



ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΤΜΗΜΑ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ
ΥΠΟΛΟΓΙΣΤΩΝ

ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ
ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ ΚΑΙ ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΜΩΝ (Co.Re.Lab.)

**Mechanism Design, Social Choice and Equilibrium
Computation in Restricted Domains**

ΔΙΔΑΚΤΟΡΙΚΗ ΔΙΑΤΡΙΒΗ

ΤΟΥ

Παναγιώτη
Πασιλινάκου

Αθήνα, Σεπτέμβριος 2023



ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΤΜΗΜΑ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ
ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ ΚΑΙ ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΜΩΝ

Mechanism Design, Social Choice and Equilibrium Computation in Restricted Domains

Διδακτορική Διατριβή
του

**Παναγιώτη
Πατσιλινάκου**

Συμβουλευτική Επιτροπή: Δημήτριος Φωτάκης (Επιβλέπων)
Αριστείδης Παγουρτζής
Αντώνιος Συμβώνης

Εγκρίθηκε από την επαμελή εξεταστική επιτροπή την 8η Σεπτεμβρίου 2023.

.....
Δημήτριος Φωτάκης
Καθηγητής
ΕΜΠ

.....
Αριστείδης Παγουρτζής
Καθηγητής
ΕΜΠ

.....
Αντώνιος Συμβώνης
Καθηγητής
ΕΜΠ

.....
Γεώργιος Χριστοδούλου
Αναπληρωτής Καθηγητής
ΑΠΘ

.....
Ιωάννης Καραγιάννης
Καθηγητής
Aarhus University

.....
Ευάγγελος Μαρχάκης
Αναπληρωτής Καθηγητής
ΟΠΑ

.....
Εμμανουήλ Βαρβαρίος
Καθηγητής
ΕΜΠ

Αθήνα, Σεπτέμβριος 2023.

.....

Παναγιώτης Πατσιλινάκος

Διδάκτωρ Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών Ε.Μ.Π.

Copyright © Παναγιώτης Σ. Πατσιλινάκος 2023.

Με επιφύλαξη παντός δικαιώματος. All rights reserved.

Η Εργασία διατίθεται με άδεια Creative Commons Αναφορά Δημιουργού 4.0 Διεθνή. Για να δείτε ένα αντίγραφο αυτής της άδειας, επισκεφθείτε το <http://creativecommons.org/licenses/by/4.0/> ή στείλετε επιστολή στο Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Οι απόψεις και τα συμπεράσματα που περιέχονται σε αυτό το έγγραφο εκφράζουν τον συγγραφέα και δεν πρέπει να ερμηνευθεί ότι αντιπροσωπεύουν τις επίσημες θέσεις του Εθνικού Μετσόβιου Πολυτεχνείου.

στους γονείς και τον αδερφό μου

στην μνήμη του παππού Παναγιώτη, που καθοδηγεί τις πράξεις μου

Abstract

The work in this thesis primarily revolves around efficient algorithmic frameworks for settings where information is not readily available. Specifically, we look at limitations of provided information from three main angles: (1) Information is difficult to quantify. In this line of work we focused on distortion in voting (JAIR'22, AAI'22), which is the notion that quantifies the impact of being able to use only limited information on the social welfare of the outcome (i.e. in terms of approximation). Here we study both the effects of various forms of limited information on metric distortion and also the distortion of a very popular mechanism, STV, in relation to the dimensionality of the underlying metric space. (2) Information is private to strategic agents and needs to be revealed to the algorithm through properly designed incentives. This area is commonly referred to as mechanism design and my related work focuses on fighting strong impossibility results by restricting our analysis in “natural” sub-classes of the general instance space (WINE'21). In this setting we have studied the approximability of the facility location problem by truthful mechanisms, whose allocation is aligned with the agent incentives. (3) Communication is expensive. Combining this restriction along with the strategic environment described previously, we show that known mechanisms have implementations with asymptotically optimal communication complexity (SAGT'20, full version under minor revision TOCS'23). In most of our works, our objective is to maximize the social welfare. Furthermore, some work has been focused on a classical aspect of algorithmic game theory, that of equilibrium computation, where we study the complexity of computing a Pure Nash Equilibrium in linear weighted congestion games and also show an efficient algorithm for computing approximate equilibria in a natural subclass of Max-Cut games (ICALP'20).

Keywords

Algorithmic Game Theory, Algorithmic Social Choice, Mechanism Design, Algorithm Design, Beyond Worst Case Analysis, Approximation Algorithms, Computational Complexity, Communication Complexity, Equilibrium Computation, Metric Distortion.

Αφαιρετική περίληψη

Η παρούσα διδακτορική διατριβή μελετάει αποδοτικά αλγοριθμικά πλαίσια για περιβάλλοντα στα οποία η πληροφορία δεν είναι άμεσα διαθέσιμη. Πιο συγκεκριμένα μελετάει περιορισμούς στην πρόσβαση στην πληροφορία από τρεις διαφορετικές γωνίες: (1) Η πληροφορία είναι προσωπική (και ιδιωτική) σε στρατηγικούς παίχτες και χρειάζεται να αποκαλυφθεί στον αλγόριθμο μέσα από κατάλληλα σχεδιασμένα κίνητρα: Αυτή η περιοχή συνήθως αναφέρεται ως «αλγοριθμικός σχεδιασμός μηχανισμών». Η έρευνα στα πλαίσια της διατριβής επικεντρώνεται στην αντιμετώπιση ισχυρών αρνητικών αποτελεσμάτων περιορίζοντας την ανάλυση σε «φυσιολογικά» υποσύνολα στιγμιότυπων του προβλήματος, μια πρακτική από την περιοχή της ανάλυσης πέραν της χειρότερης περίπτωσης στη θεωρία αλγορίθμων. Συγκεκριμένα, αναλύεται η προσεγγισμότητα του προβλήματος χωροθέτησης εγκαταστάσεων από φιλαλήθεις μηχανισμούς. (2) Η επικοινωνία είναι ακριβή: Μελετώντας αυτόν τον περιορισμό σε περιβάλλοντα με στρατηγικούς παίχτες αποδεικνύεται ότι απλοί μηχανισμοί για κλασικά προβλήματα μπορούν να υλοποιηθούν με ασυμπτωτικά βέλτιστη πολυπλοκότητα επικοινωνίας (ανάμεσα στον μηχανισμό και τους παίχτες). (3) Η πληροφορία που χρησιμοποιείται από τον μηχανισμό είναι περιορισμένη: η διατριβή εστιάζει στην παραμόρφωση στο πρόβλημα της ψηφοφορίας, που είναι η έννοια που ποσοτικοποιεί την επίπτωση της πρόσβασης σε περιορισμένη πληροφορία στην κοινωνική ωφέλεια του αποτελέσματος ενός αλγορίθμου (σε όρους προσέγγισης της βέλτιστης λύσης). Εδώ μελετώνται και οι επιπτώσεις διάφορων μορφών περιορισμένης πληροφορίας στην μετρική παραμόρφωση αλλά και η παραμόρφωση ενός πολύ διαδεδομένου μηχανισμού, του STV, σε σχέση με την διαστασιμότητα του σχετικού μετρικού χώρου. Επιπλέον, στα πλαίσια της διατριβής διερευνάται η κλασική πτυχή της αλγοριθμικής θεωρίας παιγνίων που αφορά την πολυπλοκότητα υπολογισμού αμιγών ισορροπιών. Μελετάται το πρόβλημα αυτό στο πλαίσιο των βεβαρυμένων παιγνίων συμφόρησης με γραμμικές συναρτήσεις καθυστέρησης, και παρουσιάζεται ένας αποδοτικός αλγόριθμος για τον υπολογισμό προσεγγιστικών ισορροπιών σε μια ενδιαφέρουσα κλάση των Max-Cut παιγνίων.

Λέξεις Κλειδιά

Αλγοριθμική Θεωρία Παιγνίων, Αλγοριθμική Κοινωνική Επιλογή, Σχεδιασμός Μηχανισμών, Σχεδιασμός Αλγορίθμων, Ανάλυση Πέραν της Χειρότερης Περίπτωσης, Προσεγγιστικοί Αλγόριθμοι, Υπολογιστική Πολυπλοκότητα, Πολυπλοκότητα Επικοινωνίας, Υπολογισμός Ισορροπιών, Μετρική Παραμόρφωση.

Ευχαριστίες

Η εκπόνηση της διατριβής μου κατέστη δυνατή χάρη στην συνδρομή πολλών ανθρώπων για τους οποίους νιώθω βαθεία ευγνωμοσύνη.

Πρωτίστως θα ήθελα να ευχαριστήσω τον επιβλέποντα μου, Δημήτρη Φωτάκη, για την εμπιστοσύνη, την κατανόηση και την αμέριστη συμπαράσταση του όλα αυτά τα χρόνια. Η διατριβή μου, στην εκπόνηση της οποίας πρωτοστάτησε, θα ήταν αδύνατη χωρίς εκείνον και την συνδρομή του σε όλα τα επίπεδα, ακαδημαϊκό, επιστημονικό και φιλικό. Τον ευχαριστώ πραγματικά από καρδιάς.

Πολλές ευχαριστίες οφείλω επίσης και στους καθηγητές, Στάθη Ζάχο και Αριστείδη Παγουρτζή μέσα από τους οποίους γνώρισα για πρώτη φορά το εργαστήριο CoReLab το οποίο αισθάνομαι πλέον ως ακαδημαϊκό σπίτι μου. Επιπλέον θα ήθελα να ευχαριστήσω όλα τα μέλη της επταμελούς μου επιτροπής, Δημήτριο Φωτάκη, Αριστείδη Παγουρτζή, Αντώνιο Συμβώνη, Γεώργιο Χριστοδούλου, Ιωάννη Καραγιάννη, Ευάγγελο Μαρκάκη και Εμμανουήλ Βαρβαρίγο για την βοήθεια τους στην ολοκλήρωση της διατριβής μου και για την τιμή που μου έκαναν με την συμμετοχή τους. Θα ήθελα επίσης να πω και ένα ιδιαίτερο ευχαριστώ στον Ιωάννη Καραγιάννη η διδασκαλία και επίβλεψη του οποίου στην προπτυχιακή διπλωματική εργασία μου με ώθησε με τόσο ενδιαφέρον στην Αλγοριθμική Θεωρία Παιγνίων.

Βαθιές ευχαριστίες θέλω να απευθύνω και στους επιστημονικούς μου συνεργάτες, Δημήτρη Φωτάκη, Ιωάννη Αναγνωστίδη, Στρατή Σκουλάκη, Νίκο Μουζάκη, Άνθιμο-Βαρδή Κανδήρο, Θανάση Λιανέα και Άλκη Καλαβάση για την υπέροχες επιστημονικές συζητήσεις και την κομβική για την έρευνα της διατριβής μου συνεργασία.

Επιπλέον, θέλω ευχαριστήσω όλα τα παιδιά με τα οποία συμπορευτήκαμε στο CoReLab για τις αναμνήσεις και τις φιλίες που θα κρατήσουν για πάντα, Στρατή Σκουλάκη, Τζέλα Μαθιουδάκη, Άλκη Καλαβάση, Αγάπη Ρισσάκη, Λουκά Κάβουρα, Ελένη Ψαρουδάκη, Αγγέλα Χαλκή, Αντώνη Αντωνόπουλο, Θανάση Λιανέα, Μίλτο Στούρα, Σπύρο Δραγάζη, Κατερίνα Νικολιδάκη, Ναταλία Κωτσάνη, Pouran Behrouz, Έλλη Αναστασιάδη, Σωτήρη Δήμο και όλα τα υπόλοιπα μέλη που περάσαμε χρόνο μαζί από τα οποία απολογούμαι καθώς σίγουρα αξίζουν μια θέση στην παραπάνω λίστα ονομαστικά. Ευχαριστώ επίσης και τα παιδιά που δουλέψαμε μαζί, Ιωάννη Αναγνωστίδη, Ερμή Σουμαλιά, Αθηνά Τερζόγλου, Μιχάλη Ξεφτέρη, Γιώργο Χιονά, Ηλία Παπανικολάου και Μάνο Παντουβά που πάντα κινούσαν το ενδιαφέρον μου ένα βήμα παραπέρα. Ευχαριστώ και την Δώρα Σούλιου και τον Πέτρο Ποτίκα για όλη την συνεργασία και την βοήθεια τους.

Τέλος, θα ήθελα πάνω απ' όλα να ευχαριστήσω ολόψυχα την οικογένεια μου. Τους γονείς μου, Ζωή και Σωτήρη η στήριξη και αγάπη των οποίων όλα αυτά τα χρόνια είναι ανείπωτη και έχει καταστήσει τα πάντα δυνατά. Τον αδερφό μου, Κωνσταντίνο που είναι πάντα εκεί για εμένα σε όλα τα εύκολα και τα δύσκολα. Η συμπαράσταση τους υπήρξε το παν σε όλη μου την πορεία και το παράδειγμα τους φωτίζει τον δρόμο μου. Τον θείο μου, Γιάννη, που ήταν πάντα το ακαδημαϊκό μου πρότυπο, την γιαγιά μου Μαρία, την Μαργαρίτα, την Πέννυ, την Μαρία, την Αλίχη. Τους παππούδες και γιαγιά που με μεγάλωσαν αλλά δεν πρόλαβαν να δουν αυτήν την στιγμή από κοντά,

Κωνσταντίνο, Παναγιώτη και Μαρία. Την Άννα-Αγάπη, που ανέχτηκε όλες τις παραξενιές μου αυτό το διάστημα και με έσπρωχνε πάντα μπροστά με τόση αγάπη, δίνοντας μου την ώθηση να αγωνίζομαι για το καλύτερο. Ευχαριστώ επίσης τους φίλους μου και υπερπολύτιμους συμπορευτές μου, Δημήτρη, Νίκο, Βαγγέλη, Αλέξανδρο και Αμβρόσιο, καθώς και όλους τους κοντινούς ανθρώπους και φίλους μου, από τους οποίους ζητώ επίσης συγγνώμη που δεν αναφέρονται προσωπικά παρότι το αξίζουν, που ήταν δίπλα μου σε όλη αυτήν την πορεία. Τους χιλιοευχαριστώ όλους για την αγάπη που μου έδειξαν.

Παναγιώτης Πατσιλινάκος
Αθήνα, Σεπτέμβριος 2023

Contents

1	Extended Abstract in Greek	18
1.1	Εισαγωγή	18
1.2	Σχεδιασμός μηχανισμών για στρατηγικά περιβάλλοντα: Ανάλυση πέραν της χειρότερης περίπτωσης και πολυπλοκότητα επικοινωνίας.	25
1.2.1	Σχεδιασμός Φιλαλήθων Μηχανισμών για το Πρόβλημα της Χωροθέτησης σε Στιγμιότυπα Ευσταθή σε Διαταραχές.	25
1.2.2	Δειγματοληψία και Βέλτιστη Εξαγωγή Προτιμήσεων σε Απλοούς Μηχανισμούς.	27
1.3	Σχεδιασμός μηχανισμών για μη στρατηγικά περιβάλλοντα: Περιβάλλοντα με περιορισμένη πληροφορία.	30
1.3.1	Μετρική Παραμόρφωση υπό Περιορισμένες Πληροφορίες	30
1.3.2	Διαστασιμότητα και Συντονισμός στην Ψηφοφορία: Η παραμόρφωση του STV	32
1.4	Υπολογισμός ισορροπιών.	35
1.5	Το πρόβλημα Node Max-Cut και ο Υπολογισμός Ισορροπιών σε Βεβαρημένα Παίγνια με Γραμμικές Συναρτήσεις	35
1.6	Βιβλιογραφικές πληροφορίες	37
2	Introduction	38
2.1	Strategic mechanism design: Beyond worst case and communication complexity.	43
2.1.1	Strategyproof Facility Location in Perturbation Stable Instances.	43
2.1.2	Sampling and Optimal Preference Elicitation in Simple Mechanisms.	46
2.2	Non strategic mechanism design: Domains with limited information.	49
2.2.1	Metric-Distortion Bounds under Limited Information.	49
2.2.2	Dimensionality and Coordination in Voting: The Distortion of STV	50
2.3	Computation of equilibria.	53

2.4	Node Max-Cut and Computing Equilibria in Linear Weighted Congestion Games	53
2.5	Bibliographic information	55
3	Facility location in perturbation stable instances	56
3.1	Our contribution	57
3.2	Notation, Definitions and Preliminaries	60
3.3	Perturbation Stability on the Line: Definition and Properties . . .	62
3.4	The Optimal Solution is Strategyproof for $(2 + \sqrt{3})$ -Stable Instances	67
3.5	A Deterministic Mechanism Resistant to Singleton Deviations . . .	70
3.6	Low Stability and Inapproximability by Deterministic Mechanisms	73
3.6.1	Image Sets, Holes and Well-Separated Instances	73
3.6.2	Nonexistence of strategyproof mechanisms, with bounded approximation in stable instances.	74
3.6.3	The Proof of the Impossibility Result	75
3.7	A Randomized Mechanism with Constant Approximation	77
3.A	Proof of lemma 3.3.8.	81
3.B	Proofs of Auxiliary Lemmas Used in the Proof of Theorem 3.6.5 .	87
3.B.1	The Proof of Lemma 3.6.1	88
3.B.2	The Proof of Lemma 3.6.2	88
4	Sampling and Optimal Preference Elicitation in Simple Mechanisms	91
4.1	Contributions	92
4.1.1	Overview of Techniques	93
4.2	Preliminaries	95
4.3	Facility Location Games and Approximation of the Median Mechanism	97
4.3.1	Median on the Line	97
4.3.2	High-Dimensional Median	104
4.3.3	Median on Trees	105
4.3.4	Sampling with Multiple Facilities	106
4.4	Auctions	108
4.4.1	Single-Item Auction	108
4.4.2	Multi-Item Auction with Additive Valuations	114
4.4.3	Multi-Unit Auction with Unit Demand	115
4.A	Sampling Approximation of the Plurality Rule	119
5	Metric-Distortion Bounds under Limited Information	122
5.1	Contributions	123
5.2	Preliminaries	127
5.2.1	Instance-Optimal Voting	128
5.3	Sequence of Pairwise Comparisons	129
5.4	Distortion of Deterministic Rules under Incomplete Orders	132

5.4.1	Missing Voters	133
5.4.2	Top Preferences	135
5.5	Randomized Preference Elicitation and Sampling	139
5.5.1	Approximating Copeland	140
5.5.2	Approximating Plurality Matching	141
5.A	Plurality Matching vs Minimax-LP	144
5.A.1	Instance Optimality	144
5.A.2	Decisive Metrics	145
6	Dimensionality and Coordination in Voting: The Distortion of STV	148
6.1	Contributions	149
6.2	Preliminaries	153
6.3	STV in Doubling Metrics	154
6.3.1	STV on the Line	154
6.3.2	Main Result	156
6.3.3	The Lower Bound	160
6.4	Coordination Dynamics	162
6.4.1	The Model	162
6.4.2	A Greedy Approach	163
6.4.3	Exploration/Exploitation	163
6.5	Intrinsic Robustness of Plurality Matching	165
6.5.1	Ultra-Metrics	165
6.5.2	Approximate Metrics	168
7	Node-Max-Cut and the Complexity of Equilibrium in Linear Weighted Congestion Games	171
7.1	Contributions	172
7.2	Basic Definitions and Notation	173
7.3	Hardness of Computing Equilibria in Weighted Congestion Games on Series-Parallel Networks	175
7.4	Computing Approximate Equilibria for Node-Max-Cut	178
7.A	The Proof of Theorem 7.3.1	179
7.B	Missing Technical Details from the Analysis of BRIDGEGAPS	183

Chapter 1

Extended Abstract in Greek

1.1 Εισαγωγή

Η αλγοριθμική θεωρία παιγνίων *Algorithmic Game Theory* αποτελεί ένα ισχυρό πλαίσιο που βρίσκεται στην τομή της επιστήμης των υπολογιστών, των μαθηματικών και της μικροοικονομίας. Ο κύριος της στόχος είναι να αντιμετωπίσει τις προκλήσεις που προκύπτουν από στρατηγικές αλληλεπιδράσεις μεταξύ παικτών (πχ. ατόμων ή πρακτόρων) που επιδιώκουν το δικό τους συμφέρον ή μεταξύ αυτών και μιας αρχής. Ως αναδυόμενο πεδίο μελέτης, έχει λάβει μεγάλο ερευνητικό ενδιαφέρον και προβολή τα τελευταία χρόνια, με πολλά ισχυρά κίνητρα από προβλήματα που συναντώνται στον πραγματικό κόσμο. Ορισμένα παραδείγματα είναι τα εξής:

- *Μοντελοποίηση πολύπλοκων αλληλεπιδράσεων*: Σε πολλά σενάρια, οι αποφάσεις του κάθε ατόμου σε περιβάλλοντα πολλαπλών παικτών είναι αλληλεξαρτημένες, με την έννοια ότι η προτιμώμενη ενέργεια κάθε παίκτη εξαρτάται από την επιλογή των ενεργειών των άλλων. Η αλγοριθμική θεωρία παιγνίων παρέχει τα μαθηματικά εργαλεία για τη μοντελοποίηση και ανάλυση τέτοιων αλληλεπιδράσεων, μελετώντας τη συμπεριφορά παικτών με στρατηγική συμπεριφορά μέσω της ανάλυσης της πολυπλοκότητας και της ποιότητας σχετικών ισορροπιών χρησιμοποιώντας κατάλληλες έννοιες όπως η τιμή της αναρχίας κ.λπ...
- *Χρήση δημόσιων πόρων* : Ένα πολύ συνηθισμένο πρόβλημα στην πράξη προκύπτει από την ανάγκη πολλών παικτών να ανταγωνίζονται για πόρους. Ένα πολύ καλά μελετημένο σχετικό παράδειγμα είναι τα *παίγνια συμφόρησης*, δηλαδή περιβάλλοντα στα οποία οι παίκτες πρέπει να αποφασίσουν ποιους πόρους θα χρησιμοποιήσουν για να επιτύχουν τους στόχους τους, λαμβάνοντας υπόψη ότι το κόστος κάθε πόρου εξαρτάται από το πόσοι παίκτες τον χρησιμοποιούν. Ένα πραγματικό παράδειγμα εφαρμογής αυτού του τύπου των παιγνίων είναι η μοντελοποίηση της κυκλοφορίας σε δημόσια οδικά δίκτυα.
- *Σχεδιασμός μηχανισμών συμβατών με στρατηγικά κίνητρα (incentive-compatible)*: Σε πολλά στρατηγικά περιβάλλοντα, η συμπεριφορά των παικτών ε-

πηρεάζεται από ιδιωτικές (προσωπικές) πληροφορίες. Οι παίχτες έχουν τη δυνατότητα να παραποιήσουν τις ιδιωτικές τους πληροφορίες που αναφέρουν προκειμένου να επηρεάσουν τον μηχανισμό αν αυτό τους προσφέρει μια πιο επιθυμητή έκβαση. Σε πολλές εφαρμογές, είναι επιθυμητό ή ακόμη και απαραίτητο να υπάρχει εγγύηση ότι ο χρησιμοποιούμενος μηχανισμός δεν επιτρέπει τέτοια συμπεριφορά, δηλαδή ότι ο μηχανισμός σχεδιάζεται με τέτοιο τρόπο ώστε να αναγκάζει τους παίχτες να παρέχουν αληθείς πληροφορίες. Τέτοιοι μηχανισμοί ονομάζονται *incentive-compatible* (μηχανισμοί συμβατοί με στρατηγικά κίνητρα) και αποτελούν επίσης ένα μεγάλο αντικείμενο της αλγοριθμικής θεωρίας παιγνίων. Επιπλέον, οι μηχανισμοί μπορεί να χρειάζεται να μεγιστοποιήσουν άλλα κριτήρια, όπως τη κοινωνική δικαιοσύνη ή το κέρδος, λαμβάνοντας υπόψη τις προτιμήσεις και τους περιορισμούς του περιβάλλοντος των παικτών. Ορισμένες εφαρμογές μπορούν να βρεθούν σε δημοπρασίες, συστήματα ψηφοφορίας, χωροθέτηση δημόσιων εγκαταστάσεων, κατανομή περιορισμένων πόρων μεταξύ των παικτών κ.λπ...

Ο κύριος στόχος αυτής της εργασίας επικεντρώνεται στην πτυχή του σχεδιασμού μηχανισμών στις περιοχές της αλγοριθμικής θεωρίας παιγνίων και της κοινωνικής επιλογής *Social Choice*, όπου μελετάμε επίσης περιβάλλοντα πολλαπλών παικτών από μη στρατηγική σκοπιά. Συγκεκριμένα, επικεντρωνόμαστε σε αλγοριθμικά πλαίσια για προβλήματα πολλαπλών παικτών σε περιβάλλοντα όπου οι πληροφορίες δεν είναι εύκολα προσβάσιμες, από δύο σκοπιές:

- Οι πληροφορίες είναι *ιδιωτικές*: Σε αυτήν την περίπτωση, χρειάζεται να δημιουργήσουμε *φιλαλήθεις* μηχανισμούς για περιβάλλοντα με στρατηγικούς παίχτες.
- Οι πληροφορίες είναι *περιορισμένες*: Οι παίχτες δεν μπορούν ή δεν είναι διατεθειμένοι να παράσχουν πλήρεις πληροφορίες, ή η ανταλλαγή πληροφοριών είναι δαπανηρή.

Ο στόχος της παρούσας διατριβής είναι να σχεδιαστούν αλγόριθμοι ή κανόνες και πρωτόκολλα (μηχανισμοί) που εγγυώνται επιθυμητούς λόγους προσέγγισης σε σχέση με τις βέλτιστες λύσεις, όταν οι παίχτες έχουν περιορισμένες ή ιδιωτικές πληροφορίες και παράλληλα ενθαρρύνουν επιθυμητή συμπεριφορά, όταν οι παίχτες έχουν ιδιωτικές πληροφορίες και είναι στρατηγικοί. Επιπλέον, ενδιαφερόμαστε να κατανοήσουμε τους περιορισμούς στο σχεδιασμό τέτοιων συστημάτων, δηλαδή κάποιους κάτω φράγματα για τα βέλτιστα δυνατά αποτελέσματα, σε κάθε πεδίο. Στο πεδίο του σχεδιασμού μηχανισμού για στρατηγικούς παίχτες επικεντρωνόμαστε στο παράδειγμα του προβλήματος χωροθέτησης εγκαταστάσεων (*Facility Location Problem*) και εισάγουμε την έννοια της ανάλυσης πέρα από τη χειρότερη περίπτωση στον σχεδιασμό μηχανισμών. Για την περίπτωση της περιορισμένη πληροφορία, σε μη στρατηγικό περιβάλλον, επικεντρωνόμαστε στην παραμόρφωση στο πρόβλημα της ψηφοφορίας (*Distortion in Voting*), μια σχετικά πρόσφατα εισαγόμενη έννοια που μετρά την επίδραση της δυνατότητας πρόσβασης μόνο σε μερική πληροφορία στη δυνατότητά μας να προσεγγίσουμε τη βέλτιστη κοινωνική έκβαση. Επιπλέον, συνδυάζοντας τις έννοιες των στρατηγικών

κινήτρων και των περιορισμένων πληροφοριών, μελετούμε την πολυπλοκότητα επικοινωνίας (*communication complexity*) σε βασικά, γνωστά, μονο-παραμετρικά (*single parameter*) περιβάλλοντα, όπως δημοπρασίες και το πρόβλημα της χωροθέτησης μίας δομής, και αναπτύσσουμε μηχανισμούς που είναι ασυμπτωτικά βέλτιστοι από πλευράς επικοινωνίας. Τέλος, δίνουμε έμφαση στην κλασική έννοια του υπολογισμού ισορροπιών και μελετούμε την πολυπλοκότητα του υπολογισμού των Αμιγών Ισορροπιών (*Nash Pure Nash Equilibria*) ή προσεγγιστικών Αμιγών Ισορροπιών Nash σε βεβαρυμένα παιχνίδια συμφόρησης (*Weighted Congestion Games*).

Σχεδιασμός μηχανισμών για στρατηγικά περιβάλλοντα με ιδιωτικές πληροφορίες:

Ο πρώτος τύπος προβλημάτων που πρόκειται να παρουσιάσουμε ανήκει στον τομέα της μελέτης και σχεδιασμού φιλαλήθων ή ινσεντιε-σομπατιβλε μηχανισμών. Τα κύρια παραδείγματα που μελετούνται σε αυτόν τον τομέα έχουν την ακόλουθη μορφή: αναλαμβάνουμε το ρόλο μιας αρχής που προσπαθεί να κατανείμει διάφορα αγαθά στους παίκτες ή σε ένα περιβάλλον στο οποίο ανήκουν οι παίκτες. Σε αυτήν την περίπτωση, οι παίκτες έχουν ιδιωτικές αποτιμήσεις επί των πιθανών αποτελεσμάτων και μπορεί να συμπεριφέρονται στρατηγικά προκειμένου να επηρεάσουν τον μηχανισμό, αν είναι δυνατόν να προκύψει ένα καλύτερο αποτέλεσμα για τον εαυτό τους με αυτόν τον τρόπο. Μια βασική ιδιότητα που θέλουμε πάντα οι μηχανισμοί μας να έχουν εδώ είναι η φιλαλήθεια (*truthfulness*). Δηλαδή, να εγγυώνται ότι κανένας παίκτης δεν μπορεί να επιτύχει καλύτερο τελικό αποτέλεσμα για τον εαυτό του με το να παραποιήσει τις πληροφορίες που δηλώνει στον μηχανισμό. Τα δύο βασικά παραδείγματα που μελετούνται σε αυτό το πλαίσιο είναι τα προβλήματα χωροθέτησης εγκαταστάσεων και οι δημοπρασίες.

Το πρόβλημα της χωροθέτησης εγκαταστάσεων ([177]) μελετά την αντιστάθμιση μεταξύ της προσβασιμότητας των υπηρεσιών και της αποδοτικότητας τους ως προς το κόστος. Στην έκδοση που μελετάμε, θέλουμε να τοποθετήσουμε ένα συγκεκριμένο αριθμό εγκαταστάσεων με στόχο την ελαχιστοποίηση του κοινωνικού κόστους, δηλαδή των συνολικών αποστάσεων των παικτών προς τις πλησιέστερες εγκαταστάσεις σε αυτούς. Σε αυτή την έκδοση του προβλήματος οι εγκαταστάσεις δεν έχουν κόστος για το άνοιγμά τους ή τη λειτουργία τους με την πάροδο του χρόνου. Σε αυτήν τη μορφή, αυτό είναι ένα πολύ καλά μελετημένο πρόβλημα ([116, 136, 166, 186, 111, 167]...), του οποίου η ντετερμινιστική έκδοση του έχει ουσιαστικά χαρακτηριστεί πλήρως ([112]) και η τυχαιοποιημένη έκδοση του έχει επίσης μελετηθεί εκτενώς.

Τα προβλήματα δημοπρασιών αποτελούν επίσης έναν πολύ μεγάλο τομέα μελέτης της αλγοριθμικής θεωρίας παιχνιδιών. Γενικά, η ιδέα είναι ότι ο "πωλητής" (μηχανισμός) θέλει να πουλήσει ένα ή πολλά αγαθά σε παίκτες που συμμετέχουν στη δημοπρασία. Συνήθως οι παίκτες έχουν ιδιωτικές αξιολογήσεις που αντιπροσωπεύουν πόσο θέλουν κάθε αντικείμενο και προσπαθούν να μεγιστοποιήσουν το όφελός τους προσπαθώντας να μεγιστοποιήσουν τη διαφορά μεταξύ της αξιολόγησής τους και της τιμής που τελικά πληρώνουν για τα αγαθά που αποκτούν. Υπάρχουν πολλές σχετικές μορφές

δημοπρασιών, όπως δημοπρασίες με μονό αντικείμενο ([220]), διάφορες μορφές δημοπρασιών με πολλά αντικείμενα, όπως δημοπρασίες πολλαπλών αντικειμένων όπου πωλούνται πολλαπλές αντιγραφές του ίδιου αντικειμένου, συνδυαστικές δημοπρασίες όπου οι παίχτες θέλουν συνδυασμούς αντικειμένων, ειδικοί τομείς αξιολόγησης όπως sub-modular, sub-additive αξιολογήσεις κ.λπ. Έχουν μελετηθεί επίσης πολλές μορφές δημοπρασιών, όπως οι δημοπρασίες με σφραγισμένες προσφορές (*Sealed Bid*), όπου οι παίχτες υποβάλλουν τις πλήρεις αξιολογήσεις τους στον μηχανισμό σε σφραγισμένες προσφορές, ή οι αύξουσες δημοπρασίες, όπου η τιμή του αγαθού ανεβαίνει και οι παίχτες αποφασίζουν εάν θα συνεχίσουν να συμμετέχουν ή όχι.

Σε αυτό το πλαίσιο του σχεδιασμού μηχανισμών σε στρατηγικά περιβάλλοντα, μελετάμε δύο βασικές ερωτήσεις:

Ερώτημα 1. Μπορούμε να εισάγουμε έννοιες από την ανάλυση πέρα από τη χειρότερη περίπτωση στον σχεδιασμό μηχανισμών προκειμένου να δημιουργήσουμε μηχανισμούς με καλές ιδιότητες για “ρεαλιστικά στην πράξη” στιγμιότυπα προβλημάτων, για τα οποία δεν μπορούν να υπάρχουν τέτοιοι μηχανισμοί στη γενική περίπτωση; Μπορούμε επίσης να χρησιμοποιήσουμε αυτήν την ανάλυση για να κατανοήσουμε τη φύση της φιλαλήθειας σε αυτά τα προβλήματα;

Για αυτήν την ερώτηση επικεντρωνόμαστε στο πρόβλημα χωροθέτησης εγκαταστάσεων, για το οποίο έχουν δειχθεί πολύ ισχυρά όρια όσον αφορά την ύπαρξη τέτοιων επιθυμητών μηχανισμών. Για παράδειγμα, μπορούμε να έχουμε καλούς, δηλαδή με περιορισμένη προσέγγιση και φιλαλήθεις μηχανισμούς, μόνο όταν προσπαθούμε να εκχωρήσουμε έως δύο εγκαταστάσεις σε μια γραμμή ([112]). Με αυτήν τη σκέψη, δημιουργήθηκε η έννοια της ανάλυσης “πέρα από τη χειρότερη περίπτωση” ανάλυση στον σχεδιασμό μηχανισμών.

Η ανάλυση “Πέρα από τη Χειρότερη Περίπτωση” εισήχθη για πρώτη φορά ως μαθηματικό πλαίσιο για την ανάλυση προβλημάτων που έχουν δείξει να είναι υπολογιστικά δύσκολα στη χειρότερη περίπτωση, αλλά φαίνεται να μην είναι τόσο δύσκολα στα περισσότερα ρεαλιστικά στην πράξη παραδείγματα. Αυτή η μορφή ανάλυσης ορίζει έννοιες που αιχμαλωτίζουν αυτήν την έννοια των “πραγματικών” ή “μέσων” περιπτώσεων, όπως η “ομαλότητα” (*smoothness*) ή η “ευστάθεια με μικρές διαταραχές” (*perturbation stability*) και δείχνουν ότι μέσα σε αυτές τις έννοιες η επίλυση του προβλήματος γίνεται εφικτή ([43, 31]). Για το πρόβλημα της χωροθέτησης εγκαταστάσεων, ειδικά, επιλέχθηκε η έννοια της “ευστάθειας σε διαταραχές”, λόγω της στενής σχέσης της με το πρόβλημα της συσταδοποίησης (*clustering*) για το οποίο αυτή η έννοια ανάλυσης πέρα από τη χειρότερη περίπτωση έχει επιδείξει μεγάλη επιτυχία ([13, 33, 35, 34, 202, 203]). Ουσιαστικά, ένα πρόβλημα συσταδοποίησης είναι ευσταθές με μικρές διαταραχές εάν μικρές διαταραχές σε αυτό (ή στο μετρικό χώρο στον οποίο ανήκει) δεν επηρεάζουν τη βέλτιστη συσταδοποίηση του προβλήματος. Στην περίπτωσή μας, αντί να δείξουμε ότι η περιορισμένη ευστάθεια επιτρέπει τη δημιουργία γρήγορων αλγορίθμων, θέλουμε να δείξουμε ότι μπορούμε να δημιουργήσουμε μηχανισμούς με επιθυμητές ιδιότητες, όπως φιλαλήθεια και καλή ή φραγμένη

προσέγγιση που δεν μπορούν να υπάρχουν στη γενική περίπτωση.

Ερώτημα 2. Ποιο είναι το βέλτιστο ποσό πληροφοριών που απαιτείται προκειμένου να εφαρμόσουμε μηχανισμούς για γνωστά, βασικά προβλήματα χωρίς να θυσιάζουμε τις επιθυμητές ιδιότητες του μηχανισμού;

Η αποτελεσματική επικοινωνία αποτελεί πάντα έναν από τους κύριους στόχους του σχεδιασμού μηχανισμών. Υπάρχουν πολλοί λόγοι γι' αυτό το ενδιαφέρον, όπως το γεγονός ότι η εξαγωγή δεδομένων από πολλά καταναμημένα μέρη μπορεί να είναι πολύ κοστοβόρα, ή το γεγονός ότι η αποκάλυψη προσωπικής πληροφορίας σε μεγάλο βαθμό μπορεί να μην θεωρείται θεμιτή ή ασφαλής, ή ακόμη η αναγνώριση ότι η άμεση αποκάλυψη ολόκληρης της πληροφορίας στον μηχανισμό είναι ανεπιθύμητη δεδομένου ότι επιβάλλει υψηλό γνωστικό κόστος από τους παίχτες, κ.λπ. ([47, 4, 190, 164, 215])...

Μελετάμε αυτή την έννοια για δύο κύριες κατηγορίες προβλημάτων: Δημοπρασίες και χωροθέτηση εγκαταστάσεων. Συγκεκριμένα επικεντρωνόμαστε στις δημοπρασίες ενός αντικειμένου (single item) και στις δημοπρασίες πολλαπλών μονάδων (multi unit) με παίχτες που θέλουν ο καθένας το πολύ μια μονάδα αντικειμένου, από την οπτική γωνία της προσπάθειας ελαχιστοποίησης της επικοινωνίας μεταξύ των παιχτών και του μηχανισμού. Για το πρόβλημα της χωροθέτησης εγκαταστάσεων, μελετούμε τη χωροθέτηση μίας εγκατάστασης, το οποίο μπορεί επίσης να θεωρηθεί ως μοντελοποίηση σχημάτων ψηφοφορίας, και επικεντρωνόμαστε στην εύρεση ενός προσεγγιστικού μέσου (approximate median), δηλαδή της βέλτιστης λύσης, μέσω δειγματοληψίας.

Σχεδίαση αλγορίθμων για μη στρατηγικά περιβάλλοντα με περιορισμένη πληροφορία:

Σε πολλές περιπτώσεις χρειάζεται να σχεδιάσουμε και να μελετήσουμε μηχανισμούς σε μη-στρατηγικά περιβάλλοντα με στόχους διαφορετικούς από την ανθεκτικότητα στη στρατηγική. Ένα προεξέχον παράδειγμα που περιέχει τέτοιου είδους ερωτήσεις είναι η ψηφοφορία από τον τομέα της κοινωνικής επιλογής. Σε ένα πρόβλημα ψηφοφορίας, οι παίχτες θέλουν να εκλέξουν έναν επιθυμητό υποψήφιο ή ένα σύνολο υποψηφίων. Κάθε παίχτης έχει μια προτίμηση για τους υποψηφίους και θέλουμε να δημιουργήσουμε μηχανισμούς που θα επιτρέψουν την εκλογή των κορυφαίων υποψηφίων που ικανοποιούν την πλειοψηφία όσο το δυνατόν καλύτερα. Υπάρχουν πολλές έννοιες από την κοινωνική επιλογή που μετρούν την ποιότητα ενός νικητή, όπως η έννοια του ωφέλιμιστικού νικητή (*utilitarian winner*), που μεγιστοποιεί την κοινωνική ευημερία (*social welfare*), ή ενός νικητή *Condorcet*, ο οποίος κερδίζει σε κάθε μεμονωμένη σύγκριση με οποιονδήποτε άλλο υποψήφιο κ.λπ. ([222, 223, 36, 124, 161])...

Ακόμα κι όταν οι παίχτες δεν είναι στρατηγικοί και είναι πρόθυμοι να μοιραστούν τις πληροφορίες τους, η δημιουργία μηχανισμών που εγγυώνται τέτοιες ιδιότητες, όταν αυτό είναι εφικτό, εξακολουθεί να είναι μια πρόκληση. Ένα σχετικά πρόσφατο ερευνητικό ενδιαφέρον έχει επικεντρωθεί στις *utilitarian* ψηφοφορίες, όπου κάθε ψηφοφόρος έχει ένα υποκείμενο αριθμητικό όφελος για κάθε υποψήφιο, που αναπαρι-

στά την προτίμηση του σε αυτόν, στην περίπτωση όπου δεν είναι δυνατή η απόκτηση της πλήρους πληροφορίας από τους ψηφοφόρους. Στην κύρια περίπτωση λαμβάνεται υπόψη ότι ακόμη και αν για κάθε ψηφοφόρο υπάρχουν ακριβή αριθμητικά ωφέλη για κάθε υποψήφιο, λόγω του υψηλού γνωστικού κόστους της αναγνώρισης αυτών των οφελών ή λόγω της υψηλής αξίας της ιδιωτικής πληροφορίας, οι ψηφοφόροι μπορούν να παρέχουν μόνο μια κατάταξη προτίμησης των υποψηφίων που είναι συνεπής με τις σχετικές προσωπικές τους αποτιμήσεις. Η έννοια της παραμόρφωσης μετρά την επίδραση της χρήσης περιορισμένων πληροφοριών στο πρόβλημα της ψηφοφορίας στον βαθμό που το αποτέλεσμα επιτυγχάνει κοινωνική ευημερία ([194]).

Όπως και προηγουμένως, εστιάζουμε σε δυο ευρεία ερωτήματα σε αυτήν την κατεύθυνση:

Ερώτημα 3. Ποιά είναι η επίδραση στην μετρική παραμόρφωση (*metric distortion*) σε περιβάλλοντα ψηφοφορίας με έναν νικητή όταν η πρόσβαση στην πληροφορία περιοριστεί περαιτέρω;

Η *μετρική παραμόρφωση* (*metric distortion*) αναφέρεται στην περίπτωση όπου οι ωφέλειες των παικτών είναι ενσωματωμένες σε ένα μετρικό χώρο ([18]) (δηλαδή είναι οι "αποστάσεις" τους από τους υποψήφιους σε αυτόν τον μετρικό χώρο). Αυτό το πρόβλημα έχει μελετηθεί εκτενώς και έχει καθοριστεί ένα κάτω όριο για την παραμόρφωση ([149]) το οποίο επιτυγχάνεται από έναν ντετερμινιστικό μηχανισμό ([131]). Η ερώτησή μας επεκτείνει περαιτέρω την έννοια ότι η πληροφορία μπορεί να είναι δύσκολο να αποκτηθεί και ρωτάμε ερωτήσεις όπως: ποια είναι τα όρια παραμόρφωσης όταν οι παίχτες παρέχουν ατελείς κατατάξεις μόνο των κορυφαίων υποψηφίων που προτιμούν, πόσες διαδοχικές συγκρίσεις απαιτούνται για να επιτευχθεί σχεδόν βέλτιστη παραμόρφωση και τι μπορούμε να κάνουμε με χρήση δειγματοληψίας.

