unit auction with unit-demand bidders that obtains the same bound of $1 + \epsilon$ bits – on average – per agent for any $\epsilon > 0$. In particular, in a typical ascending auction the auctioneer maintains a single price per round; in contrast, we propose a natural variant in which we announce two separate prices. The agents that are above the high price are automatically declared winners and guarantee obtaining the good at the end of the auction, while the agents that lie below the low price will withdraw from the remainder of the auction. We remark that the price of the good will be determined only at the end of the auction, (and lies between the two announced prices), a feature which is essential in order to guarantee the incentive compatibility of the mechanism. Thus, we simply recurse on the agents that lie in-between. Our main observation is that if the announced prices are "tight", these agents will constitute only a small fraction of the total participation. This would allow us to introduce a very efficient encoding scheme. More precisely, the agents that are above the high price will have to send a bit of 0, the agents that are below the low price a bit of 1 – or vice-versa, while the in-between agents some arbitrary 2-bit code, so that the encoding is non-singular. Notice that the overhead incurred from using 2-bits will be negligible, as long as the fraction of the agents in-between in small. It should be noted that this idea is reminiscent to standard techniques in Information Theory, such as the Huffman coding [Knu85].

The crux of the aforementioned method is to determine the two prices so that they are tight upper and lower bounds with respect to the market clearing price. To this end, we propose a very natural algorithm. In particular, consider a binary tree that represents the discretized valuation space. On a high level, we commence from the root of the tree and at each level we "learn" an additional bit through sampling. In particular, we simply query on whether their valuations exceed the price that corresponds to the current node. A subtle point here is that the sample could provide a rather "ambiguous" answer, in the sense that roughly half the agents in the sample prefer each choice. However, we notice that in every such node it suffices to let the upper price follow the upper path, while the lower price the lower path. This process will terminate at two separate leaves, so that with high probability the fraction of the agents in-between is arbitrarily small. More precisely, we prove this through standard Chernoff bounds. However, we remark that unlike our proposed single item auction, this mechanism sacrifices the dominant strategy equilibria due to the sampling phase; indeed, we show a weaker incentive compatibility guarantee through the notion of ex-post Nash equilibria. Finally, although we introduced this method in the context of multi-unit auctions, our algorithm determines rankings of an arbitrary unsorted list with optimal communication [Hoa61], and could be of independent interest.

Next, we consider the setting of multi-item auctions with additive valuations. If we let $m$ denote the number of items, it is clear that we can obtain a bound of $m(1+\epsilon)$ bits per bidder if we simply perform $m$ sequential auctions according to the previously introduced format. Yet, our main insight here is that we can substantially truncate this complexity through a simultaneous implementation. In particular, we design a very efficient encoding scheme that employs a very simple property of our single item auction, namely, a random agent will most likely withdraw from a given round of the auction. In light of this, a sequential implementation corresponds to a very inefficient coding scheme as we map events with very different measures to codes with the same length. In contrast, a simultaneous implementation of the $m$ auctions, where all of the $m$ rounds are executed in parallel, allows us to design an encoding scheme that exploits this asymmetry. Thus, we show that when the number of items $m$ is a constant, we can again guarantee the bound of $1 + \epsilon$ bits per bidder.

As a result, we design auctions with asymptotically optimal communication for the following exemplar settings:

- Single item auction,
- Multi-unit auction with unit demand bidders.
- Multi-item auction with additive valuations.

It should be noted that in light of well-established lower bounds in combinatorial auctions [NS06], there are natural impediments when one studies more involved settings. Yet, we believe that our results are of practical significance due to their simplicity and their communication efficiency.

### 2.2.2 Information Requirements of the Median

Turning our attention to Social Choice, we consider the generalized median mechanism in the context of facility location games. This mechanism has receive substantial attention given that it circumvents Gibbard-Satterthwaite impossibility theorem for the well-motivated setting of single-peaked preferences. Our main contribution is to show that a sampling approximation of the median obtains almost the same social cost guarantee.

To be more precise, we first analyze the one-dimensional median in which the allocated facility is placed to the median of the reported instance. Our first observation is that unlike the median itself, the social cost of the median presents a surprising sensitivity. Indeed, we prove that a unilateral deviation from a single agent can only alter the social cost by a factor of $\mathcal{O}(1/n)$, where $n$ is the number of agents. Thus, as the number of agents increases the impact of any player – with respect to the cumulative distances from the allocation – will be gradually negligible. Although properties of this kind have been observed in simple voting schemes in prior works (e.g., see [DB15]), we are the first make such as a connection in facility location games.

Next, we consider the distribution that arises from estimating the rank – with respect to the entire population – of the sample's median. The previous sensitivity argument reduces obtaining a near-optimal guarantee to showing that this distribution is concentrated around the actual median. However, standard techniques appear to be of no use in this context. For this reason, we establish concentration by showing that when the participation is large, the aforementioned distribution behaves as a transformed beta distribution. As such, we are able to provide a very precise characterization of its concentration. We believe that this result is of independent interest as it supplements several characterizations in the regime of ranking distributions.

As a result, we show that for any $\epsilon > 0$, a sample of size $c = c(\epsilon)$ guarantees a $1+\epsilon$ social cost with respect to the full information mechanism. We stress that we do use any prior knowledge on the agents' valuations. We also establish the same characterization in every metric space endowed with the $L^1$ norm. An intriguing question that arises from our results relates to the guarantees of a sampling approach in more general metric spaces where the median may not be defined (e.g. a circle). It appears to us that our insights can be extended to broader settings. Indeed, our approach fits very naturally to the recently introduced framework of distortion [ABP15; SE17; MW19; Kem20; GHS20].

**Broader Context** More broadly, communication complexity has been a primary consideration in Game Theory. A series of works have established tractable communication procedures in order to reach an approximate Nash equilibrium in two-player games [GP12; Czu+18; BR17]. Moreover, important work by Nisan and Segal [NS06] has illustrated the limitations, and in particular the exponential communication requirements in the domain of submodular bidders, as well as in combinatorial allocation problems – even when 2 players compete for $m$ indivisible items. There has been also extensive research devoted in designing incentive compatible and efficient preference elicitation mechanisms in combinatorial auctions [HS04; CS01; Blu+04].

**Roadmap** In chapter 2 we introduce some basic notions from Game Theory, as well as several fundamental mechanisms from Mechanism Design and Social Choice that are related to our work. In chapter 4 we introduce the framework of Communication Complexity, before we present our results in Auction Theory and facility location games in chapters 5 and 6 respectively.

# Chapter 3

# Algorithmic Mechanism Design

## 3.1 Introduction

Game Theory is the field that studies mathematical models of strategic interaction between rational and selfish agents; as such, it provides a rigorous framework to understand the phenomena we observe when decision-makers interact. The main assumption that characterizes the theory is that decision-makers pursue some well-defined objective – *rationality* hypothesis – based on their knowledge or expectations of the behavior of the other players. Ever since the inception of the field by Von Neumann's and Morgenstern [NMR44], Game Theory has played a vital role in Economics, Political Science, and more recently in Computer Science, underlined with the advent of the Internet.

The models employed in Game Theory are abstractions of a wide range of real-life phenomena. Indeed, the development of the theory was motivated by issues that arose in practical applications. Typical examples include oligopolistic and political competition, market equilibria, stability in large-scale economies, and many others. The fundamental question in Game Theory is how to formalize the meaning of "rationality" and "rational behavior", and how to argue in a rigorous sense about interacting and competing agents in the context of a "game". The first consideration has been addressed by the theory of Rational Behavior, through notions such as *utility* and *preference*, while for the latter issue various models have been proposed in the literature in order to capture different types of strategic situations.

Of course, we have extensive experience from participating in games, either directly or in many cases indirectly. An instance of the former scenario would be a game of chess, while indirect participation occurs, for example, when one has to decide on the most efficient route to reach a certain destination, which crucially depends on the amount of traffic congestion – i.e. the number of other commuters whose select among interfering routes. Before we proceed with a formal treatment of the fundamental notions of Game Theory, we present a simple and familiar example.

**Example 3.1.1 (Rock – Paper – Scissors).** *This game typically consists of* 2 *players with identical strategies. In each round, the players simultaneously form one of* 3 *possible shapes, namely rock, paper and scissors, each corresponding to an action available to each agent. According to the rules of the game, a player who decides to play rock will beat anyone who chose scissors, but will lose to anyone who has played paper. Furthermore, a player who chose scissors will beat a player who formed paper. On the other hand, if both players select the same shape, the round is considered to be tied.*

**Cooperative and Non-Cooperative Games** In every game-theoretic model the basic entity is the *player*. A player may be construed as a single individual or as a group of individuals making collective decisions. Based on this dichotomy, we distinguish two types of models, namely "non-cooperative" and "cooperative" respectively. It should be noted that in recent years, the overwhelming research focus – and indeed the topic of this thesis – has been on non-cooperative models.

In the following sections we are closely following the notation of Osborne and Rubinstein [OR94].

### 3.1.1 Rational Behavior

The models we consider in Game Theory assert that every agent acts *rationally* in the sense that based on the possible alternatives or actions at her disposal and the expectations about any unknown information, she selects her strategy after some optimization process. In the absence of uncertainty, rational choice consists of the following ingredients:

- A set of *actions* $A$ from which the agent makes a decision.
- A set of possible *consequences* $C$ for each action.
- A *consequence function* $g : A \mapsto C$ that maps every action to a consequence.
- A *preference relation* $\succsim$ over the consequences $C$.

In many cases the agents' preferences are specified through a *utility function* $U : C \mapsto \mathbb{R}$, which induces a preference relation $\succsim$ by the condition that $x \succsim y$ if and only if $U(x) \geq U(y)$. We remark that the preference relation is assumed to be complete, transitive, and reflexive.

For a given feasible set of actions $B \subseteq A$, a rational decision-maker selects an action $a^*$ that is feasible – i.e. $a^* \in B$ – and optimal in the sense that $g(a^*) \succsim g(a)$ for all $a \in B$. Equivalently, every rational agent solves the optimization problem $\max_{a \in B} U(g(a))$. However, in many settings individuals have to make decisions under conditions of uncertainty; typical causes of uncertainty include the following:

- Lack of information about the reasoning of the other players.
- Uncertainty due to the randomization in their environment or in the actions of other players.
- Imperfect knowledge about the parameters of the game, such as the imposed rules and the parameters of the environment.

The standard approach to address uncertainty or randomization is to assume that agents are maximizing over the expected value (von Neumann-Morgenstern utility) according to some prior beliefs or distribution over the unknown information. We remark that the rationality hypothesis has been extensively questioned by experimental psychologists.

### 3.1.2 Modeling a Game

In this section we provide a rigorous definition of a game, a central object in Game Theory that captures the model of interaction between competing agents. In particular, the *strategic form* is appropriate as a model in situations where the players perform their moves simultaneously. In such games, players select their strategies without knowing those of the other players. The canonical example in this context is the Rock – Papers – Scissors game we illustrated in our introduction. Formally, a simultaneous-move, strategic form game consists of the following:

1. A set of $n$ players $\{1, 2, \ldots, n\}$.
2. A set of possible strategies for each player $i$.
3. A utility function $u_i : S \mapsto \mathbb{R}$, where $S = S_1 \times S_2 \cdots \times S_n$.

In some cases instead of employing a utility function we model a player's preferences through a preference relation $\succsim_i$, such that $\forall s, s' \in S$ it follows that $s \succsim_i s' \iff u_i(s) \geq u_i(s')$. We will represent a strategic form game as the triple $\langle n, (S_i), (u_i) \rangle$.

**Representing Strategic Form Games** A crucial consideration in Game Theory is how to encode the aspects of the game through an efficient and succinct representation. In a so-called *matrix form*, we represent explicitly the value of the utility function for every possible strategy vector. Although this representation is quite standard, it is unsuited when the number of possible actions is exponential. Instead, application-specific representations such as *graphical* games or *congestion* games are typically employed.

### 3.1.3 Nash Equilibrium

The most commonly used solution concept in Game Theory is the notion of Nash equilibrium. This concept captures a steady state of the play of a strategic game in which every player holds the correct expectation about the other players' behavior. Yet, it does not attempt to consider the process by which a steady state is reached.

#### Pure Nash Equilibrium

**Definition 3.1.1** (Pure Nash Equilibrium)**.** *Let $\langle n, (S_i), (u_i)\rangle$ be a strategic form game. A strategy vector $s^*$ is said to be a pure Nash equilibrium of the game if for all $i \in [n]$,*

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \forall s_i \in S_i. \tag{3.1}$$

In words, a strategy vector constitutes a pure Nash equilibrium if no player has an incentive to unilaterally deviate from her strategy. Unfortunately, a pure Nash equilibrium may not exist. Consider for example the *matching-pennies* game, in which every player chooses either "head" or "tail". If the choices are different, the first player pays the second player one dollar; otherwise, the second player pays the first player one dollar. Even in this simple example it is easy to see that no pure Nash equilibrium exists. We remark that a matching pennies game is an example of a *strictly competitive* game given that the interests of the players are diametrically opposed, while in the assumed form it also constitutes a *zero-sum* game – the sum of every entry in the matrix of the game is zero.

**Mixed Nash Equilibrium** Let us denote with $\Delta(S_i)$ the set of probability distributions over the strategies in $S_i$, and $U_i$ the expected value of the utility of player $i$ given the *mixed strategies* of all players. For a given strategic game $\langle n, (S_i), (u_i)\rangle$, we define its mixed extension as the game $\langle n, (\Delta(S_i)), (U_i)\rangle$.

**Definition 3.1.2** (Mixed Nash Equilibrium)**.** *A mixed Nash equilibrium of a strategic form game is a pure Nash equilibrium of its mixed extension.*

Going back to the matching pennies game, notice that if both players randomize uniformly over their pure strategies the game reaches a mixed Nash equilibrium. In fact, Nash's celebrated work [Nas50] established that every finite game has an equilibrium point.

**Theorem 3.1.1** ([Nas50])**.** *Any game with a finite set of players and finite set of strategies has a mixed Nash equilibrium.*

Nash presented a very elegant proof of this theorem through an application of a *fixed point* theorem due to Kakutani [Kak41][1]. A question that arises immediately

[1]Similarly, Von Neumann's proof of the min-max theorem employed Brouwer's fixed point theorem.

from Nash's theorem is whether efficient algorithms for finding the equilibrium exist. This issue is of central importance given that the predictive power of an equilibrium concept crucially depends on that the agents may actually reach it after a reasonable amount of time. Unfortunately, there is strong evidence that computing a Nash equilibrium is computationally intractable even for a two-player general games, and more precisely it was shown to be PPAD-complete [DGP06].

**Mechanism Design** Mechanism Design [Hur45; Hur73; HR06; Arr12; Hur77; LLR89; MDH79] is a subfield of Game Theory that endeavors to design systems that implement desired social choices with good performance guarantees in strategic environments – assuming that every participant will act *rationally* in a game-theoretic sense. Naturally, taking into account incentive issues is crucial given that the agents' preferences are private, i.e. unknown to the mechanism. Typical environments of interest include the following:

- **Political Elections** [Mou16]: Every voter has her own preferences among a set of candidates, and the mechanism has to *aggregate* the agents' preferences in order to determine a set of winners.
- **Judgment Aggregation** [Lis12]: The goal is to determine the truth-value of logically related propositions that depend on individual *judgments* or *opinions*; as such, it is applicable in consensus reaching problems in multi-agent systems.
- **Auctions** [Mas92; Mye81; RS81]: An auctioneer has to dispose or *sell* a set of items to competing bidders.

Mechanism Design has emerged as the primary tool for rigorously studying and predicting the behavior of multi-agent systems, and has found a myriad of applications on the Internet where multiple parties with widely different goals operate and interact. In this context, the focus is to aggregate the preferences of the different participants into a social choice, while the main desideratum is that the algorithm will perform well assuming strategic and selfish behavior from each agent; in other words, the underlying mechanism should be *incentive-compatible* – non-manipulable – to potential misreports from the agents.

## 3.2 Social Choice

Social Choice is concerned with the evaluation of alternative methods of collective decision-making. The inception of this field can be traced back to antiquity, and the first organized communities where multiple individuals had to make decisions for their common cause. Indeed, a study of the principles of collective decision-making was already articulated by Aristotle in his book entitled *Politics*, dating back almost two and a half millenniums ago. Yet, a rigorous theoretical and mathematical investigation had to wait until the European Enlightenment, where the pioneering contributions of Marquis de Condorcet and Jean-Charles de Borda laid the foundations of modern Social Choice theory.

### 3.2.1 Condorcet's Paradox

Perhaps the most natural and common *voting rule* is the so-called *majority rule*, where each voter selects a single candidate among a set of alternatives. However, in 1785 Condorcet observed that the majority rule is problematic in the presence of three or more candidates. More precisely, let us denote with $a, b$, and $c$ the alternatives, and consider three voters $\{1, 2, 3\}$ with the following preferences:

(i) $a \succ_1 b \succ_1 c$

(ii) $b \succ_2 c \succ_2 a$

(iii) $c \succ_3 a \succ_3 b$,

where the notation $a \succ_i b$ implies that agent $i$ prefers candidate $a$ to $b$. In this setting, note that a majority of voters (1 and 3) prefers candidate $a$ to candidate $b$, a majority (1 and 2) prefers $b$ to $c$, and finally a majority (2 and 3) prefers $c$ to $a$. Thus, the "joint majority" choice is $a \succ b \succ c \succ a$, which is inconsistent. In other words, the method of pairwise majority voting may yield a *social preference cycle*. Indeed, for any chosen candidate in Condorcet's example there will be a majority of voters who would favor altering the outcome of the election. One of the logical implications of Condorcet's paradox is that when a majority cycle occurs in the set of social alternatives – there exists no *Cordorcet winner*, i.e. a feasible alternative which is undefeated by any other feasible alternative, the possibility of basing the social choice to a simple majority rule is excluded. It should be noted that Condorcet's paradox was taken from voting on economic policy, and it seems to have been inspired by an earlier work of Borda (1781)[2], who proposed what came to be known as the *Borda count*. In particular, this method assings for each voter a score of zero to the last ranked alternative, a score of one to the penultimate alternative, and so on until the top ranked candidate. These individuals scores are added for each candidate over all voters, and the candidate who earned the largest total score becomes the overall winner in the contest. Unfortunately, this procedure suffers from *strategic vulnerability*, since an elector has an incentive to place the strongest opponents to their favorite candidates at the foot of their submitted list. In fact, when Borda was confronted with this shortcoming he replied by saying that his scheme is "only intended for honest men" ([Bla58]).