Ερώτημα 4. Ποια είναι η σχέση μεταξύ της μετρικής παραμόρφωσης και της διαστασιμότητας (*dimensionality*) του υποκείμενου μετρικού χώρου. Επίσης, μπορούμε να χρησιμοποιήσουμε φυσικούς κανόνες μάθησης για να δημιουργήσουμε μηχανισμούς με χαμηλή παραμόρφωση;

Σε αυτήν την περίπτωση επικεντρωθήκαμε στην ανάλυση ενός κυρίαρχου μηχανισμού ψηφοφορίας: του "STV". Ο στόχος μας είναι να φράξουμε την παραμόρφωση του STV με βάση τη διάσταση του υποκείμενου μετρικού χώρου και να μελετήσουμε επίσης μηχανισμούς που πλησιάζουν αυτό το κάτω όριο. Επιπλέον, αναλύουμε την αποδοτικότητα κατανεμημένων, αποκεντρωμένων δυναμικών μάθησης από τους παίχτες και τη σύγκλισή τους σε έναν υποψήφιο με χαμηλή παραμόρφωση.

Υπολογισμός ισορροπιών:

Όπως αναφέρθηκε παραπάνω, μία από τις πιο βασικές λειτουργίες της αλγοριθμικής

θεωρίας παιγνίων είναι η μελέτη της αλληλεπίδρασης μεταξύ στρατηγικών παικτών με εγωιστικά κίνητρα. Ο βασικός και πιο γνωστός όρος που σχετίζεται με αυτήν την έννοια είναι το *Σημείο Ισορροπίας κατά NASH (Nash Equilibrium)* ([184, 185]). Μια κατάσταση του παιχνιδιού, δηλαδή η επιλογή μιας στρατηγικής για κάθε παίκτη, είναι σημείο ισορροπίας κατά Nash αν κανένας παίκτης δεν μπορεί να αυξήσει το κέρδος του αλλάζοντας μονομερώς τη στρατηγική του. Μια στρατηγική για κάθε παίκτη μπορεί να είναι είτε αμιγής (ντετερμινιστική), όπου παίζει ντετερμινιστικά μία από τις διαθέσιμες στρατηγικές του, είτε μεικτή (τυχαίος) όπου παίζει μια κατανομή των στρατηγικών του. Μία ισορροπία που αποτελείται από καθαρές στρατηγικές ονομάζεται *Αμιγής ισορροπία Nash (PNE)*. Σε αυτό το περιβάλλον επικεντρωνόμαστε στα παίγνια συμφόρησης, στα οποία οι παίκτες πρέπει να χρησιμοποιήσουν κοινούς πόρους. Συγκεκριμένα, έχουμε πεπερασμένα σύνολα παικτών και πόρων για τους οποίους ανταγωνίζονται οι παίκτες. Οι στρατηγικές των παικτών είναι υποσύνολα των πόρων, και κάθε παίκτης θέλει να ελαχιστοποιήσει τον συνολικό κόστος των πόρων στην επιλεγμένη στρατηγική του. Κάθε πόρος συσχετίζεται με μια συνάρτηση καθυστέρησης που καθορίζει το κόστος της χρήσης αυτού του πόρου ως συνάρτηση της *συμφόρησης* του, δηλαδή του αριθμού των παικτών που χρησιμοποιούν αυτόν τον πόρο στη στρατηγική τους.

Η κύρια ερώτηση στην οποία επικεντρωνόμαστε σε αυτήν την κατεύθυνση είναι η εξής:

Ερώτημα 5. Ποια είναι η πολυπλοκότητα εύρεσης Αμιγών Ισορροπιών Nash σε βεβαρημένα παίγνια συμφόρησης;

Τα παίγνια συμφόρησης ([201]) ανήκουν σε μια ευρύτερη κατηγορία παιγνίων που ονομάζονται *παιγνία δυναμικού (Potential Games)* ([178]) τα οποία πάντα διαθέτουν τουλάχιστον μία Αμιγή Ισορροπία κατά Nash. Έτσι, οι ερευνητές έχουν μελετήσει εκτενώς την πολυπλοκότητα εύρεσης τέτοιων ισορροπιών σε παίγνια συμφόρησης, με πολλά αποτελέσματα για πολλαπλές παραλλαγές, όπως συμμετρικά ή ασύμμετρα παίγνια, όπου όλοι οι παίκτες έχουν ή δεν έχουν τα ίδια σύνολα στρατηγικών αντίστοιχα, παίγνια συμφόρησης δικτύου, όπου οι στρατηγικές αντιστοιχούν σε διαδρομές σε ένα υποκείμενο δίκτυο και άλλα. Ωστόσο, λιγότερη προσοχή έχει δοθεί στη μελέτη των *βεβαρυμμένων* παιγνίων δυναμικού, όπου κάθε παίκτης συμβάλλει στη συμφόρηση των πόρων που χρησιμοποιεί με διαφορετικό βάρος ([95, 113, 118, 189, 63, 64, 153]).

Το επιχείρημα που χρησιμοποιούμε για να δείξουμε ότι κάθε παίγνιο δυναμικού διαθέτει Αμιγείς Ισορροπίες κατά Nash είναι ένα επιχείρημα τοπικής αναζήτησης, που καθιστά αυτήν την κατηγορία παιγνίων να ανήκει σε μια κλάση πολυπλοκότητας που ονομάζεται “Polynomial time Local Search” (τοπική αναζήτηση πολυωνυμικού χρόνου), ή εν συντομία PLS. Όταν ένα πρόβλημα ανήκει στις πιο δύσκολες κατηγορίες προβλημάτων μέσα σε αυτήν την τάξη, το ονομάζουμε *PLS-πλήρες (PLS-complete)* και θεωρείται ότι είναι υπολογιστικά δύσκολο. Βλέπουμε εδώ ότι παραλλαγές του παιχνιδιού που διαθέτουν γρήγορους αλγόριθμους στην μη βεβαρυμμένη περίπτωση,

γίνονται αμέσως δύσκολες όταν εισάγονται βεβαρημένοι παίχτες. Επίσης, μελετάμε την πολυπλοκότητα εύρεσης προσεγγιστικών ισορροπιών για τέτοιες περιπτώσεις.

1.2 Σχεδιασμός μηχανισμών για στρατηγικά περιβάλλοντα: Ανάλυση πέραν της χειρότερης περίπτωσης και πολυπλοκότητα επικοινωνίας.

1.2.1 Σχεδιασμός Φιλαλήθων Μηχανισμών για το Πρόβλημα της Χωροθέτησης σε Στιγμιότυπα Ευσταθή σε Διαταραχές.

Σε αυτήν την εργασία εξετάζουμε τα k -*Facility Location games* (παίγνια χωροθέτησης k εγκαταστάσεων), όπου $k \geq 2$ εγκαταστάσεις τοποθετούνται στην ευθεία (των πραγματικών αριθμών) βάσει των προτιμήσεων n στρατηγικών παικτών. Τέτοια προβλήματα ενδιαφέρουν διάφορα σενάρια στην Κοινωνική Επιλογή, όπου μια τοπική αρχή σχεδιάζει να κατασκευάσει έναν σταθερό αριθμό δημόσιων εγκαταστάσεων σε μια περιοχή (βλ. για παράδειγμα, [177]). Η επιλογή των τοποθεσιών βασίζεται στις προτιμήσεις των τοπικών κατοίκων, ή παικτών. Κάθε παίκτης αναφέρει την ιδανική του τοποθεσία, και η τοπική αρχή εφαρμόζει έναν (ντετερμινιστικό ή τυχαίο) μηχανισμό, ο οποίος απεικονίζει τις προτιμήσεις των παικτών σε k τοποθεσίες εγκαταστάσεων.

Κάθε παίκτης αξιολογεί το αποτέλεσμα του μηχανισμού με βάση το κόστος σύνδεσής του, δηλαδή την απόσταση της ιδανικής (πραγματικής) του τοποθεσίας προς την πλησιέστερη εγκατάσταση. Οι παίχτες επιδιώκουν να ελαχιστοποιήσουν το κόστος σύνδεσής τους και μπορεί να δηλώσουν λανθασμένες (ψευδείς) ιδανικές τοποθεσίες στην προσπάθειά τους να επηρεάσουν τον μηχανισμό. Συνεπώς, ο μηχανισμός πρέπει να είναι *φιλαλήθης*, δηλαδή πρέπει να εξασφαλίζει ότι κανένας παίκτης δεν μπορεί να επωφεληθεί από τη ψευδή αναφορά της τοποθεσίας του, ή ακόμη και *ομαδικά φιλαλήθης* (*group strategyproof*), δηλαδή ανθεκτικός σε συνεργατικές συμπεριφορές μεταξύ των στρατηγικών παικτών.

Ο στόχος της τοπικής αρχής (δηλ. του μηχανισμού) είναι να ελαχιστοποιήσει το *κοινωνικό κόστος*, δηλαδή το άθροισμα των κοστών σύνδεσής των παικτών. Εκτός από την τοποθέτηση των εγκαταστάσεων με έναν τρόπο που είναι συμβατός με τις στρατηγικές των παικτών, που ορίζεται με την (ομαδική) φιλαλήθεια, ο μηχανισμός πρέπει να οδηγεί σε ένα κοινωνικά επιθυμητό αποτέλεσμα, το οποίο καθορίζεται από τον λόγο προσέγγισης του μηχανισμού ως προς το βέλτιστο κοινωνικό κόστος.

Λόγω του σημαντικού ερευνητικού ενδιαφέροντος για το θέμα, η θεμελιώδης και πλέον βασική ερώτηση σχετικά με την προσέγγιση του βέλτιστου κοινωνικού κόστους από μηχανισμούς που είναι φιλαλήθεις για το k -*Facility Location* πρόβλημα στην ευθεία έχει κατανοηθεί αρκετά καλά. Για μία εγκατάσταση ($k = 1$), η τοποθέτηση της

εγκατάστασης στη θέση της του διάμεσου (median) παίκτη είναι ομαδικά φιλαλήθης και βελτιστοποιεί τον κοινωνικό κόστος. Για δύο εγκαταστάσεις ($k = 2$), ο καλύτερος δυνατός λόγος προσέγγισης είναι $n - 2$ και επιτυγχάνεται από έναν φυσικό μηχανισμό που είναι ομαδικά φιλαλήθης και τοποθετεί τις εγκαταστάσεις στην αριστερότερη και τη δεξιότερη θέση [112, 192]. Ωστόσο, για τρεις ή περισσότερες εγκαταστάσεις ($k \geq 3$), δεν υπάρχει κανένας ντετερμινιστικός, ανώνυμος¹ μηχανισμός, που να είναι φιλαλήθης με φραγμένο λόγο προσέγγισης (όσον αφορά το n και το k) [112]. Ωστόσο, υπάρχει ένας πιθανοτικός (randomized) ανώνυμος μηχανισμός που είναι ομαδικά φιλαλήθης με λόγο προσέγγισης n [116] (βλ. επίσης το Κεφάλαιο 3 για μια επιλεκτική λίστα με πρόσθετες αναφορές).

Ευστάθεια σε Διαταραχές σε k -Facility Location Games. Η εργασία μας στοχεύει στο να παρακάμψει το ισχυρό αρνητικό αποτέλεσμα του [112] και είναι εμπνευσμένη από την πρόσφατη επιτυχία στον σχεδιασμό αλγορίθμων πολυωνυμικού χρόνου για στιγμιότυπα συσταδοποίησης που είναι γ -ευσταθή (βλ. [13, 33, 35, 34, 202, 203]). Ένα στιγμιότυπο ενός προβλήματος συσταδοποίησης, όπως το k -Facility Location (αντίστοιχα k -median τιμή στη βιβλιογραφία των αλγορίθμων βελτιστοποίησης και των αλγορίθμων προσέγγισης), είναι γ -ευσταθές σε διαταραχές (ή απλά γ -ευσταθές), για κάποιο $\gamma \geq 1$, αν η βέλτιστη συσταδοποίηση δεν επηρεάζεται από την μείωση της κλίμακας κάποιου υποσυνόλου των εγγραφών του πίνακα αποστάσεων κατά έναν παράγοντα το πολύ γ (δηλ. την διαίρεση οποιονδήποτε αποστάσεων ανάμεσα σε τοποθεσίες παικτών κατά το πολύ έναν παράγοντα γ) [43, 31].

Η έννοια της ευστάθειας σε διαταραχές εισήχθη αρχικά από τους Bilu και Linial [43] και τους Awasthi, Blum και Sheffet [31] και έχει ενθαρρύνει μεγάλο όγκο επιπρόσθετης έρευνας (δείτε, για παράδειγμα, [13, 33, 34, 203] και τις αναφορές τους) με σκοπό την απόκτηση θεωρητικής κατανόησης της ανώτερης απόδοσης στην πράξη αρκετά απλών αλγορίθμων συσταδοποίησης για γνωστά δύσκολα προβλήματα συσταδοποίησης στην κλάση NP-hard (όπως η τοποθέτηση k εγκαταστάσεων σε γενικούς μετρικούς χώρους). Εν γένει, οι βέλτιστες συστάδες ενός γ -ευσταθούς στιγμιότυπου είναι κάπως καλά διαχωρισμένες μεταξύ τους και, συνεπώς, σχετικά εύκολο να αναγνωριστούν (δείτε επίσης τις κύριες ιδιότητες των ευσταθών σε διαταραχές στιγμιότυπων στο υποκεφάλαιο 3.3). Ως αποτέλεσμα, φυσικές επεκτάσεις απλών αλγορίθμων, όπως ο αλγόριθμος single linkage (αντίστοιχα αλγόριθμος του Kruskal), μπορούν να βρουν τη βέλτιστη συσταδοποίηση σε πολυωνυμικό χρόνο, υπό την προϋπόθεση ότι $\gamma \geq 2$ [13], και τυπικές αλγοριθμικές προσεγγίσεις, όπως ο δυναμικός προγραμματισμός (αντίστοιχα, η τοπική αναζήτηση), λειτουργούν σε σχεδόν γραμμικό χρόνο για $\gamma > 2 + \sqrt{3}$ (αντίστοιχα, $\gamma > 5$) [3].

Σε αυτήν την εργασία, ερευνούμε κατά πόσον ο περιορισμός της ανάλυσης σε ευσταθή στιγμιότυπα επιτρέπει βελτιωμένους μηχανισμούς που είναι φιλαλήθεις, με φραγμένους (και ιδανικά σταθερούς) λόγους προσέγγισης για την τοποθέτηση k εγκαταστάσεων στην ευθεία, με $k \geq 2$. Σημειώνουμε ότι τα αρνητικά αποτελέσματα του [112] εξαρτώνται καιρία από το γεγονός ότι η συσταδοποίηση (και η αντίστοιχη

¹Ένας μηχανισμός είναι *ανώνυμος* εάν το αποτέλεσμά του εξαρτάται μόνο από τις θέσεις των παικτών, χωρίς να λαμβάνει υπόψη του την ταυτότητά τους.

τοποθέτηση των εγκαταστάσεων) που παράγεται από οποιονδήποτε ντετερμινιστικό μηχανισμό με φραγμένο λόγο προσέγγισης πρέπει να είναι ευαίσθητη στις λανθασμένες αναφορές τοποθεσίας από συγκεκριμένους παίκτες (δείτε επίσης το υποκεφάλαιο 3.6). Επομένως, είναι πολύ φυσικό να εξετάσουμε αν ο περιορισμός σε γ -ευσταθή στιγμιότυπα επιτρέπει την δημιουργία ντετερμινιστικών ή πιθανοτικών μηχανισμών με μη τετριμμένους λόγους προσέγγισης για το πρόβλημα της τοποθέτησης k εγκαταστάσεων k στην ευθεία (k -Facility Location Game).

Μια κεντρική, ενδιαφέρουσα παρατήρηση, είναι ότι για ένα επαρκώς μεγάλο γ , τα γ -ευσταθή στιγμιότυπα του k -Facility Location διαθέτουν επιπλέον δομή που θα μπορούσε κανείς να εκμεταλλευθεί για τον σχεδιασμό φιλαλήθων μηχανισμών με καλούς λόγους προσέγγισης (βλέπε επίσης το υποκεφάλαιο 3.3). Για παράδειγμα, η τοποθεσία κάθε παίκτη είναι τουλάχιστον $\gamma - 1$ φορές πιο κοντά στην πλησιέστερη εγκατάσταση απ' ό,τι σε οποιαδήποτε άλλη τοποθεσία παίκτη που ανήκει σε διαφορετική συστάδα (Πρόταση 3.3.2). Επιπλέον, για $\gamma \geq 2 + \sqrt{3}$, η απόσταση μεταξύ δύο διαδοχικών συστάδων είναι μεγαλύτερη από τη διάμετρό τους (Λήμμα 3.3.4).

Επιγραμματικά, σε αυτό το πλαίσιο δείχνουμε ότι η βέλτιστη λύση είναι φιλαλήθης σε $(2 + \sqrt{3})$ -ευσταθή στιγμιότυπα, όταν η βέλτιστη λύση στο πραγματικό στιγμιότυπο δεν περιλαμβάνει συστάδες που αποτελούνται μόνο από έναν παίκτη, και ότι η τοποθέτηση της εγκατάστασης στον παίκτη δίπλα από τον δεξιότερο σε κάθε συστάδα της βέλτιστης συσταδοποίησης (ή στον μοναδικό παίκτη, για συστάδες που αποτελούνται από έναν παίκτη) είναι φιλαλήθης και $(n - 2)/2$ -προσεγγιστική για 5-ευσταθή στιγμιότυπα (ακόμη κι αν η βέλτιστη λύση περιλαμβάνει συστάδες με μόνο έναν παίκτη). Από την αρνητική πλευρά, δείχνουμε ότι για κάθε $k \geq 3$ και κάθε $\delta > 0$, δεν υπάρχει ντετερμινιστικός ανώνυμος μηχανισμός που να επιτυγχάνει φραγμένο λόγο προσέγγισης και να είναι φιλαλήθης σε $(\sqrt{2} - \delta)$ -ευσταθή στιγμιότυπα. Αποδεικνύουμε επίσης ότι η τοποθέτηση της εγκατάστασης σε έναν τυχαίο παίκτη σε κάθε μία συστάδα της βέλτιστης συσταδοποίησης είναι φιλαλήθης και 2-προσεγγιστική σε 5-ευσταθή. Εξ όσων γνωρίζουμε, αυτή είναι η πρώτη φορά που αποδεικνύεται η ύπαρξη ντετερμινιστικών (αντίστοιχα πιθανοτικών) φιλαλήθων μηχανισμών με περιορισμένο (αντίστοιχα σταθερό) λόγο προσέγγισης για μια μεγάλη και φυσική κλάση στιγμιότυπων του k -Facility Location προβλήματος.

1.2.2 Δειγματοληψία και Βέλτιστη Εξαγωγή Προτιμήσεων σε Απλούς Μηχανισμούς.

Η αποτελεσματική εξαγωγή προτιμήσεων (*preference elicitation*) έχει αποτελέσει ένα κεντρικό θέμα και μια μεγάλη πρόκληση από το ξεκίνημα της έρευνας στον σχεδιασμό μηχανισμών, με πολλές εφαρμογές σε περιβάλλοντα με πολλαπλούς παίκτες και σύγχρονα συστήματα τεχνητής νοημοσύνης. Πράγματι, η απαίτηση από κάθε παίκτη να επικοινωνήσει όλες τις προτιμήσεις του θεωρείται ευρέως μη πρακτική, και ένας σημαντικός όγκος εργασίας έχει εξερευνήσει εναλλακτικές προσεγγίσεις για την περικοπή της ποσότητας των εξαγόμενων πληροφοριών, ερωτώντας ακολουθιακά μια σειρά φυσικών ερωτημάτων για να αποκτήσει μόνο τα σχετικά (ή απαραίτητα) μέρη

πληροφορίας από κάθε παίκτη.

Αυτή η έμφαση έχει υποστηριχτεί σθεναρά για αρκετούς λόγους. Πρώτον, έχει αναγνωριστεί από τους συμπεριφορικούς οικονομολόγους ότι η απόκτηση πληροφοριών επιβάλλει υψηλό γνωστικό κόστος [190, 164], και οι παίχτες μπορεί ακόμη να διστάζουν να αποκαλύψουν την πλήρη (ιδιωτική) αξιολόγησή τους. Ως αποδεικτικό παράδειγμα, συχνά αναφέρεται η υπεροχή των *indirect* (μη άμεσων) μηχανισμών [29]. Πράγματι, σε ορισμένους τομείς με σοβαρούς περιορισμούς επικοινωνίας [4, 180], ένας *direct revelation* μηχανισμός (απευθείας αποκάλυψης πληροφορίας) - όπου κάθε παίκτης πρέπει να αποκαλύψει τις πλήρεις προτιμήσεις του - θεωρείται ακόμη και ανεφάρμοστος. Για παράδειγμα, η επικοινωνία αναγνωρίζεται τυπικά ως ο κύριος περιορισμός σε κατανεμημένα περιβάλλοντα [170]. Όπως εξηγούν οι Blumrosen και Feldman [46], οι παίχτες λειτουργούν συνήθως σε έναν περικομμένο χώρο δράσης λόγω τεχνικών, συμπεριφορικών ή ρυθμιστικών λόγων. Τέλος, ένας μηχανισμός που επιτυγχάνει αποτελεσματική εξαγωγή προτιμήσεων θα προσέφερε ισχυρότερες εγγυήσεις για τον έλεγχο της ιδιωτικότητας των πληροφοριών [216].

Η γενική ερώτηση για το αν ένας μηχανισμός κοινωνικής επιλογής μπορεί να προσεγγιστεί ακριβώς από ένα μικρότερο υποσύνολο του πλήρους πληθυσμού των παικτών αποτελεί ένα από τα κύρια θέματα στην υπολογιστική κοινωνική επιλογή. Ίσως η πιο καθιερωμένη προσέγγιση για την περικοπή των εξαγόμενων πληροφοριών είναι η χρήση *δειγματοληψίας*. Συγκεκριμένα, δεδομένου ότι σε πολλές εφαρμογές το ζήτημα της συγκέντρωσης προτιμήσεων από όλους τους παίχτες μπορεί να είναι αδύνατο, ο σχεδιαστής πραγματοποιεί μια συγκέντρωση προτιμήσεων επιλέγοντας τυχαία ένα μικρό υποσύνολο του συνολικού πληθυσμού [58, 83]. Αυτή η προσέγγιση είναι ιδιαίτερα γνωστή στο πλαίσιο της *ψηφοφορίας*, όπου ο στόχος είναι συνήθως να 'προβλεφθεί' το αποτέλεσμα του μηχανισμού που έχει πρόσβαση σε πλήρεις πληροφορίες χωρίς να διεξαχθεί πραγματικά η ψηφοφορία για ολόκληρο τον πληθυσμό [82]. Συγκεκριμένα παραδείγματα των παραπάνω σεναρίων περιλαμβάνουν τις δημοσκοπήσεις για τις εκλογές, τα exit poll και τις διαδικτυακές δημοσκοπήσεις.

Στο πρώτο μέρος της εργασίας μας, ακολουθούμε αυτήν τη μακρά γραμμή έρευνας στην υπολογιστική κοινωνική επιλογή. Συγκεκριμένα, αναλύουμε την *δηγματική πολυπλοκότητα* (sample complexity) του διάσημου μηχανισμού του *median* (του διάμεσου) στο πλαίσιο των *facility location games*, όπου κάθε παίκτης συσχετίζεται με ένα σημείο - που αντιστοιχεί στην προτιμώμενη του τοποθεσία - σε κάποιο υποκείμενο μετρικό χώρο. Ο μηχανισμός του διάμεσου έχει ιδιαίτερη σημασία στην κοινωνική επιλογή. Πράγματι, το διάσημο θεώρημα αδυναμίας των Gibbard-Satterthwaite [130, 207] δηλώνει ότι για κάθε κανόνα ψηφοφορίας (*voting rule*) που είναι προς (*onto*) - για κάθε εναλλακτικό αποτέλεσμα (δηλ. νικητή της ψηφοφορίας) υπάρχει ένα προφίλ ψηφοφόρων που θα έκανε αυτήν την εναλλακτική να κερδίσει - και μη-δικτατορικός (*non-dictatorial*), υπάρχουν περιπτώσεις όπου ένας παίκτης έχει συμφέρον να ψηφίσει μια εναλλακτική που δεν αντιστοιχεί στις πραγματικές του προτιμήσεις - δηλαδή ο κανόνας δεν είναι *φιλαλήθης*. Σημαντικό είναι ότι αυτό το εμπόδιο εξαφανίζεται όταν οι προτιμήσεις των ψηφοφόρων περιορίζονται. Ίσως ο πιο γνωστός τέτοιος περιορισμός

προτιμήσεων είναι αυτός των *single-peaked* (μονο-κορυφωμένων) προτιμήσεων², που εισήγαγε ο Black [45], για τις οποίες ο μηχανισμός του διάμεσου του Moulin [181] είναι πράγματι φιλαλήθης.

Το κύριο αποτέλεσμα μας σε αυτήν την κατεύθυνση είναι να δείξουμε ότι για κάθε $\epsilon > 0$, ένα δείγμα μεγέθους $c(\epsilon) = \Theta(1/\epsilon^2)$ παράγει αναμενόμενο λόγο προσέγγισης $1 + \epsilon$ ως προς τον βέλτιστο κοινωνικό κόστος του γενικευμένου μηχανισμού διαμέσου στο μετρικό χώρο $(\mathbb{R}^d, |\cdot|_1)$, όσο ο αριθμός των παιχτών $n \rightarrow \infty$.

Στο δεύτερο μέρος της εργασίας μας, προσπαθούμε να σχεδιάσουμε δημοπρασίες με ελάχιστη πολυπλοκότητα επικοινωνίας. Ως ένα συγκεκριμένο κίνητρο, εξετάζουμε μια δημοπρασία ενός αντικειμένου (*single-item*), και υποθέτουμε ότι κάθε αποτίμηση των παιχτών μπορεί να εκφραστεί με k bits. Μια σημαντική παρατήρηση είναι ότι τα πιο κυρίαρχα μοντέλα δημοπρασιών είναι πολύ μη αποδοτικά από πλευράς επικοινωνίας. Πράγματι, η *Vickrey* - ή δημοπρασία σφραγισμένης προσφοράς (*sealed bid*)-δημοπρασία [221] είναι ένας μηχανισμός άμεσης αποκάλυψης, που απαιτεί ολόκληρες τις ιδιωτικές πληροφορίες από τους n παίκτες, οδηγώντας σε μια πολυπλοκότητα επικοινωνίας των $n \cdot k$ bits. Πρέπει όμως να σημειωθεί ότι, παρόλο που η δημοπρασία του *Vickrey* διαθέτει πολλές θεωρητικά ελκυστικές ιδιότητες, η αντίστοιχη ανοδική (*ascending*) μορφή της εμφανίζει καλύτερη απόδοση στην πράξη [28, 29, 145, 144], για λόγους που σχετίζονται κυρίως με την απλότητα, τη διαφάνεια και τις εγγυήσεις απορρήτου της τελευταίας μορφής. Δυστυχώς, η υλοποίηση του κανόνα του *Vickrey* μέσω μιας δημοπρασίας *English* (δλδ *ascending*) απαιτεί - στη χειρότερη περίπτωση - εκθετική επικοινωνία τάξης $n \cdot 2^k$, καθώς ο δημοπράτης πρέπει να καλύψει ολόκληρο το χώρο αποτιμήσεων. Συνεπώς, το μοτίβο εξαγωγής προτιμήσεων σε μια δημοπρασία τύπου *English* είναι ευρέως ανεπιτυχές και η έλλειψη προγνώσεων για τις αποτιμήσεις των παιχτών υπονομεύει σημαντικά την απόδοσή της.

Σε αυτό το πλαίσιο, δείχνουμε ότι ο κανόνας του *יצחק* μπορεί να εφαρμοστεί με αναμενόμενη επικοινωνία $1 + \epsilon$ bits από έναν μέσο παίκτη, για κάθε $\epsilon > 0$, επιτυγχάνοντας ασυμπτωτικά το προφανές κάτω όριο. Ως συμπέρασμα, παρέχουμε έναν ενδιαφέροντα τρόπο αύξησης της τιμής σε μια δημοπρασία τύπου *English*. Χρησιμοποιούμε επίσης τη μορφή των δημοπρασιών ενός αντικειμένου με ένα αποδοτικό σχήμα κωδικοποίησης για να αποδείξουμε ότι τα ίδια όρια επικοινωνίας μπορούν να επιτευχθούν σε δημοπρασίες με αθροιστικές αποτιμήσεις (*additive valuations*) μέσω ταυτόχρονων αναδυόμενων δημοπρασιών, υποθέτοντας ότι ο αριθμός των αντικειμένων είναι σταθερός. Τέλος, προτείνουμε μια αύξουσα δημοπρασία πολλαπλών μονάδων αντικειμένου (*multi-unit auction*) με παίκτες που ο καθένας θέλει το πολύ μια μονάδα.

²Συγκεκριμένα, υποθέτουμε ότι οι εναλλακτικές είναι ταξινομημένες σε μια γραμμή, αντιπροσωπεύοντας τις θέσεις τους· όπως περιγράφεται στη θεωρία της χωρικής ψηφοφορίας [90, 24], είναι συχνά λογικό να υποθέτουμε ότι οι εναλλακτικές μπορούν να αντιπροσωπευθούν ως σημεία σε μια γραμμή (π.χ. στις πολιτικές εκλογές η θέση ενός υποψηφίου μπορεί να υποδεικνύει εάν ανήκει στην αριστερή ή δεξιά πτέρυγα). Οι προτιμήσεις ενός παίκτη είναι *single-peaked* αν προτιμά εναλλακτικές που είναι πιο κοντά στο κορυφαίο σημείο του. Σημειώνουμε ότι σε περιπτώσεις *single-peaked* προτιμήσεων είναι γνωστό ότι δεν μπορεί να υπάρχει κύκλος *Condorcet*.

ο μηχανισμός μας ανακοινώνει σε κάθε γύρο δύο ξεχωριστές τιμές και βασίζεται σε έναν αλγόριθμο δειγματοληψίας που εκτελεί προσεγγιστική επιλογή με περιορισμένη επικοινωνία, οδηγώντας πάλι σε ασυμπτωτικά βέλτιστη επικοινωνία. Τα αποτελέσματά μας δεν απαιτούν καμία προηγούμενη γνώση για τις αποτιμήσεις των παικτών, και κυρίως προέρχονται από φυσικές τεχνικές δειγματοληψίας.

1.3 Σχεδιασμός μηχανισμών για μη στρατηγικά περιβάλλοντα: Περιβάλλοντα με περιορισμένη πληροφορία.

1.3.1 Μετρική Παραμόρφωση υπό Περιορισμένες Πληροφορίες

Η συγκέντρωση των προτιμήσεων των ατόμων σε μια συλλογική απόφαση βρίσκεται στην καρδιά της θεωρίας της κοινωνικής επιλογής και έχει πρόσφατα βρει πολλές εφαρμογές σε πεδία όπως η ανάκτηση πληροφοριών, τα συστήματα συστάσεων και η μηχανική μάθηση [222, 223, 36, 124, 161]. Η κλασική θεωρία του ον Νευμαν ανδ Μοργενστερν [224] υποθέτει ότι οι ατομικές προτιμήσεις αντιπροσωπεύονται μέσω μιας *συνάρτησης ωφέλειας*, που αναθέτει αριθμητικές (ή *cardinal*) τιμές σε κάθε εναλλακτική. Ωστόσο, στη θεωρία των ψηφοφοριών, καθώς και στις περισσότερες πρακτικές εφαρμογές, οι μηχανισμοί συνήθως συλλέγουν μόνο *διατακτικές* (*ordinal*) πληροφορίες από τους ψηφοφόρους, που υποδεικνύουν μια σειρά προτιμήσεων για τους υποψήφιους. Παρόλο που αυτό φαίνεται ασύμβατο με μια *utilitarian* αναπαράσταση, έχει αναγνωριστεί ότι μπορεί να είναι δύσκολο για έναν ψηφοφόρο να καθορίσει μια ακριβή αριθμητική τιμή για κάθε εναλλακτική του (δλ. για κάθε υποψήφιο), και η παροχή μόνο διατακτικών πληροφοριών μειώνει σημαντικά το γνωστικό φορτίο. Αυτό θέτει το ερώτημα: Ποια είναι η απώλεια στην αποτελεσματικότητα ενός μηχανισμού που εξάγει μόνο διατακτικές πληροφορίες σε σχέση με τη *κοινωνική ευημερία*, δηλαδή το άθροισμα των ατομικών ωφελειών πάνω σε μια επιλεγμένη εναλλακτική. Το πλαίσιο της *παραμόρφωσης*, που εισήγαγαν οι Procaccia και Rosenstein [194], μετρά ακριβώς αυτήν την απώλεια από την κατεύθυνση των προσεγγιστικών αλγορίθμων, και έχει λάβει σημαντική προσοχή τα τελευταία χρόνια.

Όπως αποδεικνύεται, οι εγγυήσεις προσέγγισης που μπορούμε να ελπίζουμε ότι μπορούμε να επιτύχουμε εξαρτώνται καιρία από τις υποθέσεις που κάνουμε για τις ωφέλειες των ατόμων. Για παράδειγμα, στην απουσία οποιασδήποτε δομής, στο [194] παρατηρήθηκε ότι *κανένας* διατακτικός μηχανισμός δεν μπορεί να επιτύχει περιορισμένη παραμόρφωση. Σε αυτό το έργο, επικεντρωνόμαστε στο πλαίσιο της *μετρικής παραμόρφωσης*, που εισήγαγαν οι Anshelevich et al. [18], όπου οι ψηφοφόροι και οι υποψήφιοι θεωρούνται σημεία σε ένα αυθαίρετο μετρικό χώρο· αυτό είναι παρόμοιο με τα μοντέλα στη θεωρία της χωρικής ψηφοφορίας (*spatial voting theory*) [74]. Σε αυτό το πλαίσιο, οι προτιμήσεις των ψηφοφόρων μετρώνται μέσω της "κοντινότητάς" τους από

κάθε υποψήφιο, και ο στόχος είναι να διαλέξουμε έναν υποψήφιο που (περίπου) ελαχιστοποιεί το *κοινωνικό κόστος*, δηλαδή τη συνολική απόσταση από τους ψηφοφόρους. Μια αρκετά απλοϊκή μορφή αυτού του πλαισίου εκδηλώνεται όταν οι ψηφοφόροι και οι υποψήφιοι ενσωματώνονται σε μια μονοδιάστατη γραμμή, και η τοποθεσία τους υποδηλώνει αν βρίσκονται “αριστερά” ή “δεξιά” στο πολιτικό φάσμα. Ωστόσο, το πλαίσιο της μετρικής παραμόρφωσης έχει πολύ μεγαλύτερη εμβέλεια, καθώς δεν γίνονται καθόλου υποθέσεις για τον υποκείμενο μετρικό χώρο.

Σημαντικό είναι ότι αυτό το παράδειγμα προσφέρει έναν πειστικό τρόπο για να συγκρίνουμε ποσοτικά διάφορους κανόνες ψηφοφορίας που χρησιμοποιούνται συχνά στην πράξη [214, 149, 132, 18], ενώ λειτουργεί και ως ένα σημείο αναφοράς για τον σχεδιασμό νέων μηχανισμών που αναζητούν καλύτερα όρια παραμόρφωσης [131, 183]. Μια κοινή υπόθεση που χρησιμοποιείται σε αυτήν τη γραμμή έρευνας είναι ότι ο αλγόριθμος έχει πρόσβαση στις πλήρεις σειρές προτεραιότητας των ψηφοφόρων. Ωστόσο, υπάρχουν πολλά πρακτικά σενάρια όπου θα ήταν επιθυμητό να περικόψουμε τις διατακτικές πληροφορίες που αποκομίζονται από τον μηχανισμό. Για παράδειγμα, η απαίτηση για παροχή μόνο των κορυφαίων προτιμήσεων θα μπορούσε να ανακουφίσει περαιτέρω το γνωστικό κόστος, καθώς ενδέχεται να είναι δύσκολο για έναν ψηφοφόρο να συγκρίνει εναλλακτικές που βρίσκονται στο κάτω μέρος της λίστας προτιμήσεών του (για επιπρόσθετη κινητοποίηση για την εξέταση των ατελών ή μερικών διατάξεων παραπέμπουμε στα [123, 70, 39], και τις αναφορές τους), ενώ οποιαδήποτε περικοπή της αποκομιζόμενης πληροφορίας θα μεταφραστεί επίσης σε πιο αποτελεσματική επικοινωνία. Αυτοί οι λόγοι έχουν κατευθύνει αρκετούς συγγραφείς να μελετήσουν την υποβάθμιση της παραμόρφωσης υπό απουσία πληροφοριών [150, 14, 138, 98, 50], επιτρέποντας πιθανώς μια πιθανοτική ανάλυση (δείτε την υποενότητα μας με τα σχετικά έργα). Σε αυτό το έργο, ακολουθούμε αυτήν τη γραμμή έργου, προσφέροντας νέες αντιλήψεις και βελτιωμένα όρια σε σχέση με τα προηγούμενα αποτελέσματα.

Η κύρια συνεισφορά μας σε αυτήν την κατεύθυνση είναι τριπλή: Πρώτον, εξετάζουμε μηχανισμούς που πραγματοποιούν μια ακολουθία διαδοχικών συγκρίσεων μεταξύ των υποψηφίων. Δείχνουμε ότι ένας δημοφιλής ντετερμινιστικός μηχανισμός που λειτουργεί σε πολλές φάσεις αποκλεισμού υποψηφίων οδηγεί σε παραμόρφωση $\mathcal{O}(\log m)$ ενώ κάνει μόνο $m - 1$ από τις $\Theta(m^2)$ δυνατές διαδοχικές συγκρίσεις, όπου m αντιπροσωπεύει τον αριθμό των υποψηφίων. Η ανάλυσή μας για αυτόν τον μηχανισμό εκμεταλλεύεται ένα ισχυρό τεχνικό λήμμα που αναπτύχθηκε από τον Kempe (AAAI '20). Παρέχουμε επίσης ένα αντίστοιχο κάτω όριο στην παραμόρφωση του. Αντίθετα, αποδεικνύουμε ότι οποιοσδήποτε μηχανισμός που πραγματοποιεί λιγότερες από $m - 1$ διαδοχικές συγκρίσεις οδηγεί αναπόφευκτα σε μη περιορισμένη παραμόρφωση. Επιπλέον, μελετούμε τη δύναμη των ντετερμινιστικών μηχανισμών υπό ατελείς διατάξεις. Ειδικότερα, όταν οι παίχτες παρέχουν τις k κορυφαίες τους προτιμήσεις, δείχνουμε ένα άνω όριο $6m/k + 1$ για την παραμόρφωση, για κάθε $k \in 1, 2, \dots, m$. Επομένως, βελτιώνουμε σημαντικά το προηγούμενο όριο $12m/k$ που είχε θεσπίσει ο Kempe (AAAI '20), και πλησιάζουμε περισσότερο το καλύτερο γνωστό κάτω όριο. Τέλος, ασχολούμαστε με την πολυπλοκότητα δειγματοληψίας για να εξασφαλίσουμε σχεδόν βέλτιστη παραμόρφωση σε σχέση με τον βέλτιστο κοινωνικό κόστος με υψη-

λή πιθανότητα. Η κύρια συνεισφορά μας είναι να δείξουμε ότι ένα τυχαίο δείγμα από $\Theta(m/\epsilon^2)$ ψηφοφόρους είναι αρκετό για να διασφαλίσει παραμόρφωση $3 + \epsilon$ με υψηλή πιθανότητα, για κάθε αρκετά μικρό $\epsilon > 0$. Αυτό το αποτέλεσμα βασίζεται στην ανάλυση της ευαισθησίας του ντετερμινιστικού μηχανισμού που εισήγαγαν οι Gkatzelis, Halpern και Shah (FOCS '20). Σημαντικό είναι ότι όλα τα όρια της δειγματοληψίας μας είναι ανεξάρτητα από την κατανομή.

1.3.2 Διαστασιμότητα και Συντονισμός στην Ψηφοφορία: Η παραμόρφωση του STV

Σε αυτήν την εργασία συνεχίζουμε την έρευνα στην περιοχή της μετρικής παραμόρφωσης. Μια σημαντική οπτική, είναι ότι αυτό το πλαίσιο προσφέρει ένα ποσοτικό “σημείο αναφοράς” για τη σύγκριση διαφόρων κανόνων ψηφοφορίας που χρησιμοποιούνται συχνά στην πράξη. Πράγματι, μία από τις κύριες σχέψεις της εργασίας μας αφορά στον χαρακτηρισμό της απόδοσης του μηχανισμού *Single Transferable Vote*³, η αλλιώς STV.

Ο STV είναι ένα δημοφιλές σύστημα ψηφοφορίας που χρησιμοποιείται στις εθνικές εκλογές πολλών χωρών, συμπεριλαμβανομένων της Αυστραλίας, της Ιρλανδίας και της Ινδίας, καθώς και σε πολλές άλλες διαδικασίες συγκέντρωσης προτιμήσεων, όπως για παράδειγμα στα Βραβεία Ακαδημίας Oscar. Για να είμαστε πιο ακριβείς, ο STV προχωρά με επαναληπτικό τρόπο: Σε κάθε γύρο, οι ψηφοφόροι ψηφίζουν τον προτιμώμενο τους υποψήφιο - ανάμεσα στους ενεργούς, ενώ ο υποψήφιος που έλαβε την λιγότερη υποστήριξη στον τρέχοντα γύρο αποκλείεται. Αυτή η διαδικασία επαναλαμβάνεται για $m - 1$ γύρους, όπου το m αντιπροσωπεύει τον αριθμό των (αρχικών) υποψηφίων, και ο τελευταίος υποψήφιος που επιζεί κηρύσσεται νικητής του STV. Επιπλέον, πρέπει να σημειωθεί ότι αυτή η διαδικασία είναι γενικά μη-ντετερμινιστική λόγω της ανάγκης για μηχανισμό λήψης αποφάσεων σε περίπτωση ισοπαλίας: όπως στο [214], θα εργαστούμε με το *parallel universe* μοντέλο των “ονιτζερ κ.ά. [79], όπου ένας υποψήφιος θεωρείται νικητής του STV αν επιβιώνει μετά από κάποια ακολουθία αποκλεισμών.


Σε αυτό το πλαίσιο, οι Skowron και Elkind [214] ήταν οι πρώτοι που ανέλυσαν τη μετρική παραμόρφωση του STV υπό μετρικές προτιμήσεις. Συγκεκριμένα, έδειξαν ότι η παραμόρφωση του STV σε γενικούς μετρικούς χώρους είναι πάντα $O(\log m)$, ενώ έδωσαν επίσης ένα σχεδόν tight κάτω όριο της μορφής $\Omega(\sqrt{\log m})$. Ενδιαφέρον παρουσιάζει μια προσεκτική εξέταση του κάτω ορίου τους, αποκαλύπτοντας την ύπαρξη ενός μετρικού υπό-χώρου υψηλών διαστάσεων, όπως απεικονίζεται στο Σχήμα 1.1, και είναι γνωστό γεγονός στη θεωρία των μετρικών απεικονίσεων (metric embeddings) ότι δεν μπορούν να απεικονιστούν *ισομετρικά* (*isometrically*) σε χαμηλών διαστάσεων⁴ Ευκλείδειους χώρους [173]. Ως αποτέλεσμα, οι Skowron και Elkind

³Για συνέπεια με την προηγούμενη έρευνα, το STV θα αντιπροσωπεύει σε όλη αυτή την εργασία την μορφή του προβλήματος με *μοναδικού νικητή*

⁴Λέμε ότι ένας Ευκλείδειος χώρος είναι χαμηλής διάστασης αν η διάστασή του d είναι φραγμένη από μια “μικρή” παγκόσμια σταθερά, δηλαδή $d = O(1)$.

[214] αφήνουν ανοικτό το ακόλουθο ενδιαφέρον ερώτημα:

Ερώτημα Α. Ποια είναι η παραμόρφωση του STV σε Ευκλείδειους χώρους χαμηλής διάστασης;



Σχήμα 1.1: Μια μετρική υψηλών-διαστάσεων σε μορφή γράφου ‘άστέρα’.

Βεβαίως, η απόδοση των κανόνων ψηφοφορίας σε χώρους χαμηλής διάστασης έχει αποτελέσει αντικείμενο εντατικής μελέτης στη θεωρία της χωρικής ψηφοφορίας, υπό την υπόθεση ότι οι ψηφοφόροι και οι υποψήφιοι είναι συνήθως ενσωματωμένοι σε υποχώρους μικρής διάστασης [24, 90]. Επομένως, είναι φυσικό να αναρωτηθούμε εάν μπορούμε να εστιάσουμε την ανάλυση του STV σε χώρους χαμηλής διάστασης. Στην πραγματικότητα, ως μέρος μιας ευρύτερης ατζέντας, μπορούν να τεθούν ανάλογα ερωτήματα και για άλλους μηχανισμούς. Ωστόσο, είναι ενδιαφέρον να σημειώσουμε ότι για πολλούς κανόνες ψηφοφορίας που αναλύονται μέσα στο πλαίσιο της μετρικής παραμόρφωσης υπάρχουν κάτω όρια για χαμηλές-διαστάσεις. Μερικά παραδείγματα αυτών δίνονται στον πίνακα 1.1. Σε αντίθεση, η εργασία μας θα εστιάσει στον STV από τους μηχανισμούς στον πίνακα 1.1, απαντώντας αποτελεσματικά στο Ερώτημα Α. Σημαντικό είναι ότι θα παράσχουμε έναν χαρακτηρισμό που πηγαίνει πολύ πέρα από τους ευκλείδειους χώρους, σε μετρικές με ‘έγγενώς’ χαμηλή διάσταση (low intrinsic dimension).

Η επόμενη στόχευση της εργασίας μας ενθαρρύνεται απευθείας από την αποδοτικότητα του κανόνα STV σε σύγκριση με τον κανόνα πλειοψηφίας (plurality rule), και ιδιαίτερα τις στρατηγικές συνέπειες αυτής της ασυνέπειας. Ένα καλό σημείο εκκίνησης για αυτήν τη συζήτηση πηγάζει από το γεγονός ότι σε πολλές θεμελιώδεις περιπτώσεις συγκέντρωσης προτιμήσεων οι υποψήφιοι επιλέγονται από μη αποδοτικούς μηχανισμούς, και σε πολλές περιπτώσεις οποιαδήποτε αλλαγή αντιμετωπίζει ανυπερβλήτα εμπόδια. Για παράδειγμα, στις πολιτικές εκλογές ο μηχανισμός ψηφοφορίας καθορίζεται συνήθως από τους εκλογικούς νόμους, ή ακόμη και από το σύνταγμα [165].

Mechanism	Lower Bound	Dimension
Plurality	$2m - 1$	1
Borda	$2m - 1$	1
Copeland	5	2
Veto	$2n - 1$	1
Approval	$2n - 1$	1

Table 1.1: Η Ευκλείδεια διάσταση που απαιτείται για να κατασκευαστεί ένα (ισχυρό) κάτω όριο για διάφορους κοινούς κανόνες ψηφοφορίας, αυτά τα αποτελέσματα περιλαμβάνονται στο [18]. Να σημειώσουμε ότι για τον Copeland, η μετρική που κατασκευάστηκε στο [18] δεν είναι Ευκλείδεια, αλλά μπορεί εύκολα να τροποποιηθεί σε μια.

Ως αποτέλεσμα, η κατανόηση της συμπεριφοράς στρατηγικών πρακτόρων (παικτών) όταν αντιμετωπίζουν ανεπαρκείς μηχανισμούς είναι απολύτως σημαντική [56, 225]. Ένας αρκετά ανεξάρτητος τρόπος να το δούμε αυτό είναι εάν αυτόνομοι παίκτες μπορούν να συγχλίνουν σε αποδεκτές κοινωνικές επιλογές μέσω φυσικών κανόνων μάθησης. Αυτό θέτει το ερώτημα:

Ερώτημα Β. Σε τι βαθμό μπορεί η στρατηγική συμπεριφορά να βελτιώσει την αποδοτικότητα στο πρόβλημα της ψηφοφορίας;

Τονίζουμε ότι, αν και στην απουσία οποιασδήποτε πληροφορίας μπορεί να είναι ασαφές πώς μπορούν οι παίκτες να εμπλακούν σε στρατηγική συμπεριφορά, στις περισσότερες περιπτώσεις που ενδιαφέρουν τους παίκτες, αυτοί έχουν πολλές προγενέστερες πληροφορίες πριν ψηφίσουν, π.χ. μέσω δημοσκοπήσεων, έρευνας αγοράς, προγνώσεων, προηγούμενων εκλογών ή ακόμα και πρώιμης ψηφοφορίας. Πράγματι, υπάρχει μια εκτενής σειρά εργασιών που μελετά τη δυναμική του πληθυσμού για παίκτες που ψηφίζουν ανταποκρινόμενοι στις πληροφορίες που διαθέτουν (βλ. [200] και τις αναφορές του), καθώς και το ρόλο της πληροφορίας στον προσδιορισμό της δημόσιας πολιτικής [158].

Για να μελετήσουμε αυτήν την κατεύθυνση, προτείνουμε ένα φυσικό μοντέλο, σύμφωνα με το οποίο οι παίκτες ενεργούν επαναληπτικά βάσει ορισμένων μερικών ανατροφοδοτήσεων σχετικά με τις προτιμήσεις των άλλων ψηφοφόρων. Εξηγούμε πώς ο STV μπορεί να ενταχθεί πολύ φυσικά σε αυτό το πλαίσιο, ενώ καθιερώνουμε την ύπαρξη απλών και αποκεντρωμένων δυναμικών συντονισμού που συγχλίνουν σε μια κοντινά βέλτιστη εναλλακτική λύση.

Το τελευταίο μέρος της εργασίας μας προσφέρει ορισμένες βελτιώσεις και επεκτάσεις προηγούμενων εργασιών, κυρίως κινούμενες από ορισμένες θεμελιώδεις εφαρμογές στο πλαίσιο των facility location παιχνιδιών. Συγκεκριμένα, επικεντρωθήκαμε κυρίως στον βέλτιστο - υπό μετρικές προτιμήσεις - ντετερμινιστικό μηχανισμό που προτάθηκε πρόσφατα από τους Gkatzelis, Halpern και Shah [131]: δείχνουμε ότι επι-

τυγχάνει βέλτιστη παραμόρφωση σε ultra-metrics και σχεδόν βέλτιστη παραμόρφωση για αποστάσεις που ικανοποιούν τις τριγωνικές ανισότητες (προσεγγιστικά).

1.4 Υπολογισμός ισορροπιών.

1.5 Το πρόβλημα Node Max-Cut και ο Υπολογισμός Ισορροπιών σε Βεβαρημένα Παίγνια με Γραμμικές Συναρτήσεις

Με κίνητρο την εντυπωσιακή επιτυχία της τοπικής αναζήτησης στη συνδυαστική βελτιστοποίηση, όταν οι Johnson et al. εισήγαγαν την κλάση πολυπλοκότητας τοπικής αναζήτησης (Polynomial Local Search - PLS) [143], που αποτελείται από προβλήματα τοπικής αναζήτησης με πολυωνυμικά επαληθεύσιμα τοπικά βέλτιστα. Η PLS περιλαμβάνει πολλά φυσικά πλήρη προβλήματα (δείτε π.χ., [176]), με το CIRCUIT-FLIP [143] και το MAX-CUT [208] να αποτελούν τα πλέον γνωστά, και θέτει τα θεμέλια για την μελέτη της πολυπλοκότητας του υπολογισμού τοπικών βέλτιστων. Τα τελευταία 15 χρόνια, ένα σημαντικό μέρος της έρευνας στην PLS-πληρότητα κινητοποιήθηκε από το πρόβλημα υπολογισμού μιας αμιγούς ισορροπίας κατά Nash παιγνίων δυναμικού (potential games) (δείτε π.χ., [2, 213, 127] και τις αναφορές που αναφέρονται εκεί). Σε αυτά τα παίγνια, κάθε βελτιωτική αλλαγή στρατηγικής από έναν παίκτη μειώνει μια συνάρτηση δυναμικού και τα τοπικά βέλτιστα αντιστοιχούν σε αμιγείς ισορροπίες κατά Nash (pure Nash equilibria) [178].

Ο υπολογισμός ενός τοπικού βέλτιστου του MAX-CUT με την γειτονιά FLIP (γνωστό και ως LOCAL-MAX-CUT) έχει εξεταστεί εκτενώς ως ένα από τα πιο διαδεδομένα προβλήματα στην PLS. Δεδομένου ενός γράφου με βάρη στις ακμές, μια τομή θεωρείται τοπικά βέλτιστη εάν δεν μπορούμε να αυξήσουμε το βάρος της μετακινώντας έναν κόμβο από τη μία πλευρά της τομής στην άλλη. Από τότε που αποδείχθηκε η PLS-πληρότητά του από τους Schaffer και Yannakakis [208], οι ερευνητές έχουν δείξει ότι το LOCAL-MAX-CUT παραμένει PLS-πλήρες για γράφους με μέγιστο βαθμό πέντε [89], είναι πολυωνυμικά επιλύσιμο για κυβικούς γράφους [191], και η *smoothed* πολυπλοκότητα του είναι είτε πολυωνυμική σε πλήρεις γράφους [12] και αραιούς γράφους [89], είτε σχεδόν πολυωνυμική σε γενικούς γράφους [71, 92]. Επιπλέον, λόγω της απλότητάς του και της ευελιξίας του, το MAX-CUT έχει χρησιμοποιηθεί ευρέως σε αναγωγές PLS (βλ. για παράδειγμα [2, 127, 213]). Το LOCAL-MAX-CUT μπορεί επίσης να χαρακτηριστεί ως ένα παίγνιο, όπου κάθε κόμβος στοχεύει να μεγιστοποιήσει το συνολικό βάρος των ακμών που διασχίζουν την τομή. Τα παιχνίδια τομής είναι παίγνια δυναμικού (η τιμή της τομής είναι η συνάρτηση δυναμικού), κάτι που έχει ενθαρρύνει την έρευνα στον αποτελεσματικό υπολογισμό προσεγγιστικών ισορροπιών για το LOCAL-MAX-CUT [41, 64]. Κατά την αντίληψή μας, πέρα από την έρευνα για την *smoothed* πολυπλοκότητα του LOCAL-MAX-CUT (και ίσως το LOCAL-MAX-CUT είναι P-πλήρες για μη-βεβαρημένους γράφους [208, Θεώρημα 4.5]), δεν έχει

υπάρξει μέχρι στιγμής έρευνα για το κατά πόσο (και σε τι βαθμό) επιπρόσθετη δομή στα βάρη των ακμών επηρεάζει την πολυπλοκότητα του LOCAL-MAX-CUT.

Μια στενά συνδεδεμένη κατεύθυνση έρευνας ασχολείται με την πολυπλοκότητα του υπολογισμού μιας αμιγούς ισορροπίας Nash (ή PNE, για συντομία) σε παίγνια συμφόρησης [201], ένα τυπικό παράδειγμα παιγνίων δυναμικού [178] και μία από τις πιο εκτενώς μελετημένες κλάσεις παιχνιδιών στη περιοχή της Αλγοριθμικής Θεωρίας Παιγνίων (δείτε π.χ., [113] για μια σύντομη ανάλυση των προηγούμενων ερευνών). Στα παίγνια συμφόρησης (ή CGs, για συντομία), ένα πεπερασμένο σύνολο παικτών ανταγωνίζονται για ένα πεπερασμένο σύνολο πόρων. Οι στρατηγικές αντιστοιχούν σε υποσύνολα πόρων, και οι παίχτες στοχεύουν στο να ελαχιστοποιήσουν τον συνολικό κόστος των πόρων στις στρατηγικές τους. Κάθε πόρος e συσχετίζεται με μια (μη αρνητική και μη φθίνουσα) συνάρτηση καθυστέρησης, που καθορίζει το κόστος χρήσης του e ως συνάρτηση της συμφόρησης του e (δηλαδή, του αριθμού παικτών που περιλαμβάνουν το e στη στρατηγική τους). Οι ερευνητές έχουν εκτενώς μελετήσει τις ιδιότητες ειδικών περιπτώσεων και παραλλαγών των παιγνίων συμφόρησης. Στο πλαίσιο αυτής της εργασίας οι ποιό σχετικές εστιάζουν σε *συμμετρικά* (αντίστοιχα, *μη-συμμετρικά*) CGs, όπου οι παίχτες μοιράζονται το ίδιο σύνολο στρατηγικών (αντίστοιχα, έχουν διαφορετικά σύνολα στρατηγικών), παίγνια συμφόρησης σε *δίκτυα*, όπου οι στρατηγικές αντιστοιχούν σε μονοπάτια σε ένα υποκείμενο δίκτυο (γράφο), και στα *βεβαρημένα* CGs, όπου οι παίχτες συνεισφέρουν στη συμφόρηση με διαφορετικά βάρη ο καθένας.

Οι Fabrikant et al. [95] απέδειξαν ότι ο υπολογισμός ενός PNE για μη-συμμετρικά CGs δικτύου ή συμμετρικά CGs είναι PLS-πλήρης, και ανάγεται σε ένα πρόβλημα ελάχιστης ροής min-cost-flow για συμμετρικά CGs δικτύου. Περίπου την ίδια περίοδο, οι [118, 189] απέδειξαν ότι τα παίγνια συμφόρησης με βάρη έχουν μια (βεβαρημένη) συνάρτηση δυναμικού, και συνεπώς έχουν τουλάχιστον ένα PNE, εάν οι συναρτήσεις καθυστέρησης είναι είτε αφινικές είτε εκθετικές (και οι [140, 141] απέδειξαν ότι από μια άποψη, αυτός ο περιορισμός είναι απαραίτητος). Στη συνέχεια, οι Ackermann et al. [2] χαρακτήρισαν τα σύνολα στρατηγικών των CGs που εγγυώνται αποτελεσματικό υπολογισμό ισορροπιών. Χρησιμοποίησαν επίσης μια παραλλαγή του LOCAL-MAX-CUT, που ονομάζεται *threshold games* (παίγνια κατωφλίου), για να απλοποιήσουν την απόδειξη της PLS-πληρότητας του [95] και για να δείξουν ότι ο υπολογισμός ενός PNE για ασύμμετρα δίκτυα CGs με (εκθετικά απότομες) γραμμικές συναρτήσεις καθυστέρησης είναι PLS-πλήρης.

Από την άλλη πλευρά, η πολυπλοκότητα του υπολογισμού ισορροπίας για τα CGs με βάρη δεν είναι καλά κατανοημένη μέχρι στιγμής. Όλα τα αποτελέσματα υπολογιστικής πολυπλοκότητας που προαναφέρθηκαν ισχύουν και για τα CGs με βάρη, καθώς γενικεύουν τα τυπικά CGs (όπου οι παίχτες έχουν μοναδιαίο βάρος). Αλλά από τη θετική πλευρά, γνωρίζουμε μόνο πώς να υπολογίσουμε αποτελεσματικά ένα PNE για τα CGs με βάρη σε parallel links με γενικές συναρτήσεις καθυστέρησης [120] και για τα CGs με βάρη σε parallel links με ταυτοτικές συναρτήσεις καθυστέρησης και μη-συμμετρικές στρατηγικές [126]. Παρά το σημαντικό ενδιαφέρον για τον υπολογισμό (ακριβούς ή προσεγγιστικού) του PNE για τα CGs (βλ. για παράδειγμα [63, 64, 153]

και τις αναφορές που περιέχονται εκεί), δεν κατανοούμε πώς (και σε ποιο βαθμό) η πολυπλοκότητα του υπολογισμού ισορροπίας επηρεάζεται από τα βάρη των παικτών. Αυτό ισχύει ιδιαίτερα για τα CGs με βάρη με γραμμικές συναρτήσεις καθυστέρησης, τα οποία έχουν συνάρτηση δυναμικού και ο υπολογισμός της ισορροπίας βρίσκεται στην PLS.

Σε αυτήν την κατεύθυνση, αρχικά δείχνουμε ότι ο υπολογισμός μιας PNE με γραμμικές συναρτήσεις βαρών είναι PLS-πλήρης για πολύ περιορισμένους χώρους στρατηγικών, δηλαδή όταν οι στρατηγικές των παικτών είναι διαδρομές σε ένα δίκτυο σειριακό-παράλληλο (series parallel) με μοναδικούς (κοινούς για όλους τους παίκτες) αρχικό κόμβο και κόμβο προορισμού. Επιπλέον, στην ίδια εργασία αποδεικνύεται ότι και για πολύ περιορισμένες συναρτήσεις καθυστέρησης, δηλαδή των η καθυστέρηση σε κάθε πόρο είναι ίση με τη συμφόρηση, σε γενικά δίκτυα το πρόβλημα είναι επίσης PLS-πλήρες, μέσω μιας αναγωγής σε ένα ειδικά ορισμένο νέο πρόβλημα, το *Node-Max-Cut*. Τα αποτελέσματά μας αποκαλύπτουν ένα εντυπωσιακό χάσμα ανάμεσα στην πολυπλοκότητα υπολογισμού μιας PNE στα παίγνια συμφόρησης με παίκτες που έχουν βάρη και στα παίγνια χωρίς βάρη, καθώς στην περίπτωση των μη βεβαρημένων παικτών, το PNE μπορεί να υπολογιστεί εύκολα είτε από ένα απλό προσεκτικό αλγόριθμο (για δίκτυα σειριακά-παράλληλα) είτε από οποιαδήποτε διαδικασία βέλτιστης απόκρισης (best response) (όταν η καθυστέρηση είναι ίση με τη συμφόρηση). Τέλος δείχνουμε επίσης πώς να υπολογίσουμε αποτελεσματικά ένα $(1 + \epsilon)$ -προσεγγιστικό PNE για το Node-Max-Cut, εάν ο αριθμός των διαφορετικών βαρών των κορυφών του γράφου είναι σταθερός.

1.6 Βιβλιογραφικές πληροφορίες

Τα αποτελέσματα που παρουσιάζονται σε αυτήν τη διατριβή έχουν ήδη εμφανιστεί σε δημοσιεύσεις, με την παρακάτω δομή.

Το Κεφάλαιο 3 βασίζεται στο [114], το οποίο εμφανίστηκε στο WINE 2021. Το Κεφάλαιο 4 βασίζεται στο [7], το οποίο εμφανίστηκε στο SAGT 2020, και η πλήρης έκδοση βρίσκεται υπό (minor) αναθεώρηση για το TOCS 2023. Το Κεφάλαιο 5 βασίζεται στο [11], το οποίο εμφανίστηκε στο JAIR 2022. Μια συντομότερη έκδοση εμφανίστηκε στο ICALP 2020. Το Κεφάλαιο 6 βασίζεται στο [10], το οποίο εμφανίστηκε στο AAAI 2022. Το Κεφάλαιο 7 βασίζεται στο [122], το οποίο εμφανίστηκε στο ICALP 2020.

Chapter 2

Introduction

Algorithmic game theory is a powerful framework that lies in the intersection of computer science, mathematics and micro-economics. It's main focus is to address the challenges arising in strategic interactions among self-interested agents or between them and some authority. As an emerging field of study, it has received much research interest in the past years, motivated also by several real-world key problems. Some examples are:

- *Modeling complex interactions*: In many scenarios, the decisions of agents in multi-agent environments are interdependent, in the sense that each agent's preferred action depends on the choice of actions of others. Algorithmic game theory provides the mathematical tools to model and analyze such interactions, studying the behavior of such strategic agents by analyzing the complexity and quality of related equilibria using notions such as price of anarchy etc...
- *Resource usage* : A very common family of real-world problems is the need to have multiple agents compete over resources. A very well studied related example is *congestion games*, i.e. settings in which agents have to decide which resources to use in order to achieve their goals, keeping in mind that the cost of each resource depends on how many agents use it. A real world application of this type of games is modeling traffic on public road networks.
- *Incentive-compatible mechanism design*: In many strategic settings agents' behaviour is influenced by private information. Agents have the ability to misreport their private information in an attempt to manipulate the mechanism if that would grant them a more desirable outcome. In many applications a guarantee that the used mechanism does not allow such behaviour i.e. that the mechanism is designed in such a way that it elicits truthful information from the agents is desirable or even essential. Such mechanisms are called incentive-compatible and are also a big focus of algorithmic game theory. Furthermore, mechanisms might need to maximize other criteria

such as fairness or profit while taking into account the agents' preferences and constraints. Some applications can be found in auctions, voting systems, facility location, allocation of limited resources among agents etc...

The main focus of this work lies in the mechanism design aspect of algorithmic game theory and social choice, where we also study multi-agent environments from a non strategic point of view. Specifically, we focus on algorithmic frameworks for multi-agent problems in settings with not readily available information, from two perspectives:

- Information is *private*: In which case we need to create *truthful* mechanisms for environments with strategic agents.
- Information is *limited*: Where agents cannot or will not provide their full information, or information exchange is expensive.

The goal is to design algorithms or rules and protocols (mechanisms) that guarantee desirable approximation ratios over the optimal solutions when agents have limited or private information and also incentivize desirable behaviour, when agents hold private information. Furthermore, we are interested in understanding the limitations in designing such systems, i.e. some lower bounds on what the best possible achievable outcome could be, in each domain. On the incentivized version of mechanism design we focus on the paradigm of the facility location problem and introduce the notion of beyond worst case analysis in mechanism design. For the limited information, non-incentivized version, we focus on distortion in voting, a relatively newly introduced notion that measures the effect of having incomplete information in our ability to approximate the optimal social outcome. Furthermore, combining the notions of strategic incentives and limited information we study the communication complexity of well-studied single parameter environments, such as auctions and single-facility location and develop communication-wise asymptotically optimal mechanisms. Finally, dedicating some focus to the classical notion of equilibrium computation, we study the complexity of computing Pure Nash Equilibria, or approximate Pure Nash Equilibria in weighted congestion games.

Mechanism design for strategic domains with private information:

The first type of problems we are going to present belong in the area of studying and designing truthful, or incentive-compatible, mechanisms. The main paradigms studied in this field have the following form: we take the role of an authority that tries to allocate various goods to agents or on some underlying environment the agents belong in. In this setting agents have private valuations over outcomes and they may behave strategically in order to manipulate the mechanism, if it is possible to elicit a better outcome for themselves in that manner. A key property we always want our mechanisms to have here is *strategyproofness*. I.e. to guarantee that no agent can be better off by misreporting their private information to the

mechanism. The two main paradigms we study in this setting are the facility location problems and auctions.

The *facility location problem* ([177]) embodies the trade-off between service accessibility and cost efficiency. In the version we study, a specific number of facilities need to be allocated with the objective of minimizing the social cost, i.e. the aggregate distances of agents to their closest facilities. This version of the problem is offline and facilities don't have a cost for being opened or operated over time. In this form, this is a very well studied ([116, 136, 166, 186, 111, 167]...) problem whose deterministic version has essentially been fully characterized ([112]) and whose randomized version has also been well studied.

Auction settings are also a very big field of study of algorithmic game theory. In general the idea is that the "seller" (mechanism) wants to sell single or multiple goods to agents participating in the auction. Typically agents have a private valuation that represents how much they want each item and they try to maximize their utility by trying to maximize the difference between their valuation and the price they finally pay for acquired items. There are many related settings, such as single item auctions ([220]), various forms of multi-item auctions such as multi-item auctions where multiple instances of the same item are sold, combinatorial auctions where agents want combinations of items, specific valuation domains such as sub-modular, sub-additive valuations etc... Many auction formats have also been studied like sealed bid auctions, where the agents submit their full valuations to the mechanism in sealed bids or ascending auctions where the price of the good ascends and bidders decide whether to keep participating or not.

In this setting of strategic mechanism design we study two main broad questions:

Question 1. Can we introduce beyond worst case analysis concepts in mechanism design in order to create mechanisms with good properties for "real-world" instances for problems for which no such mechanisms can exist in the general case? Can we also use this analysis to understand the nature of strategyproofness in said problems?

For this question we focus on the facility location problem for which very hard limitations have been shown regarding the existence of desired mechanisms. For example, we can have nice, i.e. bounded approximation and truthful mechanisms, only when we try to allocate up to two facilities on the line ([112]). With that in mind the idea of introduced the logic of *beyond worst case analysis* in mechanism design was conceived.

Beyond worst case analysis was first introduced as a mathematical framework to analyze problems that have been shown to be computationally hard in the worst case but seem to be not as hard to solve in most realized examples. This form of analysis defines concepts that capture this notion of "real-world" or "average" instances such as *smoothness* or *perturbation stability* and show that within these

notions the underlying problem becomes tractable ([43, 31]). For the facility location problem specifically, the notion of *perturbation stability* was chosen, due to its close relation to the clustering problem for which this beyond worst case notion has shown great success ([13, 33, 35, 34, 202, 203]). Essentially a clustering instance is perturbation stable if making small perturbations to it (or the metric space it resides on) does not affect the optimal clustering of the instance. In our case, instead of showing that restricting the instance space to perturbation stable instances allows us to create fast algorithms we want to show that we are able to create mechanisms with desired properties such as strategyproofness and good or bounded approximation that cannot exist in the general case.

Question 2. What is the optimal amount of information required in order to implement mechanisms for key problems without sacrificing the desired properties of the mechanism?

Efficient communication has always been a primary desideratum in mechanism design. There are many reasons for this interest such as the fact that extracting data from multiple distributed parties can be expensive, or the fact that communication serves as a proxy for information considerations, or even the realization that direct revelation is undesirable in the sense that it induces a high cognitive cost and bidders may be hesitant to completely disclose their valuation etc ([47, 4, 190, 164, 215])...

We study this concept for two main categories of problems: Auctions and facility location. Specifically we focus on single item auctions and multi-unit auctions with bidders with unit demands from the perspective of trying to minimize the communication between agents and the mechanism. For the facility location problem, we study 1-facility location which can also be considered as a modeling for voting schemes and focus on finding an approximate median, i.e. optimal solution, through sampling.

Non strategic algorithm design in domains with limited information:

In many cases we need to design and study mechanisms in non strategic environments with objectives other than strategyproofness. A prominent example containing such questions is *voting* from the *social choice* domain. In voting people want to elect a desired candidate or a set of candidates. Every agent has a preference over the candidates and we want to create mechanisms that elicit top candidates that satisfy the majority as good as possible. There are many notions from social choice quantifying the quality of a winner such as the notion of a utilitarian winner, that maximizes the social welfare, or a Condorcet winner which is a candidate that wins when pairwise compared with any other candidate etc ([222, 223, 36, 124, 161])...

Even when agents are not strategic and willing to share their information it is

still a challenge to create mechanisms that guarantee such properties when they are applicable. A recent related research interest has focused in utilitarian voting, where there is an underlying numerical utility of each candidate for each voter, in the case where requiring full information from the voters is not possible. The main setting considers that even though each voter has exact utilities (numerical) for each candidate due to the high cognitive cost of realizing these utilities or due to the high value of private information they can only provide a ranking over the candidates that is consistent with their utilities. The notion of distortion then quantifies the impact of utilizing only limited information in the voting process on the social welfare of the outcome ([194]).

Again, we focus on two broad questions in this domain:

Question 3. What is the effect in metric distortion for single winner settings when information is further truncated?

Metric distortion refers to the setting where the utilities of the agents are embedded in a metric space ([18]) (i.e. are their “distances” to the candidates in that metric space). This problem has been studied extensively and a lower bound in distortion has been established ([149]) that has also later been matched by a deterministic mechanism ([131]). Our question further extends the notion that information may be difficult to attain and we ask questions such as: what are the distortion bounds when incomplete top rankings are provided by the agents, how many pairwise comparisons are required to achieve almost optimal distortion and what can we do with sampling.

Question 4. What is the relation between metric distortion and the dimensionality of the underlying metric space? Also, can we use natural learning rules to create mechanisms with low distortion?

In this case we focus on the analysis of a prominent voting mechanism: *STV*. Our goal is to bound the distortion of *STV* by the dimensionality of the underlying instance and also study mechanisms that approach this lower bound. Furthermore we analyse the efficiency of deterministic decentralized dynamics by the agents and their convergence to a candidate with low distortion.

Computation of equilibria:

As mentioned above, one of the most basic functions of algorithmic game theory is studying the interaction among strategic players with self-serving incentives. The main and most well known notion related to this concept is the famous *Nash Equilibrium* ([184, 185]). A state of the game, i.e. a selection of a strategy for each agent, is a Nash Equilibrium if no agent can increase their utility by unilaterally changing their strategy. A strategy for each player can either be pure (determin-

istic), where they deterministically play one of their available strategies, or mixed (randomized) where they play a distribution over their strategies. An equilibrium consisting of pure strategies is called a *Pure Nash Equilibrium* (PNE). In this environment we focus on congestion games which are games in which agents have to use shared resources. Specifically we have finite sets of agents and resources over which agents compete. Agent strategies are subsets of the resources and each agent wants to minimize the total cost of the resources in their chosen strategy. Each resource is associated with a latency function which determines the cost of using that resource as a function of its *congestion*, which is the number of players using that resource in their strategy.

The main question we focus on in this setting is the following:

Question 5. What is the complexity of finding *Pure Nash Equilibria* in weighted congestion games?

Congestion games ([201]) belong in a wider class of games called *potential games* ([178]) that always admit at least one PNE. Thus, researchers have studied extensively the complexity of finding such equilibria in congestion games, with many results for many variations such as symmetric or asymmetric games, where all agents have or don't have the same sets of strategies respectively, network congestion games, where strategies correspond to paths in an underlying network and more. Still, less focus has been devoted to the study of *weighted* potential games, where each agent contributes to the congestion with a different weight ([95, 113, 118, 189, 63, 64, 153]).

The argument that we use to show that every potential game admits PNE is a local search argument, which makes this class of games belong in a complexity class called *Polynomial time Local Search*, or PLS for short. When a problem belongs to the hardest sets of problems within this class, *PLS-complete*, is considered to be computationally hard. We see here that variations of the game that admit fast algorithms in the unweighted case, instantly become hard when introducing weighted players. We also study the complexity of finding approximate equilibria for such cases.

2.1 Strategic mechanism design: Beyond worst case and communication complexity.

2.1.1 Strategyproof Facility Location in Perturbation Stable Instances.

We consider *k-Facility Location games*, where $k \geq 2$ facilities are placed on the real line based on the preferences of n strategic agents. Such problems are motivated by natural scenarios in Social Choice, where a local authority plans to

build a fixed number of public facilities in an area (see e.g., [177]). The choice of the locations is based on the preferences of local people, or *agents*. Each agent reports her ideal location, and the local authority applies a (deterministic or randomized) *mechanism* that maps the agents' preferences to k facility locations.

Each agent evaluates the mechanism's outcome according to her *connection cost*, i.e., the distance of her ideal location to the nearest facility. The agents seek to minimize their connection cost and may misreport their ideal locations in an attempt of manipulating the mechanism. Therefore, the mechanism should be *strategyproof*, i.e., it should ensure that no agent can benefit from misreporting her location, or even *group strategyproof*, i.e., resistant to coalitional manipulations. The local authority's objective is to minimize the *social cost*, namely the sum of agent connections costs. In addition to allocating the facilities in a incentive compatible way, which is formalized by (group) strategyproofness, the mechanism should result in a socially desirable outcome, which is quantified by the mechanism's approximation ratio to the optimal social cost.

Since Procaccia and Tennenholtz [192] initiated the research agenda of *approximate mechanism design without money*, k -Facility Location has served as the benchmark problem in the area and its approximability by deterministic or randomized strategyproof mechanisms has been studied extensively in virtually all possible variants and generalizations. For instance, previous work has considered multiple facilities on the line (see e.g., [112, 116, 136, 166, 186]) and in general metric spaces [111, 167]), different objectives (e.g., social cost, maximum cost, the L_2 norm of agent connection costs [102, 192, 116]), restricted metric spaces more general than the line (cycle, plane, trees, see e.g., [5, 85, 108, 133, 175]), facilities that serve different purposes (see e.g., [156, 163, 212]), and different notions of private information about the agent preferences that should be declared to the mechanism (see e.g., [72, 101, 174] and the references therein).

Due to the significant research interest in the topic, the fundamental and most basic question of approximating the optimal social cost by strategyproof mechanisms for k -Facility Location on the line has been relatively well-understood. For a single facility ($k = 1$), placing the facility at the median location is group strategyproof and optimizes the social cost. For two facilities ($k = 2$), the best possible approximation ratio is $n - 2$ and is achieved by a natural group strategyproof mechanism that places the facilities at the leftmost and the rightmost location [112, 192]. However, for three or more facilities ($k \geq 3$), there do not exist any deterministic anonymous strategyproof mechanisms for k -Facility Location with a bounded (in terms of n and k) approximation ratio [112]. On the positive side, there is a randomized anonymous group strategyproof mechanism with an approximation ratio of n [116] (see also Section 3 for a selective list of additional references).

Perturbation Stability in k -Facility Location Games. Our work aims to circumvent the strong impossibility result of [112] and is motivated by the recent success on the design of polynomial-time exact algorithms for perturbation sta-

ble clustering instances (see e.g., [13, 33, 35, 34, 202, 203]). An instance of a clustering problem, like k -Facility Location (a.k.a. k -median in the optimization and approximation algorithms literature), is γ -*perturbation stable* (or simply, γ -*stable*), for some $\gamma \geq 1$, if the optimal clustering is not affected by scaling down any subset of the entries of the distance matrix by a factor at most γ . Perturbation stability was introduced by Bilu and Linial [43] and Awasthi, Blum and Sheffet [31] (and has motivated a significant volume of followup work since then, see e.g., [13, 33, 34, 203] and the references therein) in an attempt to obtain a theoretical understanding of the superior practical performance of relatively simple clustering algorithms for well known NP-hard clustering problems (such as k -Facility Location in general metric spaces). Intuitively, the optimal clusters of a γ -stable instance are somehow well separated, and thus, relatively easy to identify (see also the main properties of stable instances in Section 3.3). As a result, natural extensions of simple algorithms, like single-linkage (a.k.a. Kruskal’s algorithm), can recover the optimal clustering in polynomial time, provided that $\gamma \geq 2$ [13], and standard approaches, like dynamic programming (resp. local search), work in almost linear time for $\gamma > 2 + \sqrt{3}$ (resp. $\gamma > 5$) [3].

In this work, we investigate whether restricting our attention to stable instances allows for improved strategyproof mechanisms with bounded (and ideally, constant) approximation guarantees for k -Facility Location on the line, with $k \geq 2$. We note that the impossibility results of [112] crucially depend on the fact that the clustering (and the subsequent facility placement) produced by any deterministic mechanism with a bounded approximation ratio must be sensitive to location misreports by certain agents (see also Section 3.6). Hence, it is very natural to investigate whether the restriction to γ -stable instances allows for some nontrivial approximation guarantees by deterministic or randomized strategyproof mechanisms for k -Facility Location on the line.

To study the question above, we adapt to the real line the strict notion of γ -*metric stability* [13], where the definition also requires that the distances form a metric after the γ -perturbation. In our notion of *linear γ -stability*, the instances should retain their linear structure after a γ -perturbation. Hence, a γ -perturbation of a linear k -Facility Location instance is obtained by moving any subset of pairs of consecutive agent locations closer to each other by a factor at most $\gamma \geq 1$. We say that a k -Facility Location instance is γ -stable, if the original instance and any γ -perturbation of it admit the same unique optimal clustering

Interestingly, for γ sufficiently large, γ -stable instances of k -Facility Location have additional structure that one could exploit towards the design of strategyproof mechanisms with good approximation guarantees (see also Section 3.3). E.g., each agent location is $\gamma - 1$ times closer to the nearest facility than to any location in a different cluster (Proposition 3.3.2). Moreover, for $\gamma \geq 2 + \sqrt{3}$, the distance between any two consecutive clusters is larger than their diameter (Lemma 3.3.4).

From a conceptual viewpoint, our work is motivated by a reasoning very similar to that discussed by Bilu, Daniely, Linial and Saks [44] and summarized in

“*clustering is hard only when it doesn’t matter*” by Roughgarden [205]. In a nutshell, we expect that when k public facilities (such as schools, libraries, hospitals, representatives) are to be allocated to some communities (e.g., cities, villages or neighborhoods, as represented by the locations of agents on the real line) the communities are already well formed, relatively easy to identify and difficult to radically reshape by small distance perturbations or agent location misreports. Moreover, in natural practical applications of k -Facility Location games, agents tend to misreport “locally” (i.e., they tend to declare a different ideal location in their neighborhood, trying to manipulate the location of the local facility), which usually does not affect the cluster formation. In practice, this happens because the agents do not have enough knowledge about locations in other neighborhoods, and because “large non-local” misreports are usually easy to identify by combining publicly available information about the agents (e.g., occupation, address, habits, lifestyle). Hence, we believe that the class of γ -stable instances, especially for relatively small values of γ , provides a reasonably accurate abstraction of the instances of k -Facility Location games that a mechanism designer is more likely to deal with in practice. Thus, we feel that our work takes a small first step towards justifying that (not only clustering but also) strategyproof facility location is hard only when it doesn’t matter.

In this line of work, we show that the optimal solution is strategyproof in $(2 + \sqrt{3})$ -stable instances whose optimal solution does not include any singleton clusters, and that allocating the facility to the agent next to the rightmost one in each optimal cluster (or to the unique agent, for singleton clusters) is strategyproof and $(n - 2)/2$ -approximate for 5-stable instances (even if their optimal solution includes singleton clusters). On the negative side, we show that for any $k \geq 3$ and any $\delta > 0$, there is no deterministic anonymous mechanism that achieves a bounded approximation ratio and is strategyproof in $(\sqrt{2} - \delta)$ -stable instances. We also prove that allocating the facility to a random agent of each optimal cluster is strategyproof and 2-approximate in 5-stable instances. To the best of our knowledge, this is the first time that the existence of deterministic (resp. randomized) strategyproof mechanisms with a bounded (resp. constant) approximation ratio is shown for a large and natural class of k -Facility Location instances.

2.1.2 Sampling and Optimal Preference Elicitation in Simple Mechanisms.

Efficient *preference elicitation* has been a central theme and a major challenge from the inception of *mechanism design*, with a myriad of applications in multi-agent environments and modern artificial intelligence systems. Indeed, requesting from every agent to communicate all of her preferences is considered widely impractical, and a substantial body of work has explored alternative approaches to truncate the elicited information, sequentially asking a series of natural queries in order to elicit only the relevant parts of the information.

This emphasis has been strongly motivated for a number of reasons. First, it has been acknowledged by behavioral economists that soliciting information induces a high *cognitive cost* [190, 164], and agents may even be reluctant to reveal their complete (private) valuation; as pertinent evidence, the superiority of *indirect* mechanisms is often cited [29]. In fact, in certain domains with severe communication constraints [4, 180] a *direct revelation* mechanism—one in which every agent has to disclose her entire preferences—is considered even infeasible. Indeed, communication is typically recognized as the main bottleneck in distributed environments [170]. As explained by Blumrosen and Feldman [46], agents typically operate with a truncated action space due to technical, behavioral or regulatory reasons. Finally, a mechanism with efficient preference elicitation would provide stronger information-privacy guarantees [216].

The general question of whether a social choice function can be accurately approximated by less than the full set of agents constitutes one of the main themes in computational social choice. Perhaps the most standard approach to truncate the elicited information consists of *sampling*. More precisely, given that in many real-world applications it might be infeasible to gather preferences from all the agents, the designer performs preference aggregation by randomly selecting a small subset of the entire population [58, 83]. This approach is particularly familiar in the context of *voting*, where the goal is typically to *predict* the outcome of the full information mechanism without actually holding the election for the entire population [82]; concrete examples of the aforementioned scenarios include election polls, exit polls, as well as online surveys.

In the first part of our work, we follow this long line of research in computational social choice. Specifically, we analyze the *sample complexity* of the celebrated *median* mechanism in the context of *facility location games*, where every agent is associated with a point—corresponding to her preferred location—on some underlying metric space. The median mechanism is of particular importance in social choice. Indeed, the celebrated *Gibbard-Satterthwaite* impossibility theorem [130, 207] states that for any *onto*—for every alternative there exists a voting profile that would make that alternative prevail—and non-dictatorial voting rule, there are instances for which an agent is better off casting a vote that does not correspond to her true preferences—i.e. the rule is not *strategy-proof*. Importantly, this impediment disappears when the agents’ preferences are restricted. Arguably the most well-known such restriction is that of the *single-peaked preferences*,¹ introduced by Black [45], for which Moulin’s [181] median mechanism is

¹More precisely, suppose that the alternatives are ordered on a line, representing their *positions*; as argued in *spatial voting theory* [90, 24], it is often reasonable to assume that the alternatives can be represented as points on a line (e.g., in political elections a candidate’s position may indicate whether she is a “left-wing” or a “right-wing” candidate). An agent’s preferences are *single-peaked* if she prefers alternatives which are closer to her peak. We remark that in single-peaked domains it is known that there can be no *Condorcet cycles*.

indeed strategy-proof.

Our key result here is to show that for any $\epsilon > 0$, a sample of size $c(\epsilon) = \Theta(1/\epsilon^2)$ yields in expectation a $1 + \epsilon$ approximation with respect to the optimal social cost of the generalized median mechanism on the metric space $(\mathbb{R}^d, \|\cdot\|_1)$, while the number of agents $n \rightarrow \infty$.

In the second part of our work, we endeavor to design auctions with minimal *communication complexity*. As a concrete motivating example, we consider a *single-item* auction, and we assume that every valuation can be expressed with k bits. An important observation is that the most dominant formats are very inefficient from a communication standpoint. Indeed, *Vickrey's*—or *sealed-bid*—auction [221] is a direct revelation mechanism, eliciting the entire private information from the n agents, leading to a communication complexity of $n \cdot k$ bits; in fact, it should be noted that although Vickrey's auction possesses many theoretically appealing properties, its ascending counterpart exhibits superior performance in practice [28, 29, 145, 144], for reasons that mostly relate to the simplicity, the transparency, as well as the privacy guarantees of the latter format. Unfortunately, implementing Vickrey's rule through an *English* auction requires—in the worst case—exponential communication of the order $n \cdot 2^k$, as the auctioneer has to cover the entire valuation space. Thus, the elicitation pattern in an English auction is widely inefficient, and the lack of prior knowledge on the agents' valuations would dramatically undermine its performance.

In this direction, we study a series of exemplar environments from auction theory through a communication complexity framework, measuring the expected number of bits elicited from the agents; we posit that any valuation can be expressed with k bits, and we mainly assume that k is independent of the number of agents n . In this context, we show that Vickrey's rule can be implemented with an expected communication of $1 + \epsilon$ bits from an average bidder, for any $\epsilon > 0$, asymptotically matching the trivial lower bound. As a corollary, we provide a compelling method to increment the price in an English auction. We also leverage our single-item format with an efficient encoding scheme to prove that the same communication bound can be recovered in the domain of additive valuations through simultaneous ascending auctions, assuming that the number of items is a constant. Finally, we propose an ascending-type multi-unit auction under unit demand bidders; our mechanism announces at every round two separate prices and is based on a sampling algorithm that performs approximate selection with limited communication, leading again to asymptotically optimal communication. Our results do not require any prior knowledge on the agents' valuations, and mainly follow from natural sampling techniques.