[2]It should be noted that the same scheme was independently studied by Pierre-Simon de Laplace.

### 3.2.2 Impossibility Results

A large number of different voting rules have been suggested throughout the years by the pioneering studies of Marquis De Condorcet, Jean-Charles Borda, Charles Dodgson (better known by his pen name Lewis Carroll), Duncan Black, and many others. Important these celebrated works are, they were concerned exclusively with some specific voting scheme. In sharp contrast, Arrow developed a unified framework which allowed him to characterize all conceivable voting schemes. In particular, Arrow pioneered an axiomatic approach in Social Choice by imposing a set of natural axioms. Then, he showed that these axioms are logically incompatible.

Moreover, a central difficulty that arises in social choice is *strategic voting.* For instance, consider a voter $i$ with preferences $a \succ_i b \succ_i c$ who knows that candidate $a$ is unpopular and hence will not be selected; it is clear that such a voter has an incentive to strategically vote for $b$ instead of $a$. Unfortunately, strategic vulnerability is unavoidable for every reasonable voting rule, as stated by the celebrated *Gibbard-Satterthwaite* theorem.

Formally, consider a set of candidates $A$ and a set of $n$ voters $I$. Let us denote with $L$ the permutations – or *linear orderings* – of $A$, so that $\prec \in L$ is a total order – antisymmetric and transitive – on $A$. The preferences of a single voter $i$ are given by $\succ_i \in L$, where recall that $a \succ_i b$ implies that agent $i$ prefers $a$ to $b$.

**Definition 3.2.1.** *A mapping $F : L^n \mapsto L$ is called a social welfare function, while a mapping $f : L^n \mapsto A$ is called a social choice function.*

Thus, a social welfare function aggregates the preferences of all voters into a common preference, while a social choice function accumulates the preferences of all voters into a single candidate.

#### Arrow's Impossibility Theorem

The first axiom that will be implied in the forthcoming analysis is that every participant can express – or is *free* to express – any preference ordering, while the social welfare function has to be able to aggregate the profile of any set of profiles into a social preference ordering. The second axiom requires that the social welfare function faithfully reflects the unanimous preference expressed by all the individuals, as stated in the following definition:

**Definition 3.2.2.** *A social welfare function $F$ satisfies unanimity if for every $\prec \in L$, $F(\prec, \dots, \prec) = \prec$.*

The next axiom requires that two social states can be compared based solely on the preferences of the individuals on the two alternatives.

**Definition 3.2.3.** *A social welfare function satisfies independence of irrelevant alternatives if the social preferences between any two alternatives $a$ and $b$ depends* *only on the voters' preferences between $A$ and $B$. More precisely, for every $a, b \in A$ and every $\prec_1, \ldots, \prec_n, \prec'_1, \ldots, \prec'_n \in L$, if we denote $\prec = F(\prec_1, \ldots, \prec_n)$ and $\prec' = F(\prec'_1, \ldots, \prec'_n)$ then $a \prec_i b \iff a \prec'_i b$ for all $i$ implies that $a \prec b \iff a \prec' b$.*

The final axiom is that there should be no *dictator* in the society. In particular, we introduce the following definition:

**Definition 3.2.4.** *A voter $i$ is a dictator in a social welfare function $F$ if for all $\prec_1, \ldots, \prec_n \in L$, $F(\prec_1, \ldots, \prec_n) = \prec_i$.*

Before we prove Arrow's theorem we state a preliminary lemma.

**Lemma 3.2.1.** *Let $\succ_1, \ldots, \succ_n$ and $\succ'_1, \ldots, \succ'_n$ be two agents' profiles such that for every player $i$, $a \succ_i b \iff c \succ'_i d$. Then, $a \succ b \iff c \succ' d$, where $\succ = F(\succ_1, \ldots, \succ_n)$ and $\succ' = F(\succ'_1, \ldots, \succ'_n)$.*

**Theorem 3.2.1** ([MAS+14])**.** *Every social welfare function over a set of more than 2 candidates ($|A| \geq 3$) that satisfies unanimity and independence of irrelevant alternatives is a dictatorship.*

*Proof.* Fix some $F$ that satisfies unanimity and independence of irrelevant alternatives. Take any $a, b \in A$ with $a \neq b$, and for every $i \in \{0, 1, \ldots, n\}$ define a preference profile $\pi^i$ in which exactly the first $i$ players rank $a$ above $b$. By unanimity, in $F(\pi^0)$ we have $b \succ a$, while in $F(\pi^n)$ we have $a \succ b$. In the sequence of profiles $\pi^0, \pi^1, \ldots, \pi^n$ the ranking between $a$ and $b$ flips, so for some $j$ we have that $F(\pi^{j-1})$, $b \succ a$, while in $F(\pi^j)$, $a \succ b$. It suffices to show that $j$ is a dictator.

Take any $c, d \in A$ with $c \neq d$. We will show that if $c \succ_j d$ then $c \succ d$, where $\succ = F(\succ_1, \ldots, \succ_n)$. Indeed, consider some alternative $e$ which is different from $c$ and $d$. For $i < j$ move $e$ to the top in $\succ_i$, for $i > j$ move $e$ to the bottom in $\succ_i$, and for $j$ move $e$ so that $c \succ_j e \succ_j d$. It follows from independence of irrelevant alternatives that we have not changed the social ranking between $c$ and $d$. Moreover, notice that the players' preferences for the ordered pair $(c, e)$ are identical to their preferences for $(a, b)$ in $\pi^j$, but the preferences for $(e, d)$ are identical to the preferences for $(a, b)$ in $\pi^{j-1}$ and hence, using Lemma 3.2.1 we obtain that $c \succ e$ and $e \succ d$, implying that $c \succ d$ (by transitivity).[3] ■

We remark that the method of simple majority voting we discussed previously satisfies all of these conditions, except that the generated social preference relation lacks the general assurance of transitivity by virtue of the Condorcet paradox. Arrow's impossibility theorem serves as a clear indicator of the need to the rigorous scrutiny in search of resolutions of the identified contradiction. A notable sufficient condition on the agents' preferences was proposed by Black [Bla48], in the form of *single-peaked* preferences. This assumption has a simple geometric representation to the effect that the social alternatives can be represented by a one-dimensional

[3] This proof is due to Geanakoplos [Gea05].

variable. Black's theorem is the first possibility result of this nature in Social Choice, and it served as the prelude of the modern development of the theory of voting.

### Gibbard-Satterthwaite Theorem

It turns out that Arrow's theorem dramatically impedes the design of strategy-proof mechanisms. Let us first formally define the notion of strategic manipulation.

**Definition 3.2.5.** *A social choice function $F$ can be strategically manipulated by voter $i$ if for some $\prec_1, \ldots, \prec_n \in L$ and some $\prec'_i \in L$ we have that $a \prec_i a'$, where $a = F(\prec_1, \ldots, \prec_i, \ldots, \prec_n)$ and $a' = F(\prec_1, \ldots, \prec'_i, \ldots, \prec_n)$. $F$ is called incentive compatible if it cannot be manipulated.*

**Definition 3.2.6.** *A social choice function $F$ is monotone if $F(\prec_1, \ldots, \prec_i, \ldots, \prec_n) = a \neq a' = F(\prec_1, \ldots, \prec'_i, \ldots, \prec_n)$ implies that $a' \prec_i a$ and $a \prec'_i a'$.*

**Proposition 3.2.1.** *A social choice function is incentive compatible if and only if it is monotone.*

The obvious example of an incentive compatible social choice function over two candidates is the majority rule. However, when the number of alternatives is larger than 2, only trivial social functions are incentive compatible.

**Definition 3.2.7.** *Voter $i$ is a dictator in social choice function $F$ if for all $\prec_1, \ldots, \prec_n \in L$, $\forall b \neq a, a \succ_i b \implies F(\prec_1, \ldots, \prec_n) = a$. $F$ is called a dictatorship if there exists a dictator $i$.*

**Theorem 3.2.2** (Gibbard-Satterhwaite [Gib73; Sat75])**.** *Let $F$ be an incentive compatible social choice function onto $A$, where $|A| \geq 3$; then, $F$ is a dictatorship.*

Given that the validity of the Gibbard-Satterthwaite theorem on the ubiquity of strategic manipulation in voting schemes, a vast literature has been developed in search for either an escape route from the Gibbard-Satterthwaite impossibility theorem, or directions in which the theorem may be generalized. Indeed, Mechanism Design was funded exactly in order to devise rules so that individuals will actually express their true preferences, even when they are acting rationally.

## 3.3 Mechanisms with Money

In the previous section, we modeled an agent's preference as an ordering of the possible candidates. In particular, $a \succ_i b$ implies that $i$ prefers $a$ to $b$, but it does not capture the degree that $a$ is preferred to $b$. In this section, we redefine our setting, and we assume that the preference of an agent $i$ is given by a *valuation function* $v_i : A \mapsto \mathbb{R}$, where $v_i(a)$ represents the value that $i$ assigns to alternative $a$; this value can be thought in terms of some currency. Then, if player $i$ additionally commits some quantity of money $m$, $i$'s *utility* is $u_i = v_i(a) - m$. Utilities of this form are called *quasilinear preferences*.

### 3.3.1 Vickrey's Second Price Auction

Before we proceed to the general setting, we study a simple example, namely a single item auction. More precisely, consider that a single (indivisible) item to be disposed for sale to one of $n$ players. Every player $i$ has a scalar value $v_i$ which represents her maximum willingness to pay for the item. Thus, if $i$ wins the item at a price of $p$, her utility for obtaining the item will be $v_i - p$; naturally, the utility of every player that did not obtain the item is 0. If we place this scenario into the terms of our general setting, the set of alternatives $A$ here is the set of possible winners. The only natural social choice would be to allocate the item to the player who values it the most. However, the main challenge is that the auctioneer does not know in advance the valuations of the players and hence, the payments we impose should incentivize *truthful* behavior by the agents. Let us first focus on a simple auction format, namely a *sealed-bid auction*, which occurs as follows:

(1) Every bidder $i$ privately communicates a bid $b_i$ to the auctioneer

(2) The auctioneer determines the agent that obtains the item

(3) The auctioneer determines the payment, i.e. the selling price

Given that we allocate the item to the highest bidder, the only degree of freedom we have is the payment we impose. In this context, let us first consider some unsuitable choices of payment.

- **No payment**: We simply allocate the item for free to the player that declared the highest bid $b_i$. It is clear that this method is susceptible to strategic behavior, as every player has an incentive to exaggerate her actual valuation.
- **Pay you bid**: The winner of the auction will simply pay the declared bid. This system is also open to manipulation, since a player who sincerely reports her valuation will always obtain a total utility of 0. Thus, every agent will attempt to declare a somewhat lower value. In particular, if $i$ knows the value of the second highest bid, her optimal response is to declare a value just above it.

Our final observation motivates the following format:

**Definition 3.3.1** ([Vic61]). *Vickrey's second price auction: Allocate the item to the highest bidder $i$, and let her pay $p = \max_{j \neq i} b_j$.*

**Proposition 3.3.1** ([Vic61]). *Let $b_1, \dots v_b, b'_i \in \mathbb{R}$, $u_i$ and $u'_i$ the utility of player $i$ for a bid of $v_i$ and $b_i$ respectively in a Vickrey's auction. Then, $u_i \geq u'_i$, that is misreporting can only reduce her utility.*

*Proof.* Assume that player $i$ wins the item with a report of $v_i$, and the second highest report valuation is $p^*$. Then, $u_i = v_i - p^* \geq 0$. For any attempted manipulation $b_i \geq p^*$, it follows that $u_i' = u_i$. On the other hand, if $b_i < p^*$, then $i$ would lose the item and $u_i' = 0 \leq u_i$.

Moreover, if $i$ loses the item by bidding $v_i$, then $u_i = 0$. Let $j \neq i$ represent the winner of the auction, with $b_j \geq v_i$. If $b_i < b_j$, $i$ would still lose the item and hence $u_i' = u_i$. Likewise, for $b_i \geq b_j$, $i$ would win but overpay for the item, and her utility would be $u_i' = v_i - b_j \leq 0 = u_i$. ∎

### 3.3.2 Incentive Compatible Mechanisms

A mechanism has to select an alternative from a set $A$, as well as decide on payments. The preference of each player $i$ is modeled by a valuation function $v_i : A \mapsto \mathbb{R}$, where $v_i \in V_i$ for commonly known set of possible valuations functions for player $i$.

**Definition 3.3.2.** *A mechanism is a social choice function $f : V_1 \times \cdots \times V_n \mapsto A$ and a vector of payment functions $p_1, \ldots, p_n$, where $p_i : V_1 \times V_n \mapsto \mathbb{R}$ is the amount that players $i$ pays to the mechanism.*

**Definition 3.3.3.** *A mechanism $(f, p_1, \ldots, p_n)$ is called dominant strategy incentive compatible if for every player $i$, every $v_1 \in V_1, \ldots, v_n \in V_n$ and every $v_i' \in V_i$, if $a = f(v_i, v_{-i})$ and $a' = f(v_i', v_{-i})$, then $v_i(a) - p_i(v_i, v_{-i}) \geq v_i(a') - p_i(v_i', v_{-i})$.*

Intuitively, any player $i$ prefers telling the truth to the mechanism rather than any possible misreport.

### 3.3.3 Vickrey-Clarke-Groves Mechanisms

The main result in this section is an incentive compatible mechanism for the most natural social choice function, namely the social welfare. More precisely, the social welfare of an alternative $a \in A$ is defined as $\sum_i v_i(a)$.

**Definition 3.3.4** (VCG Mechanism)**.** *A mechanism $(f, p_1, \ldots, p_n)$ is called a Vickrey-Clarke-Groved (VCG) mechanism if $f$ maximizes the social welfare, i.e. $f(v_1, \ldots, v_n) \in \arg\max_{a \in A} \sum_i v_i(a)$, and for some functions $h_1, \ldots, h_n$, where $h_i : V_{-i} \mapsto \mathbb{R}$, we have that for all $v_1 \in V_1, \ldots, v_n \in V_n$, $p_i(v_1, \ldots, v_n) = h_i(v_{-i}) - \sum_{j \neq i} v_j(f(v_1, \ldots, v_n))$.*

Thus, in a VCG mechanism every player pays an amount equal to the sum of the values of all other players. As a result, when this term is added to her own value, the sum becomes exactly the social welfare. Intuitively, this mechanism aligns all players' incentives with the goal of maximizing social welfare. The other term in the payment has no implication for player $i$ since it does not depend on her actions.

**Theorem 3.3.1** ([Vic61; Cla71; Gro73])**.** *Every VCG mechanism is dominant strategy incentive compatible.*

*Proof.* Fix $i, v_{-i}, v_i$ and $v_i'$. It suffices to show that for player $i$ with valuation $v_i$, the obtained utility when declaring $v_i$ is not less than the utility when declaring $v_i'$. Let $a = f(v_i, v_{-i})$ and $a' = f(v_i', v_{-i})$. The utility of $i$ when declaring $v_i$ is $v_i(a) + \sum_{j \neq i} v_j(a) - h_i(v_{-i})$; similarly, the utility when declaring $v_i'$ is $v_i(a') + \sum_{j \neq i} v_j(a') - h_i(v_{-i})$. But, given that $a = f(v_i, v_{-i})$ maximizes the social welfare over all alternatives, $v_i(a) + \sum_{j \neq i} v_j(a) \geq v_i(a') + \sum_{j \neq i} v_j(a')$, concluding the proof. ■

### 3.3.4 Clarke Pivot Rule

A fundamental question in the VCG mechanism is how to select the functions $h_i$. In particular, we impose the following natural conditions:

**Definition 3.3.5.** *A mechanism is individually rational if for every $v_i, \ldots, v_n$ we have that $v_i(f(v_1, \ldots, v_n)) - p_i(v_1, \ldots, v_n) \geq 0$, i.e. players always get non-negative utility from the mechanism.*

**Definition 3.3.6.** *A mechanism has no positive transfers if for every $v_1, \ldots, v_n$ and every $i$, $p_i(v_1, \ldots, v_n) \geq 0$, i.e. no player is ever paid.*

In this context, the following choice of $h_i$'s provides the aforementioned properties:

**Definition 3.3.7** (Clarke's pivot rule)**.** *The choice $h_i(v_{-i}) = \max_{b \in A} \sum_{j \neq i} v_i(b)$ is called the Clarke pivot payment. Under this rule the payment of player $i$ is $p_i(v_1, \ldots, v_n) = \max_b \sum_{j \neq i} v_i(b) - \sum_{j \neq i} v_i(a)$, where $a = f(v_1, \ldots, v_n)$.*

Intuitively, every player $i$ pays an amount equal to the total social welfare loss that incurs to the other agents.

**Proposition 3.3.2.** *A VCG mechanism with Clarke pivot payments makes no positive transfers. Moreover, if $v_i(a) \geq 0$ for every $v_i \in V_i$ and $a \in A$, then it is also individually rational.*

*Proof.* Let $a = f(v_1, \ldots, v_n)$ be the alternative maximizing $\sum_j v_j(a)$ and $a'$ be the alternative maximizing $\sum_{j \neq i} v_j(a')$. The utility of a player $i$ is $v_i(a) + \sum_{j \neq i} v_j(a) - \sum_{j \neq i} v_j(a') \geq \sum_j v_j(a) - \sum_j v_j(b) \geq 0$; thus, the individual rationality follows. To show no positive transfers, notice that $\sum_{j \neq i} v_i(a') \geq \sum_{j \neq i} v_i(a)$. ■

### 3.3.5 Examples

**Single Item Auction** The Vickrey auction is a special case of a VCG mechanism with the Clarke pivot rule. Indeed, finding the player with the highest value is exactly equivalent to maximizing $\sum_i v_i(i)$, since only a single player gets non-zero value. Moreover, the Clarke pivot rule gives exactly Vickrey's second price auction.

**Multi-unit Auctions** In a multi-unit auction, $m$ identical units of some good are sold. We consider the simple case where every bidder is interested in obtaining a single unit. Maximizing social welfare means allocating the items to the $m$ highest bidders, and in the VCG mechanism with the pivot rule, each of them should pay the $m+1$-th highest reported price.