2.2 Non strategic mechanism design: Domains with limited information.

2.2.1 Metric-Distortion Bounds under Limited Information.

Aggregating the preferences of individuals into a collective decision lies at the heart of social choice, and has recently found numerous applications in areas such as information retrieval, recommender systems, and machine learning [222, 223, 36, 124, 161]. The classic theory of von Neumann and Morgenstern [224] postulates that individual preferences are captured through a *utility function*, assigning numerical (or *cardinal*) values to each alternative. Yet, in voting theory, as well as in most practical applications, mechanisms typically elicit only *ordinal* information from the voters, indicating an order of preferences over the candidates. Although this might seem at odds with a utilitarian representation, it has been recognized that it might be hard for a voter to specify a precise numerical value for an alternative, and providing only ordinal information substantially reduces the cognitive burden. This begs the question: What is the loss in efficiency of a mechanism extracting only ordinal information with respect to the *utilitarian social welfare*, i.e., the sum of individual utilities over a chosen candidate? The framework of *distortion*, introduced by Procaccia and Rosenschein [194], measures exactly this loss from an approximation-algorithms standpoint, and has received considerable attention in recent years.

As it turns out, the approximation guarantees we can hope for crucially depend on the assumptions we make on the utility functions. For example, in the absence of any structure Procaccia and Rosenschein [194] observed that *no* ordinal deterministic mechanism can obtain bounded distortion. In this work, we focus on the *metric distortion* framework, introduced by Anshelevich et al. [18], wherein voters and candidates are thought of as points in some arbitrary metric space; this is akin to models in spatial voting theory [74]. In this context, the voters' preferences are measured by means of their "proximity" from each candidate, and the goal is to output a candidate who (approximately) minimizes the *social cost*, i.e., the total distance to the voters. A rather simplistic view of this framework manifests itself when agents and candidates are embedded into a one-dimensional line, and their locations indicate whether they are "left" or "right" on the political spectrum. However, the metric distortion framework has a far greater reach since no assumptions whatsoever are made for the underlying metric space.

Importantly, this paradigm offers a compelling way to quantitatively compare different voting rules commonly employed in practice [214, 149, 132, 18], while it also serves as a benchmark for designing new mechanisms in search of better distortion bounds [131, 183]. A common assumption made in this line of work is that the algorithm has access to the *entire* total orders of the voters. However, there are many practical scenarios in which it might be desirable to truncate the

ordinal information elicited by the mechanism. For example, requesting only the top preferences could further relieve the cognitive burden since it might be hard for a voter to compare alternatives which lie on the bottom of her preferences’ list (for additional motivation for considering incomplete or partial orderings we refer to the works of Fotakis et al. [123], Chen et al. [70], Benferhat et al. [39], and references therein), while any truncation in the elicited information would also translate to more efficient communication. These reasons have driven several authors to study the decay of distortion under missing information [150, 14, 138, 98, 50], potentially allowing some randomization (see our related work subsection). In this work, we follow that line of research, offering several new insights and improved bounds over prior results.

In this direction, our primary contribution is threefold. First, we consider mechanisms that perform a sequence of pairwise comparisons between candidates. We show that a popular deterministic mechanism employed in many knockout phases yields distortion $\mathcal{O}(\log m)$ while eliciting only $m - 1$ out of the $\Theta(m^2)$ possible pairwise comparisons, where m represents the number of candidates. Our analysis for this mechanism leverages a powerful technical lemma developed by Kempe (AAAI ‘20). We also provide a matching lower bound on its distortion. In contrast, we prove that any mechanism which performs fewer than $m - 1$ pairwise comparisons is destined to have unbounded distortion. Moreover, we study the power of deterministic mechanisms under incomplete rankings. Most notably, when agents provide their k -top preferences we show an upper bound of $6m/k + 1$ on the distortion, for any $k \in \{1, 2, \dots, m\}$. Thus, we substantially improve over the previous bound of $12m/k$ established by Kempe (AAAI ‘20), and we come closer to matching the best-known lower bound. Finally, we are concerned with the sample complexity required to ensure near-optimal distortion with high probability. Our main contribution is to show that a random sample of $\Theta(m/\epsilon^2)$ voters suffices to guarantee distortion $3 + \epsilon$ with high probability, for any sufficiently small $\epsilon > 0$. This result is based on analyzing the sensitivity of the deterministic mechanism introduced by Gkatzelis, Halpern, and Shah (FOCS ‘20). Importantly, all of our sample-complexity bounds are distribution-independent.

2.2.2 Dimensionality and Coordination in Voting: The Distortion of STV

In this work we continue to explore the notion of *metric-distortion*. Importantly, this framework of distortion offers a quantitative “benchmark” for comparing different voting rules commonly employed in practice. Indeed, one of the primary considerations of our work lies in characterizing the performance of the *single transferable vote*² mechanism (henceforth STV).

²For consistency with prior work STV will represent throughout this paper the *single-winner* variant of the system, which is sometimes referred to as *instant-runoff voting* (IRV)

STV is a widely-popular iterative voting system employed in the national elections of several countries, including Australia, Ireland, and India, as well as in many other preference aggregation tasks; e.g., in the Academy Awards. To be more precise, STV proceeds in an iterative fashion: In each round, agents vote for their most preferred candidate—among the *active* ones, while the candidate who enjoyed the least amount of support in the current round gets eliminated. This process is repeated for $m - 1$ rounds, where m represents the number of (initial) alternatives, and the last surviving candidate is declared the winner of STV. As an aside, notice that this process is generally non-deterministic due to the need for a tie-breaking mechanism; as in [214], we will work with the *parallel universe model* of Conitzer et al. [79], wherein a candidate is said to be an STV winner if it survives under *some* sequence of eliminations.

In this context, Skowron and Elkind [214] were the first to analyze the distortion of STV under metric preferences. Specifically, they showed that the distortion of STV in general metric spaces is always $O(\log m)$, while they also gave a nearly-matching lower bound in the form of $\Omega(\sqrt{\log m})$. Interestingly, a careful examination of their lower bound reveals the existence of a high-dimensional submetric, as depicted in Figure 2.1, and it is a well-known fact in the theory of metric embeddings that such objects cannot be *isometrically* embedded into low-dimensional³ Euclidean spaces [173]. As a result, Skowron and Elkind [214] left open the following intriguing question:

Question 1. *What is the distortion of STV under low-dimensional Euclidean spaces?*




Figure 2.1: A high-dimensional metric in the form of a “star” graph.

in the literature.

³We say that a Euclidean space is low-dimensional if its dimension d is bounded by a “small” universal constant, i.e. $d = O(1)$.

Needless to say that the performance of voting rules in low-dimensional spaces has been a subject of intense scrutiny in spatial voting theory, under the premise that voters and candidates are typically embedded in subspaces with small dimension [24, 90]. For example, recent experimental work by Elkind et al. [88] evaluates several voting rules in a 2-dimensional Euclidean space, motivated by the fact that preferences are typically crystallized on the basis of a few crucial dimensions; e.g., economic policy and healthcare. Indeed, in the so-called Nolan Chart—a celebrated political spectrum diagram—political views are charted along two axes, expanding upon the traditional one-dimensional representation; to quote from the work of Elkind et al. [88]:

“...the popularity of the Nolan Chart [...] indicates that two dimensions are often sufficient to provide a good approximation of voters’ preferences.”

Thus, it is natural to ask whether we can refine the analysis of STV under low-dimensional spaces. In fact, as part of a broader agenda analogous questions can be raised for other mechanisms as well. However, it is interesting to point out that for many voting rules analyzed within the framework of distortion there exist low-dimensional lower bounds; some notable examples are given in Table 2.1. In contrast, our work will separate STV from the mechanisms in Table 2.1, effectively addressing Question 1. Importantly, we shall provide a characterization well-beyond Euclidean spaces, to metrics with “intrinsically” low dimension.

Mechanism	Lower Bound	Dimension
Plurality	$2m - 1$	1
Borda	$2m - 1$	1
Copeland	5	2
Veto	$2n - 1$	1
Approval	$2n - 1$	1

Table 2.1: The Euclidean dimension required to construct a (tight) lower bound for several common voting rules; these results appear in [18]. We should note that for Copeland the metric constructed in [18] is not Euclidean, but can be easily modified to be one.

The next consideration of our work is directly motivated by the efficiency of STV compared to the plurality rule, and in particular the strategic implications of this discrepancy. A good starting point for this discussion stems from the fact that in many fundamental preference aggregation settings alternatives are chosen by inefficient mechanisms, and in many cases any reform faces insurmountable impediments. For example, in political elections the voting mechanism is typically

dictated by electoral laws, or even the constitution [165]. As a result, understanding the behavior of strategic agents when faced with inefficient mechanisms is of paramount importance [56, 225]. A rather orthogonal way of viewing this is whether autonomous agents can converge to admissible social choices through natural learning rules; this begs the question:

Question 2. *To what extent can strategic behavior improve efficiency in voting?*

We stress that although in the absence of any information it might be unclear how agents can engage in strategic behavior, in most applications of interest agents have plenty of prior information before they cast their votes, e.g. through polls, surveys, forecasts, prior elections, or even early voting. Indeed, there is a prolific line of work which studies population dynamics for agents that cast their votes in response to the information they possess (see [200], and references therein), as well as the role of information in shaping public policy [158].

To address such considerations we propose a natural model wherein agents act iteratively based on some partial feedback on the other voters' preferences. We explain how STV can be very naturally cast in this framework, while we establish the existence of simple and decentralized *coordination* dynamics converging to a near-optimal alternative.

The final theme of our work offers certain refinements and extensions of prior works, mostly driven by some fundamental applications in the context of facility location games. Specifically, we primarily focus on the optimal—under metric preferences—deterministic mechanism recently introduced by Gkatzelis, Halpern, and Shah [131]; we show that it recovers the optimal bound under ultra-metrics, and near-optimal distortion under distances satisfying approximate triangle inequalities.

2.3 Computation of equilibria.

2.4 Node Max-Cut and Computing Equilibria in Linear Weighted Congestion Games

Motivated by the remarkable success of local search in combinatorial optimization, Johnson et al. introduced [143] the complexity class Polynomial Local Search (PLS), consisting of local search problems with polynomially verifiable local optimality. PLS includes many natural complete problems (see e.g., [176]), with CIRCUIT-FLIP [143] and MAX-CUT [208] among the best known ones, and lays the foundation for a principled study of the complexity of local optima computation. In the last 15 years, a significant volume of research on PLS-completeness was motivated by the problem of computing a pure Nash equilibrium of potential games (see e.g., [2, 213, 127] and the references therein), where any improving deviation

by a single player decreases a potential function and its local optima correspond to pure Nash equilibria [178].

Computing a local optimum of MAX-CUT under the FLIP neighborhood (a.k.a. LOCAL-MAX-CUT) has been one of the most widely studied problems in PLS. Given an edge-weighted graph, a cut is locally optimal if we cannot increase its weight by moving a vertex from one side of the cut to the other. Since its PLS-completeness proof by Schäffer and Yannakakis [208], researchers have shown that LOCAL-MAX-CUT remains PLS-complete for graphs with maximum degree five [89], is polynomially solvable for cubic graphs [191], and its smoothed complexity is either polynomial in complete [12] and sparse [89] graphs, or almost polynomial in general graphs [71, 92]. Moreover, due to its simplicity and versatility, MAX-CUT has been widely used in PLS reductions (see e.g., [2, 127, 213]). LOCAL-MAX-CUT can also be cast as a game, where each vertex aims to maximize the total weight of its incident edges that cross the cut. Cut games are potential games (the value of the cut is the potential function), which has motivated research on efficient computation of approximate equilibria for LOCAL-MAX-CUT [41, 64]. To the best of our knowledge, apart from the work on the smoothed complexity of LOCAL-MAX-CUT (and may be that LOCAL-MAX-CUT is P-complete for unweighted graphs [208, Theorem 4.5]), there has not been any research on whether (and to which extent) additional structure on edge weights affects hardness of LOCAL-MAX-CUT.

A closely related research direction deals with the complexity of computing a pure Nash equilibrium (equilibrium or PNE, for brevity) of *congestion games* [201], a typical example of potential games [178] and among the most widely studied classes of games in Algorithmic Game Theory (see e.g., [113] for a brief account of previous work). In congestion games (or CGs, for brevity), a finite set of players compete over a finite set of resources. Strategies are resource subsets and players aim to minimize the total cost of the resources in their strategies. Each resource e is associated with a (non-negative and non-decreasing) latency function, which determines the cost of using e as a function of e 's *congestion* (i.e., the number of players including e in their strategy). Researchers have extensively studied the properties of special cases and variants of CGs. Most relevant to this work are *symmetric* (resp. *asymmetric*) CGs, where players share the same strategy set (resp. have different strategy sets), *network* CGs, where strategies correspond to paths in an underlying network, and *weighted* CGs, where player contribute to the congestion with a different weight.

Fabrikant et al. [95] proved that computing a PNE of asymmetric network CGs or symmetric CGs is PLS-complete, and that it reduces to min-cost-flow for symmetric network CGs. About the same time [118, 189] proved that weighted congestion games admit a (weighted) potential function, and thus a PNE, if the latency functions are either affine or exponential (and [140, 141] proved that in a certain sense, this restriction is necessary). Subsequently, Ackermann et al. [2] characterized the strategy sets of CGs that guarantee efficient equilibrium compu-

tation. They also used a variant of LOCAL-MAX-CUT, called *threshold games*, to simplify the PLS-completeness proof of [95] and to show that computing a PNE of asymmetric network CGs with (exponentially steep) linear latencies is PLS-complete.

On the other hand, the complexity of equilibrium computation for weighted CGs is not well understood. All the hardness results above carry over to weighted CGs, since they generalize standard CGs (where the players have unit weight). But on the positive side, we only know how to efficiently compute a PNE for weighted CGs on parallel links with general latencies [120] and for weighted CGs on parallel links with identity latency functions and asymmetric strategies [126]. Despite the significant interest in (exact or approximate) equilibrium computation for CGs (see e.g., [63, 64, 153] and the references therein), we do not understand how (and to which extent) the complexity of equilibrium computation is affected by player weights. This is especially true for weighted CGs with linear latencies, which admit a potential function and their equilibrium computation is in PLS.

In this work, we seek a more refined understanding of the complexity of local optimum computation for Max-Cut and pure Nash equilibrium (PNE) computation for congestion games with weighted players and linear latency functions. We show that computing a PNE of linear weighted congestion games is PLS-complete for very restricted strategy spaces, namely when player strategies are paths on a series-parallel network with a single origin and destination. Furthermore, in the same work we also show that the problem remains PLS-complete even for very restricted latency functions, namely when the latency on each resource is equal to the congestion, for general networks, through a reduction from a newly defined appropriate local search problem we call *Node-Max-Cut*. In this light we also show how to compute efficiently a $(1 + \varepsilon)$ -approximate equilibrium for Node-Max-Cut, if the number of different vertex weights is constant.

2.5 Bibliographic information

The results presented in this thesis have already appeared in publications, following the structure below.

Chapter 3 is based on [114] which appeared in WINE 2021. Chapter 4 is based on [7] which appeared in SAGT 2020 and full version is under minor revision for TOCS 2023. Chapter 5 is based on [11] which appeared in JAIR 2022. A shorter version has appeared in ICALP 2020. Chapter 6 is based on [10] which appeared in AAAI 2022. Chapter 7 is based on [122] which appeared in ICALP 2020.

Chapter 3

Facility location in perturbation stable instances

Approximate mechanism design without money for variants and generalizations of Facility Location games on the line has been a very active and productive area of research in the last decade.

Previous work has shown that deterministic strategyproof mechanisms can only achieve a bounded approximation ratio for k -Facility Location on the line, only if we have at most 2 facilities [112, 192]. Notably, stable (called *well-separated* in [112]) instances with $n = k + 1$ agents play a key role in the proof of inapproximability of k -Facility Location by deterministic anonymous strategyproof mechanisms [112, Theorem 3.7]. On the other hand, randomized mechanisms are known to achieve a better approximation ratio for $k = 2$ facilities [166], a constant approximation ratio if we have $k \geq 2$ facilities and only $n = k + 1$ agents [91, 116], and an approximation ratio of n for any $k \geq 3$ [116]. Fotakis and Tzamos [111] considered winner-imposing randomized mechanisms that achieve an approximation ratio of $4k$ for k -Facility Location in general metric spaces. In fact, the approximation ratio can be improved to $\Theta(\ln k)$, using the analysis of [25].

For the objective of maximum agent cost, Alon et al. [5] almost completely characterized the approximation ratios achievable by randomized and deterministic strategyproof mechanisms for 1-Facility Location in general metrics and rings. Fotakis and Tzamos [116] presented a 2-approximate randomized group strategyproof mechanism for k -Facility Location on the line and the maximum cost objective. For 1-Facility Location on the line and the objective of minimizing the sum of squares of the agent connection costs, Feldman and Wilf [102] proved that the best approximation ratio is 1.5 for randomized and 2 for deterministic mechanisms. Golomb and Tzamos [136] presented tight (resp. almost tight) additive approximation guarantees for locating a single (resp. multiple) facilities on the line and the objectives of the maximum cost and the social cost.

Regarding the application of perturbation stability, we follow the approach of

beyond worst-case analysis (see e.g., [202, 203]), where researchers seek a theoretical understanding of the superior practical performance of certain algorithms by formally analyzing them on practically relevant instances. The beyond worst-case approach is not anything new for Algorithmic Mechanism Design. *Bayesian* analysis, where the bidder valuations are drawn as independent samples from a distribution known to the mechanism, is standard in revenue maximization when we allocate private goods (see e.g., [204]) and has led to many strong and elegant results for social welfare maximization in combinatorial auctions by truthful posted price mechanisms (see e.g., [86, 103]). However, in this work, instead of assuming (similar to Bayesian analysis) that the mechanism designer has a relatively accurate knowledge of the distribution of agent locations on the line (and use e.g., an appropriately optimized percentile mechanism [217]), we employ a deterministic restriction on the class of instances (namely, perturbation stability), and investigate if deterministic (resp. randomized) strategyproof mechanisms with a bounded (resp. constant) approximation ratio are possible for locating any number $k \geq 2$ facilities on such instances. To the best of our knowledge, the only previous work where the notion of perturbation stability is applied to Algorithmic Mechanism Design (to combinatorial auctions, in particular) is [107] (but see also [32, 94] where the similar in spirit assumption of endowed valuations was applied to combinatorial markets).

3.1 Our contribution

Our conceptual contribution is that we initiate the study of efficient (wrt. their approximation ratio for the social cost) strategyproof mechanisms for the large and natural class of γ -stable instances of k -Facility Location on the line. Our technical contribution is that we show the existence of deterministic (resp. randomized) strategyproof mechanisms with a bounded (resp. constant) approximation ratio for 5-stable instances and any number of facilities $k \geq 2$. Moreover, we show that the optimal solution is strategyproof for $(2 + \sqrt{3})$ -stable instances, if the optimal clustering does not include any singleton clusters (which is likely to be the case in virtually all practical applications). To provide evidence that restriction to stable instances does not make the problem trivial, we strengthen the impossibility result of Fotakis and Tzamos [112], so that it applies to γ -stable instances, with $\gamma < \sqrt{2}$. Specifically, we show that for any $k \geq 3$ and any $\delta > 0$, there do not exist any deterministic anonymous strategyproof mechanisms for k -Facility Location on $(\sqrt{2} - \delta)$ -stable instances with bounded (in terms of n and k) approximation ratio.

At the conceptual level, we interpret the stability assumption as a prior on the class of true instances that the mechanism should be able to deal with. Namely, we assume that the mechanism has only to deal with γ -stable true instances, a restriction motivated by (and fully consistent with) how the stability assumption is used in the literature on efficient algorithms for stable clustering (see e.g., [13, 33, 34, 43], where the algorithms are analyzed for stable instances only). More

specifically, our mechanisms expect as input a declared instance such that in the optimal clustering, the distance between any two consecutive clusters is at least $\frac{(\gamma-1)^2}{2\gamma}$ times larger than the diameters of the two clusters (a.k.a. *cluster-separation* property, see Lemma 3.3.5). This condition is necessary (but not sufficient) for γ -stability and can be easily checked. If the declared instance does not satisfy the cluster-separation property, our mechanisms do not allocate any facilities. Otherwise, our mechanisms allocate k facilities (even if the instance is not stable). We prove that for all γ -stable true instances (with the exact stability factor γ depending on the mechanism), if agents can only deviate so that the declared instance satisfies the cluster-separation property (and does not have singleton clusters, for the optimal mechanism), our mechanisms are strategyproof and achieve the desired approximation guarantee. Hence, if we restrict ourselves to γ -stable true instances and to agent deviations that do not obviously violate γ -stability, our mechanisms should only deal with γ -stable declared instances, due to strategyproofness. On the other hand, if non-stable true instances may occur, the mechanisms cannot distinguish between a stable true instance and a declared instance, which appears to be stable, but is obtained from a non-stable instance through location misreports.

The restriction that the agents of a γ -stable instance are only allowed to deviate so that the declared instance satisfies the cluster-separation property (and does not have any singleton clusters, for the optimal mechanism) bears a strong conceptual resemblance to the notion of strategyproof mechanisms with *local verification* (see e.g., [27, 23, 61, 68, 121, 117, 137]), where the set of each agent's allowable deviations is restricted to a so-called *correspondence set*, which typically depends on the agent's true type, but not on the types of the other agents. Instead of restricting the correspondence set of each individual agent independently, we impose a structural condition on the entire declared instance, which restricts the set of the agents' allowable deviations, but in a global and observable sense. As a result, we can actually implement our notion of verification, by checking some simple properties of the declared instance, instead of just assuming that any deviation outside an agent's correspondence set will be caught and penalized (which is the standard approach in mechanisms with local verification [23, 68, 61], but see e.g., [27, 111] for noticeable exceptions).

On the technical side, we start, in Section 3.3, with some useful properties of stable instances of k -Facility Location on the line. Among others, we show (i) the *cluster-separation* property (Lemma 3.3.4), which states that in any γ -stable instance, the distance between any two consecutive clusters is at least $\frac{(\gamma-1)^2}{2\gamma}$ times larger than their diameters; and (ii) the so-called *no direct improvement from singleton deviations* property (Lemma 3.3.7), i.e., that in any 3-stable instance, no agent who deviates to a location, which becomes a singleton cluster in the optimal clustering of the resulting instance, can improve her connection cost through the facility of that singleton cluster.

In Section 3.4, we show that for $(2 + \sqrt{3})$ -stable instances whose optimal clustering does not include any singleton clusters, the optimal solution is strategyproof

(Theorem 3.4.1). For the analysis, we observe that since placing the facility at the median location of any fixed cluster is strategyproof, a misreport cannot be profitable for an agent, unless it results in a different optimal clustering. The key step is to show that for $(2 + \sqrt{3})$ -stable instances without singleton clusters, a profitable misreport cannot change the optimal clustering, unless the instance obtained from the misreport violates the cluster-separation property. To the best of our knowledge, the idea of penalizing (and thus, essentially forbidding) a whole class of potentially profitable misreports by identifying how they affect a key structural property of the original instance, which becomes possible due to our restriction to stable instances, has not been used before in the design of strategyproof mechanisms for k -Facility Location (see also the discussion above about resemblance to mechanisms with verification).

We should also motivate our restriction to stable instances without singleton clusters in their optimal clustering. So, let us consider the rightmost agent x_j of an optimal cluster C_i in a γ -stable instance \vec{x} . No matter the stability factor γ , it is possible that x_j performs a so-called *singleton deviation*. Namely, x_j deviates to a remote location x' (potentially very far away from any location in \vec{x}), which becomes a singleton cluster in the optimal clustering of the resulting instance. Such a singleton deviation might cause cluster C_i to merge with (possibly part of the next) cluster C_{i+1} , which in turn, might bring the median of the new cluster much closer to x_j (see also Fig. 3.1 in Section 3.3). It is not hard to see that if we stick to the optimal solution, where the facilities are located at the median of each optimal cluster, there are γ -stable instances¹, with arbitrarily large $\gamma \geq 1$, where some agents can deviate to a remote location and gain, by becoming singleton clusters, while maintaining the desirable stability factor of the declared instance (see also Fig. 3.1).

To deal with singleton deviations², we should place the facility either at a location close to an extreme one, as we do in Section 3.5 with the ALMOSTRIGHTMOST mechanism, or at a random location, as we do in Section 3.7 with the RANDOM mechanism. More specifically, in Section 3.5, we show that the ALMOSTRIGHTMOST mechanism, which places the facility of any non-singleton optimal cluster at the location of the second rightmost agent, is strategyproof for 5-stable instances of k -Facility Location (even if their optimal clustering includes singleton clusters) and achieves an approximation ratio at most $(n - 2)/2$ (Theorem 3.5.1). Moreover, in Section 3.7, we show that the RANDOM mechanism, which places the facility

¹E.g., let $k = 2$ and consider the $\Theta(\gamma)$ -stable instance $(0, 1 - \varepsilon, 1, 6\gamma, 6\gamma + \varepsilon, 6\gamma + 1, 6\gamma + 1 + \varepsilon, 6\gamma + 2)$, for any $\gamma \geq 1$. Then, the agent at location 6γ can decrease its connection cost (from 1) to ε by deviating to location $(6\gamma)^2$.

²Another natural way to deal efficiently with singleton deviations is through some means of *location verification*, such as winner-imposing verification [111] or ε -symmetric verification [117, 121]. Adding e.g., winner-imposing verification to the optimal mechanism, discussed in Section 3.4, results in a strategyproof mechanism for $(2 + \sqrt{3})$ -stable instances whose optimal clustering may include singleton clusters.

of any optimal cluster at a location chosen uniformly at random, is strategyproof for 5-stable instances (again even if their optimal clustering includes singleton clusters) and achieves an approximation ratio of 2 (Theorem 3.7.1).

To obtain a deeper understanding of the challenges behind the design of strategyproof mechanisms for stable instances of k -Facility Location on the line, we strengthen the impossibility result of [112, Theorem 3.7] so that it applies to γ -stable instances with $\gamma < \sqrt{2}$ (Section 3.6). Through a careful analysis of the image sets of deterministic strategyproof mechanisms, we show that for any $k \geq 3$, any $\delta > 0$, and any $\rho \geq 1$, there do not exist any ρ -approximate deterministic anonymous strategyproof mechanisms for $(\sqrt{2} - \delta)$ -stable instances of k -Facility Location on the line (Theorem 3.6.5). The proof of Theorem 3.6.5 requires additional ideas and extreme care (and some novelty) in the agent deviations, so as to only consider stable instances, compared against the proof of [112, Theorem 3.7]. Interestingly, singleton deviations play a crucial role in the proof of Theorem 3.6.5.

3.2 Notation, Definitions and Preliminaries

We let $[n] = \{1, \dots, n\}$. For any $x, y \in \mathbb{R}$, we let $d(x, y) = |x - y|$ be the distance of locations x and y on the real line. For a tuple $\vec{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$, we let \vec{x}_{-i} denote the tuple \vec{x} without coordinate x_i . For a non-empty set S of indices, we let $\vec{x}_S = (x_i)_{i \in S}$ and $\vec{x}_{-S} = (x_i)_{i \notin S}$. We write (\vec{x}_{-i}, a) to denote the tuple \vec{x} with a in place of x_i , $(\vec{x}_{-\{i,j\}}, a, b)$ to denote the tuple \vec{x} with a in place of x_i and b in place of x_j , and so on. For a random variable X , $\mathbb{E}(X)$ denotes the expectation of X . For an event E in a sample space, $\Pr(E)$ denotes the probability that E occurs.

Instances. We consider k -Facility Location with $k \geq 2$ facilities and $n \geq k + 1$ agents on the real line. We let $N = \{1, \dots, n\}$ be the set of agents. Each agent $i \in N$ resides at a location $x_i \in \mathbb{R}$, which is i 's private information. We usually refer to a locations profile $\vec{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$, $x_1 \leq \dots \leq x_n$, as an *instance*. By slightly abusing the notation, we use x_i to refer both to the agent i 's location and sometimes to the agent i (i.e., the strategic entity) herself.

Mechanisms. A *deterministic mechanism* M for k -Facility Location maps an instance \vec{x} to a k -tuple $(c_1, \dots, c_k) \in \mathbb{R}^k$, $c_1 \leq \dots \leq c_k$, of facility locations. We let $M(\vec{x})$ denote the outcome of M in instance \vec{x} , and let $M_j(\vec{x})$ denote c_j , i.e., the j -th smallest coordinate in $M(\vec{x})$. We write $c \in M(\vec{x})$ to denote that $M(\vec{x})$ places a facility at location c . A *randomized mechanism* M maps an instance \vec{x} to a probability distribution over k -tuples $(c_1, \dots, c_k) \in \mathbb{R}^k$.

Connection Cost and Social Cost. Given a k -tuple $\vec{c} = (c_1, \dots, c_k)$, $c_1 \leq \dots \leq c_k$, of facility locations, the connection cost of agent i wrt. \vec{c} , denoted $d(x_i, \vec{c})$, is $d(x_i, \vec{c}) = \min_{1 \leq j \leq k} |x_i - c_j|$. Given a deterministic mechanism M and an instance \vec{x} , $d(x_i, M(\vec{x}))$ denotes the connection cost of agent i wrt. the outcome of $M(\vec{x})$. If M is a randomized mechanism, the expected connection cost

of agent i is $\mathbb{E}_{\vec{c} \sim M(\vec{x})}(d(x_i, \vec{c}))$. The *social cost* of a deterministic mechanism M for an instance \vec{x} is $cost(\vec{x}, M(\vec{x})) = \sum_{i=1}^n d(x_i, M(\vec{x}))$. The social cost of a facility locations profile $\vec{c} \in \mathbb{R}^k$ is $cost(\vec{x}, \vec{c}) = \sum_{i=1}^n d(x_i, \vec{c})$. The *expected social cost* of a randomized mechanism M on instance \vec{x} is

$$cost(\vec{x}, M(\vec{x})) = \sum_{i=1}^n \mathbb{E}_{\vec{c} \sim M(\vec{x})}(d(x_i, \vec{c})).$$

The *optimal social cost* for an instance \vec{x} is $cost^*(\vec{x}) = \min_{\vec{c} \in \mathbb{R}^k} \sum_{i=1}^n d(x_i, \vec{c})$. For k -Facility Location, the optimal social cost (and the corresponding optimal facility locations profile) can be computed in $O(kn \log n)$ time by standard dynamic programming.

Approximation Ratio. A mechanism M has an approximation ratio of $\rho \geq 1$, if for any instance \vec{x} , $cost(\vec{x}, M(\vec{x})) \leq \rho cost^*(\vec{x})$. We say that the approximation ratio ρ of M is *bounded*, if ρ is bounded from above either by a constant or by a (computable) function of n and k .

Strategyproofness. A deterministic mechanism M is *strategyproof*, if no agent can benefit from misreporting her location. Formally, M is strategyproof, if for all location profiles \vec{x} , any agent i , and all locations y , $d(x_i, M(\vec{x})) \leq d(x_i, M(\vec{x}_{-i}, y))$. Similarly, a randomized mechanism M is strategyproof (in expectation), if for all location profiles \vec{x} , any agent i , and all locations y , $\mathbb{E}_{\vec{c} \sim M(\vec{x})}(d(x_i, \vec{c})) \leq \mathbb{E}_{\vec{c} \sim M(\vec{x}_{-i}, y)}(d(x_i, \vec{c}))$.

Clusterings. A *clustering* (or k -clustering, if k is not clear from the context) of an instance \vec{x} is any partitioning $\vec{C} = (C_1, \dots, C_k)$ of \vec{x} into k sets of consecutive agent locations. We index clusters from left to right. I.e., $C_1 = \{x_1, \dots, x_{|C_1|}\}$, $C_2 = \{x_{|C_1|+1}, \dots, x_{|C_1|+|C_2|}\}$, and so on. We refer to a cluster C_i that includes only one agent (i.e., with $|C_i| = 1$) as a *singleton* cluster. We sometimes use (\vec{x}, \vec{C}) to highlight that we consider \vec{C} as a clustering of instance \vec{x} .

Two clusters C and C' are identical, denoted $C = C'$, if they include the exact same locations. Two clusterings $\vec{C} = (C_1, \dots, C_k)$ and $\vec{Y} = (Y_1, \dots, Y_k)$ of an instance \vec{x} are the same, if $C_i = Y_i$, for all $i \in [k]$. Abusing the notation, we say that a clustering \vec{C} of an instance \vec{x} is identical to a clustering \vec{Y} of a γ -perturbation \vec{x}' of \vec{x} (see also Definition 3.3.1), if $|C_i| = |Y_i|$, for all $i \in [k]$.

We let $x_{i,l}$ and $x_{i,r}$ denote the leftmost and the rightmost agent of each cluster C_i . Under this notation, $x_{i-1,r} < x_{i,l} \leq x_{i,r} < x_{i+1,l}$, for all $i \in \{2, \dots, k-1\}$. Exploiting the linearity of instances, we extend this notation to refer to other agents by their relative location in each cluster. Namely, $x_{i,l+1}$ (resp. $x_{i,r-1}$) is the second agent from the left (resp. right) of cluster C_i . The *diameter* of a cluster C_i is $D(C_i) = d(x_{i,l}, x_{i,r})$. The distance of clusters C_i and C_j is $d(C_i, C_j) = \min_{x \in C_i, y \in C_j} \{d(x, y)\}$, i.e., the minimum distance between a location $x \in C_i$ and a location $y \in C_j$.

A k -facility locations (or k -centers) profile $\vec{c} = (c_1, \dots, c_k)$ induces a clustering $\vec{C} = (C_1, \dots, C_k)$ of an instance \vec{x} by assigning each agent / location x_j to the

cluster C_i with facility c_i closest to x_j . Formally, for each $i \in [k]$, $C_i = \{x_j \in \vec{x} : d(x_j, c_i) = d(x_j, \vec{c})\}$. The *optimal clustering* of an instance \vec{x} is the clustering of \vec{x} induced by the facility locations profile with minimum social cost.

The social cost of a clustering \vec{C} induced by a k -facility locations profile \vec{c} on an instance \vec{x} is simply $cost(\vec{x}, \vec{c})$, i.e., the social cost of \vec{c} for \vec{x} . We sometimes refer to the social cost $cost(\vec{x}, \vec{C})$ of a clustering \vec{C} for an instance \vec{x} , without any explicit reference to the corresponding facility locations profile. Then, we refer to the social cost $cost(\vec{x}, \vec{c})$, where each facility c_i is located at the median location of C_i (the left median location of C_i , if $|C_i|$ is even).

We often consider certain structural changes in a clustering due to agent deviations. Let \vec{C} be a clustering of an instance \vec{x} , which due to an agent deviation, changes to a different clustering \vec{C}' . We say that cluster C_i is *split* when \vec{C} changes to \vec{C}' , if not all agents in C_i are served by the same facility in \vec{C}' . We say that C_i is *merged* in \vec{C}' , if all agents in C_i are served by the same facility, but this facility also serves in \vec{C}' some agents not in C_i .

3.3 Perturbation Stability on the Line: Definition and Properties

Next, we introduce the notion of γ -(linear) stability and prove some useful properties of γ -stable instances of k -Facility Location, which are repeatedly used in the analysis of our mechanisms.

Definition 3.3.1 (γ -Perturbation and γ -Stability). *Let $\vec{x} = (x_1, \dots, x_n)$ be a locations profile. A locations profile $\vec{x}' = (x'_1, \dots, x'_n)$ is a γ -perturbation of \vec{x} , for some $\gamma \geq 1$, if $x'_1 = x_1$ and for every $i \in [n - 1]$, $d(x_i, x_{i+1})/\gamma \leq d(x'_i, x'_{i+1}) \leq d(x_i, x_{i+1})$. A k -Facility Location instance \vec{x} is γ -perturbation stable (or simply, γ -stable), if \vec{x} has a unique optimal clustering (C_1, \dots, C_k) and every γ -perturbation \vec{x}' of \vec{x} has the same unique optimal clustering (C_1, \dots, C_k) .*

Namely, a γ -perturbation \vec{x}' of an instance \vec{x} is obtained by moving a subset of pairs of consecutive locations closer by a factor at most $\gamma \geq 1$. A k -Facility Location instance \vec{x} is γ -stable, if \vec{x} and any γ -perturbation \vec{x}' of \vec{x} admit the same unique optimal clustering (where clustering identity for \vec{x} and \vec{x}' is understood as explained in Section 3.2). We consistently select the optimal center c_i of each optimal cluster C_i with an even number of points as the left median point of C_i .

Our notion of linear perturbation stability naturally adapts the notion of metric perturbation stability [13, Definition 2.5] to the line. We note, the class of γ -stable linear instances, according to Definition 3.3.1, is at least as large as the class of metric γ -stable linear instances, according to [13, Definition 2.5]. Similarly to [13, Theorem 3.1] (see also [205, Lemma 7.1] and [31, Corollary 2.3]), we can show that for all $\gamma \geq 1$, every γ -stable instance \vec{x} , which admits an optimal clustering C_1, \dots, C_k with optimal centers c_1, \dots, c_k , satisfies the following γ -center proximity

property: For all cluster pairs C_i and C_j , with $i \neq j$, and all locations $x \in C_i$, $d(x, c_j) > \gamma d(x, c_i)$.

We repeatedly use the following immediate consequence of γ -center proximity (see also [205, Lemma 7.2]). The proof generalizes the proof of [205, Lemma 7.2] to any $\gamma \geq 2$.

Proposition 3.3.2. *Let $\gamma \geq 2$ and let \vec{x} be any γ -stable instance, with unique optimal clustering C_1, \dots, C_k and optimal centers c_1, \dots, c_k . Then, for all clusters C_i and C_j , with $i \neq j$, and all locations $x \in C_i$ and $y \in C_j$, $d(x, y) > (\gamma - 1)d(x, c_i)$.*

The following observation, which allows us to treat stability factors multiplicatively, is an immediate consequence of Definition 3.3.1.

Observation 3.3.3. *Every α -perturbation followed by a β -perturbation of a locations profile can be implemented by a $(\alpha\beta)$ -perturbation and vice versa. Hence, a γ -stable instance remains (γ/γ') -stable after a γ' -perturbation, with $\gamma' < \gamma$, is applied to it.*

We next show that for γ large enough, the optimal clusters of a γ -stable instance are well-separated, in the sense that the distance of two consecutive clusters is larger than their diameters.

Lemma 3.3.4 (Cluster-Separation Property). *For any γ -stable instance on the line with optimal clustering (C_1, \dots, C_k) and all clusters C_i and C_j , with $i \neq j$, $d(C_i, C_j) > \frac{(\gamma-1)^2}{2\gamma} \max\{D(C_i), D(C_j)\}$.*

The cluster-separation property of Lemma 3.3.4 was first obtained in [3] as a consequence of γ -cluster proximity. For completeness, we present a different proof that exploits the linear structure of the instance.

Proof. It suffices to establish the lemma for two consecutive clusters C_i and C_{i+1} . We recall that $d(C_i, C_{i+1}) = d(x_{i,r}, x_{i+1,l})$. Moreover, by symmetry, we can assume wlog. that $D(C_i) \geq D(C_{i+1})$.

If C_i is a singleton, $D(C_i) = 0$ and the lemma holds trivially. If $|C_i| = 2$, wlog. we can only consider the case where $x_{i,l}$ is C_i 's center. Otherwise, i.e., if $x_{i,r}$ is C_i 's center in optimal clustering $(C_1, \dots, C_i, \dots, C_k)$ with centers $(c_i, \dots, x_{i,r}, \dots, c_j)$, the same clustering $(C_1, \dots, C_i, \dots, C_k)$ with centers $(c_1, \dots, x_{i,l}, \dots, c_j)$ is also optimal for the γ -stable instance \vec{x} (and should still be optimal after a γ perturbation of \vec{x} , due to the stability of the instance). We then have:

$$\begin{aligned} D(C_i) = d(x_{i,l}, x_{i,r}) = d(c_i, x_{i,r}) &< \frac{1}{(\gamma-1)} d(x_{i,r}, x_{i+1,r}) = \frac{1}{(\gamma-1)} d(C_i, C_{i+1}) \Rightarrow \\ d(C_i, C_{i+1}) &> (\gamma-1)D(C_i) \end{aligned}$$

where the first inequality follows from Proposition 3.3.2. The lemma then follows by noticing that for any $\gamma \geq 1$:

$$\gamma - 1 \geq \frac{\gamma^2 + 1}{2\gamma} - 1$$

The most interesting case is where $|C_i| \geq 3$ and $x_{i,l} < c_i \leq x_{i,r}$. Suppose $d(x_{i,l}, c_i) = \beta D(C_i)$, for some $\beta \in (0, 1]$ and hence $d(c_i, x_{i,r}) = (1 - \beta)D(C_i)$ (i.e., β quantifies how close c_i is to C_i 's extreme points and to the closest point of C_{i+1}). We recall that $d(C_i, C_{i+1}) = d(x_{i,r}, x_{i+1,l})$.

We start with a tighter analysis of the equivalent of Proposition 3.3.2 for $x_{i,l}$ and $x_{i+1,l}$, taking into account their specific ordering on the line:

$$\begin{aligned} d(x_{i,l}, x_{i+1,l}) &\geq d(x_{i,l}, c_{i+1}) - d(x_{i+1,l}, c_{i+1}) \\ &> \gamma d(x_{i,l}, c_i) - \frac{d(x_{i+1,l}, c_i)}{\gamma} \\ &= \gamma d(x_{i,l}, c_i) - \frac{d(x_{i+1,l}, x_{i,l}) - d(x_{i,l}, c_i)}{\gamma} \Rightarrow \\ d(x_{i,l}, x_{i+1,l}) &> \frac{\gamma^2 + 1}{\gamma + 1} d(x_{i,l}, c_i) \end{aligned}$$

Where the second inequality stands due to the γ -center proximity property of γ stable instances and the equality stands because $x_{i,l} < c_i < x_{i+1,l}$. Since $d(C_i, C_{i+1}) = d(x_{i,r}, x_{i+1,l}) = d(x_{i,l}, x_{i+1,l}) - D(C_i)$, and by $d(x_{i,l}, c_i) = \beta D(C_i)$, we get that:

$$d(C_i, C_{i+1}) > \left(\frac{\beta(\gamma^2 + 1)}{\gamma + 1} - 1 \right) D(C_i) \quad (3.1)$$

Furthermore, by Proposition 3.3.2, we have that $d(x_{i,r}, x_{i+1,l}) > (\gamma - 1)d(x_{i,r}, c_i)$. Hence, by $d(c_i, x_{i,r}) = (1 - \beta)D(C_i)$, we get that:

$$d(C_i, C_{i+1}) > (1 - \beta)(\gamma - 1)D(C_i) \quad (3.2)$$

So, by (3.1) and (3.2) we have that it must be:

$$d(C_i, C_{i+1}) > \max \left\{ \frac{\beta(\gamma^2 + 1)}{\gamma + 1} - 1, (1 - \beta)(\gamma - 1) \right\} D(C_i) \quad (3.3)$$

We now observe that for any fixed $\gamma > 1$, the first term of the max in (3.3), $\frac{\beta(\gamma^2 + 1)}{\gamma + 1} - 1$, is increasing for all $\beta > 0$, while the second term, $(1 - \beta)(\gamma - 1)$, is decreasing for all $\beta \in (0, 1]$. Hence, for any fixed $\gamma > 1$, the minimum value of the max in (3.3) is achieved when β satisfies:

$$\frac{\beta(\gamma^2 + 1)}{\gamma + 1} - 1 = (1 - \beta)(\gamma - 1)$$

Solving for β , we get that:

$$\beta = \frac{1}{2} + \frac{1}{2\gamma}, \quad (3.4)$$

with $\beta \in (1/2, 1]$, when $\gamma \geq 1$.

We conclude the proof by substituting the value of β in (3.4) to (3.3). \square

Setting $\gamma \geq 2 + \sqrt{3}$, we get that:

Corollary 3.3.5. *Let $\gamma \geq 2 + \sqrt{3}$ and let \vec{x} be any γ -stable instance with unique optimal clustering (C_1, \dots, C_k) . Then, for all clusters C_i and C_j , with $i \neq j$, $d(C_i, C_j) > \max\{D(C_i), D(C_j)\}$.*

The following is an immediate consequence of the cluster-separation property in Lemma 3.3.4.

Observation 3.3.6. *Let \vec{x} be a k -Facility Location with a clustering $\vec{C} = (C_1, \dots, C_k)$ such that for any two clusters C_i and C_j , $\max\{D(C_i), D(C_j)\} < d(C_i, C_j)$. Then, if in the optimal clustering of \vec{x} , there is a facility at the location of some $x \in C_i$, no agent in C_i is served by a facility at $x_j \notin C_i$.*

(!!!What does this say?? Remove!!!)Next, we establish the so-called *no direct improvement from singleton deviations* property, used to show the strategyproofness of the ALMOSTRIGHTMOST and RANDOM mechanisms. Namely, we show that in any 3-stable instance, no agent deviating to a singleton cluster in the optimal clustering of the resulting instance can improve her connection cost through the facility of that singleton cluster.

Lemma 3.3.7. *Let \vec{x} be a γ -stable instance with $\gamma \geq 3$ and optimal clustering $\vec{C} = (C_1, \dots, C_k)$ and cluster centers (c_1, \dots, c_k) , and let an agent $x_i \in C_i \setminus \{c_i\}$ and a location x' such that x' is a singleton cluster in the optimal clustering of the resulting instance (\vec{x}_{-i}, x') . Then, $d(x_i, x') > d(x_i, c_i)$.*

Proof. We establish the lemma for the leftmost agent $x_{i,l}$ as the deviating agent. Specifically, we show that $x_{i,l}$ needs to move by at least $d(x_{i,l}, c_i)$ to the left in order to become a singleton cluster. The property then follows for the rest of the agents.

Suppose $x_{i,l}$ can create a singleton cluster by deviating less than $d(x_{i,l}, c_i)$ to the left. I.e., for some x' such that $d(x', x_{i,l}) < d(x_{i,l}, c_i)$ the optimal clustering of $\vec{x}' = (\vec{x}_{-x_{i,l}}, x')$ is such that the agent location at x' becomes a singleton cluster. We call this clustering (that is optimal for \vec{x}') \vec{C}' . Notice that since $d(x', x_{i,l}) < d(x_{i,l}, c_i)$, x' is in the gap between clusters C_{i-1} and C_i as by 3-perturbation stability we have $d(x_{i-1,r}, x_{i,l}) > 2d(x_{i,l}, c_i)$. This means that in order for this case to be feasible, no agents from C_{i-1} can be clustered together with agents in C_i in (\vec{x}', \vec{C}') , because x' lies between them and is a singleton cluster.

Consider now the instance $\vec{x}_{-x_{i,l}}$. We know that $\text{cost}(\vec{x}_{-x_{i,l}}, \vec{C}') \geq \text{cost}(\vec{x}_{-x_{i,l}}, \vec{C})$. That is, since otherwise the optimal clustering for \vec{x} would make $x_{i,l}$ a singleton cluster and serve the rest of the agents as in \vec{C}' . Let diff be the difference in the total cost agents in $\vec{x}_{-x_{i,l}}$ experience between clusterings \vec{C} and \vec{C}' . I.e.

$\text{diff} = \text{cost}(\vec{x}_{-x_{i,l}}, \vec{C}') - \text{cost}(\vec{x}_{-x_{i,l}}, \vec{C})$. As before, since $x_{i,l}$ is not a singleton cluster in (\vec{x}, \vec{C}) we know that $d(x_{i,l}, c_i) < \text{diff}$ (or else setting $x_{i,l}$ as a singleton would have a lower cost in \vec{x} than \vec{C}).

But we can perform a 3-perturbation in \vec{x} in the following way: Scale down all distances between agents from x_1 up to $x_{i-1,r}$ and all distances between agents from $x_{i,l+1}$ to x_n (x_n being the rightmost agent of the instance) by 3. Call this instance \vec{x}_{per} . Since agents of clusters C_{i-1} and C_i are not clustered together neither in \vec{C} nor in \vec{C}' we have that

$$\text{diff}_{per} \leq \frac{\text{cost}(\vec{x}_{-x_{i,l}}, \vec{C}') - \text{cost}(\vec{x}_{-x_{i,l}}, \vec{C})}{3}.$$

So $\text{diff}_{per} \leq \text{diff}/3$. Since $d(x_{i,l}, c_i)$ is unaffected in the perturbation and by stability the optimal clustering of \vec{x}_{per} must remain the same (as \vec{x}) we have that it must be $d(x_{i,l}, c_i) < \text{diff}/3$ (1).

Finally, the least amount of extra social cost suffered between $\text{cost}(\vec{x}, \vec{C})$ and the case of setting $x_{i,l}$ as a center that serves only itself and serve the remaining agents of the instance as on \vec{C}' (i.e. as they would be served should x' gets a facility that served only herself), will be $\text{diff} - d(x_{i,l}, c_i)$. This means that the optimal clustering algorithm would only choose this solution when $d(x', c_i) > \text{diff} - d(x, c_i)$. So the agent must deviate by at least $\text{diff} - 2d(x, c_i)$. But from (1) we have

$$\text{diff} - 2d(x, c_i) > 3d(x, c_i) - 2d(x_{i,l}, c_i) = d(x_{i,l}, c_i),$$

which concludes the proof of the lemma. \square

The following shows that for 5-stable instances \vec{x} , an agent cannot form a singleton cluster, unless she deviates by a distance larger than the diameter of her cluster in \vec{x} 's optimal clustering.

Lemma 3.3.8. *Let \vec{x} be any γ -stable instance with $\gamma \geq 5$ and optimal clustering $\vec{C} = (C_1, \dots, C_k)$. Let $x_i \in C_i \setminus \{c_i\}$ be any agent and x' any location such that x' is a singleton cluster in the optimal clustering of instance $\vec{x}' = (\vec{x}_{-i}, x')$, where x_i has deviated to x' . Then, $d(x', x_i) > D(C_i)$.*

Sketch. Initially, we show that a clustering \vec{C}' of instance $\vec{x}' = (\vec{x}_{-i}, x')$, with $d(x', x_i) \leq D(C_i)$, cannot be optimal and contain x' as a singleton cluster, unless some agent $\vec{x} \setminus C_i$ is clustered together with some agent in C_i . To this end, we use the lower bound on the distance between difference clusters for 5-stable instances show in Lemma 3.3.4. Then, using stability arguments, i.e. that the optimal clustering should not change for instance \vec{x} , even when we decrease, by a factor of 4, the distances between consecutive agents in $\vec{x} \setminus C_i$, we show that in \vec{C}' agents in $\vec{x} \setminus C_i$ experience an increase in cost of at least $2D(C_i)$ (notice that $\vec{x} \setminus C_i = \vec{x}' \setminus (C_i \cup \{x'\})$). But the additional cost of serving x' from c_i in clustering \vec{C} is at most $2D(C_i)$, since $d(x', x_i) \leq D(C_i)$ and $d(x_i, c_i) \leq D(C_i)$.




Figure 3.1: An example of a so-called *singleton deviation*. The deviating agent (grey) declares a remote location, becomes a singleton cluster, and essentially turns the remaining agents into a $(k - 1)$ -Facility Location instance. Thus, the deviating agent can benefit from her singleton deviation, due to the subsequent cluster merge.

Mechanism 1: OPTIMAL

Result: An allocation of k -facilities

Input: A k -Facility Location instance \vec{x} .

- 1 Compute the optimal clustering (C_1, \dots, C_k) . Let c_i be the left median point of each cluster C_i .
 - 2 **if** $(\exists i \in [k]$ with $|C_i| = 1)$ **or** $(\exists i \in [k - 1]$ with $\max\{D(C_i), D(C_{i+1})\} \geq d(C_i, C_{i+1}))$ **then**
| **Output:** “FACILITIES ARE NOT ALLOCATED”.
 - 3 **else**
| **Output:** The k -facility allocation (c_1, \dots, c_k)
-

Hence retaining clustering \vec{C} and serving location x' from c_i would have a smaller cost than the supposedly optimal clustering \vec{C}' . The complete proof follows by a careful case analysis. \square

The proof is deferred to appendix 3.A of the chapter.

3.4 The Optimal Solution is Strategyproof for $(2 + \sqrt{3})$ -Stable Instances

We next show that the OPTIMAL mechanism, which allocates the facilities optimally, is strategyproof for $(2 + \sqrt{3})$ -stable instances of k -Facility Location whose optimal clustering does not include any singleton clusters. More specifically, in this section, we analyze Mechanism 1.

In general, due to the incentive compatibility of the median location in a single cluster, a deviation can be profitable only if it results in a k -clustering different

from the optimal clustering (C_1, \dots, C_k) of \vec{x} . For γ is sufficiently large, γ -stability implies that the optimal clusters are well identified so that any attempt to alter the optimal clustering (without introducing singleton clusters and without violating the cluster separation property, which is necessary of stability) results in an increased cost for the deviating agent. We should highlight that Mechanism 1 may also “serve” non-stable instances that satisfy the cluster separation property. We next prove that the mechanism is strategyproof if the true instance is $(2 + \sqrt{3})$ -stable and its optimal clustering does not include any singleton clusters, when the agent deviations do not introduce any singleton clusters and not result in instances that violate the cluster separation property (i.e. are served by the mechanism) .

Theorem 3.4.1. *The OPTIMAL mechanism applied to $(2 + \sqrt{3})$ -stable instances of k -Facility Location without singleton clusters in their optimal clustering is strategyproof and minimizes the social cost.*

Proof. We first recall some of the notation about clusterings, introduced in Section 3.2. Specifically, for a clustering $\vec{C} = (C_1, \dots, C_k)$ of an instance \vec{x} with centers $\vec{c} = (c_1, \dots, c_k)$, the cost of an agent (or a location) x is $d(x, \vec{C}) = \min_{j \in [k]} \{d(x, c_j)\}$. The cost of a set of agents X in a clustering \vec{C} is $cost(X, \vec{C}) = \sum_{x \in X} d(x, \vec{C})$. Finally, the cost of an instance \vec{x} in a clustering \vec{C} is $cost(\vec{x}, \vec{C}) = \sum_{x_j \in \vec{x}} d(x_j, \vec{C})$. This general notation allows us to refer to the cost of the same clustering for different instances. I.e, if \vec{C} is the optimal clustering of \vec{x} , then $cost(\vec{y}, \vec{C})$ denotes the cost of instance \vec{y} in clustering \vec{C} (where we select the same centers as in clustering \vec{C} for \vec{x}).

The fact that if OPTIMAL outputs k facilities, they optimize the social cost is straightforward. So, we only need to establish strategyproofness. To this end, we show the following: Let \vec{x} be any $(2 + \sqrt{3})$ -perturbation stable k -Facility Location instance with optimal clustering $\vec{C} = (C_1, \dots, C_k)$. For any agent i and any location y , let \vec{Y} be the optimal clustering of the instance $\vec{y} = (\vec{x}_{-i}, y)$ resulting from the deviation of i from x_i to y . Then, if y does not form a singleton cluster in (\vec{y}, \vec{Y}) , either $d(x_i, \vec{C}) < d(x_i, \vec{Y})$, or there is an $i \in [k - 1]$ for which $\max\{D(Y_i), D(Y_{i+1})\} \geq d(Y_i, Y_{i+1})$.

So, we let $x_i \in C_i$ deviate to a location y , resulting in $\vec{y} = (\vec{x}_{-i}, y)$ with optimal clustering \vec{Y} . Since y is not a singleton cluster, it is clustered with agents belonging in one or two clusters of \vec{C} , say either in cluster C_j or in clusters C_{j-1} and C_j . By optimality of \vec{C} and \vec{Y} , the number of facilities serving $C_{j-1} \cup C_j \cup \{y\}$ in (\vec{y}, \vec{Y}) is no less than the number of facilities serving $C_{j-1} \cup C_j$ in (\vec{x}, \vec{C}) . Hence, there is at least one facility in either C_{j-1} or C_j .

Wlog., suppose that a facility is allocated to an agent in C_j in (\vec{y}, \vec{Y}) . By Corollary 3.3.5 and Observation 3.3.6, no agent in C_j is served by a facility in $\vec{x} \setminus C_j$ in \vec{Y} . Thus we get the following cases about what happens with the optimal clustering \vec{Y} of instance $\vec{y} = (\vec{x}_{-i}, y)$:

Case 1: y is not allocated a facility in \vec{Y} : This can happen in one of two ways:

Case 1a: y is clustered together with some agents from cluster C_j and no facility placed in C_j serves agents in $\vec{x} \setminus C_j$ in \vec{Y} .

Case 1b: y is clustered together with some agents from a cluster C_j and at least one of the facilities placed in C_j serve agents in $\vec{x} \setminus C_j$ in \vec{Y} .

Case 2: y is allocated a facility in \vec{Y} . This can happen in one of two ways:

Case 2a: y only serves agents that belong in C_j (by optimality, y must be the median location of the new cluster, which implies that either $y < x_{i,l}$ and y only serves $x_{i,l}$ or $x_{j,l} \leq y \leq x_{j,r}$).

Case 2b: In \vec{Y} , y serves agents that belong in both C_{j-1} and C_j .

We next show that the cost of the original clustering \vec{C} is less than the cost of clustering \vec{Y} in \vec{y} . Hence, mechanism OPTIMAL would also select clustering \vec{C} for \vec{y} , which would make x_i 's deviation to y non-profitable. In particular, it suffices to show that:

$$\begin{aligned} \text{cost}(\vec{y}, \vec{C}) &< \text{cost}(\vec{y}, \vec{Y}) \Leftrightarrow \\ \text{cost}(\vec{x}, \vec{C}) + d(y, \vec{C}) - d(x_i, \vec{C}) &< \text{cost}(\vec{x}, \vec{Y}) + d(y, \vec{Y}) - d(x_i, \vec{Y}) \Leftrightarrow \\ d(y, \vec{C}) - d(y, \vec{Y}) &< \text{cost}(\vec{x}, \vec{Y}) - \text{cost}(\vec{x}, \vec{C}) + d(x_i, \vec{C}) - d(x_i, \vec{Y}) \end{aligned}$$

Since x_i 's deviation to y is profitable, $d(x_i, \vec{C}) - d(x_i, \vec{Y}) > 0$. Hence, it suffices to show that:

$$\begin{aligned} d(y, \vec{C}) - d(y, \vec{Y}) &\leq \text{cost}(\vec{x}, \vec{Y}) - \text{cost}(\vec{x}, \vec{C}) \\ &= \text{cost}(C_j, \vec{Y}) - \text{cost}(C_j, \vec{C}) + \text{cost}(\vec{x} \setminus C_j, \vec{Y}) - \text{cost}(\vec{x} \setminus C_j, \vec{C}) \end{aligned} \quad (3.5)$$

We first consider Case 1a and Case 2a, i.e., the cases where \vec{Y} allocates facilities to agents of C_j (between $x_{j,l}$ and $x_{j,r}$) that serve only agents in C_j . Note that in case 2a, y can also be located outside of C_j and serve only $x_{i,l}$. We can treat this case as Case 1a, since it is equivalent to placing the facility on $x_{i,l}$ and serving y from there.

In Case 1a and Case 2a, we note that (3.5) holds if the clustering \vec{Y} allocates a single facility to agents in $C_j \cup \{y\}$, because the facility is allocated to the median of $C_j \cup \{y\}$, hence $d(y, \vec{C}) - d(y, \vec{Y}) = \text{cost}(C_j, \vec{Y}) - \text{cost}(C_j, \vec{C})$, while $\text{cost}(\vec{x} \setminus C_j, \vec{Y}) - \text{cost}(\vec{x} \setminus C_j, \vec{C}) \geq 0$, since \vec{C} is optimal for \vec{x} . So, we focus on the most interesting case where the agents in $C_j \cup \{y\}$ are allocated at least two facilities. We observe that (3.5) follows from:

$$d(y, \vec{C}) - d(y, \vec{Y}) \leq \frac{1}{\gamma} \left(\text{cost}(\vec{x} \setminus C_j, \vec{Y}) - \text{cost}(\vec{x} \setminus C_j, \vec{C}) \right) \quad (3.6)$$

$$\text{cost}(C_j, \vec{C}) - \text{cost}(C_j, \vec{Y}) \leq \left(1 - \frac{1}{\gamma}\right) \left(\text{cost}(\vec{x} \setminus C_j, \vec{Y}) - \text{cost}(\vec{x} \setminus C_j, \vec{C}) \right) \quad (3.7)$$

To establish (3.6) and (3.7), we first consider the valid γ -perturbation of the original instance \vec{x} where all distances between consecutive agent pairs to the left of C_j (i.e. agents $\{x_1, x_2, \dots, x_{j-1, r}\}$) and between consecutive agent pairs to the right of C_j (i.e. agents $\{x_{j+1, l}, \dots, x_{k, r}\}$) are scaled down by γ . By stability, the clustering \vec{C} remains the unique optimal clustering for the perturbed instance \vec{x}' . Moreover, since agents in $\vec{x} \setminus C_j$ are not served by a facility in C_j in \vec{C} and \vec{Y} , and since all distances outside C_j are scaled down by γ , while all distances within C_j remain the same, the cost of the clusterings \vec{C} and \vec{Y} for the perturbed instance \vec{x}' is $\text{cost}(C_j, \vec{C}) + \text{cost}(\vec{x} \setminus C_j, \vec{C})/\gamma$ and $\text{cost}(C_j, \vec{Y}) + \text{cost}(\vec{x} \setminus C_j, \vec{Y})/\gamma$, respectively. Using $\text{cost}(\vec{x}', \vec{C}) < \text{cost}(\vec{x}', \vec{Y})$ and $\gamma \geq 2$, we obtain:

$$\text{cost}(C_j, \vec{C}) - \text{cost}(C_j, \vec{Y}) < \frac{1}{\gamma} \left(\text{cost}(\vec{x} \setminus C_j, \vec{Y}) - \text{cost}(\vec{x} \setminus C_j, \vec{C}) \right) \quad (3.8)$$

$$\leq \left(1 - \frac{1}{\gamma} \right) \left(\text{cost}(\vec{x} \setminus C_j, \vec{Y}) - \text{cost}(\vec{x} \setminus C_j, \vec{C}) \right) \quad (3.9)$$

Moreover, if $C_j \cup \{y\}$ is served by at least two facilities in \vec{Y} , the facility serving y (and some agents of C_j) is placed at the median location of \vec{Y} 's cluster that contains y . Wlog., we assume that y lies on the left of the median of C_j . Then, the decrease in the cost of y due to the additional facility in \vec{Y} is equal to the decrease in the cost of $x_{i, l}$ in \vec{Y} , which bounds from below the total decrease in the cost of C_j due to the additional facility in \vec{Y} . Hence,

$$d(y, \vec{C}) - d(y, \vec{Y}) \leq \text{cost}(C_j, \vec{C}) - \text{cost}(C_j, \vec{Y}) \quad (3.10)$$

We conclude Case 1a and Case 2a, by observing that (3.6) follows directly from (3.10) and (3.8).

Finally, we study Case 1b and Case 2b, i.e. the cases where some agents of C_j are clustered with agents of $\vec{x} \setminus C_j$ in \vec{Y} . Let C'_{j1} and C'_{j2} denote the clusters of (\vec{y}, \vec{Y}) including all agents of C_j (i.e., $C_j \subseteq C'_{j1} \cup C'_{j2}$). By hypothesis, at least one of C'_{j1} and C'_{j2} contains an agent $z \in \vec{x} \setminus C_j$. Suppose this is true for the cluster C'_{j1} . Then, $D(C'_{j1}) > D(C_j)$, since by Corollary 3.3.5, for any $\gamma \geq (2 + \sqrt{3})$, the distance of any agent z outside C_j to the nearest agent in C_j is larger than C_j 's diameter. But since both C'_{j1} and C'_{j2} contain agents of C_j , we have that $d(C'_{j1}, C'_{j2}) < D(C_j)$. Therefore, $D(C'_{j1}) > d(C'_{j1}, C'_{j2})$ and the cluster-separation property is violated. Hence the resulting instance \vec{y} is not γ -stable and Mechanism 1 does not allocated any facilities for it. \square

3.5 A Deterministic Mechanism Resistant to Singleton Deviations

Next, we present a deterministic strategyproof mechanism for 5-stable instances of k -Facility Location whose optimal clustering may include singleton clusters. To make singleton cluster deviations non profitable, cluster merging has to

be discouraged by the facility allocation rule. So, we allocate facilities near the edge of each optimal cluster, ending up with a significantly larger approximation ratio and a requirement for larger stability, in order to achieve strategyproofness. Specifically, we now need to ensure that no agent can become a singleton cluster close enough to her original location. Moreover, since agents can now gain by splitting their (true) optimal cluster, we need to ensure that such deviations are either non profitable or violate the cluster-separation property.

Theorem 3.5.1. *ALMOSTRIGHTMOST (Mechanism 2) is strategyproof for 5-stable instances of k -Facility Location and achieves an approximation ratio of $(n - 2)/2$.*

Proof. The approximation ratio of $(n - 2)/2$ follows directly from the fact that the mechanism allocates the facility to the second rightmost agent of each non-singleton optimal cluster.

As for strategyproofness, let \vec{x} denote the true instance and $\vec{C} = (C_1, \dots, C_k)$ its optimal clustering. We consider an agent $x_i \in C_j$ deviating to location y , resulting in an instance $\vec{y} = (\vec{x}_{-i}, y)$ with optimal clustering \vec{Y} . Agent x_i 's cost is at most $D(C_j)$. Agent x_i could profitably declare false location y in the following ways:

Mechanism 2: ALMOSTRIGHTMOST

Result: An allocation of k -facilities

Input: A k -Facility Location instance \vec{x} .

- 1 Find the optimal clustering $\vec{C} = (C_1, \dots, C_k)$ of \vec{x} .
- 2 **if** there are two consecutive clusters C_i and C_{i+1} with
 $\max\{D(C_i), D(C_{i+1})\} \geq d(C_i, C_{i+1})$ **then**
| **Output:** "FACILITIES ARE NOT ALLOCATED".
- 3 **for** $i \in \{1, \dots, k\}$ **do**
- 4 **if** $|C_i| > 1$ **then**
- 5 | Allocate a facility to the location of the second rightmost agent of C_i ,
| i.e., $c_i \leftarrow x_{i,r-1}$.
- 6 **else**
- 7 | Allocate a facility to the single agent location of C_i : $c_i \leftarrow x_{i,l}$
- 8 **end**
- 9 **end**

Output: The k -facility allocation $\vec{c} = (c_1, \dots, c_k)$.

Case 1: The agents in C_j are clustered together in \vec{Y} and y is allocated a facility with $d(y, x_i) < d(x_i, x_{i,r-1}) \leq D(C_j)$ ($x_{i,r-1}$ is the location of x_i 's facility, when she is truthful).

Case 1a: y is a singleton cluster and $d(y, x_i) < D(C_j)$. For 5-stable instances, Lemma 3.3.8 implies that $x_i \in C_j$ has to move by at least $D(C_j)$ to become a singleton cluster, a contradiction.

Case 1b: y is the second rightmost agent of a cluster C'_j in (\vec{y}, \vec{Y}) . Then, the agent x_i can gain only if $d(y, x_i) < D(C_j)$. In Case 1, the agents in C_j are clustered together in \vec{Y} . If $y < x_i$, y must be the second rightmost agent of a cluster on the left of $x_{j,l}$ and by Lemma 3.3.4, $d(x_i, y) \geq d(x_{j,l}, x_{j-1,r}) > D(C_j)$. Hence, such a deviation cannot be profitable for x_i (note how this case crucially uses the facility allocation to the second rightmost agent of a cluster). If $y > x_i$, x_i can only gain if y is the second rightmost agent of a cluster including $C_j \cup \{y, x_{j+1,l}\}$ and possibly some agents on the left of C_j , which is treated below.

Case 2: The agents in C_j are clustered together in \vec{Y} and C_j is merged with some agents from C_{j+1} and possibly some other agents to the left of $x_{j,l}$ (note that merging C_j only with agents to the left of $x_{j,l}$ does not change the facility of x_i). Then, we only need to consider the case where the deviating agent x_i is $x_{j,r}$, since any other agent to the left of $x_{j,r-1}$ cannot gain, because cluster merging can only move their serving facility further to the right. As for $x_{j,r}$, we note that by optimality and the hypothesis that agents in C_j belong in the same cluster of \vec{Y} , $x_{j,r}$ cannot cause the clusters C_j and C_{j+1} to merge in \vec{Y} by deviating in the range $[x_{j,r}, x_{j+1,l}]$. The reason is that the set of agents $(C_j \setminus \{x_{j,r}\}) \cup \{y\} \cup C_{j+1}$ cannot be served optimally by a single facility, when the set of agents $C_j \cup C_{j+1}$ requires two facilities in the optimal clustering \vec{C} . Hence, unless C_{j+1} is split in \vec{Y} (which is treated similarly to Case 3a), $x_{j,r}$ can only move her facility to C_{j+1} , which is not profitable for her, due to Lemma 3.3.4.

Case 3: C_j is split into two clusters in \vec{Y} . Hence, the leftmost agents, originally in C_j , are served by a different facility than the rest of the agents originally in C_j . We next show that in any profitable deviation of x_i where C_j is split, either the deviation is not feasible or the cluster-separation property is violated. The case analysis below is similar to the proof of Theorem 3.4.1.

Case 3a: Agents in C_j are clustered together with some agents of $\vec{x} \setminus C_j$ in \vec{Y} . By hypothesis, there are agents $z, w \in C_j$ placed in different clusters of \vec{Y} , and at least one of them, say z , is clustered together with an agent $p \in C_\ell$, with $\ell \neq j$, in \vec{Y} . For brevity, we refer to the (different) clusters in which z and w are placed in clustering \vec{Y} as C'_z and C'_w , respectively. Then, $D(C'_z) \geq d(p, z) > D(C_j)$, by Lemma 3.3.4. But also $d(C'_z, C'_w) < d(z, w) \leq D(C_j)$, and consequently, $D(C'_z) > d(C'_z, C'_w)$, which implies that the cluster-separation property is violated and Mechanism 2 does not allocate any facilities in this case.

Case 3b: Agents in C_j are split and are not clustered together with any agents of $\vec{x} \setminus C_j$ in \vec{Y} . Hence, y is not clustered with any agents in $\vec{x} \setminus C_j$

in \vec{Y} . Otherwise, i.e., if y is not clustered with agents of C_j in \vec{Y} , it would be suboptimal for clustering \vec{Y} to allocate more than one facility to agents of $C_j \setminus \{x_i\}$ and at most $k - 2$ facilities to $(\vec{x} \cup \{y\}) \setminus C_j$, while the optimal clustering \vec{C} allocates a single facility to C_j and $k - 1$ facilities to $\vec{x} \setminus C_j$. But again if y is only clustered with agents of C_j , it is suboptimal for clustering \vec{Y} to allocate more than one facility to agents of $(C_j \cup \{y\}) \setminus \{x_i\}$ and at most $k - 2$ facilities to $\vec{x} \setminus C_j$, while the optimal clustering \vec{C} allocates a single facility to C_j and $k - 1$ facilities to $\vec{x} \setminus C_j$, as shown in the proof of Theorem 3.4.1.

□

3.6 Low Stability and Inapproximability by Deterministic Mechanisms

We next extend the impossibility result of [112, Theorem 3.7] to $\sqrt{2}$ -stable instances of k -Facility Location on the line, with $k \geq 3$. Thus, we provide strong evidence that restricting our attention to stable instances does not make strategyproof mechanism design trivial.

3.6.1 Image Sets, Holes and Well-Separated Instances

We start with some basic facts about strategyproof mechanisms and by adapting the technical machinery of well-separating instances from [112, Section 2.2] to stable instances.

Image Sets and Holes. Given a mechanism M , the *image set* $I_i(\vec{x}_{-i})$ of an agent i with respect to an instance \vec{x}_{-i} is the set of facility locations the agent i can obtain by varying her reported location. Formally, $I_i(\vec{x}_{-i}) = \{a \in \mathbb{R} : \exists y \in \mathbb{R} \text{ with } M(\vec{x}_{-i}, y) = a\}$.

If M is strategyproof, any image set $I_i(\vec{x}_{-i})$ is a collection of closed intervals (see e.g., [209, p. 249]). Moreover, a strategyproof mechanism M places a facility at the location in $I_i(\vec{x}_{-i})$ nearest to the declared location of agent i . Formally, for any agent i , all instances \vec{x} , and all locations y , $d(y, M(\vec{x}_{-i}, y)) = \inf_{a \in I_i(\vec{x}_{-i})} \{d(y, a)\}$.

Some care is due, because we consider mechanisms that need to be strategyproof only for γ -stable instances (\vec{x}_{-i}, y) . The image set of such a mechanism M is well defined (possibly by assuming that all facilities are placed to essentially $+\infty$), whenever (\vec{x}_{-i}, y) is not γ -stable. Moreover, the requirement that M places a facility at the location in $I_i(\vec{x}_{-i})$ nearest to the declared location y of agent i holds only if the resulting instance (\vec{x}_{-i}, y) is stable. We should underline that all instances considered in the proof of Theorem 3.6.5 are stable (and the same holds for the proofs of the propositions adapted from [112, Section 2.2]).

Any (open) interval in the complement of an image set $I \equiv I_i(\vec{x}_{-i})$ is called a *hole* of I . Given a location $y \notin I$, we let $l_y = \sup_{a \in I} \{a < y\}$ and $r_y = \inf_{a \in I} \{a > y\}$ be the locations in I nearest to y on the left and on the right, respectively. Since I is a collection of closed intervals, l_y and r_y are well-defined and satisfy $l_y < y < r_y$. For convenience, given a $y \notin I$, we refer to the interval (l_y, r_y) as a y -hole in I .

Well-Separated Instances. Given a deterministic strategyproof mechanism M with a bounded approximation $\rho \geq 1$ for k -Facility Location, an instance \vec{x} is $(x_1 | \dots | x_{k-1} | x_k, x_{k+1})$ -*well-separated* if $x_1 < \dots < x_k < x_{k+1}$ and $\rho d(x_{k+1}, x_k) < \min_{i \in \{2, \dots, k\}} \{d(x_{i-1}, x_i)\}$. We call x_k and x_{k+1} the *isolated pair* of the well-separated instance \vec{x} .

Hence, given a ρ -approximate mechanism M for k -Facility Location, a well-separated instance includes a pair of nearby agents at distance to each other less than $1/\rho$ times the distance between any other pair of consecutive agents. Therefore, any ρ -approximate mechanism serves the two nearby agents by the same facility and serve each of the remaining “isolated” agents by a different facility. We remark that well-separated instances are also ρ -stable.

We are now ready to adapt some useful properties of well-separated instances from [112, Section 2.2]. It is not hard to verify that the proofs of the auxiliary lemmas below apply to $\sqrt{2}$ -stable instances, either without any change or with some minor modifications (see also [112, Appendix A]). For completeness, we give the proofs of the lemmas below in Appendix 3.B.

Lemma 3.6.1 (Proposition 2.2, [112]). *Let M be any deterministic strategyproof mechanism with a bounded approximation ratio $\rho \geq 1$. For any $(x_1 | \dots | x_{k-1} | x_k, x_{k+1})$ -well-separated instance \vec{x} , $M_k(\vec{x}) \in [x_k, x_{k+1}]$.*

Lemma 3.6.2 (Proposition 2.3, [112]). *Let M be any deterministic strategyproof mechanism with a bounded approximation ratio $\rho \geq 1$, and let \vec{x} be a $(x_1 | \dots | x_{k-1} | x_k, x_{k+1})$ -well-separated instance with $M_k(\vec{x}) = x_k$. Then, for every $(x_1 | \dots | x_{k-1} | x'_k, x'_{k+1})$ -well-separated instance \vec{x}' with $x'_k \geq x_k$, $M_k(\vec{x}') = x'_k$.*

Lemma 3.6.3 (Proposition 2.4, [112]). *Let M be any deterministic strategyproof mechanism with a bounded approximation ratio $\rho \geq 1$, and let \vec{x} be a $(x_1 | \dots | x_{k-1} | x_k, x_{k+1})$ -well-separated instance with $M_k(\vec{x}) = x_{k+1}$. Then, for every $(x_1 | \dots | x_{k-1} | x'_k, x'_{k+1})$ -well-separated instance \vec{x}' with $x'_{k+1} \leq x_{k+1}$, $M_k(\vec{x}') = x'_{k+1}$.*

3.6.2 Nonexistence of strategyproof mechanisms, with bounded approximation in stable instances.

In this subsection we show that our restriction to deterministic mechanisms that only guarantee strategyproofness in stable instances is imposed by the nature of the problem and is actually the only way that bounded approximation strategyproof mechanisms can be created for this family of instances. More specifically, a natural question to arise after observing the negative result of [112] and our

positive results about creating good mechanisms that only work when restricted to the domain of strategyproof instances would be: "Can we create mechanisms that are strategyproof (not necessarily bounded approximation) for all instances but also bounded approximation for stable instances?". Here we show that the answer to this question is no.

The idea behind this observation is the following: In the proof of the negative result in [112, Theorem 3.7] we see that the authors start with an original instance that due to bounded approximation has to have a specific allocation. Then they argue that because of the mechanism's behaviour in that first instance the mechanism would have to follow a specific behaviour in a slightly altered instance than the original for one of two reasons: either to maintain strategyproofness or to maintain bounded approximation. Building on this logic then we follow a series of intermediate instances along with the equivalent "forced" behaviour of the mechanism, in order to guarantee these properties, to finally end up in a contradictory allocation (i.e. recreate the original instance showing that the mechanism would allocate two facilities over the well separated pair, violating the bounded approximation property). All we need to do to reach the required conclusion then, is to recreate a similar (or the same proof) showing that each intermediate instance in which we determine the mechanism's behaviour by the fact that it should maintain a bounded approximation ratio is stable. Thus we allow the mechanism to only be strategyproof in general instances but only require it to be bounded approximation in stable ones. Reaching the same contradiction as in the original proof we conclude that such a mechanism doesn't exist.

Theorem 3.6.4. *For every $k \geq 3$, any deterministic strategyproof mechanism for k -facility location on the real line with $n \geq k + 1$ agents has unbounded approximation ratio within any domain of γ -stable instances, for any $\gamma \geq 1$.*

Proof. We only consider the case where $k = 3$ and $n = 4$. It is then not hard to verify that the result stands for any $k \geq 3$ and $n \geq k + 1$. To reach a contradiction, let M be any deterministic anonymous mechanism □

3.6.3 The Proof of the Impossibility Result

We are now ready to establish the main result of this section. The proof of the following builds on the proof of [112, Theorem 3.7]. However, we need some additional ideas and to be way more careful with the agent deviations used in the proof, since our proof can only rely on $\sqrt{2}$ -stable instances.

Theorem 3.6.5. *For every $k \geq 3$ and any $\delta > 0$, any deterministic anonymous strategyproof mechanism for $(\sqrt{2} - \delta)$ -stable instances of k -Facility Location on the real line with $n \geq k + 1$ agents has an unbounded approximation ratio.*

Proof. We only consider the case where $k = 3$ and $n = 4$. It is not hard to verify that the proof applies to any $k \geq 3$ and $n \geq k + 1$. To reach a contradiction, let

M be any deterministic anonymous strategyproof mechanism for $(\sqrt{2} - \delta)$ -stable instances of 3-Facility Location with $n = 4$ agents and with an approximation ratio of $\rho \geq 1$.

We consider a $(x_1|x_2|x_3, x_4)$ -well-separated instance \vec{x} . For a large enough $\lambda \gg \rho$ and a very large (practically infinite) $B \gg 6\rho\lambda$, we let $\vec{x} = (0, \lambda, 6B + \lambda, 6B + \lambda + \varepsilon)$, for some small enough $\varepsilon > 0$ ($\varepsilon \ll \lambda/\rho$). By choosing λ and ε appropriately, becomes the instance \vec{x} γ -stable, for $\gamma \gg \sqrt{2}$.

By Lemma 3.6.1, $M_3(\vec{x}) \in [x_3, x_4]$. Wlog, we assume that $M_3(\vec{x}) \neq x_3$ (the case where $M_3(\vec{x}) \neq x_4$ is fully symmetric and requires Lemma 3.6.2). Then, by moving agent 4 to $M_3(\vec{x})$, which results in a well-separated instance and, by strategyproofness, requires that M keeps a facility there, we can assume wlog. that $M_3(\vec{x}) = x_4$.

Since \vec{x} is well-separated and M is ρ -approximate, both x_3 and x_4 are served by the facility at x_4 . Hence, there is a x_3 -hole $h = (l, r)$ in the image set $I_3(\vec{x}_{-3})$. Since $M(\vec{x})$ places a facility at x_4 and not in x_3 , the right endpoint r of h lies between x_3 and x_4 , i.e. $r \in (x_3, x_4]$. Moreover, since M is ρ -approximate and strategyproof for $(\sqrt{2} - \delta)$ -stable instances, agent 3 should be served by a facility at distance at most $\rho\lambda$ to her, if she is located at $4B$. Hence, the left endpoint of the hole h is $l > 3B$. We distinguish two cases based on the distance of the left endpoint l of h to x_4 .

Case 1: $x_4 - l > \sqrt{2}\lambda$. We consider the instance $\vec{y} = (\vec{x}_{-3}, a)$, where $a > l$ is arbitrarily close to l (i.e., $a \gtrsim l$) so that $d(a, x_4) = \sqrt{2}\lambda$. Since $d(x_1, x_2) = \lambda$, $d(x_2, a)$ is quite large, and $d(a, x_4) = \sqrt{2}\lambda$, the instance \vec{y} is $(\sqrt{2} - \delta)$ -stable, for any $\delta > 0$. By strategyproofness, $M(\vec{y})$ must place a facility at l , since $l \in I_3(\vec{x}_{-3})$.

Now, we consider the instance $\vec{y}' = (\vec{y}_{-4}, l)$. Since we can choose $a > l$ so that $d(l, a) \ll \lambda$, the instance \vec{y}' is $(x_1|x_2|l, a)$ -well-separated and $(\sqrt{2} - \delta)$ -stable. Hence, by strategyproofness, $M(\vec{y}')$ must keep a facility at l , because $l \in I_4(\vec{y}_{-4})$.

Then, by Lemma 3.6.3, $y'_4 = a \in M(\vec{y}')$, because for the $(x_1|x_2|x_3, x_4)$ -well-separated instance \vec{x} , $M_3(\vec{x}) = x_4$, and \vec{y}' is a $(x_1|x_2|l, a)$ -well-separated instance with $y'_4 \leq x_4$. Since both $l, a \in M(\vec{y}')$, either agents 1 and 2 are served by the same facility of $M(\vec{y}')$ or agent 2 is served by the facility at l . In both cases, the social cost of $M(\vec{y}')$ becomes arbitrarily larger than $a - l$, which is the optimal social cost of the 3-Facility Location instance \vec{y}' .

Case 2: $x_4 - l \leq \sqrt{2}\lambda$.

Let $m = (r + l)/2$ be the midpoint of the x_3 -hole (l, r) in $I_3(\vec{x}_{-3})$. We consider the instance $\vec{y} = (\vec{x}_{-3}, a)$, where $a < m$ is arbitrarily close to m (i.e., $a \lesssim m$) so that $a - l < r - a$ and $d(a, x_4) \lesssim \sqrt{2}\lambda/2$. The latter is possible since x_3 is already arbitrarily close to x_4 and the right endpoint r of the hole $h = (l, r)$ lies in $(x_3, x_4]$. Since $d(x_1, x_2) = \lambda$, $d(x_2, a)$ is quite large, and $d(a, x_4) \lesssim \sqrt{2}\lambda/2$, the instance \vec{y} is $(\sqrt{2} - \delta)$ -stable, for any $\delta > 0$. By strategyproofness, $M(\vec{y})$ must place a facility at l , since $l \in I_3(\vec{x}_{-3})$ and l is the nearest endpoint of the hole $h = (l, r)$ to a .

As before, we now consider the instance $\vec{y}' = (\vec{y}_{-4}, l)$. Since $d(x_1, x_2) = \lambda$, $d(x_2, a)$ is quite large, and $d(a, l) < d(a, r) \leq \sqrt{2}\lambda/2$, the instance \vec{y}' is $(\sqrt{2} - \delta)$ -

stable, for any $\delta > 0$. Hence, by strategyproofness, $M(\vec{y}')$ must keep a facility at l , because $l \in I_4(\vec{y}'_{-4})$.

To conclude the proof, we need to construct a $(x_1|x_2|l', l' + \varepsilon)$ -well-separated instance \vec{z} with $l' \in M(\vec{z})$. Then, we can reach a contradiction to the hypothesis that M has a bounded approximation ratio, by applying Lemma 3.6.3, similarly to Case 1.

To this end, we consider the image set $I_4(\vec{y}'_{-4})$ of agent 4 in $\vec{y}'_{-4} = (x_1, x_2, a)$. Since $l \in M(\vec{y}')$, $l \in I_4(\vec{y}'_{-4})$. If $a - \varepsilon \in I_4(\vec{y}'_{-4})$, the instance $\vec{z} = (\vec{y}'_{-4}, a - \varepsilon)$ is $(x_1|x_2|a - \varepsilon, a)$ -well-separated (and thus, $(\sqrt{2} - \delta)$ -stable, for any $\delta > 0$). Moreover, by strategyproofness, $M(\vec{z})$ must place a facility at $a - \varepsilon$, because $a - \varepsilon \in I_4(\vec{y}'_{-4})$. Otherwise, there must be a hole $h' = (l', r')$ in the image set $I_4(\vec{y}'_{-4})$, with $l' > l$ (because $l \in I_4(\vec{y}'_{-4})$) and $r' < a - \varepsilon$ (because of the hypothesis that $a - \varepsilon \notin I_4(\vec{y}'_{-4})$). We consider the instance $\vec{z}' = (\vec{y}'_{-4}, l' + \varepsilon) = (x_1, x_2, l' + \varepsilon, a)$. Since $l' + \varepsilon \in (l, a)$, $d(a, l' + \varepsilon) < d(a, l) < \sqrt{2}\lambda/2$ and the instance \vec{z}' is $(\sqrt{2} - \delta)$ -stable, for any $\delta > 0$. Therefore, by strategyproofness and since $l' \in I_4(\vec{y}'_{-4})$, $M(\vec{z}')$ must place a facility at l' . We now consider the instance $\vec{z} = (\vec{z}'_{-3}, l') = (x_1, x_2, l', l' + \varepsilon)$, which is $(x_1|x_2|l', l' + \varepsilon)$ -well-separated (and thus, $(\sqrt{2} - \delta)$ -stable, for any $\delta > 0$). Moreover, by strategyproofness and since $l' \in M(\vec{z}')$, and thus, $l' \in I_3(\vec{z}'_{-3})$, $M(\vec{z})$ must place a facility at l' .