**Public Project** Consider a government project – e.g. infrastructure such as libraries, schools, roads, and so on – that has a *commonly* known cost $C$, and is valued by every citizen $i$ at a *privately* known value $v_i$. The project will be undertook if and only if $\sum_i v_i > C$, the constituents value the project more than its actual cost. According to the VCG mechanism with the Clarke pivot rule a player $i$ – with a positive value for the project $v_i \geq 0$ – will pay a non-zero amount only if $i$ is pivotal, i.e. $\sum_{j \neq i} v_j \leq C$ but $\sum_j v_j > C$, in which case $i$ will pay $p_i = C - \sum_{j \neq i} v_j$. It should be noted that $\sum_i p_i < C$ and hence, the collected payments do not cover the project's cost.

### 3.3.6 The Revelation Principle

Our focus so far has been on *direct revelation* protocols, that is every agent reveal the entire private valuation to the mechanism. This emphasis in Mechanism Design is usually justified by the revelation principle, as stated in the following theorem.

**Theorem 3.3.2** (Revelation Principle)**.** *For every mechanism $\mathcal{M}$ in which every participant has a dominant strategy, there is an equivalent direct revelation dominant strategy incentive compatible mechanism $\mathcal{M}'$.*

*Proof.* For the proof we use a simulation argument (Figure 3.1). By assumption, every participant with any private valuation $v_i$ has a dominant strategy $s_i(v_i)$ in the given mechanism. Consider a direct revelation mechanism $\mathcal{M}'$ to which the agents delegate the responsibility of playing their dominant strategies. In particular, $\mathcal{M}'$ accepts sealed bids $b_1, \ldots, b_n$ from the players, and it submits the bids $s_1(b_1), \ldots, s_n(b_n)$ to mechanism $\mathcal{M}$. Finally, $\mathcal{M}'$ outputs the outcome of mechanism $\mathcal{M}$. It is clear that $\mathcal{M}'$ is dominant strategy incentive compatible. Indeed, if an agent $i$ submits a bid other than $v_i$, it would result in playing a strategy other than $s_i(v_i)$ in $\mathcal{M}$, which can only decrease $i$'s utility. ∎

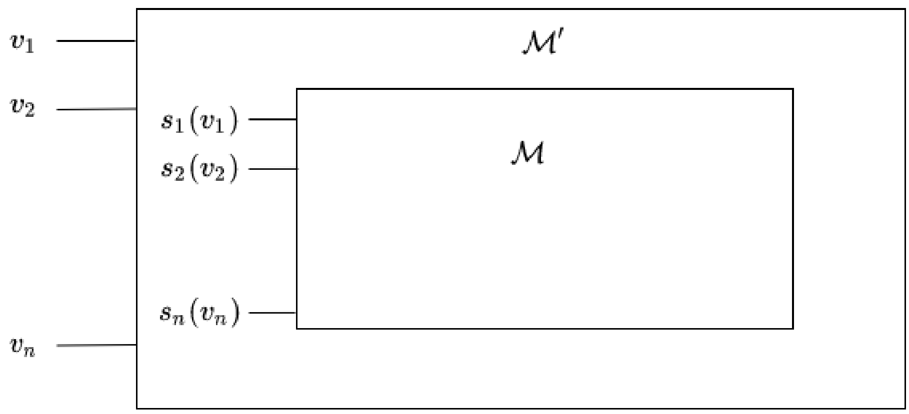


Figure 3.1: Construction of the direct revelation mechanism $\mathcal{M}$' given a mechanism $\mathcal{M}$ with dominant strategies.

## 3.4 Ascending Implementations in Auction Theory

### 3.4.1 English Auction

Let us first illustrate an ascending auction through an example. Consider that we want allocate $m$ identical items to *unit-demand* bidders. The basic idea in an ascending format is to increment the proposed selling price until the demand equals the supply. More precisely, consider some parameter $\epsilon$ that serves as the step of the ascend – typically assumed a priori fixed; the *English* auction consists of the following steps:

1. Set the initial price $p := 0$
2. Let $S_0$ the set of agents
3. For $t = 1, 2, \ldots$
    - Let $S_t$ the agents that remain active for a price of $p + \epsilon$
    - If $|S_t| > m$ increment $p$ by $\epsilon$
    - Otherwise, sell the items at price $p$ to the agents in $S_t$; if there are items leftover, sell them to an arbitrary subset of the bidders of $S_{t-1} \setminus S_t$ at price $p$

An English auction is an indirect auction, that is it does not explicitly elicit the valuation of a bidder. Naturally, there is a number of reasons for employing an indirect auction format [Li17; Aus04; AM06; MW82]:

1. **Privacy Guarantees**: The winner in an English auction does not reveal her private valuation – but only that it exceeds the second highest bidder, while parameter $\epsilon$ also conceals the private information of every other participant.

2. **Cognitive Cost**: In many cases deriving the actual valuation can be burdensome for a bidder and an ascending auction is *easier* for the participants. This point is connected to the notion of *obvious strategy-proofness*.

3. **Transparency**: Indeed, ascending auctions are typically open and a bidder is able to keep track of the progress, unlike sealed-bid auctions.

4. **Higher Revenue**: Empirical studies suggest that indirect formats induce higher revenue, a point that relates to the so-called "bidding wars".

Return to the analysis of the English auction, it is important to point that sincere bidding in an iterative auction means that a player answers to all queries truthfully. Our first proposition is that the incentive guarantee of the English auction is almost as good as the Vickrey's auction.

**Proposition 3.4.1.** *In an English auction, sincere bidding is a dominant strategy for every bidder, up to an $\epsilon$.*

Thus, as $\epsilon \to 0$, sincere bidding becomes a dominant strategy for every player. Finally, it should be clear that the social welfare of the outcome is within $m \cdot \epsilon$ of the optimal social welfare; thus, as $\epsilon \to 0$ the English auction obtains the optimal welfare.

### 3.4.2 Additive Valuations

As a second example in the regime of ascending auctions, consider a set of $m$ items – in general non-identical. Every bidder $i$ has a private valuation $v_{ij}$ for every item $j$, while our main assumption in this subsection is that the valuations are additive, that is for any non-empty bundle of items $S$, the obtained value for agent $i$ is

$$v_i(S) = \sum_{j \in S} v_{ij}. \tag{3.2}$$

The additivity assumption implies that a bidder's value for an item does not depend on the other received items; hence, there are no *substitutes* or *complements*. In this case, the direct revelation solution for this domain is straightforward: perform a separate Vickrey auction for each of the $m$ items. It is easy to show that the induced mechanism is dominant strategy incentive compatible, as an instance of the general VCG mechanism. However, the main observation here is that the incentive guarantee does not transfer for parallel English auctions. Indeed, consider the following example:

**Example 3.4.1.** *Consider two bidders and two items with* $v_{11} = 5, v_{12} = 3, v_{21} = 4$*, and* $v_{22} = 1$*. If both players bid sincerely, the first bidder will win both items at prices* 4 *and* 1 *respectively. However, it is possible that the second player adopts the following action: If the first player starts bidding on the first item, retaliate by bidding on both items until termination; otherwise, bid sincerely for the entire auction. If the first agent bids sincerely, her received utility will be* 0*, given the action of her opponent. On the other hand, if the first agent completely abandons the first item and commits only to the second item, the obtained utility is* 2*. Thus, for the first player sincere bidding is not a best response.*

As it is illustrated in the previous example, iterative auctions have a much richer action space than sealed-bid formats, necessitating a relaxed notion incentive compatibility. More precisely, consider $n$ agents with $V_1, \dots, V_n$ representing the sets of possible private valuations, and $A_1, \dots, A_n$ the sets of possible actions. A *strategy* $s_i$ is a function from $V_i$ to $A_i$; for instance, sincere bidding in an iterative auction is a potential strategy.

**Definition 3.4.1.** *A strategy profile* $(s_1, \dots, s_n)$ *is an ex-post Nash equilibrium if for every bidder* $i$ *and valuation* $v_i \in V_i$*, the action* $s_i(v_i)$ *is a best response to every action profile* $s_{-i}(v_{-i})$*.*

**Definition 3.4.2.** *We call a mechanism ex-post incentive compatible if sincere bidding constitutes an ex-post Nash equilibrium.*

**Proposition 3.4.2.** *Parallel English auctions are ex-post incentive compatible.*

### 3.4.3 Unit-Demand Bidders

Continuing our survey in ascending auctions, in this subsection we introduce a setting that strictly generalizes the multi-unit scenario with identical items. More precisely, we consider a set $U$ of $m$ non-identical items, while every bidder $i$ has a private valuation $v_{ij}$ for each item $j$. The main assumption here is that every bidder $i$ has *unit-demand*, meaning that her value for a non-empty bundle of items $S \subseteq U$ is

$$v_i(S) := \max_{j \in S} v_{ij}. \tag{3.3}$$

Notice that the case of identical items is the special case where $v_{ij} = v_i$, i.e. independent of $j$. A prominent example captured by this model are housing and automobile markets [SS71].

**Direct Revelation Solution** As a thought experiment, assume that all of the private valuations $v_{ij}$ are known to the mechanism. The main observation is that welfare-maximization in this scenario is tantamount to the maximum-weight bipartite matching of the induced graph. More precisely, consider a weighted bipartite graph $(N, U, E, w)$, where the vertex set $N$ corresponds to the $n$ bidders, the vertex set $U$ corresponds to the $m$ items, and for every $i \in N, j \in U$, there is an edge $(i, j) \in E$ with weight $v_{ij}$. It is clear that a matching in the induced graph corresponds to particular allocations; subsequently, every maximum-weight matching yields a welfare-maximizing allocation. Thus, the solution can be obtained in polynomial time via Linear Programming, or several celebrated combinatorial algorithms such as the *Hungarian method*. Next, employing the VCG mechanism yields that every bidder should be charged its *externality*, that is the welfare loss incurred to the other agents by her presence, leading to a dominant strategy incentive compatible mechanism that maximizes the social welfare in polynomial time.

In the following analysis, the goal will be to establish an ascending auction with analogously strong guarantees. In this context, first notice that the termination condition in an ascending auction is essentially that "supply equals demand". Yet, in this setting it's not trivial to even guarantee this property. For this reason, we will introduce the concept *Walrasian equilibrium*, which is in a sense the natural outcome of an ascending auction.

### Walrasian Equilibrium

**Definition 3.4.3.** *A Walrasian equilibria (WE) with unit-demand bidders is a price vector* $\mathbf{p}$ *and an allocation* $M$ *such that:*

1. *For every bidder* $i$, $M(i) \in \arg\max\{v_{ij} - \mathbf{p}(j)\}$;
2. *An item* $j \in U$ *remains unmatched in* $M$ *only if* $\mathbf{p}(j) = 0$.

It is important to point that WE is essentially the natural termination point in an ascending auction; hence, it is crucial to characterize WE and correlate them with the VCG outcome in order to understand the limits of an ascending implementation.

**Theorem 3.4.1** (First Welfare Theorem)**.** *Consider an auction with unit-demand bidders; if* $(\mathbf{p}, M)$ *is a Walrasian equilibria, then* $M$ *induces a welfare-maximizing allocation.*

*Proof.* Let $M^*$ be a welfare maximizing allocation. Given that $(\mathbf{p}, M)$ constitutes a Walrasian equilibria, it follows that for every bidder $i$,

$$v_i(M(i)) - \mathbf{p}(M(i)) \geq v_i(M^*(i)) - \mathbf{p}(M^*(i)). \tag{3.4}$$

Summing this inequality over all biders $i$ yields that

$$\sum_{i=1}^{n} v_i(M(i)) - \sum_{i=1}^{n} \mathbf{p}(M(i)) \geq \sum_{i=1}^{n} v_i(M^*(i)) - \sum_{i=1}^{n} \mathbf{p}(M^*(i)). \tag{3.5}$$

Finally, the claim follows from $\sum_{i=1}^{n} \mathbf{p}(M^*(i)) \leq \sum_{i=1}^{n} \mathbf{p}(M(i))$. ∎

The first welfare theorem can be thought of as a justification for the *efficient-markets* hypothesis, which maintains that market prices fully reflect all available information. In fact, this theorem can be extended for an $\epsilon$-WE, in which case the welfare of $M$ is within $\epsilon \cdot \min\{n, m\}$ of the maximum possible.

The next step is to correlate the payment predicted by the VCG mechanism to price vectors that participate in a Walrasian equilibria. It should be noted that while the former are uniquely defined, the latter are not. Indeed, notice that for a single-item auction, the Walrasian price vectors are $[v_2, v_1]$, where $v_1$ and $v_2$ are the first and second highest valuations among the agents respectively.

First, consider some bidder $i$, and let us denote with $M^{-i}$ the matching that leaves $i$ unmatched. Then, the payment of $i$ according to the VCG mechanism is

$$p_i = \sum_{k \neq i} v_k(M^{-i}(k)) - \sum_{k \neq i} v_k(M(k)). \tag{3.6}$$

Thus, the VCG outcome induces a price vector $\mathbf{q}$ of item prices. Indeed, if an item $j$ is sold we define its price as the VCG payment of the agent that is assigned to the item; otherwise, we let $\mathbf{q}(j) = 0$.

**Proposition 3.4.3.** *Consider an instance with unit-demand bidders. If $\mathbf{q}$ is the VCG price vector and $\mathbf{p}$ a Walrasian price vector, it follows that $\mathbf{q}(j) \leq \mathbf{p}(j)$ for every item $j$.*

*Proof.* Let us denote with $M$ the allocation computed by the VCG mechanism. If an item $j$ is not sold, it follows that $\mathbf{q}(j) = 0 \leq \mathbf{p}(j)$. Moreover, consider an item $j$ allocated to a bidder $i$, and let $M^{-i}$ represents a welfare-maximizing allocation that leaves bidder $i$ unmatched. For every bidder $k \neq i$, it follows that

$$v_k(M(k)) - \mathbf{q}(M(k)) \geq v_k(M^{-i}(k)) - \mathbf{q}(M^{-i}(k)). \tag{3.7}$$

Summing this inequality over all bidders $k \neq i$ yields that

$$\sum_{k \neq i} v_k(M(k)) - \sum_{k \neq i} \mathbf{q}(M(k)) \geq \sum_{k \neq i} v_k(M^{-i}(k)) - \sum_{k \neq i} \mathbf{q}(M^{-i}(k)). \tag{3.8}$$

Finally, rearranging these terms yields that

$$\mathbf{p}(j) \geq \sum_{k \neq i} v_k(M^{-i}(k)) - \sum_{k \neq i} v_k(M(k)) = \mathbf{q}(j). \tag{3.9}$$

■

**Proposition 3.4.4.** *Consider an instance with unit-demand bidders. If $M$ and $\mathbf{q}$ represent the allocation and the induced price vector of the VCG outcome, it follows that $(\mathbf{q}, M)$ is a Walrasian equilibria.*

As a result, we deduce that the VCG outcome coincides with the "smallest" Walrasian equilibria, i.e. the one that is component-wise in the payment vector smaller than any other. This property is crucial in order to obtain an incentive compatibility guarantee for the ascending auction, since it is necessary to terminate with the VCG payments.

**The Crawford-Knoer Auction** In this context, we analyze an auction that was proposed by Crawford and Knoer [CK81] (see also [ZSP08; DGS86; Leo83]).

1. Initialize the price of every item $j$ to $\mathbf{p}(j) = 0$
2. Query the demand of every unassigned agent $j \in D_i(\mathbf{p}+\epsilon) = \arg\max\{v_i(j) - (\mathbf{p}(j) + \epsilon)\}$.
3. If no unassigned bidder submits a bid, terminate the auction
4. Otherwise, pick some unassigned bidder $i$ with preferred item $j$, and assign item $j$ to $i$. If the item was previously assigned to some bidder $k$, mark $k$ as unassigned and set $\mathbf{p}(j) := \mathbf{p}(j) + \epsilon$.

Notice that in the Crawford-Knower (henceforth CK) auction a bidder cannot relinquish the item, unless some other agent outbids her. Thus, once an item has been requested, it will be assigned to some agent at termination. Moreover, assuming sincere reporting the CK auction terminates in a pseudo-polynomial number of iterations. Before we establish the main theorems concerning the CK auction, we first state some preliminary lemmas.

**Lemma 3.4.1.** *Assuming sincere reporting, the CK auction terminates at an $\epsilon$-WE $(\mathbf{p}, M)$.*

This lemma follows easily from the definition of the Walrasian equilibria and the termination condition of the CK auction. Moreover, recall that this lemma implies that the allocation obtains a welfare within $\epsilon \cdot m$ of the maximum possible.

**Lemma 3.4.2.** *Let $\mathbf{p}$ the price vector of the CK auction at termination with sincere reporting, and $\mathbf{q}$ the induced price vector of the VCG mechanism. Then, $\mathbf{p}(j) \leq \mathbf{q}(j) + \epsilon \cdot \min\{n, m\}$.*

Thus, the prices computed by the CK auction are no higher that the VCG prices, up to some small error.

**Lemma 3.4.3.** *Let $\mathbf{p}$ the price vector of the CK auction at termination with sincere reporting, and $\mathbf{q}$ the induced price vector of the VCG mechanism. Then, $\mathbf{p}(j) \geq \mathbf{q}(j) - \epsilon \cdot \min\{n, m\}$.*

Thus, the previous lemmas directly imply the following theorem:

**Theorem 3.4.2.** *Assuming sincere reporting, the outcome of the CK auction simulates – up to $\epsilon$ – the VCG outcome under truthful revelation.*

Finally, we state the incentive guarantee of the Crawford-Knower auction:

**Theorem 3.4.3.** *The CK auction is ex-post incentive compatible up to* $2\epsilon \cdot \min\{n, m\}$ *error.*

## 3.5 Mechanisms Without Money

The Gibbard-Satterthwaite impossibility theorem asserts that on *general domain* preferences only dictatorial rules can be implemented in dominant strategies. However, in many scenarios the agents' preferences are not completely unrestricted, as it was assumed in the previously studied voting context. For instance, in the previous section we showed that monetary transfers allows for a rich strategy-proof mechanism, namely the VCG mechanism. Nonetheless, there are many important applications where there are significant incentive issues, but monetary transfers between the mechanisms and the agents are considered infeasible, unethical or illegal. Mechanism Design without money is relevant for designing and understanding methods for environments such as voting, organ donation and school choice.

### 3.5.1 Facility Location Games

The facility location class of problems models a large number of important problems that occur in practice, ranging from traditional areas such as urban planning, to more recent ones such as computer networking. In particular, we are concerned with the placement of facilities that will supply services to clients in order to minimize some function of cost; this cost function may be defined in different ways, depending on each specific problem.

**Uncapacitated Facility Location Problem** Let $F$ be a set of facilities, $T$ a set of terminals, $c_f$ opening costs for each facility $f \in F$, and $d_{tf}$ connection costs for connecting terminal $t \in T$ to facility $f \in F$. The problem is to find a subset of facilities to open and establish connections from terminals to this subset such that the sum of all costs is minimized. An integer program formulation for this problem is presented below:

$$\text{minimize} \sum_{f \in F} c_f y_f + \sum_{f \in F} \sum_{t \in T} d_{tf} x_{tf},$$

$$\text{subject to} \sum_{f \in F} x_{tf} = 1, \forall t \in T,$$
$$y_f \leq x_{tf}, \forall f \in F, \forall t \in T,$$
$$y_f, x_{tf} \in \{0, 1\}, \forall f \in F, \forall t \in T.$$

In this formulation, $y_f$ is a boolean variable that indicates whether a facility $f$ is opened, while the boolean variable $x_{tf}$ indicates whether terminal $t$ is connected to facility $f$. Many possible variants may arise from this problem. In particular, there might be capacities or quotas associated with each facility, or the facilities can be selected in any point on a metric space [Shm00; Li13]. In Game Theory, we are interested in a scenario where the clients select their strategies independently in order to minimize their connection cost. Mechanism Design is employed in order to design games with desired properties, or be used to measure the inefficiency arising from selfish behavior (Price Of Anarchy). Indeed, there have been significant advancements in designing truthful mechanisms for such games [DMV05; LSW05; FT10; LS04; PT03; Tha10].

**The Median Mechanism** Here we analyze the median mechanism, an allocation rule proposed by Moulin [Mou80]. Taking a step back, the Gibbard-Satterthwaite theorem implies that if the preference of the agents on the set of alternatives can be any ordering (a condition usually referred to as *unrestricted domain*), then apart from dictatorial rules every decision scheme will include an incentive for strategic misreporting of preferences for at least one preference profile. Moulin investigated a relaxation of the unrestricted domain assumption when the preferences of the agents are all single-peaked along the real line. In this context, it is natural to consider only the voting schemes where every agent simply announces her peak-alternative. Moulin showed that every strategyproof, efficient and anonymous voting scheme is obtained by adding $n-1$ fixed ballots to the $n$ voters' ballots – here $n$ represents the number of agents – and then choosing the median of this larger set of ballots.

We are interested in the implications of this characterization in facility location games with strategic agents. Consider a metric space $(M, \mathrm{d}(\cdot,\cdot))$, and let $\mathbf{x}_i \in M$ denote the preferred location of agent $i$. In this case, the social cost of an allocation $F$ is defined as

$$\mathrm{SC}(F) = \sum_{i=1}^{n} \min_{f \in F} \mathrm{d}(\mathbf{x}_i, f). \tag{3.10}$$

We will first present the median mechanism and its properties for a one-dimensional instance, and we will then extend our analysis for a high-dimensional Euclidean space. For the one-dimensional case, the location of agent $i$ will be denoted with $x_i$. The median mechanism allocates a facility on the median of the reported instance. It should be clear that assuming truthful reporting this mechanism obtains optimal social cost. Moreover, we can prove the following proposition:

**Proposition 3.5.1.** *The one-dimensional median mechanism is strategyproof.*

*Proof.* Consider some agent $i$, some arbitrary vector $x'_{-i}$ indicating the reports from the other agents and $x^* = \text{median}(x'_{-i}, x_i)$ the location of the facility when $i$ reports truthfully. If $x_i < x^*$, then $i$ could only alter the allocation with a report of $x'_i > x^*$; however, this could only increase her distance from the allocated facility. A similar argument applies for $x_i \geq x^*$, concluding the proof. ■

Next, we extend the analysis on any metric space $(\mathbb{R}^d, ||\cdot||_1)$. The basic idea is to consider a set of axes that constitute a basis for the vector system; then, the generalized median allocation derives from the median projected to each of the axes.

**Proposition 3.5.2.** *The generalized median is strategyproof.*

**Proposition 3.5.3.** *The generalized median minimizes the social cost in the $L^1$ norm.*

*Proof.* If we denote with superscripts the coordinates of each vector, then

$$\begin{aligned}\min_{\mathbf{x}\in\mathbb{R}^d}\left(\sum_{i=1}^{n}||\mathbf{x}-\mathbf{x}_i||_1\right) &= \min_{(x^1,\dots,x^d)\in\mathbb{R}^d}\left(\sum_{i=1}^{n}\sum_{j=1}^{d}|x^j - x_i^j|\right)\\ &= \sum_{j=1}^{d}\left(\min_{x^j\in\mathbb{R}}\sum_{i=1}^{n}|x^j - x_i^j|\right).\end{aligned}$$

As a result, it follows that allocating $x^j = \text{median}(x_1^j, x_2^j, \dots, x_n^j), \forall j \in [d]$ minimizes the social cost. ■

# Chapter 4

# Communication Complexity

The increasing importance of distributed computing, networking, and VLSI routing have highlighted the importance of communication as a resource. Indeed, in many applications communication is the real bottleneck, as it is significantly slower and more expensive than local computation. In this context, the field of communication complexity endeavors to study the number of bits that the participants of a communication system need to exchange in order to perform certain tasks [KN96]. More precisely, Yao's pivotal work [Yao79] introduced an elegant mathematical model for studying these type of questions, in which only two parties (Alice and Bob) have to evaluate a function $f(x, y)$, where $x$ is Alice's input and $y$ is Bob's input. Despite the simplicity of this formulation, it turns out that it already captures many of the fundamental issues related to complexity of communication, and the results proven in this model can be often extended to more complicated settings. In the following sections, we closely follow the notation of Kushilevitz's survey [Kus97].

## 4.1 Two-Party Model

In this section we describe the two-party communication complexity model, as defined by Yao [Yao79]. Perhaps the most appealing feature of this model is its simplicity; indeed, it considers a scenario where only two communicating parties have to compute a two argument Boolean function, with each argument known to only a single party. It is important to point out that this model completely ignores the computational resources required by each party, and it focuses solely on the amount of communication exchanged between the parties. More precisely, consider a two-argument, Boolean function $f : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$. Alice is given an input $x \in \{0,1\}^n$ and Bob an input $y \in \{0,1\}^n$, while their goal is to compute the value of $f(x, y)$. Some examples which have attracted considerable attention in the literature include:

- Equality: $\mathrm{EQ}(x, y)$ is defined to be 1 if and only if $x = y$.

- Greater than: $\mathrm{GT}(x, y)$ is defined as 1 if and only if $x > y$, where $x$ and $y$ are viewed as the binary representation of numbers in the range $0, 1, \ldots 2^n - 1$.
- Disjointness: $\mathrm{DISJ}(x, y)$ is defined as 1 if and only if there is no index $i$ such that $x_i = y_i = 1$. One should think of $x$ and $y$ as subsets of $\{1, \ldots, n\}$, represented by their characteristic vectors. In this case, $\mathrm{DISJ}(x, y) = 1$ only if these subsets are disjoint, i.e. $x \cap y = \emptyset$.

Despite the simplicity of these functions, they represent natural functionalities that are often performed and required in distributed systems; e.g. consistency regulation typically employ the equality function.

The computation of the value of $f(x, y)$ is performed through a *communication protocol*. Specifically, during the execution of the procedure the two parties alternate roles in transmitting messages, where each message consists of a string of bits, and the protocol determines what message the sender should transmit next, as a function of its current input and of the communication performed so far.

Of course, every function can be computed by the following *trivial* protocol: Alice sends her entire input $x$ to Bob ($n$ bits of communication), and then Bob – knowing both $x$ and $y$ – computes $f(x, y)$ and sends it back to Alice. However, we general think of $n$ as a large number and hence, sending the entire input is very burdensome. Indeed, in many cases there are much more efficient protocols. In this context, the complexity measure of a protocol $\mathcal{P}$ is the minimal number of bits that must be sent in order to compute function $f$. Formally, we let $s_{\mathcal{P}}(x, y) = (m_1, m_2, \ldots, m_r)$ represent the communication exchanged during the execution of $\mathcal{P}$, where $m_i$ denotes the $i$-th message sent in the protocol. We will also use $|m_i|$ to refer to the length – i.e. the number of bits – of $m_i$. The deterministic communication complexity $D(\mathcal{P})$ of a given protocol $\mathcal{P}$ is the worst-case number of bits exchanged by the protocol. Moreover, the deterministic communication complexity $D(f)$ of a function $f$ is the communication complexity of the most efficient protocol that computes $f$. By employing the trivial protocol we described earlier it follows that for any function $f$, $D(f) \leq n + 1$.

### 4.1.1 Lower Bounds

Our main concern in this subsection is to establish *lower bounds* on the communication complexity of specific functions. Naturally, the task of proving lower bounds is usually much more involved than proving *upper bounds*, since the former task has to reference any possible solution to the problem. The main motivation for proving lower bounds is to know the limits of our communication protocols, and whether the complexity of a given procedure can be improved. In particular, we shall analyze the combinatorial structure imposed by protocols. The basic combinatorial element is called a *rectangle*.

**Definition 4.1.1** ([Kus97]). *A rectangle is a subset of $\{0, 1\}^n \times \{0, 1\}^n$ of the form $A \times B$, where each of $A$ and $B$ is a subset of $\{0, 1\}^n$. A rectangle $R = A \times B$ is* *called $f$-monochromatic if for every $x \in A$ and $y \in B$ the value of $f(x, y)$ is the same.*

As an example, consider the equality function EQ. If $A$ is the set of all strings in $\{0,1\}^n$ whose first bit is 1 and $B$ the set of all strings in $\{0,1\}^n$ whose first bit is 0, then $A \times B$ is EQ-monochromatic rectangle. Indeed, for every $x \in A$ and $y \in B$, we have that $x \neq y$ and hence $\mathrm{EQ}(x, y) = 0$.

The following lemma shows that the inputs for which the communication is the same form an $f$-monochromatic rectangle.

**Lemma 4.1.1.** *Let $\mathcal{P}$ be a protocol that computes a function $f$ and $(m_1, \dots m_r)$ be a sequence of messages. The set of inputs $(x, y)$ for which $s_{\mathcal{P}}(x, y) = (m_1, \dots, m_r)$ forms an $f$-monochromatic rectangle.*

*Proof.* First, we show by induction that the set of inputs for which the communication starts with $(m_1, \dots, m_i)$ is a rectangle. For $i = 0$ this set if $\{0,1\}^n \times \{0,1\}^n$, which is clearly a rectangle. Let $R$ be the set of inputs for which the communication starts with $(m_1, \dots, m_i)$. By the induction hypothesis, $R$ forms a rectangle $A \times B$. Let us assume – without any loss of generality – that the message $m_{i+1}$ is sent by Alice. If $A'$ is the set of inputs $x \in A$ for which given $m_1, \dots, m_i$ the message sent by Alice is $m_{i+1}$, then the set of inputs consistent with $(m_1, \dots, m_{i+1})$ is $A' \times B$, which is a rectangle. As a result, the set of inputs $(x, y)$ for which $s_{\mathcal{P}}(x, y) = (m_1, \dots, m_r)$ is rectangle, concluding the proof. ∎

For a function $f : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}^n$, we define $C^{\mathcal{P}}(f)$ as the minimum number of $f$-monochromatic rectangles that partition the space of inputs $\{0,1\}^n \times \{0,1\}^n$.

**Lemma 4.1.2.** *For every function $f : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}^n$,*

$$D(f) \geq \log_2 C^{\mathcal{P}}(f). \tag{4.1}$$

*Proof.* By Lemma 4.1.1, every protocol $\mathcal{P}$ partitions the space of inputs $\{0,1\} \times \{0,1\}^n$ into $f$-monochromatic rectangles. The number of these rectangles, that is the number of possible communications, is at most $2^{D(\mathcal{P})}$ and $D(f) \leq D(\mathcal{P})$; thus, $C^{\mathcal{P}}(f) \leq 2^{D(f)}$ and the lemma follows. ∎

Lemma 4.1.2 implies that for proving lower bounds on the communication complexity of $f$, it is sufficient to establish lower bounds on the number of $f$-monochromatic rectangles that partition the space of inputs. In this context, we shall first present the *fooling set* method, an approach for proving lower bounds on $C^{\mathcal{P}}(f)$, developed in [Yao79; LS81].

### Fooling Set Method

**Definition 4.1.2.** *A set of input pairs $\{(x_1, y_1), (x_2, y_2), \ldots, (x_\ell, y_\ell)\}$ is called a fooling set of size $\ell$ with respect to $f$ if there exists $b \in \{0, 1\}$ such that*

1. *For all $i$, $f(x_i, y_i) = b$.*
2. *For all $i \neq j$, either $f(x_i, y_j) \neq b$ or $f(x_j, y_i) \neq b$.*

**Lemma 4.1.3.** *If there exists a fooling set of size $\ell$ with respect to $f$, then*

$$D(f) \geq \log_2 \ell. \tag{4.2}$$

*Proof.* By Lemma 4.1.2 it suffices to show that $C^{\mathcal{P}}(f) \geq \ell$. In particular, we will show that in any partition of $\{0, 1\}^n \times \{0, 1\}^n$ into $f$-monochromatic rectangles, the number of rectangles is at least $\ell$. For the sake of contradiction, suppose that the number of $f$-monochromatic rectangles is smaller than $\ell$. In this case, there exist two pairs in the fooling set $(x_i, y_i)$ and $(x_j, y_j)$ that belong to the same rectangle $A \times B$. Thus $x_i, x_j \in A$ and $y_i, y_j \in B$, which means that $(x_i, y_j)$ and $(x_j, y_i)$ also belong to the rectangle $A \times B$. However, by the definition of the fooling set, $f(x_i, y_i) = f(x_j, y_j) = b$, while at least one of $f(x_i, y_j)$ and $f(x_j, y_i)$ is different from $b$; this implies that the rectangle $A \times B$ is not $f$-monochromatic. ∎

**Theorem 4.1.1.** *The deterministic communication complexity of EQ is at least $n$ bits.*

*Proof.* By Lemma 4.1.3 it suffices to find a fooling set of size $2^n$ for the function EQ. Specifically, consider the following set:

$$\{(\alpha, \alpha) : \alpha \in \{0, 1\}^n\}. \tag{4.3}$$

It is clear that this set is of size $2^n$. Moreover, note that $\mathrm{EQ}(\alpha, \alpha) = 1$ for every $\alpha$, while $\mathrm{EQ}(\alpha, \alpha') = \mathrm{EQ}(\alpha', \alpha) = 0$, for every $\alpha \neq \alpha'$. ∎

**Theorem 4.1.2.** *The deterministic communication complexity of DISJ is at least $n$ bits.*

*Proof.* For the DISJ we observe that the following is a fooling set of size $2^n$:

$$\left\{(A, \bar{A}) : A \subseteq \{1, \ldots, n\}\right\}. \tag{4.4}$$

∎

**Theorem 4.1.3.** *The deterministic communication complexity of GT is at least $n$ bits.*

*Proof.* The proof is analogous to Theorem 4.1.1. ∎

### Algebraic Method

The second method for establishing lower bounds has an algebraic flavor, and employs known results and tools from Linear Algebra. In particular, we map every function $f$ to a $2^n \times 2^n$ zero-one matrix, denoted as $M_f$, so that every row of $M_f$ is associated with a string $x \in \{0,1\}^n$ and each column of $M_f$ is associated with a string $y \in \{0,1\}^n$. Naturally, the $(x, y)$ entry of matrix $M_f$ contains the value $f(x, y)$.

**Lemma 4.1.4** ([MS]). *For any function $f : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$,*

$$D(f) \geq \log \mathrm{Rank}(M_f). \tag{4.5}$$

*Proof.* By Lemma 4.1.2, it suffices to show that $C^{\mathcal{P}}(f) \geq \mathrm{Rank}(M_f)$. Let $R_1, \ldots, R_t$ be the 1-monochromatic rectangles of an optimal cover of $\{0,1\}^n \times \{0,1\}^n$ with $f$-monochromatic rectangles. We associate every rectangle $R_i$ to a $2^n \times 2^n$ matrix $M_i$, whose $(x, y)$ entry is 1 if $(x, y) \in R_i$, and 0 otherwise. By construction, we have that $M_f = \sum_{i=1}^{t} M_i$. Moreover, given that the rank is sub-additive, we obtain that

$$\mathrm{Rank}(M_f) \leq \sum_{i=1}^{t} \mathrm{Rank}(M_f). \tag{4.6}$$

But, since $R_i$ is a 1-monochromatic rectangle it follows that $\mathrm{Rank}(M_i) = 1$, for all $i$, and the lemma follows directly from 4.6. ∎

**Theorem 4.1.4.** *The deterministic communication complexity of EQ is at least $n$ bits.*

*Proof.* The matrix $M_f$ that corresponds to the equality function has an entry 1 in its main diagonal, and 0 elsewhere. Thus, it follows that $\mathrm{Rank}(M_f) = 2^n$, and the theorem follows from Lemma 4.1.4. ∎

The algebraic method has everything one could hope for in a lower bound technique. Indeed, it frees us from studying the communication protocols and lets us consider the properties of matrix $M_f$ as a linear operator between Euclidean spaces. In this context, we should mention the *log rank* conjecture, one of the most prominent open problems in communication complexity. In particular, analogously to matrix $M_f$, the *sign* matrix $A$ of a function $f$ is a $2^n \times 2^n$ and $\{\pm 1\}$-valued matrix where its entry $(x, y)$ is $-1$ if $f(x, y) = 1$, and 0 otherwise. The log rank conjecture is stated as follows:

**Conjecture 4.1.1** ([LS88]). *There exists a constant $c$ such that for every sign matrix $A_f$,*

$$D(f) \leq (\log \mathrm{Rank}(A_f))^c + 2. \tag{4.7}$$

The additive term is needed because a rank-one sign matrix can require two bits of communication; see [NW95] for progress in the conjecture.