Therefore, starting from the $(\sqrt{2} - \delta)$ -stable instance \vec{y}' , with $l \in M(\vec{y}')$, we can construct a $(x_1|x_2|l', l' + \varepsilon)$ -well-separated instance \vec{z} with $l' \in M(\vec{z})$. Then, by Lemma 3.6.3, $z_4 = l' + \varepsilon \in M(\vec{z})$, because for the $(x_1|x_2|x_3, x_4)$ -well-separated instance \vec{x} , $M_3(\vec{x}) = x_4$, and \vec{z} is a $(x_1|x_2|l', l' + \varepsilon)$ -well-separated instance with $z_4 \leq x_4$. Since both $l', l' + \varepsilon \in M(\vec{z})$, the social cost of $M(\vec{z})$ is arbitrarily larger than ε , which is the optimal social cost of the 3-Facility Location instance \vec{z} . \square

3.7 A Randomized Mechanism with Constant Approximation

In this section, we show that for an appropriate stability, a simple randomized mechanism is strategyproof, can deal with singleton clusters and achieves an approximation ratio of 2.

The intuition is that the ALMOSTRIGHTMOST mechanism can be easily transformed to a randomized mechanism, using the same key properties to guarantee strategyproofness, but achieving an $O(1)$ -approximation, as opposed to $O(n)$ -approximation of ALMOSTRIGHTMOST. Specifically, RANDOM (see also Mechanism 3) again finds the optimal clusters, but then places a facility at the location of an agent selected uniformly at random from each optimal cluster. We again use cluster-separation property, as a necessary condition for stability of the optimal clustering. The stability properties required to guarantee strategyproofness are very similar to those required by ALMOSTRIGHTMOST, because the set of possible

profitable deviations is very similar for ALMOSTRIGHTMOST and RANDOM. Finally, notice that the cluster-separation property step of RANDOM (step 2) now makes use that due to Lemma 3.3.4, it must be $1.6 \cdot \max\{D(C_i), D(C_{i+1})\} < d(C_i, C_{i+1})$ for 5-stable instances.

Mechanism 3: RANDOM

Result: An allocation of k -facilities

Input: A k -Facility Location instance \vec{x} .

- 1 Find the optimal clustering $\vec{C} = (C_1, \dots, C_k)$ of \vec{x} .
- 2 **if** there are two consecutive clusters C_i and C_{i+1} with $1.6 \cdot \max\{D(C_i), D(C_{i+1})\} \geq d(C_i, C_{i+1})$ **then**
 - | **Output:** “FACILITIES ARE NOT ALLOCATED”.
- 3 **for** $i \in \{1, \dots, k\}$ **do**
- 4 | Allocate the facility to an agent c_i selected uniformly at random from the agents of cluster C_i
- 5 **end**

Output: The k -facility allocation $\vec{c} = (c_1, \dots, c_k)$.

Theorem 3.7.1. RANDOM (Mechanism 3) is strategyproof and achieves an approximation ratio of 2 for 5-stable instances of k -Facility Location on the line.

Sketch. We present here the outline of the proof. The full proof then follows. The approximation guarantee is straightforward to verify. As mentioned, the proof of strategyproofness is similar to the proof of Theorem 3.5.1. In general, we need to cover the key deviation cases, which include the following:

Case 1: why agent deviating agent $x \in C_i$ cannot gain by becoming a member of another cluster,

Case 2: or by becoming a self serving center,

Case 3: or by merging or splitting C_i .

Cases 2 and 3 can be immediately derived from the proof of Theorem 3.5.1.

The most interesting case is Case 1: x_i deviates to x' to be clustered together with agents from a different cluster of \vec{C} , in order to gain, without splitting C_i (again consider $\vec{C} = (C_1, \dots, C_k)$ the optimal clustering of original instance \vec{x} and \vec{C}' the optimal clustering of instance $\vec{x}' = (\vec{x}_{-i}, x')$).

By analyzing the expected value of agent x_i in both clusterings \vec{C} and \vec{C}' we show that in order for her to be able to gain from such a deviation, it must be $d(x', x_i) < D(C_i)$ and x' is clustered together with agents in C_{i-1} or C_{i+1} , suppose C_{i-1} w.l.o.g. Since agents in $C_i \setminus x_i$ are not split in clustering \vec{C}' , we know they form cluster $C'_{i'} \in \vec{C}'$. Hence, in this case $x \in C'_{i'-1}$. The key to the proof is to show that since $d(x', x_i) < D(C_i)$ then clustering \vec{C}' on instance \vec{x}' violates the cluster separation property verification step, either between clusters $C'_{i'}$ and $C'_{i'-1}$ or between clusters $C'_{i'-1}$ and $C'_{i'-2}$. This is also the reason why in this case the

cluster separation property verification step needs to be more precise, for 5-stable instances, as mentioned in the description of the algorithm. \square

Proof. The approximation guarantee easily follows from the fact that since a facility is uniformly at random placed over each optimal cluster, the expected cost of the sum of the cost of the agents in each cluster is 2 times their cost in the optimal clustering.

As is it always with our mechanisms, agent $x_i \in C_i$ cannot gain by moving within the range of C_i (this would only increase her utility).

Since the analysis of RANDOM is so similar to the analysis of the mechanism in Section 3.5, we skip the detailed case analysis and mention only the key deviation cases that need be covered. Specifically these include:

Case 1: why agent $x_i \in C_i$ cannot gain by becoming a member of another cluster,

Case 2: or by becoming a self serving center

Case 3: or by merging or splitting C_i .

Without loss of generality, consider the deviating agent to be the edge agent $x_{i,l} \in C_i$, declaring location x' creating instance $\vec{x}' = (\vec{x}_{-x_{i,l}}, x')$ with optimal clustering \vec{C}' . If our results stand for her, they easily transfer to all agents in C_i . C_i contains n agents, including $x_{i,l}$. For simplicity, without loss of generality we index these agents from left to right, excluding $x_{i,l}$, such as $x_{i,l} \leq x_{i,1} \leq \dots \leq x_{i,n-1}$, where $x_{i,1} = x_{i,l+1}$ and $x_{i,n-1} = x_{i,r}$. Now for simplicity, we represent $d(x_{i,l}, x_{i,j})$ by $d_{i,j}$. Of course $d_{i,l} = 0$. We define as X_i the discrete random variable that takes values from sample space $\{d_{i,l}, d_{i,1}, d_{i,2}, \dots, d_{i,n-1}\}$ uniformly at random. That is, X_i represents the cost agent $x_{i,l}$ experiences if she is served by the facility placed in C_i by the mechanism. Then, the expected cost of $x_{i,l}$ should she not deviate is:

$$\mathbb{E}(X_i) = \frac{0 + d_{i,1} + \dots + d_{i,n-1}}{n}$$

That is, since for any agent $x_j \notin C_i$, $d(x_j, x_{i,l}) > D(C_i) = d_{i,n-1}$ by Lemma 3.3.4.

Now, for Case 1, “why agent $x \in C_i$ cannot gain by becoming a member of another cluster”. Notice that this is the case where agents in C_i are not merged or splitted in \vec{C}' . With some abuse of notation, this allows us to refer to the cluster containing agents in $C_i \setminus x_{i,l}$ in \vec{C}' of \vec{x} as C'_i . C'_{i-1} then is the set of agents belonging to the cluster immediately to the left of C'_i (i.e. the rightmost agent of C'_{i-1} , excluding x' , is $x_{i-1,r}$). Consider a deviation x' that places the deviating agent in cluster C'_{i-1} after step 1 of the mechanism. Again for simplicity consider $d(x', x_{i,l}) = c$ and we index agents in C_{i-1} inversely, such that $x_{i-1,\hat{1}} \geq x_{i-1,\hat{2}} \geq \dots \geq x_{i-1,\hat{n}'}$ (meaning that now $x_{i-1,r} = x_{i-1,\hat{1}}$, $x_{i-1,r-1} = x_{i-1,\hat{2}}$ etc.) where $|C_{i-1}| = n'$. Equivalently we set $d(x_{i,l}, x_{i-1,\hat{j}}) = d_{i-1,j}$. By Corollary 3.3.5, we have $d_{i,1} \leq d_{i,2} \leq \dots \leq d_{i,n-1} \leq d_{i-1,1} \leq \dots \leq d_{i-1,n'}$. Now we define

uniform random variable X'_i with sample space $\{d_{i,1}, \dots, d_{i,n-1}\}$ (see that $d_{i,l}$ is now absent) and random variable X'_{i-1} with sample space $\{c, d_{i-1,1}, \dots, d_{i-1,n'}\}$. Now X'_i represents the cost of $x_{i,l}$ should she be served by the facility placed in C'_i of the changed instance (which now doesn't include her) and X'_{i-1} her cost should she be served by the facility placed at C'_{i-1} (which now includes her false declared location). The expected cost of $x_{i,l}$ now becomes $\mathbb{E}(\min\{X'_i, X'_{i-1}\})$.

But, since $d_{i,1} \leq d_{i,2} \leq \dots \leq d_{i,n-1} \leq d_{i-1,1} \leq \dots \leq d_{i-1,n'}$, unless $d(x', x_{i,l}) < d_{i,n-1} = D(C_i)$, we have that:

$$\mathbb{E}(\min\{X'_i, X'_{i-1}\}) = \mathbb{E}(X'_i) = \frac{d_{i,1} + \dots + d_{i,n-1}}{n-1} > \mathbb{E}(X_i)$$

That means that $x_{i,l}$ cannot gain by this deviation unless x' both belongs in C'_{i-1} and $d(x', x_{i,l}) < D(C_i)$. All we need to show now is that any such situation would result in a violation of the inter-cluster distance between C'_{i-1} and C'_i or between C'_{i-1} and C'_{i-2} , guaranteed by the cluster-separation property and hence it would be caught by the mechanism's cluster-separation property verification step.

Specifically consider the distance of $x_{i,l}$ to her center c_i of C_i in the optimal clustering. We know that it must be $d(C_{i-1}, C_i) \geq D(C_i) \cdot 1.6$, by Lemma 3.3.4, for the given stability factor of 5. But in order for this distance to be tight, it must be that $d(x_{i,l}, c_i) = 0.4 \cdot D(C_i)$ (see factor c of proof of Lemma 3.3.4 -due to stability properties, if $d(x_{i,l}, c_i) < 0.4 \cdot D(C_i)$ or $> 0.4 \cdot D(C_i)$, $d(C_{i-1}, C_i)$ grows larger than $D(C_i) \cdot 1.6$). Furthermore, in order for this distance to be tight, it must also be $d(c_{i-1}, x_{i-1,r}) < 0.4 \cdot D(C_i)$ (since by stability it must be $d(C_{i-1}, C_i) = d(x_{i-1,r}, x_{i,l}) > (\gamma - 1)d(x_{i-1,r}, c_{i-1})$).

Now, since it must be $d(x', x_{i,l}) < D(C_i)$ it will be $d(x', c_i) < 1.4D(C_i)$ and $d(x', x_{i-1,r}) > 0.6D(C_i)$ (since $d(C_i, C_{i-1}) > 1.6D(C_i)$ by Lemma 3.3.4). Finally we distinguish between two cases:

Case 1: $c_{i-1} \in C'_{i-1}$. Now notice that $d(x_{i,l}, c_{i-1}) > 5d(x_{i,l}, c_i)$ so $d(x_{i,l}, c_{i-1}) > 2D(C_i)$. Then $D(C'_{i-1}) \geq d(c_{i-1}, x') > D(C_i)$ (since $d(x', x_{i,l}) < D(C_i)$). But then $d(C'_{i-1}, C'_i) \leq d(x', c_i) \leq 1.4 \cdot D(C_i)$ which means that the cluster separation verification property of step 2 would be violated.

Case 2: $c_{i-1} \notin C'_{i-1}$. Then, in this edge case we notice it would be $d(C'_{i-1}, C'_{i-2}) \leq d(c_{i-1}, x_{i-1,r}) \leq 0.4D(C_i)$. But $D(C'_{i-1}) \geq d(x_{i-1,r}, x') \geq 0.6D(C_i)$. Hence the verification property of step 2 is again violated between C'_{i-1} and C'_{i-2} .

All we have to do to finish, is note that as c_i moves to the right or to the left, $d(C_{i-1}, C_i)$ grows by a multiplicative factor $\gamma - 1$ ($=4$) of $d(x_{i,l}, c_i)$ (see proof of Lemma 3.3.4) and $d(x_{i,l}, c_{i-1})$ by a multiplicative factor of 5 (remember, it must be both $d(x_{i,l}, c_{i-1}) > 5d(x_{i,l}, c_i)$ and $d(x_{i,r}, c_{i-1}) > 5d(x_{i,r}, c_i)$). Which means that the above inequalities will still hold. ³

³Notice here that while this property was a must-have for ALMOSTRIGHTMOST to work

For Case 2, why agent $x_{i,l}$ cannot gain by becoming a self serving cluster, we simply notice the following: her cost, should she not deviate, is at most $D(C_i)$ (see expected value from previous case). But, from Lemma 3.3.8 we know that $x_{i,l}$ must deviate by at-least $\geq D(C_i)$, for a stability factor of 5. So she cannot gain from this deviation⁴.

For Case 3, it is not hard to see that by merging all the agents in C_i with agents $\notin C_i$, her expected cost can only increase. Furthermore, splitting the agents in C_i would cause the cluster-separation property verification step to identify the split (see the proof of the strategyproofness of the ALMOSTRIGHTMOST mechanism, in Section 3.5) and remove all agents of C_i from the game. \square

3.A Proof of lemma 3.3.8.

We first present the outline of the proof and then the proof follows. We do this because despite the mostly relatively straight forward arguments used in the proof, due to the delicate formalization required in order to formally describe all the mentioned conditions, the proof gains a good amount of descriptive length. We consider random agent $x_i \in C_i$ of instance \vec{x} with optimal clustering $\vec{C} = (C_1, \dots, C_k)$, deviating to location x' creating instance $\vec{x}' = (x_{-i}, x')$.

Initially we show that due to the large distance between clusters C_i and C_j with $i \neq j$, guaranteed by Lemma 3.3.4 for 5-stable instances, we need only study the cases where $x' \in (x_{i-1,r}, x_{i,l})$ and $x' \in (x_{i,r}, x_{i+1,l})$ and in the optimal clustering \vec{C}' of instance \vec{x}' no agent in $\vec{x}' \setminus C_i$ is served together with any agent in C_i ⁵, as in all other cases either x' is not a singleton in \vec{C}' or $d(x', x_i) > D(C_i)$.

The rest of the proof follows the logic of the proof of Theorem 3.4.1 (which follows), tailored to this specific case. More specifically, given the observation above, we notice the following: In alternative clustering \vec{C}'' in which we forcefully place two facilities serving only agents in C_i (optimally with regards to serving agents in C_i), and serve the remaining agents, $\vec{x} \setminus C_i$, optimally with $k-2$ facilities, the cost agents in $\vec{x} \setminus C_i$ experience in clustering \vec{C}'' is the same cost agents in $\vec{x}' \setminus C_i \cup x'$ experience in clustering \vec{C} (notice that the sets $\vec{x} \setminus C_i$ and $\vec{x}' \setminus C_i \cup x'$). Now, the cost of agents in C_i in clustering \vec{C}'' is at least $D(C_i)/2$ smaller than it is in \vec{C}' (that is since we can always place the facility to the edge agent further from c_i - see proof of Theorem 3.4.1). But since \vec{C}'' is not optimal for \vec{x} this means that

i.e. the mechanism wouldn't work if x' both belongs in C'_{i-1} and $d(x', x_{i,l}) < D(C_i)$, here this might not be the case. We can easily see this guarantees strategyproofness, but it might not be necessary which means the mechanism may work for smaller stability factors.

⁴Again, while this property guarantees strategyproofness, it might not be necessary for example, we see that in one of the bad edge cases, where all agents of C_i are gathered on $x_{i,r}$, with $c_i = x_{i,r}$ a stability of 3 would suffice to guarantee that $x_{i,l}$ needs to deviate by at-least $D(C_i)$ to become a self-serving cluster.

⁵Note here that we refer to the group of agents that belong in cluster C_i of the optimal clustering of instance \vec{x} . This group is well defined for instance \vec{x}' as well.

agents in $\vec{x} \setminus C_i$ experience an increase in cost larger than $D(C_i)/2$ in clustering \vec{C}'' when compared to clustering \vec{C} . For brevity we symbolize this cost increase as cst , so we say $cst > D(C_i)/2$.

We now we consider the 4-perturbation of instance \vec{x} in which all distances among agents to the left and to the right of C_i are shrunk by a factor of 4. By stability we know that the optimal clustering of the perturbed instance should be the same as the optimal clustering of the original! But in the perturbed instance all costs of agents in $\vec{x} \setminus C_i$ are divided by 4 in both clusterings \vec{C} and \vec{C}'' while the costs of agents in C_i remain the same. So, in order for \vec{C}'' to be sub-optimal in the perturbed instance it must be $cst/4 > D(C_i)/2$ which means $cst > 2D(C_i)$. But serving agent x' of \vec{x}' by c_i has cost at most $2D(C_i)$ if $d(x', x_i) < D(C_i)$ since $d(x_i, c_i) < D(C_i)$. This means that clustering \vec{C}' cannot be optimal for \vec{x} .

Proof. We want to show the lemma for any γ -stable instance for $\gamma \geq 5$.

We prove the lemma for random agent $x_j \in C_i$ for some cluster C_i in optimal clustering $\vec{C} = (C_1, \dots, C_k)$ of the γ -stable instance \vec{x} . Consider that the agent declares false location x' providing input profile $\vec{x}' = (\vec{x}_{-j}, x')$ to the mechanism in order to become a singleton cluster. That is, if the optimal clustering of instance \vec{x}' is \vec{Y} then x' is a single agent cluster in \vec{Y} .

We first study the case where $|C_i| = 2$. But then, from Lemma 3.3.7 we know that for any γ -stable instance for $\gamma \geq 3$ agent $x_j \in C_i$ of optimal clustering \vec{C} must deviate by at least his distance to C_i 's center in order to become a singleton cluster in \vec{Y} . I.e. it must be $d(x', x_j) > d(x_j, c_i) = D(C_i)$, so the lemma stands for this case.

For the most general case, $|C_i| \geq 3$ we start with some observations. By Lemma 3.3.4 we know that for any two clusters C_i and C_j of optimal clustering $\vec{C} = (C_1, \dots, C_k)$ of \vec{x} we have $d(C_i, C_j) > \left(\frac{\gamma-1}{2\gamma}\right) \max\{D(C_i), D(C_j)\}$. For $\gamma \geq 5$ that is:

$$d(C_i, C_j) > 1.6 \max\{D(C_i), D(C_j)\}. \quad (3.11)$$

To begin, we notice the following claim:

Claim 1. *Agent x_i cannot declare a false location x' with $x_{i,l} \leq x' \leq x_{i,r}$ in such a way that x' is a singleton cluster in \vec{Y} .*

We can easily see the validity of the claim, since by optimality (also see proof of Theorem 3.4.1) $x_j \in C_i$ cannot change the optimal clustering by deviating within the bounds of cluster C_i , i.e. if $x_{i,l} \leq x' \leq x_{i,r}$. Hence it must be $x' \notin [x_{i,l}, x_{i,r}]$. Even so, for completeness, we provide a proof of the claim, tailored to the case of 5-stable instances, after the proof of the lemma.

In addition, we notice that if $x' \leq x_{i-1,r}$ or $x' \geq x_{i+1,l}$ then the lemma trivially stands, again by Equation 3.11 (I.e. in this case it would be $d(x', x_j) > 1.6D(C_i)$).

This means that we need only study the cases where $x' \in (x_{i-1,r}, x_{i,l})$ or $x' \in (x_{i,r}, x_{i+1,l})$ and $d(x', x_j) < D(C_i)$ (and show that x' cannot become a singleton cluster in \vec{Y} in these cases).

Suppose, contrary to the lemma's claim, that agent declares location $x' \in (x_{i-1,r}, x_{i,l})$ with $d(x', x_j) < D(C_i)$ such that x' is a singleton in \vec{Y} (the other case, $x' \in (x_{i,r}, x_{i+1,l})$, is symmetrical). Then we notice the following three properties for optimal clustering \vec{Y} of instance \vec{x}' :

Property 1: In \vec{Y} there is a facility among agents in $C_i \setminus x_j$.

Property 2: In \vec{Y} no agent “to the left” of cluster C_i (i.e. by an agent in some cluster C_l for $l < i$, of \vec{C}) is served by an agent in $C_i \setminus x_j$.

Property 3: In \vec{Y} no agent “to the right” of cluster C_i (i.e. by an agent in some cluster C_l for $l > i$, of \vec{C}) is served by an agent in $C_i \setminus x_j$.

The imminent conclusion from Properties 1, 2 and 3 is the following: Consider instance $\vec{x} \setminus C_i$ and its optimal $k - 2$ -clustering \vec{C}_{-2} . Then $cost(\vec{x} \setminus C_i, \vec{Y}) = cost(\vec{x} \setminus C_i, \vec{C}_{-2})$ ⁶. We provide short proofs for each one of these three properties right after the proof of the lemma.

We are now ready to complete the proof. In order to do so we bound the extra cost experienced by agents in $\vec{x} \setminus C_i$ in the possible re-clustering after x_i 's deviation, i.e. $cost(\vec{x} \setminus C_i, \vec{Y}) - cost(\vec{x} \setminus C_i, \vec{C})$. We do this by considering the following alternative clustering C' of instance \vec{x} : serve agents in C_i using two facilities, optimally and agents in $\vec{x} \setminus C_i$ using the remaining $k - 2$ facilities optimally. So in \vec{C}' we have:

$$cost(C_i, \vec{C}') \leq cost(C_i, \vec{C}) - \frac{D(C_i)}{2}, \quad (3.12)$$

since placing the second facility placed among agents in C_i to the edge-agent further away from c_i reduces the cost by at least $\frac{D(C_i)}{2}$.

But since \vec{C} is optimal in \vec{x} and hence \vec{C}' is not, it must be:

$$cost(\vec{x} \setminus C_i, \vec{C}') - cost(\vec{x} \setminus C_i, \vec{C}) > \frac{D(C_i)}{2} \quad (3.13)$$

Otherwise it would be $cost(\vec{x}, \vec{C}') < cost(\vec{x}, \vec{C})$. Now notice that properties 1, 2 and 3 mean that agents in $\vec{x} \setminus C_i$ are clustered in exactly the same way in \vec{C}' as in \vec{Y} . That means that:

$$cost(\vec{x}' \setminus \{C_i \cup x'\}, \vec{Y}) = cost(\vec{x} \setminus C_i, \vec{C}') \quad (3.14)$$

and that no agent to the left of C_i is clustered together with any agent to the right of C_i in \vec{C}' .

⁶For a description of this notation, of the form $cost(\vec{x}, \vec{C})$, see proof of Theorem 3.4.1

The last observation means if we consider a 4-perturbation of instance \vec{x} , instance \vec{x}_p , in which we divide all distances among agents between $[x_l, x_{i-1,r}]$ and agents between $[x_{i+1,l}, x_r]$, where x_l and x_r the leftmost and rightmost agents of the instance equivalently we have that:

$$\text{cost}(\vec{x}_p \setminus C_i, \vec{C}') - \text{cost}(\vec{x}_p \setminus C_i, \vec{C}) = \frac{\text{cost}(\vec{x} \setminus C_i, \vec{C}') - \text{cost}(\vec{x} \setminus C_i, \vec{C})}{4}$$

But in x_p the distances among agents in C_i remain unaffected which means that in x_p , Equation 3.12 still stands. This means, that since the instance is 5-stable, clustering \vec{C}' must still be sub-optimal in \vec{x}_p and hence it must be

$$\begin{aligned} \text{cost}(\vec{x}_p \setminus C_i, \vec{C}') - \text{cost}(\vec{x}_p \setminus C_i, \vec{C}) &> \frac{D(C_i)}{2} \Rightarrow \\ \frac{\text{cost}(\vec{x} \setminus C_i, \vec{C}') - \text{cost}(\vec{x} \setminus C_i, \vec{C})}{4} &> \frac{D(C_i)}{2} \Rightarrow \\ \text{cost}(\vec{x} \setminus C_i, \vec{C}') - \text{cost}(\vec{x} \setminus C_i, \vec{C}) &> 2D(C_i). \end{aligned} \quad (3.15)$$

Noticing again that by Equation (3.14), $\text{cost}(\vec{x}' \setminus \{C_i \cup x'\}, \vec{Y}) = \text{cost}(\vec{x} \setminus C_i, \vec{C}')$ and by Equation (3.15) and $\text{cost}(\vec{x} \setminus C_i, \vec{C}) = \text{cost}(\vec{x}' \setminus \{C_i \cup x'\}, \vec{C})$ we have

$$\text{cost}(\vec{x}' \setminus \{C_i \cup x'\}, \vec{Y}) - \text{cost}(\vec{x}' \setminus \{C_i \cup x'\}, \vec{C}) > 2D(C_i).$$

Finally, since $d(x', x_i) < D(C_i)$,

$$\text{cost}(\{C_i \cup x' \setminus x_i\}, \vec{C}) - \text{cost}(\{C_i \cup x' \setminus x_i\}, \vec{Y}) < 2D(C_i),$$

since $d(x_i, c_i) \leq D(C_i)$. By adding the last two equations we get that $\text{cost}(\vec{x}', \vec{Y}) > \text{cost}(\vec{x}', \vec{C})$ which means that \vec{Y} is not optimal. \square

We now present the proofs of **Claim 1** and **Properties 1, 2 and 3**, used in the main proof of Lemma 3.3.8.

Of Claim 1. Consider $x'_{i,l}$ and $x'_{i,r}$ to be the leftmost and rightmost agents of $C_i \setminus x_j$ (i.e. if $x_j \neq x_{i,r}, x_{i,l}$ then $x_{i,r} = x'_{i,r}$ and $x_{i,l} = x'_{i,l}$).

Contrary to the claim, suppose $x'_{i,l} \leq x' \leq x'_{i,r}$ and x' is a singleton cluster \vec{Y} . Since x' is a singleton and $x'_{i,l}$ and $x'_{i,r}$ are to her left and right side equivalently, $x'_{i,l}$ and $x'_{i,r}$ cannot be served by the same facility in \vec{Y} (since clustering \vec{Y} is optimal for \vec{x}'). This means that either $x'_{i,l}$ or $x'_{i,r}$ is served by an agent in $\vec{x} \setminus C_i$ or there are two facilities among agents in $C_i \setminus x_j$ in \vec{Y} . Both of these cases are infeasible though. For the first one, suppose that $x_{i,r}$ is not served by an agent in $C_i \setminus x_j$. By Equation 3.11 that means that the cost of serving $x_{i,r}$ is at-least $1.6D(C_i)$. But since $x'_{i,l} \leq x' \leq x'_{i,r}$ x' , $d(x', x_{i,r}) < D(C_i)$ so \vec{Y} could not be optimal in \vec{x}' . For

the latter case (\vec{Y} places two facilities among agents in $C_i \setminus x_j$) we see that if \vec{Y} is optimal for \vec{x}' then the optimal $(k-1)$ -clustering of instance (\vec{x}_{-j}) would place two facilities among agents in $C_i \setminus x_j$ (since x' is a singleton removing her and one facility from the instance should yield the exact same clustering for the rest of the agents). But then, since in \vec{C} there is only one facility among agents in C_i , \vec{C} could not be optimal for instance \vec{x} (because if the optimal $(k-1)$ -clustering of instance (\vec{x}_{-j}) places two facilities among agents in set $C_i \setminus x_j$ then the optimal k -clustering of instance \vec{x} should place at least as many among agents in C_i), which is a contradiction.

Finally we notice that if $x_j = x_{i,l}$, x' cannot become a singleton in \vec{Y} if $x' \in [x_{i,l}, x'_{i,l}]$ since the cost serving agent x_j by c_i in that interval is only decreased (in relation to the cost of serving her by c_i in \vec{x} - she's getting closer to her serving facility). Similarly for the case of $x_j = x_{i,r}$ moving in interval $[x'_{i,r}, x_{i,r}]$. The above mean that agent x_j cannot become a singleton cluster by moving within the bounds of C_i (i.e. if x' is a singleton in \vec{Y} it must be $x' \notin [x_{i,l}, x_{i,r}]$), which is the claim. □

Of Property 1. We know that $|C_i \setminus x_j| \geq 2$. Furthermore, since $d(x', x_j) < D(C_i)$ we have that $d(x', c_i) < 2D(C_i)$. But if there is no facility among agents in $C_i \setminus x_j$ that means that these agents are all served by a facility placed in a location x_l with $x_l \in C_l$ with $l \neq i$. But, again by Equation 3.11 that would mean that

$$\text{cost}(C_i \setminus x_j, \vec{Y}) > 2 * 1.6D(C_i) + \text{cost}(C_i \setminus x_j, \vec{C}) \quad (3.16)$$

(since $|C_i \setminus x_j| \geq 2$, $d(C_i, C_l) \geq 1.6D(C_i)$). Furthermore, since agents in $\vec{x} \setminus C_i$ are served by the same number of facilities in \vec{Y} as in \vec{C} , but also have to serve agents in $C_i \setminus x_j$ in \vec{Y} (i.e. the placement of the $(k-1)$ facilities among agents in $\vec{x} \setminus C_i$ is not optimal in \vec{Y} as it is in \vec{C} , for these agents) we have

$$\text{cost}(\vec{x} \setminus C_i, \vec{Y}) \geq \text{cost}(\vec{x} \setminus C_i, \vec{C}). \quad (3.17)$$

Hence, by adding (3.16) and (3.17) we have that :

$$\text{cost}(\vec{x} \setminus x_j, \vec{Y}) > 2 * 1.6D(C_i) + \text{cost}(\vec{x} \setminus x_j, \vec{C})$$

By remembering that $\vec{x}' = (\vec{x}_{-j}, x')$ and in \vec{Y} x' is a singleton cluster (i.e. has cost 0) the above becomes:

$$\text{cost}(\vec{x}', \vec{Y}) > 2 * 1.6D(C_i) + \text{cost}(\vec{x} \setminus x_j, \vec{C}) \quad (3.18)$$

But, alternative clustering \vec{C}' for \vec{x}' in which we serve all agents as we do in \vec{C} and also serve location x' by c_i has cost

$$\text{cost}(\vec{x}', \vec{C}') \leq \text{cost}(\vec{x} \setminus x_j, \vec{C}) + 2D(C_i), \quad (3.19)$$

since $d(x', c_i) < 2D(C_i)$.

This means that, by (3.18) and (3.19) $\text{cost}(\vec{x}', \vec{Y}) > \text{cost}(\vec{x}', \vec{C}')$ which means that clustering \vec{Y} would be sup-optimal for instance \vec{x}' , which is a contradiction.

Notice that by Observation 3.3.6, property 1 means that no agents in $C_i \setminus x_j$ are served by an agent not in C_i in \vec{Y} . □

Of Property 2. Property 2 is trivial: since $x' \in (x_{i-1,r}, x_{i,l})$ and x' forms a singleton cluster in \vec{Y} , by optimality no agent to the left of x' is clustered together with agents to the right of x' . □

Of Property 3. Initially, for property 3 we notice the following: At most 1 agent in C_{i+1} can be clustered together with agents in C_i in \vec{Y} . Otherwise, due to the distance between C_i and C_{i+1} , clustering \vec{Y} would be sub-optimal (using the same reasoning as for property 1). Obviously, due to optimality, this agent can only be $x_{i+1,l}$.

We now consider the structure of cluster C_i in relation to agent $x_{i+1,l}$. Specifically, by Equation (3.11) it must be

$$d(x_{i,r}, x_{i+1,l}) > 1.6 \cdot D(C_i), \quad (3.20)$$

since $d(x_{i,r}, x_{i+1,l}) = d(C_i, C_{i+1})$. By looking at the proof of Lemma 3.3.4 we see that the smallest possible distance between C_i and C_{i+1} is achieved when $d(x_{i,l}, c_i) = \frac{D(C_i)}{c}$ for $c = \frac{2\gamma^2}{\gamma^2 + \gamma} \Rightarrow \frac{1}{c} = 0.6$ for $\gamma = 5$. This means that since agent x_j deviates to the left in this case, by at most $D(C_i)$, it must be

$$d(x', c_i) \leq 1.6D(C_i), \quad (3.21)$$

in the edge case. Furthermore, by Observation 3.3.6, since $x_{i+1,l}$ is not served by an agent in C_{i+1} there is no facility among agents in C_{i+1} in \vec{Y} . I.e. all agents in $C_{i+1} \setminus x_{i+1,l}$ are served by a facility placed on $[x_{i+2,l}, x_n]$ where x_n the rightmost agent location in the instance. But, by Lemma 3.3.4, if $x_{i+1,l}$ is served by $c_{i+1} \in C_{i+1}$ in \vec{C} , $d(C_{i+1}, C_{i+2}) > 1.6D(C_{i+1}) \geq 1.6d(x_{i+1,l}, c_{i+1})$ and so, it is

$$\text{cost}(x_{i+1,o}, \vec{Y}) \geq d(x_{i+1,o}, x_{i+2,l}) \geq 1.6d(x_{i+1,l}, c_{i+1}), \quad (3.22)$$

for every $x_{i+1,o} \in C_{i+1} \setminus x_{i+1,l}$.

Now we are able to show that clustering \vec{Y} cannot be optimal for instance \vec{x}' in the edge case. We will compare it with clustering \vec{C} (where every agent is served by the same facility as in clustering \vec{C} and x' is served by c_i). We have the following:

$$\text{cost}(\vec{x}' \setminus \{x' \cup C_i \cup C_{i+1}\}, \vec{Y}) \geq \text{cost}(\vec{x}' \setminus \{x' \cup C_i \cup C_{i+1}\}, \vec{C}),$$

by optimality. Furthermore,

$$\text{cost}(C_{i+1}, \vec{Y}) \geq \text{cost}(C_{i+1}, \vec{C}) - d(x_{i+1,l}, c_i) + 1.6d(x_{i+1,l}, c_{i+1}) + 1.6D(C_i),$$

by optimality and equations (3.20) and (3.22). Also,

$$\text{cost}(C_i \setminus x, \vec{Y}) \geq \text{cost}(C_i \setminus x, \vec{C}),$$

by optimality. Finally,

$$\text{cost}(x', \vec{Y}) + 1.6D(C_i) > \text{cost}(x', \vec{C}),$$

by equation (3.21).

By adding we get $\text{cost}(\vec{x}', \vec{Y}) > \text{cost}(\vec{x}', \vec{C})$ which means that \vec{Y} is sub-optimal for instance \vec{x}' . All we need to finalize this observation is realize that as we move away from the edge case, the above inequalities become easier to satisfy. Specifically if C_i had center $c'_i < c_i$ we see that factor 1.6 of inequality (3.21) decreases while $d(C_i, C_{i+1})$ increases. If $c'_i > c_i$ the same factor of inequality (3.21) may increase by $|c'_i - c_i|$, but then $d(C_i, C_{i+1})$ increases by at least $\frac{\gamma^2+1}{\gamma+1} \cdot |c'_i - c_i| > 4.3|c'_i - c_i|$ (since $d(x_{i,l}, x_{i+1,l}) > \frac{\gamma^2+1}{\gamma+1}d(x_{i,l}, c_i)$ - see proof of Lemma 3.3.4), hence maintaining $\text{cost}(\vec{x}', \vec{Y}) > \text{cost}(\vec{x}', \vec{C})$. □

3.B Proofs of Auxiliary Lemmas Used in the Proof of Theorem 3.6.5

For completeness, we restate the proofs of the auxiliary lemmas with the properties of well-separated instances adapted from [112] and used in the proof of Theorem 3.6.5.

Before we proceed with the proofs of the auxiliary lemmas, we need the following basic fact about the facility allocation of any deterministic strategyproof mechanism.

Lemma 3.B.1 (Proposition 2.1, [112]). *Let M be a deterministic strategyproof with a bounded approximation ratio of $\rho \geq 1$ for $\sqrt{2}$ -stable instances of k -Facility location on the line. For any $(k+1)$ -location instance \vec{x} with $x_1 \leq x_2 \leq \dots \leq x_{k+1}$, $M_1(\vec{x}) \leq x_2$ and $M_k(\vec{x}) \geq x_k$.*

Proof. We show it for $M_1(\vec{x}) \leq x_2$, the other case is symmetric. Suppose $x_2 < M_1(\vec{x})$. Then the agent in x_1 has the incentive to deviate to location x_2 , since $M_1(\vec{x}_{-1}, x_2) = x_2$ due to the bounded approximation of M (i.e., in (\vec{x}_{-1}, x_2) , M allocates k facilities to k different locations). Notice that (\vec{x}_{-1}, x_2) is γ -stable for any $\gamma \geq 1$. □

3.B.1 The Proof of Lemma 3.6.1

Proof. Since M has a bounded approximation, the isolated pair x_k and x_{k+1} must be served by the same facility $M_k(\vec{x})$. By Lemma 3.B.1, we know that $M_k(\vec{x}) \geq x_k$. Then, it must also be $M_k(\vec{x}) \leq x_{k+1}$. Otherwise, like in Lemma 3.B.1, agent x_k could declare location x_{k+1} and decrease her cost, since $M_k(\vec{x}_{-k}, x_{k+1}) = x_{k+1}$ by the bounded approximation of M . Again, the instance (\vec{x}_{-k}, x_{k+1}) is arbitrarily stable. \square

3.B.2 The Proof of Lemma 3.6.2

We can now proceed to the proofs of the auxiliary lemmas, Lemma 3.6.3 and Lemma 3.6.2, which refer to the movement of isolated pairs. We only present the proof of Lemma 3.6.2 here. The proof of Lemma 3.6.3 is fully symmetric.

The proof shown here, refers to 2-Facility Location on well separated instances with 3 agents. All arguments as well as the stability factor of the instance only depend on the well separated property of the rightmost pairs of agents as well as their distance from the third agent from the right. That is, that since in all instances studied in the proof we only change distance between the agents of the isolated, rightmost pair, in the range $(0, d(x_1, x_2)/r)$ and only increase the distance between the isolated pair and the leftmost agent x_1 , any instance with a large enough distance between x_1 and x_2 , i.e. for which $d(x_1, x_2) > \gamma \cdot \rho d(x_2, x_3)$ will be γ -stable in all parts of the proof. In that way it is easy to verify that the arguments presented here extend to $(x_1 | \dots | x_{k-1} | x_k, x_{k+1})$ -well separated and stable instances of at least a specific minimum distance $d(x_{k-1}, x_k)$.

Consider M to be a deterministic, strategyproof, anonymous and bounded approximation mechanism, with approximation ration of at most ρ , for 2-facility location. We will work on instance \vec{x} with three agents $x_1 < x_2 < x_3$ which is $(x_1 | x_2, x_3)$ -well separated.

The proof of Lemma 3.6.2 directly follows from the following propositions, originally established in [112, Appendix A].

Proposition 3.B.2. *Consider $(x_1 | x_2, x_3)$ -well separated, stable instance \vec{x} for which $M_2(\vec{x}) = x_2$. Then for instance $\vec{x}' = (\vec{x}_{-2}, x'_2)$ where $x_2 \leq x'_2 \leq x_3$ it will be $M_2(\vec{x}') = x'_2$*

Proof. Notice that since $d(x'_2, x_3) < d(x_2, x_3)$ instance \vec{x}' is still $(x_1 | x_2, x_3)$ -well separated. Furthermore, since x_1 is allocated a facility (by the ρ -approximation property of the instance), \vec{x}' is at least as stable as \vec{x} since the distance between the isolated pair is shortened and their distance from x_1 has grown. All that needs to be shown is that image set $I_2(\vec{x}'_{-2})$ includes the interval $[x_2, x_3]$. Since x_2 is allocated a facility, we know $x_2 \in I_2(\vec{x}'_{-2})$. Furthermore, by the bounded approximation property of M $x_3 \in I_2(\vec{x}'_{-2})$. Assume there is a hole $(l, r) \in I_2(\vec{x}'_{-2})$ with $x_2 \leq l < r \leq x_k$. Consider location $y \in (l, r)$ such that $d(y, l) < d(y, r)$.

By strategyproofness $l \in M(\vec{x}_{-2}, y)$. But then, by Lemma 3.B.1 we have that $F_2(\vec{x}_{-j}, y) > y$ which contradicts M 's bounded approximation ratio, since the two agents of the isolated pair of (\vec{x}_{-j}, y) are served by different facilities. \square

Proposition 3.B.3. *Consider $(x_1|x_2, x_3)$ -well separated stable instance \vec{x} for which $M_2(\vec{x}) = x_2$. Then for every $(x_1|x_2, x'_3)$ -well separated instance $\vec{x}' = (\vec{x}_{-3}, x'_3)$, if \vec{x}' is also well separated, $M_2(\vec{x}') = x_2$.*

We notice that in that case, the distance between the agents of the isolated pair might grow a from \vec{x} to \vec{x}' . Since the proof of this proposition uses instances where the distance of the isolated pair varies from ϵ to $d(x_1, x_2)/\rho$ the proposition stands for stable instances only if all possible $(x_1|x_2, x'_3)$ -well separated instances $\vec{x}' = (\vec{x}_{-3}, x'_3)$ are well separated. It is easy to see, that since in all these instances it must be $d(x_2, x_3) < d(x_1, x_2)/\rho$ then for a large enough distance $d(x_1, x_2)$ (i.e. $d(x_1, x_2) > \gamma \cdot \rho d(x_2, x_3)$) \vec{x}' is always stable. We show the following proof considering that we have made this assumption.

Proof. Since $M_2(\vec{x}) < x_3$, we know that $x_3 \notin I_3(x_{-3})$. So, there is a x_3 -hole $(l, r) \in I_3(x_{-3})$. Since $M_2(\vec{x}) = x_2$, $l = x_2$ and $r > 2x_3 - x_2$ (by strategyproofness). By strategyproofness, if $x'_3 < (r + l)/2$ (for $x_2 < x'_3$ for well separated instance \vec{x}'), $M_2(\vec{x}') = x_2$.

To finish, we show that there are no $(x_1|x_2, x'_3)$ -well separated instances $\vec{x}' = (x_{-3}, x'_3)$ with $x'_3 \geq (r + l)/2$ and $M_2(\vec{x}') \neq x_2$. Again, we reach a contradiction by assuming that there is a point $y \geq (r + l)/2$ for which (x_{-3}, y) is a $(x_1|x_2, y)$ -well separated instance with $M_2((x_{-3}, y)) \neq x_2$. If such a y exists, then there exists $x'_k \in [(r + l)/2, r)$ for which $\vec{x}' = (x_{-3}, x'_3)$ is a $(x_1|x_2, x'_3)$ -well separated. But then, $M_2(\vec{x}') = r > x'_3$ (by strategyproofness, because x'_3 is closer to r than to l). Since \vec{x}' is $(x_1|x_2, x'_3)$ -well separated this contradicts lemma 3.6.1 which dictates that it must be $M_2(x_{-3}, x'_3) \in [x_2, x'_3]$. \square

Proposition 3.B.4. *Consider $(x_1|x_2, x_3)$ -well separated stable instance \vec{x} for which $M_2(\vec{x}) = x_2$. Then for every $(x_1|x'_2, x'_3)$ -well separated instance $\vec{x}' = (\vec{x}_{-\{2,3\}}, x'_2, x'_3)$, with $x_2 < x'_2 < (x_2 + x_3)/2$, if \vec{x}' is also well separated, $M_2(\vec{x}') = x_2$.*

Note that, as for proposition 3.B.3 the restriction that \vec{x}' is also γ -stable is equivalent to $d(x_1, x_2) > \gamma \cdot \rho d(x_2, x_3)$.

Proof. Since $x'_2 \in [x_2, x_3]$ we have that $M_2(x_{-2}, x'_2) = x'_2$, by proposition 3.B.2. But since $d(x'_2, x_3) < d(x_2, x_3)$, (x_{-2}, x'_2) is $(x_1|x'_2, x_3)$ -well separated. Hence, by proposition 3.B.3, for $(x_1|x'_2, x'_3)$ -well separated instance $\vec{x}' = (\vec{x}_{-\{2,3\}}, x'_2, x'_3)$, $M_2(\vec{x}') = x'_2$. \square

Proposition 3.B.5. *Consider $(x_1|x_2, x_3)$ -well separated stable instance \vec{x} for which $M_2(\vec{x}) = x_2$. Then for every $(x_1|x'_2, x'_3)$ -well separated instance $\vec{x}' = (\vec{x}_{-\{2,3\}}, x'_2, x'_3)$, with $x_2 \leq x'_2$, if \vec{x}' is also well separated, $M_2(\vec{x}') = x_2$.*

Proof. We will inductively use proposition 3.B.4 to create instance \vec{x}' . Consider $d = d(x'_2, x_2)$, $\delta = d(x_3, x_2)/2$ and $\kappa = \lceil d/\delta \rceil$. Then for every $\lambda = 1, 2, 3, \dots, \kappa$ consider instance $\vec{x}_\lambda = (\vec{x}_{-\{2,3\}}, x_2 + (\lambda-1)\delta, x_3 + (\lambda-1)\delta)$. Now observe that \vec{x}_λ is well separated since for it's rightmost pair, $x'_2 = x_2 + (\lambda-1)\delta$ and $x'_3 = x_3 + (\lambda-1)\delta$ it is $d(x'_2, x'_3) > 2\delta$ while $d(x_1, x'_2) > d(x_1, x_2)$. By iteratively applying proposition 3.B.4 to \vec{x}_λ , we have that for every $(\vec{x}_{-\{2,3\}}, y_2, y_3)$ well separated instance with $x_2 + (\lambda-1)\delta \leq y_2 \leq x_2 + \lambda\delta$, $M_2(\vec{x}_{-\{2,3\}}, y_2, y_3) = y_2$. For $\lambda = \kappa$ we get $M_2(\vec{x}_{-\{2,3\}}, x'_2, x'_3) = x'_2$. \square

Chapter 4

Sampling and Optimal Preference Elicitation in Simple Mechanisms

Preference elicitation has received considerable attention in computational social choice [76, 78, 84, 169, 168]. Segal [211] 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. The boundaries of computational tractability and the strategic issues that arise in optimal preference elicitation were investigated by Conitzer and Sandholm for a series of voting schemes [77], while the same authors established the worst-case number of bits required to execute common voting rules [78]. The trade-off between accuracy and information leakage in facility location games has been studied by Feldman et al. [105], investigating the behavior of truthful mechanisms with truncated input space—e.g., *ordinal* information models.

More recently, the trade-off between efficiency—in terms of *distortion* as introduced by Procaccia and Rosenschein [193]—and communication was addressed by Mandal et al. [171] (see also [6]). Their results were subsequently improved using machinery from streaming algorithms [172], such as *linear sketching* and *L_p -sampling* [179, 142]. In similar spirit, some works [188, 40] address efficient preference elicitation in the form of *top- ℓ elicitation*, meaning that the agents are asked to provide the length ℓ prefix of their ranking instead of their full ranking. This trade-off between efficiency and communication has also been an important consideration in the *metric distortion framework* [19]—a setting closely related to our work; e.g., see [152, 151, 51, 9], as well as [21] for a comprehensive overview of that line of work.

Another important consideration in the literature relates to the interplay between communication constraints and incentive compatibility. In particular, Van Zandt [219] articulated conditions under which communication and incentive com-

patibility can be examined separately, while Reichelstein [198], and Segal and Fadel [96] examined the communication overhead induced in truthful protocols, i.e. the communication cost of truthfulness. An additional prominent line of research studies mechanism design under communication constraints (see [180] and references therein). Specifically, in a closely related to ours work, Blumrosen, Nisan and Segal [47] considered the design of a single-item auction in a communication model in which every bidder could only transmit a limited number of bits. One of their key results was a 0.648 social welfare approximation for 1-bit auctions (every bidder could only transmit a *single* bit to the mechanism) and uniformly distributed bidders. Further, the design of optimal—w.r.t. the obtained revenue—bid levels in an English auction was considered in [80], assuming known prior distributions.

4.1 Contributions

Our work provides several new insights on sampling and preference elicitation for a series of exemplar environments from mechanism design.

Facility Location Games. First, we turn our attention to facility location games; specifically, we consider Moulin’s *median* mechanism, one of the most fundamental allocation rules in voting. We observe that unlike the median as a function, the *social cost* of the median exhibits a *sensitivity* property. Subsequently, we show that for any $\epsilon > 0$, a random sample of size $c = \Theta(1/\epsilon^2)$ suffices to recover a $1 + \epsilon$ approximation w.r.t. the optimal social cost of the full information median in the metric space $(\mathbb{R}, |\cdot|)$, while the number of agents $n \rightarrow \infty$; this guarantee is established both in terms of expectation (Theorem 4.3.8), and with high probability (Theorem 4.3.9). Consequently, it is possible to obtain a near-optimal approximation with an arbitrarily small fraction of the total information. Our analysis is quite robust, implying directly the same characterization for the median on *simple* and *open* curves. Next, we extend this result for the *generalized median* in high-dimensional metric spaces $(\mathbb{R}^d, \|\cdot\|_1)$ in Theorem 4.3.11. In contrast, the sensitivity property of the median does not extend on trees, as implied by Theorem 4.3.12. Finally, for completeness, we show that sampling cannot provide meaningful guarantees w.r.t. the expected social cost when allocating at least 2 facilities on the line through the *percentile* mechanism (Theorem 4.3.13).

These results constitute natural continuation on efficient preference elicitation and sampling in social choice [83, 82, 58, 76], and supplement the work of Feldman et al. [105]; yet, to the best of our knowledge, we are the first to investigate the performance of sampling in facility location games. We stress that our guarantees do not require any prior knowledge or any discretization assumptions. Moreover, the sensitivity property of the median’s social cost could be of independent interest, as it can be potentially employed to design *differentially private* [87] and *noise-tolerant* implementations of the median mechanism. From a technical standpoint,

one of our key contributions is a novel asymptotic characterization of the rank of the sample’s median (Theorem 4.3.3), discussed further in the subsection below.

Auctions. Next, in Section 4.4, we espouse a communication complexity framework in order to design a series of auctions with optimal preference elicitation. In particular, we measure the number of bits elicited from the agents in expectation, endeavoring to minimize it. We mainly make the natural assumption that the number of bits that can represent any valuation—expressed with parameter k —is *independent* on the number of agents n ; thus, we focus on the communication complexity while n asymptotically grows. In this context, we show that we can asymptotically match the lower bound of 1 bit per bidder for a series of fundamental settings, without possessing any prior knowledge.

- First, we propose an ascending auction in which the ascend of the price is calibrated *adaptively* through a sampling mechanism. Thus, in Theorem 4.4.9 we establish that we can implement Vickrey’s rule with only $1 + \epsilon$ bits from an average bidder, for any $\epsilon > 0$;
- Moreover, we consider the design of a multi-item auction with m items and bidders with additive valuations. Our main contribution is to design an efficient *encoding scheme* that substantially truncates the communication in a *simultaneous* implementation, asymptotically recovering the same optimal bound whenever the number of items m is a small *constant* (Theorem 4.4.12);
- Finally, we develop a novel ascending-type multi-unit auction in the domain of unit-demand bidders. Our proposed auction announces in every round two separate prices—based on a natural sampling algorithm (see Theorem 4.4.13), leading again to the optimal communication bound (Theorem 4.4.16).

Our results supplement prior work [198, 96] by showing that the incentive compatibility constraint does not augment the communication requirements of the interaction process for a series of fundamental settings. We also corroborate on one of the main observations of Blumrosen et al. [47]: *asymmetry helps*—in deriving tight communication bounds. Indeed, in our mechanisms the information elicitation is substantially asymmetrical. Finally, we believe that our results could have practical significance due to their simplicity and their communication efficiency.

4.1.1 Overview of Techniques

Facility Location Games. In Section 4.3, our first key observation in Theorem 4.3.2 is that the social cost of the median admits a *sensitivity* property. Thus, it is possible to obtain a near-optimal approximation w.r.t. the optimal social cost even if the allocated facility is very far from the median. The sensitivity of the

social cost of the median essentially reduces a near-optimal approximation to the concentration of the rank of the sample’s median. Based on this insight, we prove in Theorem 4.3.3 that when the participation is large (i.e. $n \rightarrow \infty$) the distribution of the rank of the sample’s median converges to a continuous *transformed beta* distribution. This result should not be entirely surprising given that—as is folklore in statistics—the order statistics of the uniform distribution on the unit interval have marginal distributions belonging to the beta distribution family (e.g., see [81]). Having made these connections, the rest of our guarantees in Section 4.3 for more general metrics follow rather directly.

Auctions. With regard to our results in Section 4.4, we commence our overview with the single-item auction, and we then describe our approach for several extensions. In particular, for our ascending auction we consider as a black-box an algorithm that implements a second-price auction; then, at every round of the auction we *simulate* a sub-auction on a random sample of agents, and we broadcast the “market-clearing price” in the sub-auction as the price of the round. From a technical standpoint, we show that as the size of the sample increases, the fraction of the agents that will remain active in the forthcoming round gets gradually smaller (Theorem 4.4.5), leading to Theorem 4.4.8 and Theorem 4.4.9. It should also be noted that the truncated communication does *not* undermine the incentive compatibility of our mechanism, as implied by Theorem 4.4.2 and Theorem 4.4.3. Moving on to the multi-item auction with additive valuations, we employ a basic principle from information theory: encode the more likely events with fewer bits. This simple observation along with a property of our single-item auction allow us to substantially reduce the communication complexity when the auctions are executed in *parallel*.

Finally, we alter our approach for the design of a multi-unit auction with unit demand bidders. In contrast to a standard ascending auction, our idea is to broadcast in every round *two* separate prices, a “high” price and a “low” price. Subsequently, the mechanism may simply recurse on the agents that reside between the two prices. The crux of this implementation is to design an algorithm that takes as input a *small* number of bits, and returns prices that are tight upper and lower bounds on the final VCG payment. To this end, we design a novel algorithm that essentially performs *stochastic binary search* on the tree that represents the valuation space; more precisely, in every junction, or decision, the branching is made based on a small *sample* of agents, eliminating at every step half the points on the valuation space. From an algorithmic standpoint, this gives a simple procedure performing approximate *selection* from an unordered list with very limited communication.

4.2 Preliminaries

Facility Location Games. Consider a metric space $(\mathcal{M}, \text{dis}(\cdot, \cdot))$, where $\text{dis} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ is a *metric* (or distance function) on \mathcal{M} ; i.e., for any $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathcal{M}$, dis satisfies the following: (i) identity of indiscernibles: $\text{dis}(\mathbf{x}, \mathbf{y}) = 0 \iff \mathbf{x} = \mathbf{y}$; (ii) symmetry: $\text{dis}(\mathbf{x}, \mathbf{y}) = \text{dis}(\mathbf{y}, \mathbf{x})$; and (iii) triangle inequality: $\text{dis}(\mathbf{x}, \mathbf{y}) \leq \text{dis}(\mathbf{x}, \mathbf{z}) + \text{dis}(\mathbf{z}, \mathbf{y})$. Given as input an n -tuple $\mathcal{I} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, with $\mathbf{x}_i \in \mathcal{M}$, the ℓ -facility location problem consists of allocating ℓ facilities on \mathcal{M} in order to minimize the *social cost*; more precisely, if L is the finite set of allocated facilities, the induced social cost w.r.t. the distance function dis is defined as

$$\text{cost}(L, \text{dis}) \triangleq \sum_{i=1}^n \min_{\mathbf{x} \in L} \text{dis}(\mathbf{x}, \mathbf{x}_i);$$

that is, every point is assigned to its closest (allocated) facility. For notational simplicity, we omit the distance function and we simply write $\text{cost}(L)$. With a slight abuse of notation, when $|L| = 1$, we will use $\text{cost}(\mathbf{x})$ to represent the social cost of allocating a single facility on $\mathbf{x} \in \mathcal{M}$. In a mechanism design setting every point in the instance \mathcal{I} is associated with a *strategic* agent, and \mathbf{x}_i represents her preferred *private* location (e.g. her address); naturally, every agent i endeavors to minimize her (atomic) distance from the allocated facilities. A mechanism is called *strategy-proof* if for every agent i , and for any possible valuation profile, i minimizes her distance by reporting her actual valuation; if the mechanism is *randomized*, one is typically interested in strategy-proofness in expectation.

The Median Mechanism. Posit a metric space $(\mathbb{R}^d, \|\cdot\|_1)$. The MEDIAN is a mechanism for the 1-facility location problem which allocates a single facility on the median of the reported instance. For high-dimensional spaces, the median is defined coordinate-wise; more precisely, if $\mathcal{I} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ represents an arbitrary instance on \mathbb{R}^d , and we denote with x_i^j the j^{th} coordinate of \mathbf{x}_i in some underlying coordinate system,

$$\text{median}(\mathcal{I}) \triangleq (\text{median}(x_1^1, \dots, x_n^1), \dots, \text{median}(x_1^d, \dots, x_n^d)).$$

In this context, the following properties have been well-established in social choice, and we state them without proof.

Proposition 4.2.1. *The MEDIAN mechanism is strategy-proof.*

Proposition 4.2.2. *The MEDIAN mechanism is optimal w.r.t. (i.e. minimizes) the social cost in the metric space $(\mathbb{R}^d, \|\cdot\|_1)$.*

The median can also be defined beyond Euclidean spaces [210], as it will be discussed in more detail in Section 4.3.3. It should also be noted that the median can be employed heuristically for the ℓ -facility location problem when additional separability conditions are met; for example, the instance could correspond to residents of isolated cities, and it would be natural to assign one facility to a single “cluster”.

Auctions. In the domains we are studying in Section 4.4, the valuation of an agent i for any bundle of items S is fully specified by the values $v_{i,j}$, for every item j ; if a single item—and potentially multiple units of the same item—are to be disposed, we use v_i for simplicity. Moreover, in Section 4.4 we employ a *communication complexity* framework in order to analyze the measure of information elicited from the agents; thus, we need to assume that every value $v_{i,j}$ can be expressed with k bits. We will mainly assume that k is a parameter independent of the number of agents n , and one should imagine that a small constant k (e.g. 32 or 64 bits) would suffice. In this context, we define the communication complexity of a mechanism to be the expected number of bits elicited from the participants during the interaction process, and we study the asymptotic growth of this function as $n \rightarrow \infty$. We will assume that an agent remains active in the auction only when *positive* utility can be obtained. Thus, if the (monotonically increasing) *announced price* for item j coincides with the value $v_{i,j}$ of some agent i , we presume that i will withdraw from the auction; we use this hypothesis to handle certain singular cases (e.g., all agents could have the same value for the item).

For the incentive compatibility analysis in Section 4.4 we will need to refine and extend the notion of strategy-proofness from facility location games. A mechanism will be referred to as *strategy-proof* if truthful reporting is a *universally dominant* strategy—a best response under any possible action profile and *randomized* realization—for every agent.

Obvious Strategy-Proofness. 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. Notice that OSP implies strategy-proofness, so it is a stronger notion of incentive compatibility [164].

Ex-Post Incentive Compatibility. We will also require a weaker notion of incentive compatibility; a strategy profile (s_1, \dots, s_n) constitutes an *ex-post Nash equilibrium* if the action $s_i(v_i)$ is a best response to every action profile $\mathbf{s}_{-i}(\mathbf{v}_{-i})$ —for any agent i and valuation v_i . A mechanism will be called *ex-post incentive compatible* if sincere bidding constitutes an ex-post Nash equilibrium.

Additional Notation. We will denote with $N = [n]$ the set of agents that participate in the mechanism. In a single parameter environment, the *rank* of an agent corresponds to the index of her private valuation in ascending order (ties are broken arbitrarily according to some predetermined rule; e.g. lexicographic order) and indexed from 1, unless explicitly stated otherwise. We will use the standard notation of $f(n) \sim g(n)$ if $\lim_{n \rightarrow +\infty} f(n)/g(n) = 1$ and $f(n) \lesssim g(n)$ if $\lim_{n \rightarrow +\infty} f(n)/g(n) \leq 1$, where n will always be implied as the asymptotic parameter. For notational brevity, we will let $\binom{n}{m} = 0$ when $m > n$. Finally, $\|\mathbf{x}\|_1$

denotes the L_1 norm of $\mathbf{x} \in \mathbb{R}^d$, while d will mainly represent the dimension of the underlying space.

4.3 Facility Location Games and Approximation of the Median Mechanism

The most natural approximation of the MEDIAN mechanism consists of taking a random sample of size c , and allocating a (single) facility to the median of the sample, as implemented in APPROXMEDIANVIASAMPLING (Mechanism 4). Perhaps surprisingly, we will show that this simple approximation yields a social cost arbitrarily close to the optimal for the metric space $(\mathbb{R}^d, \|\cdot\|_1)$ —for a sufficiently large sample. Our analysis commences with the median on the line, where our main contribution lies in Theorem 4.3.8. Our approach is quite robust and yields analogous guarantees for the median defined on curves (Theorem 4.3.10) and the *generalized median* on the metric space $(\mathbb{R}^d, \|\cdot\|_1)$ (Theorem 4.3.11). We conclude this section by illustrating the barriers of sampling approximations when allocating a single facility on a *tree metric*, as well as allocating *multiple* facilities on the line.

Mechanism 4: APPROXMEDIANVIASAMPLING(N, ϵ, δ)

Input: Set of agents N , accuracy parameter $\epsilon > 0$, confidence parameter $\delta > 0$

Output: $\mathbf{x} \in \mathbb{R}^d$ such that $cost(\mathbf{x}) \leq (1 + \epsilon)cost(\mathbf{x}_m)$, where $\mathbf{x}_m = \text{median}(N)$

Set $c = \Theta(1/(\epsilon\delta)^2)$ to be the size of the sample

Let S be a random sample of c agents from N

return median(S)

In this section we do not have to dwell on incentive considerations given that our sampling mechanism APPROXMEDIANVIASAMPLING directly inherits its truthfulness from the MEDIAN mechanism (recall Theorem 4.2.1)—assuming of course that the domain admits a median.

Proposition 4.3.1. APPROXMEDIANVIASAMPLING *is strategy-proof*.

4.3.1 Median on the Line

In the following, we will tacitly consider an underlying arbitrary instance $\mathcal{I} = (x_1, \dots, x_n)$, with $x_i \in \mathbb{R}$ the (private) valuation—the preferred location—of agent i . To simplify the exposition, we will assume—without any loss of generality—that the number of agents n is odd, with $n = 2\kappa + 1$ for some $\kappa \in \mathbb{N}$.

Sensitivity of the Median. The first potential obstacle in approximating the MEDIAN mechanism relates to the *sensitivity* of the median. In particular, notice that the function of the median has an *unbounded sensitivity*, that is, a unilateral deviation in the input can lead to an arbitrarily large shift in the output; more concretely, consider an instance with $n = 2\kappa + 1$ agents, and let κ agents reside at $-l$ and $\kappa + 1$ agents at $+l$ for some large $l > 0$. If an agent from the rightmost group was to switch from $+l$ to $-l$, then the median would also relocate from $+l$ to $-l$, leading to a potentially unbounded deviation. It should be noted that in the regime of statistical learning theory, one technique to circumvent this impediment and ensure differential privacy revolves around the notion of *smooth sensitivity*; e.g., see [30, 57]. Our approach is based on the observation that the *social cost* of the MEDIAN inherently presents a sensitivity property. Formally, we establish the following lemma:

Lemma 4.3.2 (Sensitivity of the MEDIAN). *Let $x_m = \text{median}(\mathcal{I})$ and $x \in \mathbb{R}$ some position such that $\epsilon \cdot n$ agents reside between x and x_m , with $0 \leq \epsilon < 1/2$. Then,*

$$\text{cost}(x) = \text{cost}(x_m) \left(1 + \mathcal{O} \left(\frac{4\epsilon}{1 - 2\epsilon} \right) \right).$$

Proof. For the sake of presentation, let us assume that $x \geq x_m$. Consider the set L containing the $\lfloor n/2 \rfloor - \epsilon \cdot n$ leftmost agents (ties are broken arbitrarily), and the set R with the $\lfloor n/2 \rfloor - \epsilon \cdot n$ rightmost agents, leading to a partition as illustrated in Figure 4.1. Now observe that if we transfer the facility from $x_m = \text{median}(\mathcal{I})$ to x the cumulative social cost of L and R remains invariant, i.e.,

$$\sum_{i \in R \cup L} |x_i - x_m| = \sum_{i \in R \cup L} |x_i - x|.$$

In other words, the increase in social cost incurred by group L is exactly the social cost decrease of group R . Thus, it follows that

$$\text{cost}(x) \leq \text{cost}(x_m) + 2\epsilon n |x - x_m|, \quad (4.1)$$

where notice that this inequality is tight when the agents in the interval (x_m, x) are accumulated arbitrarily close to x_m . Moreover, we have that

$$\text{cost}(x_m) \geq \sum_{i \in R} |x_i - x_m| \geq (\lfloor n/2 \rfloor - \epsilon n) |x - x_m|, \quad (4.2)$$

and the claim follows from bounds (4.1) and (4.2). \square

As a result, Theorem 4.3.2 implies that a unilateral deviation by a player can only lead to an increase of $\mathcal{O}(1/n)$ in the social cost. To put it differently, if an adversary corrupts arbitrarily a *constant* number of reports, the increase in social cost will be asymptotically negligible. Indeed, even if the allocated facility lies arbitrarily far from the median, the induced social cost might still be near-optimal.




Figure 4.1: Partition of the agents on the line.

Now let us assume—without loss of generality—that $c = 2\rho + 1$, for some $\rho \in \mathbb{N}$, where c is the size of the sample; also recall that $n = 2\kappa + 1$. Motivated by Theorem 4.3.2, our analysis will be *oblivious* to the agents' locations on the line, but instead, our focus will be on characterizing the *rank* of the sample's median—the relative order of the sample's median w.r.t. the entire instance; this approach will also allow us to directly obtain a guarantee in more general metric spaces. More precisely, consider a random variable X_r that represents the rank—among the entire instance—of the sample's median, normalized in the domain $[-1, 1]$; for instance, if the sample's median happens to coincide with the median of the entire instance, then $X_r = 0$. The reason behind this normalization relates to our asymptotic characterization (Theorem 4.3.3). Now fix a particular rank i/κ in $[-1, 1]$. Notice that the number of configurations that correspond to the event $\{X_r = i/\kappa\}$ is

$$\binom{\kappa - i}{\rho} \binom{\kappa + i}{\rho}.$$

As a result, the probability mass function of X_r can be expressed as follows:

$$\mathbb{P}\left[X_r = \frac{i}{\kappa}\right] = \frac{\binom{\kappa - i}{\rho} \binom{\kappa + i}{\rho}}{\binom{2\kappa + 1}{2\rho + 1}}. \quad (4.3)$$

It is interesting to notice the similarity of this expression to the probability mass function of a *hypergeometric* distribution. We also remark that the normalization constraint of the probability mass function (4.3) yields a well-known variant of the Vandermonde identity:

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

For this reason, we shall refer to the distribution of X_r as the (κ, ρ) -*Vandermonde distribution*. Importantly, Theorem 4.3.2 tells us that the concentration of the Vandermonde distribution—for sufficiently large values of parameter ρ —suffices to obtain a near-optimal approximation with respect to the social cost. However, quantifying exactly the concentration of the Vandermonde distribution appears

to be challenging.¹ In light of this, our main insight—and the main technical contribution of this section—is an asymptotic characterization of this distribution.

Theorem 4.3.3 (Convergence of the Vandermonde Distribution). *If we let $\kappa \rightarrow +\infty$ the (κ, ρ) -Vandermonde distribution converges to a continuous distribution with probability density function $\phi : [-1, 1] \rightarrow \mathbb{R}$, such that*

$$\phi(t) = C(\rho)(1 - t^2)^\rho, \quad (4.4)$$

where

$$C(\rho) = B\left(\frac{1}{2}, \rho + 1\right)^{-1} = \frac{(2\rho + 1)!}{(\rho!)^2 2^{2\rho + 1}}.$$

In the statement of the theorem, B represents the *beta function*. Recall that for $x, y \in \mathbb{R}_{>0}$, the beta function is defined as

$$B(x, y) = \int_0^1 t^{x-1}(1-t)^{y-1} dt = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}, \quad (4.5)$$

where Γ represents the *gamma function*. One can verify the normalization constraint in Theorem 4.3.3 using a quadratic transform $u = t^2$ as follows:

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

Moreover, the final term can be expressed succinctly through the following lemma:

Lemma 4.3.4. *If Γ represents the gamma function and $n \in \mathbb{N}$,*

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

Before we proceed with the proof of Theorem 4.3.3, we state an elementary result from real analysis.

Lemma 4.3.5. *Let $f : [-1, 1] \rightarrow \mathbb{R}$ be an integrable function² and x some number in $(-1, 1)$; then,*

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

¹Ariel Procaccia pointed out to us that there is a more elementary way to “upper-bound” the concentration of X_r ; see [54, Lemma 1]. Yet, we remark that their argument would only provide a guarantee with high probability, and not in expectation.

²The integrability here is implied in the standard Riemannian-Darboux sense.

Proof of Theorem 4.3.3. Take some arbitrary $x \in (-1, 1)$, let $\nu = \lfloor x\kappa \rfloor + \kappa + 1$, and consider a random variable X_r drawn from a (κ, ρ) -Vandermonde distribution. It follows that

$$\begin{aligned} \lim_{\kappa \rightarrow +\infty} \mathbb{P}[X_r \leq x] &= \lim_{\kappa \rightarrow +\infty} \sum_{i=-\kappa}^{\lfloor x\kappa \rfloor} \frac{\binom{\kappa-i}{\rho} \binom{\kappa+i}{\rho}}{\binom{2\kappa+1}{2\rho+1}} \\ &= \lim_{n \rightarrow +\infty} \sum_{i=1}^{\nu} \frac{\binom{n-i}{\rho} \binom{i-1}{\rho}}{\binom{n}{2\rho+1}}, \end{aligned} \quad (4.6)$$

where recall that $n = 2\kappa + 1$. Given that $n!/(n-j)! = \Theta(n^j), \forall j \in \mathbb{N}$, we can recast (4.6) as

$$\begin{aligned} \lim_{\kappa \rightarrow +\infty} \mathbb{P}[X_r \leq x] &= \frac{(2\rho+1)!}{(\rho!)^2} \lim_{n \rightarrow +\infty} \frac{1}{n^{2\rho+1}} \sum_{i=1}^{\nu} \frac{(n-i)!}{(n-i-\rho)!} \frac{(i-1)!}{(i-1-\rho)!} \\ &= \frac{(2\rho+1)!}{(\rho!)^2} \lim_{n \rightarrow +\infty} \frac{1}{n^{2\rho+1}} \sum_{i=1}^{\nu} (n-i)^\rho i^\rho, \end{aligned} \quad (4.7)$$

where the last derivation follows by ignoring lower order terms. Finally, from (4.7) we obtain that

$$\begin{aligned} \lim_{\kappa \rightarrow +\infty} \mathbb{P}[X_r \leq x] &= \frac{(2\rho+1)!}{(\rho!)^2} \lim_{n \rightarrow +\infty} \frac{1}{n} \sum_{i=1}^{\nu} \left(\frac{i}{n} - \left(\frac{i}{n} \right)^2 \right)^\rho \\ &= \frac{(2\rho+1)!}{(\rho!)^2} \lim_{\nu \rightarrow +\infty} \frac{x+1}{2\nu} \sum_{i=1}^{\nu} \left(i \cdot \frac{x+1}{2\nu} - \left(i \cdot \frac{x+1}{2\nu} \right)^2 \right)^\rho \\ &= \frac{(2\rho+1)!}{(\rho!)^2 2^{2\rho+1}} \lim_{\nu \rightarrow +\infty} \frac{x+1}{\nu} \sum_{i=1}^{\nu} \left(2i \cdot \frac{x+1}{\nu} - \left(i \cdot \frac{x+1}{\nu} \right)^2 \right)^\rho \\ &= \frac{(2\rho+1)!}{(\rho!)^2 2^{2\rho+1}} \lim_{\nu \rightarrow +\infty} \frac{x+1}{\nu} \sum_{i=1}^{\nu} \left(1 - \left(-1 + i \cdot \frac{x+1}{\nu} \right)^2 \right)^\rho \\ &= \frac{(2\rho+1)!}{(\rho!)^2 2^{2\rho+1}} \int_{-1}^x (1-t^2)^\rho dt, \end{aligned}$$

where in the last line we applied Theorem 4.3.5, concluding the proof. \square

Having established this asymptotic characterization, we are now ready to argue about the concentration of the induced distribution with respect to parameter ρ .

Theorem 4.3.6 (Concentration). *Consider a random variable X^3 with probability density function $\phi(t) = C(\rho)(1 - t^2)^\rho$. Then, for any $\epsilon > 0$ and $\delta > 0$, there exists some $\rho_0 = \Theta(1/(\epsilon\delta)^2)$ such that $\forall \rho \geq \rho_0$,*

$$\mathbb{P}[|X| \geq \epsilon] \leq \delta.$$

Proof. Consider some $j \in \mathbb{N}$. The j^{th} moment of $|X|$ can be expressed as

$$\mathbb{E}[|X|^j] = C(\rho) \int_{-1}^1 |t|^j (1 - t^2)^\rho dt = 2C(\rho) \int_0^1 t^j (1 - t^2)^\rho dt.$$

Applying the quadratic transformation $u = t^2$ yields

$$\mathbb{E}[|X|^j] = \frac{B\left(\frac{j}{2} + \frac{1}{2}, \rho + 1\right)}{B\left(\frac{1}{2}, \rho + 1\right)}. \quad (4.8)$$

Recall from Stirling's approximation formula that $n! = \Theta(\sqrt{2\pi n} \left(\frac{n}{e}\right)^n)$; thus, we obtain that $C(\rho)$ grows as

$$C(\rho) = \frac{(2\rho + 1)!}{(\rho!)^2 2^{2\rho+1}} = \Theta(\sqrt{\rho}).$$

In particular, this along with (4.8) imply that $\mathbb{E}[|X|] = \Theta(1/\sqrt{\rho})$. Thus, if we apply Markov's inequality we obtain that

$$\mathbb{P}[|X| \geq \epsilon] = \mathcal{O}\left(\frac{1}{\epsilon\sqrt{\rho}}\right).$$

As a result, it suffices to select $\rho = \Theta(1/(\epsilon\delta)^2)$ so that $\mathbb{P}[|X| \geq \epsilon] \leq \delta$; notice that tighter bounds w.r.t. the confidence parameter δ can be obtained if we apply Markov's inequality for higher moments of $|X|$ through (4.8). \square

We also highlight the following important byproduct, which was established above en route to proving Theorem 4.3.6.

Corollary 4.3.7. *Consider a random variable X with probability density function $\phi(t) = C(\rho)(1 - t^2)^\rho$. Then, $\mathbb{E}[|X|] = \Theta(1/\sqrt{\rho})$.*

We are now ready to analyze the approximation ratio of our sampling median, both in terms of expectation and with high probability.

³It is interesting to note that X is a *sub-Gaussian* random variable [160] with *variance proxy* $\sigma^2 = \Theta(1/\rho)$; indeed, notice that $(1 - t^2)^\rho \leq e^{-\rho t^2}$, with the bound being tight for $|t| \downarrow 0$. This observation leads to an alternative—and rather elegant—way to analyze the concentration of X .

Theorem 4.3.8 (Sampling Median on the Line). *Consider a set of agents $N = [n]$ that lie on the one-dimensional metric space $(\mathbb{R}, |\cdot|)$. Then, for any $\epsilon > 0$, $\text{APPROXMEDIANVIASAMPLING}(N, \epsilon, \delta = 1)$ takes a sample of size $c = \Theta(1/\epsilon^2)$ and yields in expectation a $1 + \epsilon$ approximation w.r.t. the optimal social cost of the full information MEDIAN , while $n \rightarrow +\infty$.*

Proof. Consider a random variable X with probability density function $\phi(t) = C(\rho)(1-t^2)^\rho$. Theorem 4.3.3 implies that X corresponds to the rank of the sample's median with sample size $c = 2\rho + 1$, while $n \rightarrow +\infty$. Let $g : (0, 1) \ni x \mapsto 2x/(1-x)$; we know from Theorem 4.3.2 that the expected approximation ratio on the social cost is $1 + \mathcal{O}(\mathbb{E}[g(|X|)])$. But, it follows that

$$\begin{aligned} \mathbb{E}[g(|X|)] &= 4C(\rho) \int_0^1 \frac{t}{1-t} (1-t^2)^\rho dt \leq 8C(\rho) \int_0^1 t(1-t^2)^{\rho-1} dt \\ &= 8 \frac{2\rho+1}{2\rho} C(\rho-1) \int_0^1 t(1-t^2)^{\rho-1} dt. \end{aligned}$$

Now notice that Theorem 4.3.7 implies that

$$C(\rho-1) \int_0^1 t(1-t^2)^{\rho-1} dt = \Theta\left(\frac{1}{\sqrt{\rho}}\right).$$

As a result, we have shown that the expected approximation ratio is $1 + \mathcal{O}(1/\sqrt{\rho})$, and taking $\rho = \Theta(1/\epsilon^2)$ concludes the proof. \square

Corollary 4.3.9. *Consider a set of agents $N = [n]$ that lie on the metric space $(\mathbb{R}, |\cdot|)$. Then, for any $\epsilon > 0$ and $\delta > 0$, $\text{APPROXMEDIANVIASAMPLING}(N, \epsilon, \delta)$ takes a sample of size $c = \Theta(1/(\epsilon\delta)^2)$ and yields with probability at least $1 - \delta$ a $1 + \epsilon$ approximation w.r.t. the optimal social cost of the full information MEDIAN , while $n \rightarrow +\infty$.*

Proof. The claim follows directly from Theorem 4.3.2, Theorems 4.3.3 and 4.3.6. \square

Extension to the Median on Curves

Here we give a slight extension of the previous guarantee when the agents lie on a *curve*. More precisely, let \mathcal{C} be a curve parameterized by a continuous function $\psi : [0, 1] \rightarrow \mathbb{R}^d$. We will assume that \mathcal{C} is *simple* and *open*, i.e. ψ is injective in $[0, 1]$; see an example in Figure 4.2. We also consider the distance between two points $A, B \in \mathcal{C}$ to be the *length* of the induced (simple) sub-curve from A to B , denoted with $\ell(A, B)$.

Notice that any simple and open curve naturally induces a ranking—a *total order*—over its domain; indeed, for any points $A, B \in \text{Im}(\psi) \equiv \mathcal{C}$, we let $A \preceq B \iff \psi^{-1}(A) \leq \psi^{-1}(B)$. Thus, we may define the median on the curve, which is strategy-proof and optimal w.r.t. the induced social cost. Importantly, our

previous analysis is robust, and our methodology for the median on the line is directly applicable.

Theorem 4.3.10 (Sampling Median on Curves). *Consider a set of agents $N = [n]$ that lie on the metric space $(\mathcal{C}, \ell(\cdot, \cdot))$, where \mathcal{C} represents a simple and open curve on some subset of a Euclidean space. Then, for any $\epsilon > 0$, $\text{APPROXMEDIANVIASAMPLING}(N, \epsilon, \delta = 1)$ takes a sample of size $c = \Theta(1/\epsilon^2)$ and yields in expectation a $1 + \epsilon$ approximation w.r.t. the optimal social cost of the full information MEDIAN , while $n \rightarrow +\infty$.*




Figure 4.2: An example of a simple and open curve \mathcal{C} . The distance $\ell(A, B)$ between two points $A, B \in \mathcal{C}$ is defined as the length from A to B .

4.3.2 High-Dimensional Median

Next, in this subsection we extend our analysis for the high-dimensional metric space $(\mathbb{R}^d, \|\cdot\|_1)$. Specifically, consider some arbitrary instance $\mathcal{I} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, where $\mathbf{x}_i \in \mathbb{R}^d$; we will denote with x_i^j the j^{th} coordinate of \mathbf{x}_i in some underlying coordinate system. Also, recall that the median of \mathcal{I} is derived through the median in every coordinate, i.e.,

$$\text{median}(\mathcal{I}) = (\text{median}(x_1^1, \dots, x_n^1), \dots, \text{median}(x_1^d, \dots, x_n^d)).$$

Naturally, the MEDIAN takes as input some instance \mathcal{I} and outputs the $\text{median}(\mathcal{I})$. We establish the following theorem:

Theorem 4.3.11 (High-Dimensional Sampling Median). *Consider a set of agents $N = [n]$ that lie on the metric space $(\mathbb{R}^d, \|\cdot\|_1)$. Then, for any $\epsilon > 0$, $\text{APPROXMEDIANVIASAMPLING}(N, \epsilon, \delta = 1)$ takes a sample of size $c = \Theta(1/\epsilon^2)$ and yields in expectation a $1 + \epsilon$ approximation w.r.t. the optimal social cost of the full information MEDIAN , while $n \rightarrow +\infty$.*

Proof. Consider some facility at $\mathbf{x} \in \mathbb{R}^d$, with $\mathbf{x} = (x^1, \dots, x^d)$; the induced social cost of \mathbf{x} can be expressed as

$$\text{cost}(\mathbf{x}) = \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{x}\|_1 = \sum_{i=1}^n \sum_{j=1}^d |x_i^j - x^j| = \sum_{j=1}^d \sum_{i=1}^n |x_i^j - x^j|.$$

Let $\mathbf{X} = (X^1, \dots, X^d)$ be the output of the APPROXMEDIANVIASAMPLING (a random variable). For any $j \in [d]$, and with $c = \Theta(1/\epsilon^2)$, Theorem 4.3.8 implies that

$$\mathbb{E} \left[\sum_{i=1}^n |x_i^j - X^j| \right] \leq (1 + \epsilon) \sum_{i=1}^n |x_i^j - x_m^j| + o_n(1),$$

where $\mathbf{x}_m = (x_m^1, \dots, x_m^d) = \text{median}(\mathcal{I})$. Thus, by linearity of expectation we obtain that for $n \rightarrow \infty$,

$$\mathbb{E}[\text{cost}(\mathbf{X})] \leq (1 + \epsilon) \text{cost}(\mathbf{x}_m) = (1 + \epsilon) \text{cost}^*.$$

□

Importantly, observe that even in the high-dimensional case it suffices to take $c = \Theta(1/\epsilon^2)$, *independently* from the dimension of the space d , in order to obtain a guarantee in expectation. We also remark that it is straightforward to extend Theorem 4.3.9 for the high-dimensional median, and recover a near-optimal allocation with high probability.

4.3.3 Median on Trees

In contrast to our previous positive results, our characterization breaks when allocating a facility on a general *network*. In particular, consider an *unweighted tree* $G = (V, E)$, and assume that every node is occupied by a *single* agent. One natural way to define the median on G is by arbitrarily choosing a *generalized median* for each path of the tree; however, as articulated by Vohra and Schummer [210], for any two paths that intersect on an interval, it is crucial that the corresponding generalized medians must not contradict each other, a condition they refer to as *consistency*. Providing a formal definition of consistency would go beyond the scope of our study; instead, we refer the interested reader to the aforementioned work.

Now imagine that the designer has *no* prior information on the topology of the network, and will have to rely solely on the information extracted by the agents. Given that the graph might be vast, we consider a sample of nodes and we then construct the *induced* graph by querying the agents in the sample⁴ (naturally, we posit that every agent knows her neighborhood). However, notice that the induced graph need not be a tree, and hence, it is unclear even how to determine the output of the sampling approximation. In fact, *any* node in the subgraph may lead to a social cost far from the optimal.

⁴This model is analogous to the standard approach in *property testing* [135].




Figure 4.3: An unweighted “star” graph G with n nodes, with every node occupied by a *single* agent. It is easy to verify that v^* constitutes the unique median on G , satisfying the *consistency* condition of Vohra and Schummer. Notice that allocating a facility on v^* yields a social cost of $n - 1$, while any other allocation leads to a social cost of $2n - 3$.

Proposition 4.3.12. *Consider an unweighted star graph $G = (V, E)$ with $|V| = n$. Then, even if we take a sample of size $c = n/2$, every node in the sample will yield an approximation of $2 - 1/(n - 1)$ w.r.t. the optimal social cost with probability $1/2$.*

An approximation ratio of 2 is trivial in the following sense: There exists a mechanism, namely RANDOMDICTATOR, which selects uniformly at random a single player and allocates the facility on her preferred position on the underlying metric space. An application of the triangle’s inequality shows that RANDOMDICTATOR yields in expectation a 2 approximation w.r.t. the optimal social cost.⁵ In that sense, augmenting the sample does not seem particularly helpful when the underlying metric space corresponds to a network.

4.3.4 Sampling with Multiple Facilities

Finally, we investigate the performance of a sampling approximation when allocating *multiple* facilities. For the sake of simplicity, we posit a metric space $(\mathbb{R}, |\cdot|)$, and we consider the PERCENTILE mechanism, an allocation rule that assigns facilities on particular *percentiles* of the input. More precisely, the PERCENTILE mechanism is parameterized by a sequence $r_1 < r_2 < \dots$, with every $r_j \in [n]$ corresponding to a rank of the input; if the instance $\mathcal{I} = (x_1, \dots, x_n)$ is given in increasing order, we allocate a facility j on x_{r_j} for every j . Naturally, the

⁵On the other hand, it is easy to see that any deterministic dictatorship yields in the worst-case an $n - 1$ approximation.

MEDIAN can be classified in this family of mechanisms. Another prominent member is the TWOEXTREMES mechanism, proposed by Procaccia and Tennenholtz [195] for the 2-facility location problem. As the name suggests, this mechanism allocates two facilities at the minimum and the maximum reports of the instance, leading to an $n - 2$ approximation ratio w.r.t. the optimal social cost (in fact, the TWOEXTREMES is the only *anonymous* and deterministic mechanism with bounded approximation ratio [115]). We remark that the PERCENTILE mechanism is always strategy-proof, while its approximation ratio w.r.t. the optimal social cost is generally unbounded.

Now consider APPROXPERCENTILEVIASAMPLING, simulating PERCENTILE on a sample of size c ; we will tacitly presume that at least 2 facilities are to be allocated. Let us assume that the leftmost percentile L contains at most $(1 - \alpha)n$ agents, for some *constant* $\alpha > 0$, and denote with l the distance between L and the complementary set of agents R ; see Figure 4.4. If we let the inner-distance between in L and R approach to 0 and $l \rightarrow \infty$, we can establish the following:

Proposition 4.3.13. *There are instances for which even with a sample of size $c = \alpha n = \Theta(n)$ the APPROXPERCENTILEVIASAMPLING has in expectation an unbounded approximation ratio w.r.t. the social cost of the full information PERCENTILE mechanism.*




Figure 4.4: An instance that corresponds to Theorem 4.3.13, as $l \rightarrow \infty$; given that at least two facilities are to be allocated, any mechanism with bounded social cost has to allocate facilities on both L and R .

Indeed, for the instance we described the social cost of the full information mechanism approaches to 0, as it always allocates a facility on L , and at least one facility on R . In contrast, there will be a positive probability—albeit exponentially small—that the approximation mechanism fails to sample an agent from L . Moreover, the same limitation (Theorem 4.3.13) applies for an additive approximation, instead of a relative one. Perhaps, it would be interesting to examine the performance of a sampling approximation if we impose additional restrictions on the instance, such as *stability* conditions.