### 4.1.2 Randomized Communication Complexity

In this subsection, we strengthen the two-party model we previously defined by allowing the two parties to make randomized choices in order to decide which messages they will send. In particular, our goal is to show that at least for specific functions, randomized protocols have a significantly smaller complexity than deterministic ones.

In this context, every message sent by Alice or Bob is a *probabilistic* function of the transmitter's input and the communication so far. We will also allow the protocol to make errors; in particular, we say that a protocol $\mathcal{P}$ computes the function $f$ if for every input $(x, y)$ the probability that the output of $\mathcal{P}$ on $(x, y)$ equals $f(x, y)$ is at least $3/4$. The randomized communication complexity of $\mathcal{P}$, denoted with $R(\mathcal{P})$ is the worst-case – over all inputs $(x, y)$ – number of bits exchanged between the parties. We illustrate the power of randomization with the following theorems:

**Theorem 4.1.5.** *The randomized communication complexity of EQ is* $\mathcal{O}(\log n)$.

*Proof.* Let us denote with $a = a_0 a_1 \dots a_{n-1}$ the input of Alice and $b_0 b_1 \dots b_{n-1}$ the input of Bob. We think of these inputs as polynomials over the field $\mathrm{GF}[p]$, consisting of the numbers $0, 1, \dots, p-1$ with the operations of addition and multiplication modulo $p$, where $p$ is a prime number such that $4n^3 < p8n^3$; that is,

$$A(x) = a_0 + a_1 x + \dots + a_{n-1} x^{n-1} \mod p, \tag{4.8}$$

$$B(x) = b_0 + b_1 x + \dots + b_{n-1} x^{n-1} \mod p. \tag{4.9}$$

Alice picks uniformly at random a number $t \in \mathrm{GF}[p]$ and sends Bob the values $t$ and $A(t)$. Bob outputs 1 if $A(t) = B(t)$ and 0 otherwise. The number of bits exchanged is $2 \cdot \log p = \mathcal{O}(\log n)$. If $a = b$ then $A(t) = B(t)$ for all $t$ and the output is always 1. On the other hand, if $a \neq b$ we have to distinct polynomials $A$ and $B$ of degree $n - 1$. Such polynomials can be equal on at most $n - 1$ elements of the field. Indeed, note that their difference is a non-zero polynomial of degree at most $n - 1$. Thus, the probability of error is at most

$$\frac{n}{p} \leq \frac{n}{4n^3} = \frac{1}{4n^2} \leq \frac{1}{4}. \tag{4.10}$$

■

**Theorem 4.1.6.** *The randomized communication complexity of GT is* $\mathcal{O}(\log n)$.

*Proof.* Let us assume that Alice and Bob operate based on two variables: $L$, the left border of search initiated as 1, and $R$, the right border of search initiated as $n$. In every step of the protocol, Alice and Bob both compute $M = \lfloor (L+R)/2 \rfloor$ and make the equality test $\mathrm{EQ}(x_L \dots X_M, y_L \dots y_M)$ based on the protocol we introduced in

Theorem 4.1.5. If the strings are equal they set $L = M + 1$; otherwise, Alice and Bob set $R = M$. If at some stage $L = R$, they simply exchange the bits $x_L$ and $y_L$ and decide on the value of GT based on the bit which is larger.

If the output of the protocol for EQ is correct in every iteration the correctness of the protocol is obvious. Thus, given that the probability of error of the EQ protocol is bounded by $1/(4n^2)$, it follows that the probability of error in the GT protocol is at most $\log n/(4n^2) \leq 1/(4n)$ (by the union bound). ■

In fact, it is known that the upper bound for the GT function can be improved to $\mathcal{O}(\log n)$. In contrast, for other functions randomized protocols are not more efficient than deterministic protocols. For example, it has been established that $R(\text{DISJ}) = \Omega(n)$; see [Raz92].

## 4.2 Applications

In this section, we show that lower and upper bounds in communication complexity can be employed to other domains. Remarkably, this approach applies for problems in which communication does not seem to appear in the problem at all.

### 4.2.1 Finite Automata

A deterministic *finite automaton* $\mathcal{A}$ has a finite set of states $Q$, an alphabet $\Sigma$, and a transition function $\delta : Q \times \Sigma \mapsto Q$ which determines for each state and for each character in the alphabet what should be the next state. An input $w = w_1 w_2 \dots w_m$ induces a sequence of $m + 1$ states as follows: commence with the initial state $q_0 \in Q$, and in the $i$-th step transition to the state defined by $\delta$ with input the current state $q$ and character $w_i$. If at the end the state of the automaton belongs to the set of *accepting* states $F \subseteq Q$, then we say that the sequence $w$ is accepted by $\mathcal{A}$, or equivalently that $w$ belongs to the *language* of $\mathcal{A}$. Otherwise, we say that $w$ is rejected by $\mathcal{A}$ or that $w$ does not belong to the language of $\mathcal{A}$.

In this context, the basic goal is to show lower bounds on the number of states required for a given language $\Sigma$. For simplicity, we assume that the alphabet is binary, i.e. $\Sigma = \{0, 1\}$. The following lemma relates this problem to an appropriate communication complexity problem.

**Lemma 4.2.1.** *Let $f : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$ be a function, $L \subseteq \Sigma^*$ a language, and $\mathcal{A}$ an automaton. If we assume that*

$$L \cap \{0,1\}^{2n} = \{xy : |x| = |y| = n, f(x,y) = 1\}, \tag{4.11}$$

*then $D(f) \leq \log |Q| + 1$.*

*Proof.* Consider the following communication protocol: On input $(x, y)$ Alice simulates the path taken by the automaton $\mathcal{A}$ on her input. Then, she transmits the last state in this path to Bob ($\log |Q|$ bits). Subsequently, Bob simulates the automaton $\mathcal{A}$ with initial state $q$ and his input string $y$. Finally, Bob simply transmits to Alice whether the final state of his sequence belongs to the set of accepting states. By construction, if $xy$ is accepted then $f(x, y) = 1$, while if $xy$ is rejected then $f(x, y) = 0$. ■

Now consider a language $L_n = \{xx : |x| = n\}$, for a fixed $n$. This language can be expressed as

$$L_n = \{xy : |x| = |y| = n, \mathrm{EQ}(x, y) = 1\}. \tag{4.12}$$

The previous lemma implies that an automaton for $L_n$ has a set of states of size

$$|Q| \geq 2^{D(\mathrm{EQ})-1}. \tag{4.13}$$

However, we have already shown that $D(\mathrm{EQ}) \geq n$, implying that every automaton for $L_n$ has size $|Q| \geq 2^n$. As a result, we conclude that the language $L = \{xx : x \in \Sigma^*\}$ has no finite automaton, i.e. $L$ is not a *regular language*. A similar argument shows that every automaton for the language $L'_n = \{xy : |x| = |y| = n, x \neq y\}$ also has size at least $2^n$. One the other hand, it is easy to construct a *non-deterministic automaton* for $L'_n$ of size $\mathcal{O}(n)$ (e.g. see [HMU07]).

### 4.2.2 Decision Trees

A *decision* tree is a binary tree such that every node is labeled by a variable from $x_1, \ldots, x_m$, and from every node there are two outgoing towards the children of the node labeled 0 and 1; the leaves of the tree are labeled either 0 or 1. A decision tree determines a function $f : \{0,1\}^m \mapsto \{0,1\}$ with the following process: for a boolean assignment to the $m$ variables, we commence from the root of tree; whenever we reach a node labeled by some variable $x_i$ we proceed by the edge whose label matches the label of variable $x_i$. Finally, when we reach a leaf, we simply return its label as the value of $f$ on the given assignment.

It is clear that for every function $f : \{0,1\}^m \mapsto \{0,1\}$ there exists a decision tree of *depth* (the longest path from the root to a leaf) $m$, where each of the $2^m$ leaves corresponds to a single assignment, and subsequently the label of the leaf is the value of $f$ on that assignment. Naturally, the goal is to find decision trees with as small depth as possible. The following lemma provides lower bounds through communication complexity lower bounds.

**Lemma 4.2.2** ([Nis94; GT91]). *Let $f : \{0,1\}^{2n} \mapsto \{0,1\}$ be a function. If $f$ has a decision tree of depth $d$, then the function $f(x_1 \ldots x_n, x_{n+1} \ldots x_{2n})$ has communication complexity $D(f) \leq d$.*

*Proof.* For a given decision tree, Alice and Bob can simulate its computation as follows: Commencing from the root of the tree, whenever they encounter a node the players examine the label $x_j$ of the node, and the player that holds the value of variable $x_j$ announces it to the other party. Then, the value of $x_j$ determines the next node. The simulation terminates when they reach a leaf of the tree, in which case the label of the corresponding leaf is the desired value of $f$. Thus, the number of bits exchanged is at most $d$. ■

### 4.2.3 VLSI chips

A VLSI chip can be viewed as an $a \times b$ grid with $n$ input ports and a single output port. On some of the vertices of the grid there are *gates*, while *wires* connect the gates with other gates on the chip or with the ports. Our main goal here is to show how communication complexity bounds can be employed to provide area-time trade-offs for VLSI chips.

The most important complexity measures for a VLSI chip are the *area* of the chip $A = a \cdot b$, and the *time* $T$, i.e. the number of steps required to determine the result in the output port from the time the input is provided. The designer of a chip that computes a function $f(x_1, \dots, x_m)$ has to decide what gates to use, and how to connect them. The following lemma shows that at least for certain functions, a VLSI chip should either have a "large" area or require "many" time steps.

**Lemma 4.2.3** ([Len90]). *Assume that we have to compute a function $f : \{0,1\}^{2n} \mapsto \{0,1\}$ through a VLSI chip with area $A$ and time $T$ to perform the computation. Then,*

$$AT^2 \geq (D(f))^2 . \tag{4.14}$$

*Proof.* For a given chip, we construct a two-party protocol as follows. Assume – without any loss of generality – that $a \leq b$. We can always split the chip into two pieces in a way that partitions the input ports of the chip into two equal size sets, and with at most $a$ wires connecting the two pieces. We assign each of Alice and Bob one of the two pieces of the chip; hence, we partition the input bits between Alice and Bob in the same way that the input bits are partitioned among the two pieces. Then, Alice and Bob will have to simulate the computation of the chip. In particular, in step $i$ of the protocol each of the two players evaluates the output of all the gates in his piece of the chip in time $i$. The parties only exchange the values which are going on the wires connecting the two pieces. Thus, since the number of wires is at most $a$ and the number of steps to compute the output is at most $T$, it follows that the total number of bits exchanged is at most $a \cdot T \leq \sqrt{A} \cdot T$, concluding the proof. ■

For more applications of communication complexity we refer to [Mur71; ROK94;

Pat96; NW93; NW94; LT80; Mil94; Mil+95; AUY83; BKL95], and referecnes therein.

# Chapter 5

# Communication in Auctions

In this chapter we present asymptotically optimal protocols for a series of exemplar auction formats. Importantly, at the same time we retain the incentive compatibility of the mechanism as well as the obtained social welfare. Our approach combines ideas employed in sampling with techniques from Information Theory. In this way, we guarantee optimal communication for the following settings:

- Single-item auction.
- Multi-unit auctions with unit-demand bidders.
- Multi-item auctions with additive valuations.

## 5.1 Preliminaries

In the following we denote with $n$ the number of participants in the game. In single parameter environments the *rank* of agent $i$ corresponds to the index of her private valuation in ascending order (and indexed from 1 unless explicitly stated otherwise). In the case of identical valuation profiles we accept some arbitrary but fixed order among the agents – e.g. lexicographic order. We also assume that an agent remains active in the auction only when positive utility can be obtained; that is, if the announced price for the item is greater or equal to the valuation of some agent $i$, then $i$ will withdraw from the forthcoming rounds of the auction, while we remark that in Mechanism 4 we will assume that the agents' valuations are distinct.

A mechanism will be referred to as *strategy-proof* or *incentive compatible* if truthful reporting is a *universally* dominant strategy – a best response under any possible action profile and randomized realization – for every agent. A strategy $s_i$ is *obviously dominant* if, for any deviating strategy $s_i'$, starting from any earliest information set where $s_i'$ and $s_i$ disagree, the best possible outcome from $s_i'$ is no better than the worst possible outcome from $s_i$. A mechanism is *obviously strategy-proof* (OSP) if it has an equilibrium in obviously dominant strategies.

We use the standard notation of $f(n) \sim g(n)$ if $\lim_{n \to +\infty} f(n)/g(n) = 1$ and $f(n) \lesssim g(n)$ if $\lim_{n \to +\infty} f(n)/g(n) \leq 1$, where $n$ will be implied as the asymptotic parameter. Moreover, in order to analyze the bit complexity the valuation space will be assumed discretized and every valuation will be represented with $k$ bits; we will mostly consider $k$ to be a constant. Finally, communication complexity is defined as the cumulative amount of bits elicited from the participants; our analysis will be worst-case with respect to the input – i.e. the agents' valuations – and average-case with respect to the introduced randomization in the procedure.

## 5.2 Single Item Auction

Our implementation is established based on a black-box algorithm; in particular, let $\mathcal{A}$ be an algorithm that faithfully *simulates* a second-price auction; that is, $\mathcal{A}$ interacts with a set of agents and returns the VCG outcome, without actually allocating items and imposing payments. However, the agents that are excluded by $\mathcal{A}$ will be also automatically eliminated from the remainder of the auction. Naturally, we assume that $c \geq 2$, so that the second-price rule is properly implemented. It should be noted that our mechanism induces a format that couples the auction that is simulated by $\mathcal{A}$ with an ascending auction.

**Algorithm 3:** Ascending Auction through Sampling

**Result:** Winner & VCG payment
**Input**: Set of agents $N$, size of sample $c$, algorithm $\mathcal{A}$
**while** $|N| > c$ **do**
  $S :=$ random sample of $c$ agents from $N$
  w := winner in $\mathcal{A}(S)$
  Announce $p :=$ payment in $\mathcal{A}(S)$
  Update the active agents: $N := \{i \in N \setminus S \mid v_i > p\} \cup \{w\}$
**end**
**if** $|N| = 1$ **then**
  return $w, p$
**else**
  return $\mathcal{A}(N)$
**end**

**Proposition 5.2.1.** *Assuming truthful bidding, Mechanism 3 implements – with probability 1 – the VCG allocation rule.*

*Proof.* First, if after the termination of some round only a single agent $i$ remains active, it follows that the announced price $p$ – that coincides with the valuation of some player – exceeds the valuation of every player besides $i$; thus, by definition, the outcome implements the VCG allocation rule. Moreover, the claim when $2 \leq |N| \leq c$ follows given that $\mathcal{A}$ faithfully simulates a second-price auction. Otherwise, in a given round – with $|N| > c$ – only agents that are below or equal to the second-highest valuation will withdraw from the auction. Thus, the allocation rule over the active players remains invariant between rounds, concluding the proof. ∎

**Proposition 5.2.2.** *If $\mathcal{A}$ simulates a sealed-bid auction, Mechanism 3 is strategy-proof.*

*Proof.* Consider some round of the auction and some agent $i$ that has been selected in the sample $S$; if we fix the reports from the agents in the sample besides $i$ we can identify the following two cases. First, if $v_i \geq x_j, \forall j \in S \setminus \{i\}$, with $x_j$ representing the report of agent $j$, then sincere bidding is a best response for $i$ given that her valuation exceeds the announced price. Indeed, note that since $\mathcal{A}$ simulates a second-price auction the winner in the sample does not have any control over the announced price of the round. In the contrary case, agent $i$ does not have an incentive to misreport and remain active in the auction given that the reserved price will be greater or equal to her valuation. Let $p$ the market clearing price in $\mathcal{A}$ and $i$ some agent that was not selected in the sample. It is clear that if $v_i \leq p$ then a best response for $i$ is to withdraw from the auction, while if $v_i > p$ then $i$'s best response is to remain active in the forthcoming round. ∎

**Proposition 5.2.3.** *If $\mathcal{A}$ simulates an English Auction, Mechanism 3 is OSP.*

*Proof.* The claim follows from the OSP property of the English auction. In particular, note that we simply perform an English auction without interacting with every active agent in each round, but instead with a small sample; when only a single player survives from the sample, we announce the price to the remainder of the agents. ∎

Before we establish the communication complexity of the induced auction, we should point out that a trivial lower bound to recover the optimal social welfare – with probability 1 – is $n$ bits. Indeed, since the information is distributed to $n$ parties and the goal is to allocate the item to the agent with the highest utility, every player has to commit at least 1 bit to the procedure. Through this prism, we will show that our mechanism reaches this lower bound with arbitrarily small error, assuming that $k$ is a constant. We should also note that the information leakage in Mechanism 3 is asymmetrical, in the sense that statistically, the agents that are closer to winning the item have to reveal relatively more bits from their private valuation. It is clear that in order to truncate the communication complexity of the mechanism, one has to guarantee small inclusion rate – in expectation – for each round of the auction; this property is implied by the following lemma:

**Lemma 5.2.1.** *Let $X_a$ the proportion of the agents that remain active in a given round of the auction; then*

$$\mathbb{E}[X_a] \lesssim \frac{2}{c+1}. \tag{5.1}$$

Before we proceed with the proof of this lemma, we first state the following auxiliary lemmas:

**Lemma 5.2.2.**

$$\binom{n}{c} \sim \frac{n^c}{c!}. \tag{5.2}$$

**Lemma 5.2.3.**

$$\sum_{i=1}^{n} i^p \sim \frac{n^{p+1}}{p+1} \sim \sum_{i=1}^{n} (i-1)^p. \tag{5.3}$$

**Lemma 5.2.4.** *Consider $n$ active players in the auction and $X_r = X_r(n, c)$ the rank - among all the active agents - of the player with the second highest valuation in the sample; then*

$$\mathbb{E}[X_r] \sim n\frac{c-1}{c+1}. \tag{5.4}$$

*Proof.* First, if we apply simple combinatorial arguments it follows that

$$\Pr[X_r = i] = \frac{\binom{n-i}{1}\binom{i-1}{c-2}}{\binom{n}{c}}. \tag{5.5}$$

As a result, we have that

$$\begin{aligned}
\mathbb{E}[X_r] = \sum_{i=1}^{n} i \Pr[X_r = i] &\sim \frac{c!}{n^c} \sum_{i=1}^{n} i(n-i)\binom{i-1}{c-2} \\
&= \frac{c!}{n^c}\left( n\sum_{i=1}^{n}(i-1)\binom{i-1}{c-2} - \sum_{i=1}^{n} i^2 \binom{i-1}{c-2}\right) \\
&\sim \frac{c!}{n^c}\left( n\sum_{i=1}^{n} i\binom{i}{c-2} - \sum_{i=1}^{n} i^2 \binom{i}{c-2}\right) \\
&\sim \frac{c!}{n^c}\left( n\sum_{i=1}^{n} \frac{i^{c-1}}{(c-2)!} - \sum_{i=1}^{n} \frac{i^c}{(c-2)!}\right) \\
&\sim \frac{c(c-1)}{n^c}\left(\frac{n^{c+1}}{c} - \frac{n^{c+1}}{c+1}\right) \\
&= n\frac{c-1}{c+1},
\end{aligned}$$

where we have applied the asymptotic results from Lemma 5.2.2 and Lemma 5.2.3. Also note that we used Cesaro's means theorem in the fourth line and we ignored the lower-order magnitude terms in the third. ■

Finally, Lemma 5.2.1 follows from $X_a \leq (n - X_r)/n$. Note that the inequality derives from the fact that multiple agents could have the same valuation with the second highest bidder in the sample and we have assumed that every such agent will withdraw from the auction.

Let us assume that $Q(n; k)$ is the (deterministic) communication complexity of $\mathcal{A}$ with $n$ players. In particular, when $\mathcal{A}$ is a sealed-bid auction it follows that $Q(n; k) = n \cdot k$. On the other hand, the worst-case communication cost of an English auction is $Q(n, k) = 2^k n$. Indeed, given that $\mathcal{A}$ faithfully simulates a second-price auction, the auctioneer has to cover every possible point on the valuation space. If $T(n; c, k)$ is the (randomized) communication complexity of the induced Mechanism 3, it follows that when $n > c$

$$\mathbb{E}[T(n; c, k)] = \mathbb{E}[T(nX_a; c, k)] + Q(c; k) + n - c. \tag{5.6}$$

Solving recursions of such form are standard in the analysis of randomized algorithms (see [Cor+09]); in particular, we can establish the following theorem.

**Theorem 5.2.1.** *Let $t(n; c, k)$ the expected communication complexity of Mechanism 3 with $k$ assumed constant; then, $\forall \epsilon > 0, \exists c_0 = c_0(\epsilon)$ such that $\forall c \geq c_0$*

$$t(n; c, k) \lesssim n(1 + \epsilon). \tag{5.7}$$

In order to simplify the exposition, we will assume that the agents' valuations are discrete; note that under this assumption we will obtain an upper bound on the communication complexity of our mechanism.

**Lemma 5.2.5.** *Let $X_a$ the proportion of the agents that remain active in a given round of the auction; then, assuming discrete valuations, it follows that*

$$\mathrm{Var}[X_a] \sim \frac{2(c-1)}{(c+2)(c+1)^2}. \tag{5.8}$$

*Proof.* The proof of this claim follows analogously to Lemma 5.2.1; in particular, if we consider the proxy random variable $X_r$, as defined in Lemma 5.2.4, we have that

$$\begin{aligned}\mathbb{E}[X_r^2] = \sum_{i=1}^{n} i^2 \Pr[X_r = i] &\sim \frac{c!}{n^c} \sum_{i=1}^{n} i^2 (n-i) \binom{i-1}{c-2} \\ &\sim \frac{c!}{n^c} \left( n \sum_{i=1}^{n} i^2 \binom{i}{c-2} - \sum_{i=1}^{n} i^3 \binom{i}{c-2} \right) \\ &\sim \frac{c!}{n^c} \left( n \sum_{i=1}^{n} \frac{i^c}{(c-2)!} - \sum_{i=1}^{n} \frac{i^{c+1}}{(c-2)!} \right) \\ &\sim \frac{c(c-1)}{n^c} \left( \frac{n^{c+2}}{c+1} - \frac{n^{c+2}}{c+2} \right) \\ &= n^2 \frac{c(c-1)}{(c+2)(c+1)}.\end{aligned}$$

As a result, the claim follows given that $\mathrm{Var}[X_r] = \mathbb{E}[X_r^2] - (\mathbb{E}[X_r])^2$ and $\mathrm{Var}[X_a] = \mathrm{Var}[X_r]/n^2$. ■

Having established this lemma, we proceed with a sketch proof of Theorem 5.2.1.

*Proof.* Consider a fixed round of the auction with $n$ active agents. If we let $\mu = \mathbb{E}[X_a]$ and $\sigma = \sqrt{\mathrm{Var}[X_a]}$, Chebyshev's inequality implies that $\Pr[|X_a - \mu| \geq \sqrt{c}\sigma] \leq 1/c$. Moreover, we have that $\mu + \sqrt{c}\sigma \leq 4/\sqrt{c}$, for $c \geq 2$. As a result, if $t(n; c, k) = \mathbb{E}[T(n; c, k)]$, it follows from (5.6) that

$$t(n; c, k) \lesssim \left(1 - \frac{1}{c}\right) t\left(\frac{4n}{\sqrt{c}}; c, k\right) + \frac{1}{c} t(n; c, k) + n + Q(c; k), \tag{5.9}$$

where we used the fact that $t(n; c, k)$ is decreasing with respect to $n$. Given that this inequality holds asymptotically, it will also hold with up to an $\epsilon$ multiplicative error, for any $\epsilon > 0$ and for sufficiently large $n \geq n_0 = n_0(\epsilon)$. Thus, solving the induced recursion completes the proof. ■

Note that our asymptotic guarantee is invariant on the communication complexity of the second-price algorithm $\mathcal{A}$, assuming that $k$ is a constant. On the other hand, if we allow $k$ to depend on $n$ our guarantee crucially depends on $\mathcal{A}$ (see Theorem 5.4.1).

## 5.3 Multi-Item Auctions with Additive Valuations

As a direct extension of the previous setting, let us assume that the auctioneer has to allocate $m$ (indivisible) items and the valuation space is *additive*, that is for every agent $i$ and for a bundle of items $S$, $v_i(S) = \sum_{j \in S} v_{ij}$. In this setting, we shall perform an auction for each item using Mechanism 3; it is clear that assuming truthful bidding, the induced auction will implement – with probability 1 – the VCG allocation rule, as implied by Proposition 5.2.1. Moreover, the following proposition holds.

**Proposition 5.3.1.** *The mechanism induced by employing $m$ auctions as described in Mechanism 3 is ex-post incentive compatible.*

However, we will illustrate that a simultaneous implementation can significantly truncate the communication exchange – relatively to a sequential format – through an efficient encoding scheme. First, we assume that $m$ is arbitrary and that we have to perform a separate and independent auction for each of the $m$ items. Under this assertion, the optimality condition yields the lower bound of $n \cdot m$ bits, which can be again asymptotically reached with arbitrarily small error:

**Proposition 5.3.2.** *Let $t(n; m, c, k)$ the expected communication complexity of implementing $m$ sequential auctions as described in Mechanism 3 with $k$ assumed constant; then, $\forall \epsilon > 0, \exists c_0 = c_0(\epsilon)$ such that $\forall c \geq c_0$*

$$t(n; m, c, k) \lesssim nm(1 + \epsilon). \tag{5.10}$$

On the other hand, if we assume that $m$ is constant and that we perform the auctions simultaneously, we will show that we can reach the bound of $n$ bits with a very simple coding scheme. To be precise, recall that – asymptotically – the expected inclusion rate in Mechanism 3 is at most $2/(c + 1)$ and thus, as the sample size increases the overwhelmingly most probable scenario is that some random agent will drop from the next round of the auction; we shall exploit this property by considering the following encoding. An agent $i$ – that remains active in at least one auction – will transmit the bit 0 in the case of withdrawal from every auction; otherwise, $i$ will transmit an $m$ bit vector that will indicate the auctions that she wishes to remain active. Although the latter part of the encoding is clearly sub-optimal, we will show that in fact, we can asymptotically obtain an optimality guarantee. In particular, consider a round of the auction with $n$ players that remain active in at least one auction and $p$ the expected probability that a player will withdraw from every auction in the current round. Since every player is active in at most $m$ auctions, it follows from the union bound that $1 - p \lesssim 2m/(c + 1)$. As a result, if $N_b$ denotes the total number of bits transmitted in the round, we have that

$$\mathbb{E}[N_b] = n\left(1 \cdot p + m \cdot (1 - p)\right) \lesssim n\left(\left(1 - \frac{2m}{c+1}\right) + m\left(\frac{2m}{c+1}\right)\right). \tag{5.11}$$

As a result, since $m$ is a constant it follows that $\mathbb{E}[N_b] \lesssim n(1+\delta)$, for any $\delta > 0$ and for a sufficiently large constant $c$. Moreover, the expected inclusion rate – the proportion of agents that remain active in at least one auction – is asymptotically at most $2m/(c + 1)$ and thus, we can establish the following theorem.

**Theorem 5.3.1.** *Let $t(n; m, c, k)$ the expected communication complexity of implementing $m$ simultaneous auctions as described in Mechanism 3 with the aforementioned encoding scheme and $k$ and $m$ assumed constant; then, $\forall \epsilon > 0, \exists c_0 = c_0(\epsilon)$ such that $\forall c \geq c_0$*

$$t(n; m, c, k) \lesssim n(1 + \epsilon). \tag{5.12}$$

## 5.4 Multi-Unit Auctions with Unit Demand

Consider that we have to allocate $m$ identical items to $n$ unit demand bidders. We are interested in the non-trivial case where $m \leq n$; in this setting, our approach will differ depending on the asymptotic value of $m$.

First, we consider the canonical case where $m$ is constant. In this setting, we can extend Mechanism 3 as follows. In each round, we invoke an algorithm $\mathcal{A}$ that simulates the VCG outcome [1] for a random sample of active agents $c = \kappa m + 1$ for $\kappa \in \mathbb{N}$. Next, the market clearing price in the sample will be announced in order to 'prune' the active agents. Through parameter $\kappa$, we are able to restrain the inclusion rate in the following rounds. As a result, we can prove statements analogous to Propositions 5.2.1, 5.2.2, 5.2.3 and Theorem 5.2.1. The analysis is very similar to the single item mechanism and therefore it can be omitted.

Next, we study the case where $m = \gamma \cdot n$ for $\gamma \in (0, 1)$. Our main goal is to design an interaction process that shrinks the remaining or active agents very rapidly; however, unlike the previously studied cases, the winners from the auction could constitute a large fraction of the participation and thus, a different approach is required. In particular, we introduce the following ingredients.

First, instead of announcing a single price – as in a standard ascending auction – we shall broadcast 2 distinct prices, $p_h$ and $p_\ell$. The agents that are willing to pay $p_h$ will guarantee to obtain an item, not for a price of $p_h$, but for a price that will be later determined in the process; this of course is essential in order to obtain the incentive compatibility guarantee. On the other hand, the agents that are not willing to bid $p_\ell$ or more will be automatically excluded from the forthcoming rounds. Hence, we are able to recurse on the agents that lie in the intermediate region between $p_\ell$ and $p_h$. As a result, an agent should be able to broadcast 3 distinct signals and hence, a single bit does not suffice. Let us assume that for bidder $i$, the bit 1 corresponds to $v_i > p_h$, bit 0 to $v_i < p_\ell$ and some 2-bit code for the complementary case. Although this encoding could augment the communication complexity, we will select the prices $p_\ell$ and $p_h$ in a way that the agents in-between are very few and as a corollary, the overhead of the 2-bit representation will be negligible. This approach resembles the techniques used in

[1] The VCG outcome for this auction consists of allocating a single unit to each of the $m$-highest bidders for a price coinciding with the $m + 1$-highest bid

Coding Theory (e.g. see Huffman coding [Huf52]). On a high level, our mechanism consists of the following steps:

---

**Algorithm 4:** $\mathcal{M}(N, m)$: Multi-Unit Auction through Sampling

---

**Result:** Winners & VCG payment
**Input**: Set of agents $N$, number of items $m := \gamma n$
Initialize the winners $W := \emptyset$ and the losers $L := \emptyset$
$p_h :=$ estimated upper bound on the price
$p_\ell :=$ estimated lower bound on the price
Announce $p_\ell$ and $p_h$
Update the winners: $W := W \cup \{i \in N \mid v_i > p_h\}$
Update the losers: $L := L \cup \{i \in N \mid v_i < p_\ell\}$
**if** $p_h = p_\ell$ **then**
  return $W, p_h$
**else**
  $N := N \setminus (W \cup L)$
  Update the number of items $m$
  return $\mathcal{M}(N, m)$
**end**

---

**Proposition 5.4.1.** *Mechanism 4 is ex-post incentive compatible.*

However, we remark that Mechanism 4 is not strategyproof and in particular, answering sincerely to the queries is not necessarily a dominant strategy for the agents in the sample due to potential retaliation strategies.

The crux of the algorithm is to determine the prices $p_h$ and $p_\ell$ such that $p_h$ is at most $\epsilon n$-ranked higher than the $m+1$-ranked player and $p_\ell$ is at least $\epsilon n$-ranked lower. Consider a perfect binary tree with $k$ levels, with each level adding an additional bit to the representation, starting from the most significant bit and leading to each of the $2^k$ leaves, each representing a single valuation profile. This structure can be seen in figure 5.1.

The basic idea is to perform stochastic binary search; to be precise, in each level of the tree we will estimate an additional bit. Formally, let $x_1, x_2, \ldots x_r$ be the predicted bits after $r$ levels. In the next level we consider a sample of size $c$ in order to estimate the proportion of the agents for which $v_i \leq \overline{x_1 x_2 \ldots x_r 011 \ldots 1}$. Let $\hat{X}_c$ denote the estimation as derived from the sample and some parameter $\epsilon$. Recall that we want to create two different estimations on the tree that correspond to $p_\ell$ and $p_h$.

For simplicity, a sample will be called $\epsilon$-ambiguous if $|\hat{X}_c - \gamma| < \epsilon$. Intuitively, the smaller the $\epsilon$ the less clear the next branching becomes. In every ambiguous branch, the "high" estimation will predict a bit of 1, whilst the "lower" estimation

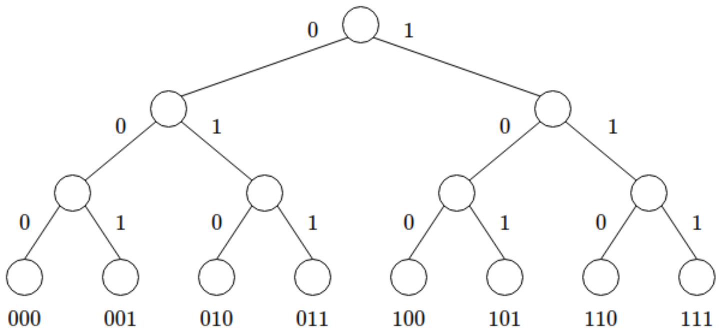


Figure 5.1: Representation of the valuation space

will predict a bit of 0. If the sample is not $\epsilon$-ambiguous, then we predict bit 1 if $\hat{X}_c < \gamma$ and bit 0 if $\hat{X}_c > \gamma$. In other words, we can imagine that the two estimations will coincide in the first levels, until they separate when a "close" decision arises. We claim that this algorithm will terminate with high probability with the desired payment range. For the purpose of the analysis we will use the following lemma:

**Lemma 5.4.1** (Chernoff-Hoeffding bound)**.** *Let $X_1, X_2, \ldots, X_c$ i.i.d. random variables with $X_i \sim Be(p)$ and $X_\mu = (X_1 + X_2 + \cdots + X_c)/c$; then*

$$\Pr(|X_\mu - p| \geq \epsilon) \leq 2e^{-2\epsilon^2 n}. \tag{5.13}$$

It is easy to see that if all of the $k$ samples have at most $\epsilon$ error, then $p_h$ will be at most $2\epsilon n$-ranked higher than the $m+1$ ranked player and $p_\ell$ will be at least $2\epsilon n$-ranked lower. Moreover, the probability $p_e$ that – for a single estimate and after $k$ rounds – there exists a sample with more than $\epsilon$ error can be upper bounded using the union bound: $p_e \leq 2ke^{-2\epsilon^2 c}$. As a result, $\forall \delta > 0$, there exists $c_0 = c_0(\delta, \epsilon, k)$ such that $p_e \leq \delta$, namely

$$c_0 = \frac{1}{2\epsilon^2} \log\left(\frac{2k}{\delta}\right). \tag{5.14}$$

Moreover, the union bound implies that the probability of error for either of the two estimates is at most $2\delta$. Let denote with $t(n; c, k)$ the expected communication complexity of the mechanism with $n$ active players. We have established the following bound:

$$\begin{aligned} t(n; c, k) \leq 2ck + (1 - 2\delta)\left(n(1 + 4\epsilon) + t(4\epsilon n; c, k)\right)) \\ + 2\delta(2n + t(n; c, k)), \end{aligned}$$

where the term $2ck$ corresponds to the communication complexity of the sampling process, the next term in the first line to the communication complexity of the round if all of the $2k$ samples have at most $\epsilon$ error and the term in the second line to the worst-case cost when at least one of the samples exceeds an $\epsilon$ error. Thus, solving the recursion and recalling that $k \in \mathcal{O}(n^{1-\ell})$ implies the following theorem:

**Theorem 5.4.1.** *Let $t(n; c, k)$ the communication complexity of Mechanism 4 with $k \in \mathcal{O}(n^{1-\ell})$ for some $\ell > 0$; then, $\forall \epsilon > 0, \exists c_0 = c_0(\epsilon, k)$ such that $\forall c \geq c_0$*

$$t(n; c, k) \lesssim n(1 + \epsilon). \tag{5.15}$$

Note that for this theorem we allowed $k$ to depend on the (initial) number of agents.

# Chapter 6

# Information Requirements of the Median

In this chapter we consider Moulin's median mechanism in the context of facility location games. We show that when the median is estimated through a random sample of constant size the obtained social cost is near optimal. Our result is inherently asymptotic and applies to settings with large participation.

**The Median Mechanism** Consider that we have to allocate a single facility on a metric space $(\mathbb{R}^d, ||\cdot||_1)$ and $n$ agents, with $\mathbf{x}_i \in \mathbb{R}^d$ the preferred location of agent $i$. The social cost of an allocation $\mathbf{x}$ is defined as $\mathrm{SC} = \sum_{i=1}^{n} \mathrm{d}(\mathbf{x}, \mathbf{x}_i)$. In this context, the generalized median [Mou80] is a strategyproof and optimal – with respect to the social cost in $L^1$ – mechanism that allocates the facility to the coordinate-wise median of the reported instance.

## 6.1 Approximating the Median

Before we proceed to our analysis, we remark that when $k$ – the number of bits that can represent any valuation – is assumed constant our result can be obtained with an iterative process – analogously to our approach in multi-unit auctions – and Chernoff bounds in order to correlate the accuracy of the approximation with the size of the sample. Nonetheless, our approach here is more robust since we do not even need the discretized valuation space hypothesis. More precisely, our mechanism will simply employ the generalized median scheme $\mathcal{M}$ for a random sample of $c$ agents.

**Proposition 6.1.1.** *The approximate median Mechanism 5 is strategyproof.*

*Proof.* The claim follows from the incentive compatibility of the median mechanism. ∎

Our analysis commences with the one-dimensional case – i.e. allocating a single facility on the line; the extension to any metric space $(\mathbb{R}^d, ||\cdot||_1)$ will then follow easily. We conclude this section by illustrating why our sampling approach cannot be efficiently applied for allocating multiple facilities. In order to make the analysis more concise – and without any loss of generality – we assume that $n = 2\kappa + 1$ and $c = 2\rho + 1$ for some $\kappa, \rho \in \mathbb{N}$. Let $X_r$ be the rank – among all of the agents – of the sample's median; in this section we shall assume that $X_r$ is normalized in the domain $[-1, 1]$. Thus, when $X_r = 0$ the median of the sample coincides with the median among the entire instance. Through this prism, we can determine the probability mass function with simple combinatorial arguments as follows:

$$\Pr\left(X_r = \frac{i}{\kappa}\right) = \frac{\binom{\kappa - i}{\rho}\binom{\kappa + i}{\rho}}{\binom{2\kappa + 1}{2\rho + 1}}. \tag{6.1}$$

As a result, we note that the normalization constraint of the probability mass function (6.1) yields a variant of the Chu-Vandermonde identity:

$$\sum_{i=-\kappa}^{\kappa} \binom{\kappa - i}{\rho}\binom{\kappa + i}{\rho} = \sum_{i=0}^{2\kappa} \binom{i}{\rho}\binom{2\kappa - i}{\rho} = \binom{2\kappa + 1}{2\rho + 1}. \tag{6.2}$$

For this reason, the distribution defined in (6.1) shall be referred to as Chu-Vandermonde distribution. One of the key aspects of our analysis is that we are oblivious to the agents' individual valuations and instead, we rely solely on their relative rank. This approach is justified by the following lemma:

**Lemma 6.1.1.** *Let $x_{opt} \in \mathbb{R}$ be the optimal location - i.e. the median of the instance - and $x \in \mathbb{R}$ some location, such that only at most $\epsilon \cdot n$ agents reside in the interval from $x$ to $x_{opt}$. Then, if $D_{opt}$ is the minimum social cost, allocating a facility on $x$ yields a social cost $D$ such that*

$$D \leq D_{opt}\left(1 + \frac{4\epsilon}{1 - 2\epsilon}\right). \tag{6.3}$$

*Proof.* Let $d = \mathrm{dis}(x, x_{opt}) = |x - x_{opt}|$; shifting the facility from $x$ to $x_{opt}$ can only reduce the social cost by at most $2\epsilon n d$, that is $D \leq D_{opt} + 2\epsilon n d$. Moreover, it is

---

**Algorithm 5:** Approximate Median through Sampling

---

**Result:** Facility's Location $\in \mathbb{R}^d$
**Input**: Set of agents $N$, size of sample $c$
$S :=$ random sample of $c$ agents from $N$
return $\mathcal{M}(S)$

---

clear that

$$D_{opt} \geq \left(\frac{n}{2} - \epsilon n\right) d \iff d \leq D_{opt} \frac{2}{n(1-2\epsilon)}. \tag{6.4}$$

As a result, if we combine the previous bounds the lemma will follow. We should mention that the analysis and subsequently the obtained bound is tight for certain instances. ■

As a corollary, obtaining a near-optimal approximation ratio is tantamount to accumulating the probability mass close to the median. The main challenge is to quantify this concentration as a function of the sample's size. To this end, we prove that for $\kappa \to +\infty$ the Chu-Vandermonde distribution converges to a continuous function, allowing for a concise characterization of the concentration.

**Theorem 6.1.1.** *If we let $\kappa \to \infty$, the Chu-Vandermonde distribution converges to a transformed beta distribution with the following probability density function:*

$$f(t) = \frac{(2\rho+1)!}{(\rho!)^2 2^{2\rho+1}} (1-t^2)^{\rho}. \tag{6.5}$$

Before we prove this theorem, we commence our analysis with the following auxiliary lemmas:

**Lemma 6.1.2.** *Let $\Gamma(\cdot)$ denote the Gamma function; then, for every $a \in \mathbb{R}$*

$$\lim_{n \to +\infty} \frac{\Gamma(n+a)}{\Gamma(n)} n^{-a} = 1. \tag{6.6}$$

**Lemma 6.1.3.** *Let $f$ be an integrable function and $x \in [-1,1]$; then*

$$\lim_{n \to +\infty} \frac{x+1}{n} \sum_{i=1}^{n} f\left(-1 + i \cdot \frac{x+1}{n}\right) = \int_{-1}^{x} f(t)dt. \tag{6.7}$$

We are now ready to prove Theorem 6.1.1.

*Proof.* Consider some arbitrary $x \in (-1,1)$ and $m = \lfloor \kappa x + \kappa \rfloor$; then

$$
\begin{aligned}
\lim_{\kappa\to+\infty} \Pr(X_r \le x) &= \lim_{\kappa\to+\infty} \sum_{i=-\kappa}^{m-\kappa-1} \frac{\binom{\kappa-i}{\rho}\binom{\kappa+i}{\rho}}{\binom{2\kappa+1}{2\rho+1}} \\
&= \lim_{n\to+\infty} \sum_{i=1}^{m} \frac{\binom{i-1}{\rho}\binom{n-i}{\rho}}{\binom{n}{2\rho+1}} \\
&= \lim_{n\to+\infty} \sum_{i=1}^{m} \frac{(2\rho+1)!}{(\rho!)^2} \frac{\frac{\Gamma(i)}{\Gamma(i-\rho)}\frac{\Gamma(n+1-i)}{\Gamma(n+1-i-\rho)}}{\frac{\Gamma(n+1)}{\Gamma(n-2\rho)}} \\
&= \frac{(2\rho+1)!}{(\rho!)^2} \lim_{n\to+\infty} \sum_{i=1}^{m} \frac{(i-\rho)^\rho (n-i-\rho)^\rho}{(n-2\rho)^{2\rho+1}} \\
&= \frac{(2\rho+1)!}{(\rho!)^2} \lim_{n\to+\infty} \frac{1}{n} \sum_{i=1}^{m} \left(\frac{i}{n} - \left(\frac{i}{n}\right)^2\right)^\rho \\
&= \frac{(2\rho+1)!}{(\rho!)^2 4^\rho} \lim_{m\to+\infty} \frac{x+1}{2m} \sum_{i=1}^{m} \left(2i\cdot\frac{x+1}{m} - \left(i\cdot\frac{x+1}{m}\right)^2\right)^\rho \\
&= \frac{(2\rho+1)!}{(\rho!)^2 2^{2\rho+1}} \lim_{m\to+\infty} \frac{x+1}{m} \sum_{i=1}^{m} \left(1 - \left(-1+i\cdot\frac{x+1}{m}\right)^2\right)^\rho \\
&= \frac{(2\rho+1)!}{(\rho!)^2 2^{2\rho+1}} \int_{-1}^{x} (1-t^2)^\rho dt,
\end{aligned}
$$

where we have applied Lemma 6.1.2 in the fourth line and Lemma 6.1.3 in the last line. Also note that we used Cesaro's means theorem in the fourth line and we ignored the lower-order magnitude terms in the fifth. ■

The distribution (6.5) derives from the Beta family, having applied a quadratic transformation. Indeed, recall that for any $x, y \in \mathbb{R}^+$, the Beta function B is defined as

$$
B(x,y) = \int_0^1 t^{x-1}(1-t)^{y-1}dt = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}. \tag{6.8}
$$

Moreover, we remark the following useful lemma:

**Lemma 6.1.4.** *For any $n \in \mathbb{N}$ we have that*

$$\Gamma\left(\frac{1}{2}+n\right) = \frac{(2n)!}{4^n n!}\sqrt{\pi}. \tag{6.9}$$

Thus, we can verify the normalization constraint using Lemma 6.1.4 and the quadratic transformation $u = t^2$ as follows.

$$\int_{-1}^{1}(1-t^2)^\rho dt = \int_0^1 u^{-\frac{1}{2}}(1-u)^\rho du = B\left(\frac{1}{2},\rho+1\right) = \frac{2^{2\rho+1}(\rho!)^2}{(2\rho+1)!}. \tag{6.10}$$

We let $X$ represent a random variable that follows distribution (6.5). Next, we correlate the concentration of the distribution with parameter $\rho$.

**Theorem 6.1.2.** *For any $\epsilon > 0$ and for any $\delta > 0$, there exists some constant $\rho_0 = \rho_0(\epsilon, \delta)$ such that $\forall \rho \geq \rho_0$*

$$\Pr(|X| \geq \epsilon) \leq \delta. \tag{6.11}$$

We commence the proof by determining the moments of $|X|$:

**Lemma 6.1.5.** *Let $j \in \mathbb{N}$; we can express the $j^{th}$ moment of $|X|$ as*

$$\mathbb{E}[|X|^j] = \frac{B\left(\frac{j}{2}+\frac{1}{2},\rho+1\right)}{B\left(\frac{1}{2},\rho+1\right)}. \tag{6.12}$$

This lemma follows from standard techniques in integration; moreover, we can establish the following trivial bound.

$$B\left(\frac{1}{2},\rho+1\right) = \frac{2^{2\rho+1}(\rho!)^2}{(2\rho+1)!} \geq \frac{2}{2\rho+1} \geq \frac{1}{\rho+1}. \tag{6.13}$$

Having established these auxiliary results we provide the proof of Theorem 6.1.2.

*Proof.* For simplicity, we consider some $j = 2\ell - 1 \geq 3$; then, it follows from Lemma 6.1.5 and (6.13) that

$$\mathbb{E}[|X|^j] = \frac{B\left(\frac{j}{2}+\frac{1}{2},\rho+1\right)}{B\left(\frac{1}{2},\rho+1\right)} \leq \frac{(\ell-1)!}{(\rho+\ell)\cdots(\rho+2)}. \tag{6.14}$$

Moreover, if we apply Markov's inequality for the $j^{\text{th}}$ moment of $X$ we have that

$$\Pr(|X| \geq \epsilon) = \Pr(|X|^j \geq \epsilon^j) \leq \frac{(\ell-1)!}{(\rho+\ell)\cdots(\rho+2)}\epsilon^{-j}. \tag{6.15}$$

As a result, for sufficiently large $\rho = \rho(\epsilon, \delta)$ the last bound will be also upper bounded by $\delta$ and the claim follows. ∎

Having established the concentration of the distribution, we apply Lemma 6.1.1 to prove the following theorem:

**Theorem 6.1.3.** *The approximate one-dimensional median Mechanism 5 obtains in expectation a $1+\epsilon$ approximation of the optimal social welfare, for any $\epsilon > 0$ and with constant input $c = c(\epsilon)$, while $n \to \infty$.*

*Proof.* Consider a random variable $X$ that follows the distribution 6.5 with some parameter $\rho$ and let $g : (0,1) \ni x \mapsto 2|x|/(1-|x|)$; we know from Lemma 6.1.1 that the approximation error in the social welfare can be upper bounded by $\mathbb{E}[g(X)]$; but, it follows that

$$\mathbb{E}[g(X)] = 4c(\rho) \int_0^1 \frac{t}{1-t}(1-t^2)^{\rho} dt \leq 8\frac{2\rho+1}{2\rho} c(\rho-1) \int_0^1 t(1-t^2)^{\rho-1} dt, \quad (6.16)$$

where $c(\rho)$ represents the normalization constant of 6.5. As a result, if $X'$ represents a random variable that follows the transformed beta distribution 6.5 with parameter $\rho - 1$, Theorem 6.1.2 implies that the mean value of $|X'|$ can be upper bounded by any $\epsilon > 0$ – for sufficiently large $\rho$; hence, the claim follows. ■

This result can be easily extended for the generalized median scheme that applies to any metric space $(\mathbb{R}^d, ||\cdot||_1)$; to be precise, let us consider some basis for the metric space. Then, we can invoke the one-dimensional sampling approximation for each of the principal axes individually. As a result, we can prove the following proposition.

**Corollary 6.1.1.** *The approximate generalized median Mechanism 5 obtains in expectation a $1+\epsilon$ approximation of the optimal social welfare, for any $\epsilon > 0$ and with constant input $c = c(\epsilon)$, while $n \to \infty$.*

Finally, we illustrate why a sampling approach – with a constant sample size – cannot provide meaningful guarantees when allocating multiple facilities. In particular, we consider the family of the *percentile* mechanisms, namely strategyproof allocation rules on the line that partition the agents' reports into particular percentiles; the median can be clearly classified in this family. We will also assume that at least 2 facilities are to be allocated and that the leftmost percentile contains at most $(1-\alpha)\cdot n$ of the agents, for some $\alpha > 0$. Let us imagine a dynamic instance where the agents from the entire leftmost percentile – including the pivotal agent – have gradually smaller valuations $x \to -\infty$, while the other agents remain fixed at a finite distance; then, any sampling approximation has in expectation an unbounded competitive ratio with respect to the full information mechanism. Indeed, there will always be a positive probability, albeit exponentially small, that we fail to sample a single agent from $-\infty$, whilst the full information percentile mechanism will allocate a facility to accommodate the divergent agents. Thus, a sampling approach cannot provide a meaningful approximation of the percentile mechanisms – at least with respect to the expected social cost. An interesting open question is whether this limitation can be overcome if we impose additional restrictions on the instance, such as stability conditions or bounded valuation space.

### 6.1.1 Chu-Vandermonde Distribution

To the best of our knowledge, the distribution defined in (6.1) – which we refer to as the Chu-Vandermonde distribution – has not been rigorously studied in the literature; one of our key results establishes an asymptotic characterization of the distribution and an inherent nexus with the family of Beta distributions. Additional analysis could be of independent interest and illuminate further interesting properties. The following figures provide a graphical representation of the distribution and illustrate its behavior for gradually extended samples, as well as the convergence that occurs when $\kappa$ gently increases.

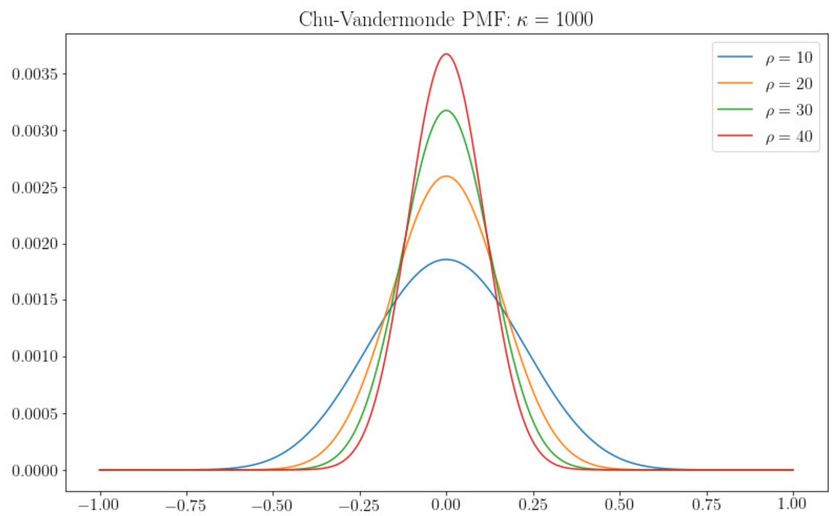


Figure 6.1: Augmenting $\rho$ accumulates the probability mass close to the median

Although our main result in Theorem 6.1.1 has an asymptotic flavour, the last figure – as well as additional experimental findings – indicate that the convergence of the distribution occurs with a substantial rate, even for relatively small values of $\kappa$.

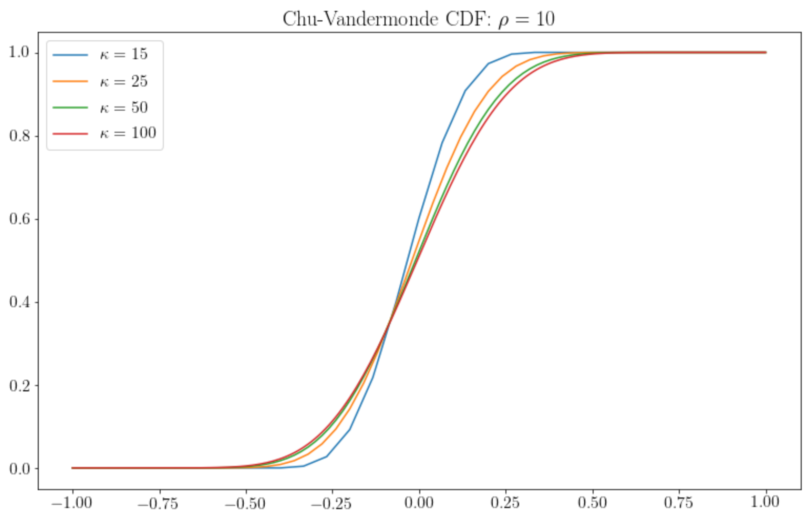


Figure 6.2: The rapid convergence of the distribution while $\kappa$ increases

# Bibliography


[1] Adrian K. Agogino and Kagan Tumer. "Handling Communication Restrictions and Team Formation in Congestion Games". In: *Auton. Agents Multi Agent Syst.* 13.1 (2006), pp. 97–115.

[2] Alfred V. Aho, Jeffrey D. Ullman, and Mihalis Yannakakis. "On Notions of Information Transfer in VLSI Circuits". In: *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*. STOC '83. Association for Computing Machinery, 1983, pp. 133–139.

[3] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. "Approximating Optimal Social Choice under Metric Preferences". In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. Ed. by Blai Bonet and Sven Koenig. AAAI Press, 2015, pp. 777–783.

[4] Kenneth J. Arrow. *Social Choice and Individual Values*. Yale University Press, 2012.

[5] Lawrence M. Ausubel. "An Efficient Ascending-Bid Auction for Multiple Objects". In: *American Economic Review* 94.5 (Dec. 2004), pp. 1452–1475.

[6] Lawrence M. Ausubel and Paul Milgrom. "The Lovely but Lonely Vickrey Auction". In: *Combinatorial Auctions, chapter 1*. MIT Press, 2006.

[7] László Babai, Peter G. Kimmel, and Satyanarayana V. Lokam. "Simultaneous Messages vs. Communication". In: *STACS 95, 12th Annual Symposium on Theoretical Aspects of Computer Science*. Ed. by Ernst W. Mayr and Claude Puech. Vol. 900. Lecture Notes in Computer Science. Springer, 1995, pp. 361–372.

[8] Yakov Babichenko and Aviad Rubinstein. "Communication Complexity of Approximate Nash Equilibria". In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2017. Montreal, Canada: Association for Computing Machinery, 2017, pp. 878–889. ISBN: 9781450345286.

[9] Duncan Black. Cambridge: Cambridge University Press, 1958.

[10] Duncan Black. "On the Rationale of Group Decision-making". In: *Journal of Political Economy* 56.1 (1948), pp. 23–34.

[11] Avrim Blum et al. "Preference Elicitation and Query Learning". In: *J. Mach. Learn. Res.* 5 (Dec. 2004), pp. 649–667. ISSN: 1532-4435.

[12] Liad Blumrosen, Noam Nisan, and Ilya Segal. "Auctions with Severely Bounded Communication". In: *J. Artif. Intell. Res.* 28 (2007), pp. 233–266.

[13] Ran Canetti, Guy Even, and Oded Goldreich. "Lower bounds for sampling algorithms for estimating the average". In: *Information Processing Letters* 53.1 (1995), pp. 17–25.

[14] Edward H. Clarke. "Multipart Pricing of Public Goods". In: *Public Choice* 11 (1971), pp. 17–33.

[15] Wolfram Conen and Tuomas Sandholm. "Preference Elicitation in Combinatorial Auctions". In: *Proceedings of the 3rd ACM Conference on Electronic Commerce*. EC '01. Tampa, Florida, USA: Association for Computing Machinery, 2001, pp. 256–259.

[16] Vincent Conitzer. "Eliciting Single-Peaked Preferences Using Comparison Queries". In: *J. Artif. Int. Res.* 35.1 (2009), pp. 161–191. ISSN: 1076-9757.

[17] Vincent Conitzer and Tuomas Sandholm. "Communication Complexity of Common Voting Rules". In: *Proceedings of the 6th ACM Conference on Electronic Commerce*. EC '05. New York, NY, USA: Association for Computing Machinery, 2005, pp. 78–87.

[18] Vincent Conitzer and Tuomas Sandholm. "Vote Elicitation: Complexity and Strategy-Proofness". In: *Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence*. AAAI Press / The MIT Press, 2002, pp. 392–397.

[19] Thomas H. Cormen et al. *Introduction to Algorithms, Third Edition*. 3rd. The MIT Press, 2009. ISBN: 0262033844.

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

[21] Artur Czumaj et al. "Distributed Methods for Computing Approximate Equilibria". In: *Algorithmica* 81.3 (June 2018), pp. 1205–1231. ISSN: 1432-0541.

[22] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. “The complexity of computing a Nash equilibrium”. In: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*. Ed. by Jon M. Kleinberg. ACM, 2006, pp. 71–78.

[23] Esther David et al. “Optimal Design Of English Auctions With Discrete Bid Levels”. In: *ACM Conference on Electronic Commerce (EC'05)*. 2005, pp. 98–107.

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

[25] Nikhil R. Devanur, Milena Mihail, and Vijay V. Vazirani. “Strategyproof cost-sharing mechanisms for set cover and facility location games”. In: *Decision Support Systems* 39.1 (2005). The Fourth ACM Conference on Electronic Commerce, pp. 11–22.

[26] Palash Dey and Arnab Bhattacharyya. “Sample Complexity for Winner Prediction in Elections”. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. AAMAS '15. International Foundation for Autonomous Agents and Multiagent Systems, 2015, pp. 1421–1430.

[27] Swapnil Dhamal and Y. Narahari. “Scalable Preference Aggregation in Social Networks”. In: *HCOMP*. 2013.

[28] Shahar Dobzinski and Jan Vondrák. “Communication Complexity of Combinatorial Auctions with Submodular Valuations”. In: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013*. Ed. by Sanjeev Khanna. SIAM, 2013, pp. 1205–1215.

[29] Ronald Fadel and Ilya Segal. “The communication cost of selfishness”. In: *Journal of Economic Theory* 144.5 (2009), pp. 1895–1920. ISSN: 0022-0531.

[30] Michal Feldman, Amos Fiat, and Iddan Golomb. “On Voting and Facility Location”. In: *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16*. Ed. by Vincent Conitzer, Dirk Bergemann, and Yiling Chen. ACM, 2016, pp. 269–286.

[31] Dimitris Fotakis and Christos Tzamos. “Winner-Imposing Strategyproof Mechanisms for Multiple Facility Location Games”. In: *Internet and Network Economics*. Ed. by Amin Saberi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 234–245.

[32] John Geanakoplos. “Three Brief Proofs of Arrow's Impossibility Theorem”. In: *Economic Theory* 26.1 (2005), pp. 211–215.

[33] Allan Gibbard. “Manipulation of Voting Schemes: A General Result”. In: *Econometrica* 41.4 (1973), pp. 587–601.

[34] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. “Resolving the Optimal Metric Distortion Conjecture”. In: *CoRR* abs/2004.07447 (2020).

[35] Paul W. Goldberg and Arnoud Pastink. “On the Communication Complexity of Approximate Nash Equilibria”. In: *Algorithmic Game Theory*. Ed. by Maria Serna. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 192–203. ISBN: 978-3-642-33996-7.

[36] Hans Dietmar Gröger and György Turán. “On Linear Decision Trees Computing Boolean Functions”. In: *Proceedings of the 18th International Colloquium on Automata, Languages and Programming*. Madrid, Spain: Springer-Verlag, 1991, pp. 707–718.

[37] Theodore Groves. “Incentives in Teams”. In: *Econometrica* 41.4 (1973), pp. 617–631.

[38] C. A. R. Hoare. “Algorithm 65: Find”. In: *Commun. ACM* 4.7 (July 1961), pp. 321–322.

[39] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation, 3rd Edition*. Pearson international edition. Addison-Wesley, 2007.

[40] Benoit Hudson and Tuomas Sandholm. “Effectiveness of Query Types and Policies for Preference Elicitation in Combinatorial Auctions”. In: *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*. AAMAS ’04. New York, New York: IEEE Computer Society, 2004, pp. 386–393.

[41] D. A. Huffman. “A Method for the Construction of Minimum-Redundancy Codes”. In: *Proceedings of the IRE* 40.9 (1952), pp. 1098–1101.

[42] Leonid Hurwicz. “Optimality and informational efficiency in resource allocation processes”. In: *Studies in Resource Allocation Processes*. Ed. by Kenneth J. Arrow and LeonidEditors Hurwicz. Cambridge University Press, 1977, pp. 393–460.

[43] Leonid Hurwicz. “The Design of Mechanisms for Resource Allocation”. In: *The American Economic Review* 63.2 (1973), pp. 1–30.

[44] Leonid Hurwicz. “The Theory of Economic Behavior”. In: *The American Economic Review* 35.5 (1945), pp. 909–925.

[45] Leonid Hurwicz and Stanley Reiter. *Designing Economic Mechanisms*. Cambridge University Press, 2006.

[46] John H. Kagel, Ronald M. Harstad, and Dan Levin. "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study". In: *Econometrica* 55.6 (1987), pp. 1275–1304. ISSN: 00129682, 14680262.

[47] John H. Kagel and Dan Levin. "Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders". In: *The Economic Journal* 103.419 (1993), pp. 868–879. ISSN: 0013-0133.

[48] Shizuo Kakutani. "A generalization of Brouwer's fixed point theorem". In: *Duke Math. J.* 8.3 (Sept. 1941), pp. 457–459.

[49] David Kempe. "An Analysis Framework for Metric Voting based on LP Duality". In: *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020.* AAAI Press, 2020, pp. 2079–2086.

[50] Donald E Knuth. "Dynamic huffman coding". In: *Journal of Algorithms* 6.2 (1985), pp. 163–180.

[51] Eyal Kushilevitz. "Communication Complexity". In: ed. by Marvin V. Zelkowitz. Vol. 44. Advances in Computers. Elsevier, 1997, pp. 331–360.

[52] Eyal Kushilevitz and Noam Nisan. *Communication Complexity.* Cambridge University Press, 1996.

[53] W. Leininger, P. B. Linhart, and R. Radner. "Equilibria of the sealed-bid mechanism for bargaining with incomplete information". In: *Journal of Economic Theory* 48.1 (1989), pp. 63–106.

[54] Thomas Lengauer. "VLSI Theory". In: *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity.* Ed. by Jan van Leeuwen. Elsevier and MIT Press, 1990, pp. 835–868.

[55] Herman B Leonard. "Elicitation of Honest Preferences for the Assignment of Individuals to Positions". In: *Journal of Political Economy* 91.3 (1983).

[56] Stefano Leonardi and Guido Schaefer. "Cross-Monotonic Cost-Sharing Methods for Connected Facility Location Games". In: *Proceedings of the 5th ACM Conference on Electronic Commerce.* EC '04. New York, NY, USA: Association for Computing Machinery, 2004, pp. 242–243.

[57] Shengwu Li. “Obviously Strategy-Proof Mechanisms”. In: *American Economic Review* 107.11 (Nov. 2017), pp. 3257–87.

[58] Shi Li. “A 1.488 approximation algorithm for the uncapacitated facility location problem”. In: *Information and Computation* 222 (2013). 38th International Colloquium on Automata, Languages and Programming (ICALP 2011), pp. 45–58.

[59] Xiang-Yang Li, Zheng Sun, and Weizhao Wang. “Cost Sharing and Strategyproof Mechanisms for Set Cover Games”. In: *STACS 2005*. Ed. by Volker Diekert and Bruno Durand. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 218–230.

[60] Richard J. Lipton and Robert Sedgewick. “Lower Bounds for VLSI”. In: *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing*. STOC ’81. Milwaukee, Wisconsin, USA: Association for Computing Machinery, 1981, pp. 300–307. ISBN: 9781450373920.

[61] Richard J. Lipton and Robert Endre Tarjan. “Applications of a Planar Separator Theorem”. In: *SIAM J. Comput.* 9.3 (1980), pp. 615–627.

[62] Christian List. “The theory of judgment aggregation: an introductory review”. In: *Synthese* 187.1 (2012), pp. 179–207.

[63] L. Lovasz and M. Saks. “Lattices, mobius functions and communications complexity”. In: *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*. 1988, pp. 81–90.

[64] Eric Maskin. “Auctions and Privatization”. In: *Privatization*. Ed. by Horst Siebert. J.C.B. Mohr Publisher, 1992, pp. 115–136.

[65] Eric Maskin, P. Dasgupta, and P. Hammond. “The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility”. In: *Review of Economic Studies* 46.2 (1979), pp. 185–216.

[66] ERIC MASKIN et al. *The Arrow Impossibility Theorem*. Columbia University Press, 2014.

[67] Kurt Mehlhorn and Erik M. Schmidt. “Las Vegas is Better than Determinism in VLSI and Distributed Computing (Extended Abstract)”. In: *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*. STOC ’82. Association for Computing Machinery, pp. 330–337.

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

[69] Peter Bro Miltersen. “Lower Bounds for Union-Split-Find Related Problems on Random Access Machines”. In: *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*. STOC ’94. Association for Computing Machinery, 1994, pp. 625–634.

[70] Peter Bro Miltersen et al. “On Data Structures and Asymmetric Communication Complexity”. In: *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing*. STOC ’95. Association for Computing Machinery, 1995, pp. 103–111.

[71] Dilip Mookherjee and Masatoshi Tsumagari. “Mechanism Design with Communication Constraints”. In: *Journal of Political Economy* 122.5 (2014), pp. 1094–1129. ISSN: 00223808, 1537534X.

[72] H. Moulin. “On Strategy-Proofness and Single Peakedness”. In: *Public Choice* 35.4 (1980), pp. 437–455. ISSN: 00485829, 15737101.

[73] Hervé Moulin. *Handbook of Computational Social Choice*. Ed. by Felix Brandt et al. Cambridge University Press, 2016.

[74] Kamesh Munagala and Kangning Wang. “Improved Metric Distortion for Deterministic Social Choice Rules”. In: *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019*. Ed. by Anna Karlin, Nicole Immorlica, and Ramesh Johari. ACM, 2019, pp. 245–262.

[75] Saburo Muroga. *Threshold logic and its applications*. Wiley, 1971.

[76] Roger B. Myerson. “Optimal Auction Design”. In: *Mathematics of Operations Research* 6.1 (1981), pp. 58–73.

[77] John F. Nash. “Equilibrium Points in n-Person Games”. In: *Proceedings of the National Academy of Science* 36.1 (Jan. 1950), pp. 48–49.

[78] John von Neumann, Oskar Morgenstern, and Ariel Rubinstein. *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press, 1944.

[79] N. Nisan and A. Wigderson. “On rank vs. communication complexity”. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 1994, pp. 831–836.

[80] Noam Nisan. “The Communication Complexity of Threshold Gates”. In: *In Proceedings of “Combinatorics, Paul Erdos is Eighty*. 1994, pp. 301–315.

[81] Noam Nisan and Ilya Segal. “The communication requirements of efficient allocations and supporting prices”. In: *Journal of Economic Theory* 129.1 (2006), pp. 192–224. ISSN: 0022-0531.

[82] Noam Nisan and Avi Wigderson. “On Rank vs. Communication Complexity”. In: *Comb.* 15.4 (1995), pp. 557–565.

[83] Noam Nisan and Avi Wigderson. “Rounds in Communication Complexity Revisited”. In: *SIAM J. Comput.* 22.1 (1993), pp. 211–219.

[84] Joel Oren, Yuval Filmus, and Craig Boutilier. “Efficient Vote Elicitation under Candidate Uncertainty”. In: *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*. IJCAI ’13. AAAI Press, 2013, pp. 309–316.

[85] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. Vol. 1. MIT Press Books. The MIT Press, 1994.

[86] Martin Pál and Éva Tardo. “Group Strategyproof Mechanisms via Primal-Dual Algorithms”. In: *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*. FOCS ’03. USA: IEEE Computer Society, 2003, p. 584.

[87] David C. Parkes, Lyle H. Ungar, and Dean P. Foster. “Accounting for Cognitive Costs in On-line Auction Design”. In: *Agent Mediated Electronic Commerce*. Ed. by Pablo Noriega and Carles Sierra. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 25–40. ISBN: 978-3-540-48835-4.

[88] R. Paturi. “Discrete Neural Computation: A Theoretical Foundation [Book Review]”. In: *IEEE Transactions on Neural Networks* 7.2 (1996), pp. 533–.

[89] A.A. Razborov. “On the distributional complexity of disjointness”. In: *Theoretical Computer Science* 106.2 (1992), pp. 385–390. ISSN: 0304-3975.

[90] Stefan Reichelstein. “Incentive compatibility and informational requirements”. In: *Journal of Economic Theory* 34.1 (1984), pp. 32–51. ISSN: 0022-0531.

[91] John G. Riley and William F. Samuelson. “Optimal Auctions”. In: *The American Economic Review* 71.3 (1981), pp. 381–392.

[92] V. P. Roychowdhury, A. Orlitsky, and Kai-Yeung Siu. “Lower bounds on threshold and related circuits via communication complexity”. In: *IEEE Transactions on Information Theory* 40.2 (1994), pp. 467–474.

[93] Mark Allen Satterthwaite. “Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions”. In: *Journal of Economic Theory* 10.2 (1975), pp. 187–217. ISSN: 0022-0531.

[94] Ilya Segal. “The communication requirements of social choice rules and supporting budget sets”. In: *Journal of Economic Theory* 136.1 (2007), pp. 341–378. ISSN: 0022-0531.

[95] L.S. Shapley and M. Shubik. In: *International Journal of Game Theory* 1 (1971), pp. 111–130.

[96] David B. Shmoys. “Approximation Algorithms for Facility Location Problems”. In: *Approximation Algorithms for Combinatorial Optimization*. Ed. by Klaus Jansen and Samir Khuller. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 27–32.

[97] Piotr Krzysztof Skowron and Edith Elkind. “Social Choice Under Metric Preferences: Scoring Rules and STV”. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. Ed. by Satinder P. Singh and Shaul Markovitch. AAAI Press, 2017, pp. 706–712.

[98] H. Smith, Tamara Dinev, and Heng Xu. “Information Privacy Research: An Interdisciplinary Review”. In: *MIS Quarterly* 35 (2011), pp. 989–1015.

[99] Nguyen Kim Thang. “On (Group) Strategy-Proof Mechanisms without Payment for Facility Location Games”. In: *Internet and Network Economics*. Ed. by Amin Saberi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 531–538.

[100] Timothy Van Zandt. “Communication Complexity and Mechanism Design”. In: *Journal of the European Economic Association* 5.2/3 (2007), pp. 543–553. ISSN: 15424766, 15424774.

[101] William Vickrey. “Counterspeculation, Auctions, and Competitive Sealed Tenders”. In: *The Journal of Finance* 16.1 (1961), pp. 8–37.

[102] Andrew Chi-Chih Yao. “Some Complexity Questions Related to Distributive Computing(Preliminary Report)”. In: *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*. STOC ’79. New York, NY, USA: Association for Computing Machinery, 1979, pp. 209–213.

[103] M. M. Zavlanos, L. Spesivtsev, and G. J. Pappas. “A distributed auction algorithm for the assignment problem”. In: *2008 47th IEEE Conference on Decision and Control*. 2008, pp. 1212–1217.