Remark. One of our insights is that unlike the *outcome* of the underlying mechanism, the social cost often presents a sensitivity property. This observation could be of independent interest in analyzing sampling approximations in voting. We explore this idea for the *plurality rule* in Section 4.A.

4.4 Auctions

In this section, we study a series of environments in the regime of auction theory. For every instance, we develop a mechanism which asymptotically minimizes the communication complexity, i.e. the numbers of bits elicited from the participants. We commence with the design of a single-item auction, and we gradually extend our techniques to cover more general domains.

4.4.1 Single-Item Auction

This subsection presents our ascending auction for disposing a *single* and *indivisible* item. On a high level, instead of updating the price in a static manner through a *fixed* increment, we propose a simple *adaptive* mechanism. Before we proceed with the description and the analysis of our mechanism, we wish to convey some intuition for our implementation. Specifically, as a thought experiment, imagine that every valuation v_i is drawn from some arbitrary distribution \mathcal{D} , and assume—rather unreasonably—that we have access to this distribution. In this context, one possible approach to minimize the communication would be to determine a threshold value T_h such that $\mathbb{P}_{v \sim \mathcal{D}}[v \geq T_h] \leq \epsilon$, for some small $\epsilon > 0$. The auctioneer could then simply broadcast at the first round of the auction the price $p := T_h$, and with high probability the agents above the threshold would constitute only a small fraction of the initial population. Moreover, the previous step could be applied recursively for the distribution \mathcal{D} conditioned on $v \geq T_h$, until only a few agents remain active. Our proposed mechanism will essentially mirror this thought experiment, but without any distributional assumptions, or indeed, any prior knowledge.

Specifically, we consider some black-box *deterministic* algorithm \mathcal{A} that *faithfully* simulates a second-price auction. Namely, \mathcal{A} takes as an input a set of agents S and returns a tuple (w, p) : w is the agent with the highest valuation among S (ties are broken arbitrarily), and p corresponds to the second highest valuation. Of course, \mathcal{A} only *simulates* a second-price auction, without actually allocating items and imposing payments.

In every round, our mechanism selects a random sample from the active agents,⁶ and simulates through algorithm \mathcal{A} a sub-auction. Then, the “market-clearing price” in the sub-auction is announced as the price of the round, and this process is then repeated iteratively. The pseudocode for our mechanism (ASCENDINGAUCTIONVIASAMPLING) is given in Mechanism 5.

Interestingly, our mechanism induces a format that couples an ascending auction with the auction simulated by \mathcal{A} . We shall establish the following properties:

Proposition 4.4.1. *Under truthful reporting, ASCENDINGAUCTIONVIASAMPLING returns—with probability 1—the VCG outcome.*

⁶Naturally, we assume that the sample contains at least 2 agents, so that the second-price rule is properly defined.

Mechanism 5: ASCENDINGAUCTIONVIASAMPLING(N, \mathcal{A}, ϵ)

Input: Set of agents N , algorithm \mathcal{A} which simulates a second-price auction, parameter $\epsilon > 0$

Output: VCG outcome (Winner & Payment)

```

while  $|N| > c$  do
  | Let  $S$  be a random sample of  $c = \Theta(1/\epsilon^2)$  agents from  $N$ 
  | Set  $w :=$  winner in  $\mathcal{A}(S)$ 
  | Announce  $p :=$  payment in  $\mathcal{A}(S)$ 
  | Update the active agents:  $N := \{i \in N \setminus S : v_i > p\} \cup \{w\}$ 
end
if  $|N| = 1$  then
  | return  $(w, p)$ 
else
  | return  $\mathcal{A}(N)$ 
end

```

Proof. First, notice that in any iteration of the while loop only agents that are below or equal to the second highest valuation will withdraw from the auction. Now consider the case where upon exiting the while loop only a *single* agent w remains in the set of active agents N . Then, it follows that the announced price p in the final round—which by construction coincides with the valuation of some player—exceeds the valuation of every player besides w . Thus, by definition, the outcome implements the VCG rule. Moreover, if after the last round $2 \leq |N| \leq c$, the claim follows given that \mathcal{A} faithfully simulates a second-price auction. \square

Proposition 4.4.2. ASCENDINGAUCTIONVIASAMPLING is *ex-post incentive compatible* if \mathcal{A} simulates a sealed-bid auction.⁷

Proof. Consider any round of the auction and some agent i that has been selected in the sample S ; we identify two cases. First, if $v_i \geq v_j, \forall j \in S \setminus \{i\}$, sincere reporting clearly constitutes a best response for i . Indeed, notice that since \mathcal{A} simulates a second-price auction, the winner in the sub-auction 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 final payment will always be greater or equal to her valuation—observe that the

⁷In fact, ASCENDINGAUCTIONVIASAMPLING (with \mathcal{A} implemented as a sealed-bid auction) is dominant strategy incentive compatible if the sequence of announced prices is non-decreasing; this property can be guaranteed if the “market-clearing price” in the sub-auction serves as a *reserved price*. Otherwise, truthful reporting is *not* necessarily a dominant strategy. For example, assume that every agent i , besides some agent j , commits to the following—rather ludicrous—strategy: i reports truthfully, unless i remains active with at most $2c$ other agents; in that case, i will act as if her valuation is 0. Then, the best response for j is to act as if her valuation is ∞ .

announced price always increases throughout the auction. Next, let p the payment in $\mathcal{A}(S)$ and $i \notin S$. It follows 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. \square

Proposition 4.4.3. *ASCENDINGAUCTIONVIASAMPLING is obviously strategy-proof (OSP) if \mathcal{A} simulates an English auction.*

Proof. Notice that the induced mechanism performs a standard English auction, but instead of interacting with *every* agent in a given round, we “ascend” in a small sample; only when a single agent remains active we broadcast the current price to the rest of the agents. Now consider some agent i that participates in the sub-auction. If the current price is below her valuation v_i , then the best possible outcome from quitting is no better than the worst possible outcome from staying in the auction. Otherwise, if the price is above v_i , then the best possible outcome from staying in the auction is no better than the worst possible outcome from withdrawing in the current round. Indeed, notice that the announced price can only increase throughout the auction. Of course, the same line of reasoning applies for a round of interaction with the entire set of active agents; essentially, the claim follows from the OSP property of the English auction. \square

Next, we analyze the communication complexity of our proposed auction. Naturally, we have to assume that the valuation space is discretized, with k bits being sufficient to express any valuation. The first thing to note is a trivial lower bound on the communication.

Fact 4.4.4 (Communication Lower Bound). *Every mechanism that determines the agent with the highest valuation—with probability 1—must elicit at least n bits.*

Importantly, we will show that this lower bound can be asymptotically recovered. In particular, observe that as the size of sample increases, the fraction of agents that will choose to remain active—at least in the forthcoming round—gradually diminishes; the following lemma makes this property precise.

Lemma 4.4.5 (Inclusion Rate). *Let X_a be a random variable representing the proportion of agents that will remain active in a given round of the ASCENDINGAUCTIONVIASAMPLING with sample size c ; then,*

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

As a result, the size of the sample c allows us to calibrate the number of agents that we wish to include in the following round—the inclusion rate. Before we proceed with the proof of Theorem 4.4.5, we first state some standard asymptotic formulas.

Fact 4.4.6.

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

Fact 4.4.7.

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

Proof of Theorem 4.4.5. Suppose that there are n active agents in some round of the ASCENDINGAUCTIONVIASAMPLING (here we slightly abuse notation given that n corresponds to the initial number of agents). Let us denote with X_r the rank—in the domain $[n]$ —of the player with the second highest valuation (recall that ties are broken arbitrarily according to some *fixed* order) in the sample. We will show that

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

Indeed, simple combinatorial arguments yield that the probability mass function of X_r can be expressed as

$$\mathbb{P}[X_r = i] = \frac{\binom{n-i}{1} \binom{i-1}{c-2}}{\binom{n}{c}}.$$

As a result, it follows that

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

where we applied the asymptotic bounds from Fact 4.4.6 and Fact 4.4.7; also note that we ignored lower order terms in the third and fourth line. Finally, the proof follows given that $X_a \leq (n - X_r)/n$; the inequality here derives from the fact that multiple agents could have the same valuation with the agent with rank X_r , and we assumed that such agents will quit. \square

Next, we are ready to analyze the communication complexity of our mechanism. In the following theorem, we implicitly assume that the agents report sincerely.

Theorem 4.4.8. *Let Q be the communication complexity of a deterministic algorithm \mathcal{A} , faithfully simulating a second-price auction, and $N = [n]$ be a set of agents. For any $\epsilon > 0$ and $c = \Theta(1/\epsilon^2)$, denote by $t(n; c, k)$ the expected communication complexity of $\text{ASCENDINGAUCTIONVIASAMPLING}(N, \mathcal{A}, \epsilon)$. Then,*

$$t(n; c, k) \lesssim (1 + \epsilon)n + Q(c; k) \log n. \quad (4.9)$$

Proof. Consider some round of the $\text{ASCENDINGAUCTIONVIASAMPLING}$ with n active agents; as in Theorem 4.4.5, let us denote with X_a the proportion of agents that will remain active in the following round of the auction. If T represents the (randomized) communication complexity of our mechanism, we obtain that

$$\mathbb{E}[T(n; c, k)] = \mathbb{E}[T(nX_a; c, k)] + Q(c; k) + n - c. \quad (4.10)$$

There are various ways to bound randomized recursions of such form; our analysis will leverage the concentration of X_a . In the sequel, we will tacitly assume that the agents' valuations are pairwise distinct, as this yields an upper bound on the actual communication complexity (in our case, ties can only truncate communication). We will first establish that

$$\mathbb{V}[X_a] \sim \frac{2(c-1)}{(c+2)(c+1)^2}. \quad (4.11)$$

Indeed, consider the random variable X_r that represents the rank of the agent with the second highest valuation in the sample. Analogously to the proof of Theorem 4.4.5, it follows that

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

As a result, (4.11) follows given that $\mathbb{V}[X_r] = \mathbb{E}[X_r^2] - (\mathbb{E}[X_r])^2$ and $\mathbb{V}[X_a] = \mathbb{V}[X_r]/n^2$; notice that under the assumption that the valuations are pairwise distinct, it follows that $X_a = (n - X_r)/n$. For notational simplicity, let us denote with $\mu = \mathbb{E}[X_a]$ and $\sigma = \sqrt{\mathbb{V}[X_a]}$. Chebyshev's inequality implies that $\mathbb{P}[|X_a - \mu| \geq \sqrt{c}\sigma] \leq 1/c$. It is also easy to see that $\mu + \sqrt{c}\sigma \leq 4/\sqrt{c}$; hence, with probability at least $1 - 1/c$, $nX_a \leq 4n/\sqrt{c}$. Consequently, (4.10) gives that

$$t(n; c, k) \lesssim \left(1 - \frac{1}{c}\right) t\left(\frac{4n}{\sqrt{c}}; c, k\right) + \frac{1}{c}t(n; c, k) + n + Q(c; k), \quad (4.12)$$

where we used the fact that $t(n; c, k)$ is decreasing w.r.t. n . Moreover, (4.12) can be recast as

$$t(n; c, k) \lesssim t\left(\frac{4n}{\sqrt{c}}; c, k\right) + \frac{c}{c-1}n + \frac{c}{c-1}Q(c; k). \quad (4.13)$$

Now consider any small $\epsilon > 0$, and let $c = \Theta(1/\epsilon^2)$;⁸ from the previous recursion we obtain that

$$t(n; c, k) \lesssim (1 + \epsilon)n + Q(c; k) \log n,$$

as desired. □

In particular, if \mathcal{A} simulates a sealed-bid auction $Q(c; k) = c \cdot k$, while if \mathcal{A} simulates an English auction $Q(c; k) = c \cdot 2^k$; indeed, implementing a *faithful* second-price auction through a standard ascending format necessitates covering the entire valuation space, i.e. 2^k potential prices. Theorem 4.4.8 implies that the size of the sample c induces a trade-off between two terms: As we augment the size of the sample c we truncate the first term in (4.9)—most agents withdraw from a given round—at the cost of increasing the simulation of the sub-auction \mathcal{A} —the second term in (4.9). Returning to our earlier thought experiment where we had access to the distribution over the valuations, the term $Q(c; k) \log n$ is essentially the overhead which we incur given that we do not possess any prior information. Yet, if k does not depend on n and we examine the asymptotic growth of the expected communication complexity w.r.t. the number of agents n , we obtain the following:

Corollary 4.4.9 (Single-Item Auction with Optimal Communication). *Let \mathcal{A} be an algorithm faithfully simulating a second-price auction, and $N = [n]$ a set of agents. For any $\epsilon > 0$ and $c = \Theta(1/\epsilon^2)$, denote by $t(n; c, k)$ the expected communication complexity of `ASCENDINGAUCTIONVIASAMPLING`(N, \mathcal{A}, ϵ). If k is a constant independent of n ,*

$$t(n; c, k) \lesssim (1 + \epsilon)n.$$

⁸We do not claim that our analysis w.r.t. the size of the sample is tight; the rather crude bound $c = \Theta(1/\epsilon^2)$ is an artifact of our analysis, but nonetheless it will suffice for our purposes.

Remark. It is important to point out that the elicitation pattern in our proposed mechanism (ASCENDINGAUCTIONVIASAMPLING) is highly *asymmetrical*. Indeed, while most of the agents will be eliminated after the first round, having only revealed a *single* bit from their valuations, the agents who are close to winning the item will have to disclose a substantial amount of information; arguably, this property is desirable. Notice that this is in stark contrast to a standard English auction in which every withdrawing bidder approximately reveals her valuation.

4.4.2 Multi-Item Auction with Additive Valuations

As an extension of the previous setting, consider that the auctioneer has to allocate m (indivisible) items to n agents, with the valuation space being *additive*; that is, for every agent i and for any bundle of items $S \neq \emptyset$,

$$v_i(S) = \sum_{j \in S} v_{i,j},$$

where recall that $v_{i,j}$ represents the value of item j for agent i . Naturally, we are going to employ an ASCENDINGAUCTIONVIASAMPLING for every item. It should be clear that—by virtue of Theorem 4.4.1—the induced mechanism implements with probability 1 the VCG outcome; every item is awarded to the agent who values it the most, and the second-highest valuation for that particular item is imposed as the payment. Moreover, Theorem 4.4.2 implies the following:

Proposition 4.4.10. *Employing for every item ASCENDINGAUCTIONVIASAMPLING yields an ex-post incentive compatible multi-item auction with additive valuations.*

Our main insight in this domain is that a *simultaneous* implementation can lead to a much more communication-efficient interaction process.

Sequential Implementation. First, assume that we have to perform an independent and separate auction for each item. Then, Fact 4.4.4 implies that our mechanism has to elicit at least $n \cdot m$ bits. As in the single-item setting, we can asymptotically match this lower bound.

Proposition 4.4.11. *Consider a set of agents $N = [n]$ with additive valuations for m (indivisible) items, and denote by $t(n; m, c, k)$ the expected communication complexity of employing for each item ASCENDINGAUCTIONVIASAMPLING. Then, for any $\epsilon > 0$ and with k assumed a constant independent of n ,*

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

Simultaneous Implementation. On the other hand, we will show that the communication complexity can be substantially reduced when the m auctions are

performed in *parallel*. Specifically, our approach employs some ideas from information theory in order to design a more efficient encoding scheme. More concretely, let us first describe the general principle. Consider a discrete random variable that has to be *encoded* and subsequently transmitted to some receiver; it is well understood in coding theory that the values of the random variable which are more likely to be observed have to be encoded with relatively fewer bits, so that the communication complexity is minimized in expectation.

Now the important link is that in the `ASCENDINGAUCTIONVIASAMPLING` with a large sample size c a random agent will most likely withdraw from a given round. Thus, we consider the following encoding scheme: An agent i —remaining active in at least one of the m auctions—will transmit a bit 0 in the case of withdrawal from *all* the auctions; otherwise, i may simply transmit an m_i -bit vector that indicates the auctions that i wishes to remain active, where $m_i \leq m$ is the number of auctions in which i is still active. Although the latter part of the encoding is clearly sub-optimal—given that we have encoded events with substantially different probabilities with the same number of bits, it will be sufficient for our argument. Consider a round of the parallel implementation with n agents and let p be the probability that a random agent will withdraw from every auction in the current round. Given that every player is active in at most m auctions, it follows from the union bound that $1 - p \lesssim 2m/(c + 1)$. Thus, if B represents the total number of bits transmitted in the round, we obtain that

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

As a result, for every $\delta > 0$ and size of sample $c = \Theta(m^2/\delta)$, we have that $\mathbb{E}[B] \leq n(1 + \delta)$. Also note that in expectation only a small fraction of agents will “survive” in a given round of the parallel auction—asymptotically at most $2m/(c + 1)$. Thus, similarly to Theorem 4.4.9, we can establish the following theorem:

Theorem 4.4.12 (Simultaneous Single-Item Auctions). *Consider some set of agents $N = [n]$ with additive valuations for m indivisible items, and assume that k and m are constants independent of n . There exists an encoding scheme such that if $t(n; m, c, k)$ is the expected communication complexity of implementing in parallel an `ASCENDINGAUCTIONVIASAMPLING` for every item, then for any $\epsilon > 0$ and for sufficiently large $c = c(\epsilon, m)$,*

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

4.4.3 Multi-Unit Auction with Unit Demand

Finally, we design a multi-unit auction where m units of the same item are to be disposed to n *unit demand* bidders; naturally, we are interested in the non-trivial case where $m \leq n$. We shall consider two canonical cases.

First, let us assume that the number of units m is a small *constant*. Then, we claim that our approach in the single-item auction can be directly applied. Indeed, we propose an ascending auction in which at every round we invoke some algorithm \mathcal{A} that simulates the VCG outcome—i.e., \mathcal{A} identifies the m agents with the highest valuations, as well as the $(m + 1)$ -highest valuation as the payment—for a random sample. Next, the “market-clearing price” in the sub-auction is announced in order to “prune” the active agents. As a result, we can establish guarantees analogous to Theorems 4.4.1 and 4.4.2 and Theorem 4.4.9; the analysis is similar to the single-item auction, and is therefore omitted.

Our main contribution in this subsection is to address the case where $m = \gamma \cdot n$, for some *constant* $\gamma \in (0, 1)$. Specifically, unlike a standard English auction, our idea is to broadcast in every round *two* separate prices; the agents who are above the high price p_h are automatically declared winners,⁹ while the agents below the lower price p_ℓ will have to quit the auction. Then, the mechanism may simply recurse on the agents that lie in the intermediate region. In this context, we consider the following encoding scheme:

- If $v_i > p_h$, then i transmits a bit of 1;
- If $v_i < p_\ell$, then i transmits a bit of 0;
- otherwise, i may transmit some arbitrary 2-bit code.

Observe that the last condition ensures that the encoding is *non-singular*. In contrast to our approach in the single-item auction, this communication pattern requires the transmission of a 2-bit code from some agents; nonetheless, we will show that this overhead can be negligible, and in particular, the fraction of agents that reside between p_h and p_ℓ can be made arbitrarily small. For simplicity in the exposition, here we will tacitly assume that the agents’ valuations are pairwise distinct. The pseudocode for our mechanism is given in Mechanism 6.

The crux of the `MULTIUNITAUCTIONVIASAMPLING` lies in the implementation of the subroutines at steps 4 and 5. This is addressed in the following theorem:

Theorem 4.4.13. *Consider a set of agents $N = [n]$ and a number of units m . There exists a sampling algorithm such that for any $\epsilon > 0$ and any $\delta > 0$ satisfies the following:*

- *It takes as input at most $4k \log(4k/\delta)/\epsilon^2$ bits.*
- *With probability at least $1 - \delta$ it returns prices p_h and p_ℓ , such that p_h is between the $(m + 1)$ -ranked player and the $(m + 1 + \lceil \epsilon n \rceil)$ -ranked player, and p_ℓ is between the $(m + 1 - \lceil \epsilon n \rceil)$ -ranked player and the $(m + 1)$ -ranked player.*

⁹However, the winners do not actually pay p_h , but rather a common price determined at the final round of the auction; this feature is necessary in order to provide any meaningful incentive compatibility guarantees.

Mechanism 6: MULTIUNITAUCTIONVIASAMPLING(N, m)

Input: Set of agents N , number of items m
Output: VCG outcome (Winners & Payment)
Initialize the winners $W := \emptyset$ and the losers $L := \emptyset$
 $p_h := \text{ESTIMATEUPPERBOUND}(N, m)$
 $p_\ell := \text{ESTIMATELOWERBOUND}(N, m)$
Announce p_h and p_ℓ
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
 | Set $m := m - |\{i \in N : v_i > p_h\}|$
 | Set $N := N \setminus (W \cup L)$
 | **return** MULTIUNITAUCTIONVIASAMPLING(N, m)
end

Proof. Consider a perfect binary tree of height k , such that each of the 2^k leaves corresponds to a point on the discretized valuation space, as illustrated in Figure 4.5. Our algorithm will essentially perform *stochastic binary search* on this tree. To be precise, beginning from the root of the tree, we will estimate an additional bit of p_h and p_ℓ in every level of the tree. Let us denote with x_1, x_2, \dots, x_r , with $x_i \in \{0, 1\}$, the predicted bits after r levels. In the current level, we take a random sample S of size c *with replacement*¹⁰ (S here is potentially a *multiset*), and we query every agent $i \in S$ on whether $v_i \leq \overline{x_1 x_2 \dots x_r 0 1 \dots 1}$, where the threshold is expressed in binary representation. Let us denote with X_μ the random estimation derived from the sample, i.e.,

$$X_\mu = \frac{\sum_{i \in S} \mathbb{1}\{v_i \leq \overline{x_1 x_2 \dots x_r 0 1 \dots 1}\}}{c}.$$

Recall that the output of the algorithm consists of two separate k -bit numbers p_h and p_ℓ . For convenience, a sample will be referred to as ϵ -*ambiguous* if $|X_\mu - \gamma| < \epsilon$, where $\gamma = m/n$ and $\epsilon > 0$ some parameter. Intuitively, whenever the sample is *unambiguous* we can branch with very high confidence; that is, we predict a bit of 1 if $X_\mu < \gamma$, and a bit of 0 if $X_\mu > \gamma$. In contrast, in every ϵ -ambiguous junction the “high” estimation—corresponding to p_h —will predict a bit of 1, whilst the “lower” estimation—corresponding to p_ℓ —will predict a bit of 0. One should imagine that the two estimators initially coincide, until they separate when a “close” decision arises (see Figure 4.5). We claim that this algorithm will terminate with high

¹⁰We assume sampling with replacement to slightly simplify the analysis; our approach is also directly applicable when the sampling occurs *without* replacement.

probability with the desired bounds. For our analysis we will employ the following standard lemma:

Lemma 4.4.14 (Chernoff-Hoeffding Bound). *Let $\{X_1, X_2, \dots, X_c\}$ be a set of i.i.d. random variables with $X_i \sim \text{Bern}(p)$ and $X_\mu = (X_1 + X_2 + \dots + X_c)/c$; then,*

$$\mathbb{P}(|X_\mu - p| \geq \epsilon) \leq 2e^{-2\epsilon^2 c}.$$

The main observation is that if for all samples X_μ has at most $\epsilon/2$ error, then the estimations p_h and p_ℓ will be within the desired range in our claim. Let us denote with p_e the probability that for a single estimate and after k levels there exists a sample with more than $\epsilon/2$ error; the union bound implies that $p_e \leq 2ke^{-\epsilon^2 c/2}$. Thus, for any $\epsilon > 0$ and $\delta > 0$, $p_e \leq \delta$ for $c \geq 2\log(2k/\delta)/\epsilon^2$. Furthermore, by the union bound, we obtain that the probability of error of either of the two estimates with input at most $2ck$ bits is at most 2δ ; rescaling $\delta := \delta/2$ concludes the proof. \square




Figure 4.5: The binary-tree representation of the valuation space. The red lines correspond to potential branching of the two estimates.

Consequently, the algorithm described in Theorem 4.4.13 will be employed to implement lines 4 and 5 in our `MULTIUNITAUCTIONVIASAMPLING`. In addition, notice that we can recognize whenever the estimated prices p_h and p_ℓ are incongruous—in the sense that either the winners are more than the available items, or that the remaining agents are less than the available items, in which case we simply repeat the estimation. Thus, we obtain the following properties:

Proposition 4.4.15. *The `MULTIUNITAUCTIONVIASAMPLING` is ex-post incentive compatible.*

Theorem 4.4.16 (Multi-Unit Auction with Optimal Communication). *Consider some set of unit demand agents $N = [n]$, and m identical units. Moreover, denote*

by $t(n; k)$ the expected communication complexity of the mechanism `MULTIUNITAUCTIONVIASAMPLING`(N, m), with steps 4 and 5 implemented through the algorithm of Theorem 4.4.13. If $k = \mathcal{O}(n^{1-\ell})$ for some $\ell > 0$, then for any $\epsilon > 0$,

$$t(n; k) \lesssim (1 + \epsilon)n.$$

Proof. Theorem 4.4.13 implies that for any $\epsilon > 0$ and $\delta > 0$,

$$t(n; k) \leq (1 - \delta)((1 + 2\epsilon)n + t(2en; k)) + \delta(2n + t(n; k)) + 4k \frac{\log(4k/\delta)}{\epsilon^2},$$

where the first term corresponds to the induced communication when the sampling algorithm of Theorem 4.4.13 succeeds, the second term is the worst-case communication whenever the sampling algorithm fails to return prices within the desired range, while the last term is the communication of the sampling algorithm. Thus, solving the induced recursion and using that $k = \mathcal{O}(n^{1-\ell})$ concludes the proof. \square

4.A Sampling Approximation of the Plurality Rule

Approximating the plurality rule is quite folklore in social choice; e.g., see [82, 58]. More recently, Bhattacharyya and Dey [82] analyzed the sample complexity of determining the outcome in several common voting rules under a γ -margin condition; that is, they assumed that the minimum number of votes that need to be modified in order to alter the winner is at least $\gamma \cdot n$. In fact, a standard result by Canetti et al. [58] establishes that at least $\Omega(\log(1/\delta)/\gamma^2)$ samples are required in order to determine the winner in the plurality rule with probability at least $1 - \delta$, even with 2 candidates. This lower bound might appear rather unsatisfactory; for one thing, the designer does not typically have any prior information on the margin γ . More importantly, in many practical scenarios we expect the margin to be negligible, leading to a substantial overhead in the sample complexity.

The purpose of this section is to show that these obstacles could be circumvented once we take a *utilitarian* approach. Indeed, instead of endeavoring to determine the winner in the election with high probability, we are interested in approximating the social welfare. More precisely, assume that n agents have to select among m alternatives or candidates. We let $u_{i,j}$ represent the score that agent $i \in [n]$ assigns to candidate $j \in [m]$. In this way, the social welfare of an outcome $j \in [m]$ is defined as

$$\text{SW}(j) = \sum_{i=1}^n u_{i,j}.$$

Notice, however, that the social welfare approximation problem through the plurality rule is hopeless under arbitrary valuations, in light of obvious information-theoretic barriers. (the framework of *distortion* introduced by Procaccia and

Rosenschein [193] quantifies exactly these limitations.) For this reason, and for the sake of simplicity, we are considering the social welfare approximation problem in a very simplistic setting.

Definition 4.A.1. *A voter i is said to be single-minded if $u_{i,r} = 1$ for some candidate $r \in [m]$, and $u_{i,j} = 0, \forall j \neq r$.*

In fact, in the simple setting where all agents are single-minded it is easy to see that the plurality rule is actually strategy-proof. Recall that in the plurality rule every agent i votes for a *single* candidate $j \in [m]$, i.e. agent i broadcasts an m -tuple $(0, 0, \dots, 1, 0)$, assigning 1 to the position that corresponds to her preferred candidate. We are now ready to analyze the sampling approximation of the plurality rule.

Mechanism 7: APPROXPLURALITYVIASAMPLING($N, [m], \epsilon, \delta$)

Input: Set of agents N , set of candidates $[m]$, accuracy parameter $\epsilon > 0$, confidence parameter $\delta > 0$

Output: Candidate $j \in [m]$ such that $\text{SW}(j) \geq (1 - \epsilon) \text{SW}(j^*)$ with probability at least $1 - \delta$, where j^* represents the optimal candidate

Set $c = 2m^2 \log(2m/\delta)/\epsilon^2$, the size of the sample

Let S be a random sample¹¹ of c agents from N

return PLURALITY(S)

Theorem 4.A.2. *Consider a set of single-minded agents N , and any number of candidates m . For any $\epsilon > 0$ and any $\delta > 0$,*

APPROXPLURALITYVIASAMPLING($N, [m], \epsilon, \delta$) yields with probability at least $1 - \delta$ an approximation ratio of $1 - \epsilon$ w.r.t. the optimal social welfare of the full information PLURALITY, for any $c \geq c_0(m, \epsilon, \delta)$, where

$$c_0(m, \epsilon, \delta) = \frac{2m^2 \log(2m/\delta)}{\epsilon^2}.$$

The proof of this theorem is simple and proceeds with a standard Chernoff bound argument; for completeness, we provide a detailed proof.

Proof of Theorem 4.A.2. First, let us denote with $s_j = (\sum_{i=1}^n u_{i,j})/n$ the score of the j^{th} candidate in the full information plurality rule. Consider a sample of size c , and let $(X_{i,1}, \dots, X_{i,m})$ represent the voting profile of the agent that was selected in the i^{th} iteration of the sampling process. It follows that $X_{i,j} \sim \text{Bern}(s_j), \forall j \in [m]$. Moreover, recall that we consider sampling with replacement, so that the set of random variables $\{X_{1,j}, \dots, X_{c,j}\}$ is pairwise independent for any $j \in [m]$; our results are also applicable when the sampling occurs without replacement given

¹¹Here we assume that we sample with replacement.

that the correlation between the random variables is negligible – for sufficient large n , although we do not formalize this here.

As a result, if we denote with $\widehat{s}_j = (\sum_{i=1}^c X_{i,j})/c$, a standard Chernoff bound argument implies that $\forall \epsilon' > 0, \forall \delta' > 0$, and $c \geq \log(2/\delta')/(2(\epsilon')^2)$, $|\widehat{s}_j - s_j| \leq \epsilon'$ with probability at least $1 - \delta'$. By the union bound, we obtain that $|\widehat{s}_j - s_j| \leq \epsilon'$ for all $j \in [m]$ and with probability at least $1 - m\delta'$. Thus, we let $\delta' = \delta/m$ for some arbitrary $\delta > 0$. We also let $s^* = s_{j^*}$ to be the score of the most preferred candidate j^* – among the entire population N . By definition, the optimal social welfare is $\text{SW}^* = \text{SW}(j^*) = \sum_{i=1}^n u_{i,j^*} = ns^*$. If $r = \arg \max_j \widehat{s}_j$ is the random output of $\text{APPROXPLURALITYVIASAMPLING}(N, [m], \epsilon', \delta)$, it follows that $s_r \geq s^* - 2\epsilon'$ with probability at least $1 - \delta$. Thus, we obtain that

$$\begin{aligned} \text{SW}(r) &= \sum_{i=1}^n u_{i,r} = ns_r \geq n(s^* - 2\epsilon') \\ &= \text{SW}^* - 2\epsilon' \frac{\text{SW}^*}{s^*} \\ &\geq \text{SW}^*(1 - 2\epsilon'm), \end{aligned}$$

where in the final inequality we used that $s^* \geq 1/m$. Finally, setting $\epsilon' = \epsilon/(2m)$ for an arbitrary $\epsilon > 0$ concludes the proof. \square

Chapter 5

Metric-Distortion Bounds under Limited Information

Research in the metric distortion framework was initiated by Anshelevich et al. [18]. Specifically, they analyzed the distortion of several common voting rules, most notably establishing that Copeland’s rule has distortion at most 5, with the bound being tight for certain instances. They also conjectured that the *ranked pairs* mechanism always achieves distortion at most 3, which is also the lower bound for any deterministic mechanism. This conjecture was disproved by Goel et al. [132],¹ while they also studied *fairness* properties of certain voting rules. Moreover, Skowron and Elkind [214] established that a popular rule named *single transferable vote* (STV) has distortion $\mathcal{O}(\log m)$, along with a nearly-matching lower bound. The barrier of 5 set out by Copeland was broken by Munagala and Wang [183], presenting a novel deterministic rule with distortion $2 + \sqrt{5}$. The same bound was obtained by Kempe [149] through an LP duality framework, who also articulated sufficient conditions for proving the existence of a deterministic mechanism with distortion 3. This conjecture was only recently confirmed by Gkatzelis et al. [131], introducing the *plurality matching* mechanism. Closely related to our study is also the work of Gross et al. [138], wherein the authors provide a near-optimal mechanism that only asks $m + 1$ voters for their top-ranked alternatives. The main difference with our setting is that we require an efficiency guarantee with high probability, and not in expectation.

Broader Context. Beyond the metric case most focus has been on analyzing distortion under a *unit-sum* assumption on the utility function, ensuring that agents have equal “weights”. In particular, Boutilier et al. [52] provide several upper and lower bounds, while they also study learning-theoretic aspects under the premise that every agent’s utility is drawn from a distribution; cf., see the

¹A tight bound of $\Theta(\sqrt{m})$ for the ranked pairs mechanism was subsequently given by Kempe [149].

work of Procaccia et al. [196]. Moreover, several multi-winner extensions have been studied in the literature. Caragiannis et al. [66] study the *committee selection* problem, which consists of selecting k alternatives that maximize the social welfare, assuming that the value of each agent is defined as the maximum value derived from the committee’s members. We also refer to Benadè et al. [37] for the *participatory budgeting* problem, and to Benadè et al. [38] when the output of the mechanism should be a total order over alternatives (instead of a single winner).

More special metric spaces have been considered by Feldman et al. [104], Fain et al. [97], Anagnostides et al. [8], strengthening some of the results we previously described. The trade-off between efficiency and communication has been addressed by Mandal et al. [171, 172], while Amanatidis et al. [6] investigated the decay of distortion under a limited amount of cardinal queries—in addition to the ordinal information. We should also note a series of works analyzing the power of ordinal preferences for some fundamental graph-theoretic problems [110, 15, 16, 17]. Finally, we point out that strategic issues are typically ignored within this line of work. We will also posit that agents provide truthfully their preferences, but we refer to Bhaskar et al. [42], Caragiannis et al. [65, 67] for rigorous considerations on the strategic issues that arise. We refer the interested reader to the excellent survey of Anshelevich et al. [22], as we have certainly not exhausted the literature.

5.1 Contributions

First, we study voting rules which perform a sequence of pairwise comparisons between two candidates, with the result of each comparison being determined by the majority rule over the entire population of voters. This class includes many common mechanisms such as Copeland’s rule [206] or the minimax scheme of Levin and Nalebuff [162], and has received considerable attention in the literature of social choice; cf., see the discussion of Lang et al. [157], and references therein. Within the framework of (metric) distortion, the following fundamental question arises:

How many pairwise comparisons between two candidates are needed to guarantee non-trivial bounds on the distortion?

For example, Copeland’s rule (and most of the common voting rules within this class) elicits all possible pairwise comparisons, i.e. $\binom{m}{2} = \Theta(m^2)$, and guarantees distortion at most 5 [18]. Thus, it is natural to ask whether we can substantially truncate the number of elicited pairwise comparisons without sacrificing too much the efficiency of the mechanism. We stress that we allow the queries to be dynamically adapted during the execution of the algorithm. In this context, we provide the following strong positive result:

Theorem 5.1.1. *There exists a deterministic mechanism that elicits only $m - 1$ pairwise comparisons and guarantees distortion $\mathcal{O}(\log m)$.*

The corresponding mechanism is particularly simple and natural: In every round we arbitrarily pair the remaining candidates and we only extract the corresponding comparisons. Next, we eliminate all the candidates who “lost” and we continue recursively until a single candidate emerges victorious. Interestingly, this mechanism is widely employed in practical applications, for example in the knockout phases of many competitions, with the difference that typically some “prior” is used in order to construct the pairings. The main technical ingredient of the analysis is a powerful lemma developed by Kempe via an LP duality argument [149]. Specifically, Kempe characterized the social cost ratio between two candidates when there exists a sequence of intermediate alternatives such that every candidate in the chain pairwise-defeats the next one. We also supplement our analysis for this mechanism with a matching lower bound on a carefully constructed instance (Theorem 5.3.6). Moreover, we show that any mechanism which performs (strictly) fewer than $m - 1$ pairwise comparisons has *unbounded* distortion (Theorem 5.3.1). This limitation applies even if we allow randomization either during the elicitation or the winner determination phase. Indeed, there are instances for which only a single alternative can yield bounded distortion, but the mechanism simply does not have enough information to identify the “right” candidate.

Next, we study deterministic mechanisms which only receive an incomplete order of preferences from every voter, instead of the entire rankings. This setting has already received attention in the literature, most notably by Kempe [150], and has numerous applications in real-life voting systems. Arguably the most important such consideration arises when every voter provides her k -top preferences, for some parameter $k \in [m]$. Kempe [150] showed that there exists a deterministic mechanism which elicits only the k -top preferences and whose distortion is upper-bounded by $79m/k$; using a powerful tool developed in [149] this bound can be improved all the way down to $12m/k$. However, this still leaves a substantial gap with respect to the best-known lower bound, which is $2m/k$ if we ignore some additive constant factors. Thus, [150] left as an open question whether the aforementioned upper bound can be improved. In our work, we make substantial progress towards bridging this gap, proving the following:

Theorem 5.1.2. *There exists a deterministic mechanism that only elicits the k -top preferences and yields distortion at most $6m/k + 1$.*

We should stress that the constant factors are of particular importance in this framework; indeed, closing the gap even for the special case of $k = m$ has received intense scrutiny in recent years [18, 183, 149, 131]. From a technical standpoint, the main technique for proving such upper bounds consists of identifying a candidate for which there exists a path to any other node such that every candidate in the path pairwise-defeats the next one by a sufficiently large margin (which depends on k). Importantly, the derived upper bound crucially depends on the length of the path. Our main technical contribution is to show that there always exists a

path of length 2 with the aforescribed property, while the previous best result by Kempe established the claim only for paths of length 3.

Although our approach can potentially bring further improvements, closing the gap inevitably requires different techniques. In particular, a promising direction appears to stem from extending some of the claims established by Gkatzelis et al. [131]. Indeed, we observe that a natural generalization of their main technical ingredient would lower the upper bound to $4m/k - 1$ (Theorem 5.4.7), which appears to be optimal when k is close to m . More precisely, Gkatzelis et al. [131] proved that a certain graph always has a perfect matching when the entire rankings are available; we conjecture that under k -top preferences there always exists a perfect matching for a subset of a k/m fraction of the voters (see Theorem 5.4.6 for a more precise statement).

We also provide some other important bounds for deterministic mechanisms under missing information. Most notably, if the voting rule performs well on an arbitrary (potentially adversarially selected) subset of the voters can we quantify its distortion over the entire population? We answer this question with a sharp upper bound in Theorem 5.4.3. In fact, we use this result as a tool for some of our other proofs, but nonetheless we consider it to be of independent interest. It should be noted that even in the realm of partial or incomplete rankings there exists an *instance-optimal* mechanism via linear programming; this was first observed by Goel et al. [132] when the total orders are available, but it directly extends in more general settings. Interestingly, we show that the recently introduced mechanism of Gkatzelis et al. [131] which always obtains distortion at most 3 can be substantially outperformed by the LP mechanism. Namely, for some instances the mechanism of Gkatzelis et al. [131] yields distortion almost 3, while the instance-optimal mechanism yields distortion close to 1.

Finally, we consider mechanisms which receive information from only a “small” *random sample* of voters; that is, we are concerned with the *sample complexity* required to ensure efficiency, which boils down to the following fundamental question:

How large should the size of the sample be in order to guarantee near-optimal distortion with high probability?

More precisely, we are interested in deriving sample-complexity bounds which are *independent* of the number of voters n . This endeavor is particularly motivated given that in most applications $n \gg m$. Naturally, sampling approximations are particularly standard in the literature of social choice. Indeed, in many scenarios one wishes to predict the outcome of an election based on a small sample (e.g. in polls or exit polls), while in many other applications it is considered even infeasible to elicit the entire input (e.g. in online surveys). In this context, we will be content with obtaining near-optimal distortion *with high probability* (e.g. 99%). This immediately deviates from the line of research studying randomized mechanisms wherein it suffices to obtain a guarantee in expectation; cf., see Anshelevich and Postl [14]. We point out that it has been well-understood that a guarantee only

in expectation might be insufficient in many cases; for example, Fain et al. [98] considered as the objective the *squared distortion* as a proxy in order to limit as much as possible the variance in the distortion. In fact, Fain et al. [98] are also concerned with sample complexity issues, but from a very different standpoint.

We stress that we only allow randomization during the preference elicitation phase; for a given random sample, which corresponds to the *entire* rankings of the voters, the mechanisms we consider act deterministically. Specifically, we analyze two main voting rules along this vein.

Theorem 5.1.3 (Approximate Copeland). *For any sufficiently small $\epsilon > 0$ there exists a mechanism that takes a sample of size $\Theta(\log(m)/\epsilon^2)$ voters and yields distortion at most $5 + \epsilon$ with probability 0.99.*

The techniques required for the proof of this theorem are fairly standard. More importantly, we analyze the sample complexity of PLURALITYMATCHING, the mechanism of Gkatzelis et al. [131] which recovers the optimal distortion bound of 3 (among deterministic mechanisms). In this context, we establish the following result:

Theorem 5.1.4 (Approximate PLURALITYMATCHING). *For any sufficiently small $\epsilon > 0$ there exists a mechanism that takes a sample of size $\Theta(m/\epsilon^2)$ voters and yields distortion at most $3 + \epsilon$ with probability 0.99.*

More precisely, the main ingredient of PLURALITYMATCHING is a maximum-matching subroutine for a certain bipartite graph. Our first observation is that the size of the maximum matching can be determined through a much smaller graph which satisfies a “proportionality” condition with respect to a maximum-matching decomposition. Although this condition cannot be explicitly met since the algorithm is agnostic to the decomposition, our observation is that sampling (with sufficiently many samples) will approximately satisfy this requirement, eventually leading to the desired conclusion.

We stress that we do *not* guarantee that the winner in our sample will coincide with that over the entire population. In fact, the sample complexity bounds for the *winner determination* problem—for virtually every reasonable voting rule—depend on the *margin* of victory; see the work of Dey and Bhattacharyya [82]. However, we argue that this feature is undesirable. For one thing, the algorithm does not have any prior information on the margin, and hence it is unclear how to tune this parameter in practice. More importantly, in many scenarios the margin might be very small, leading to a substantial overhead in the sample-complexity requirements of the mechanism. One of our conceptual contributions is to show that we can circumvent such limitations once we espouse a utilitarian framework. Indeed, all of our bounds are *distribution-independent* (and instance-oblivious).

We should also point out that, although we are emphasizing sample-complexity considerations, we believe that our results have another very clear motivation. Namely, given that in most applications $n \gg m$, it is important to provide *sublinear*

algorithms whose running time does not depend on n . In this context, we provide a Monte Carlo implementation of PLURALITYMATCHING whose time complexity scales independently of n .

5.2 Preliminaries

A *metric space* is a pair (\mathcal{M}, d) , where $d : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ constitutes a *metric* on \mathcal{M} , i.e., (i) $\forall x, y \in \mathcal{M}, d(x, y) = 0 \iff x = y$ (identity of indiscernibles), (ii) $\forall x, y \in \mathcal{M}, d(x, y) = d(y, x)$ (symmetry), and (iii) $\forall x, y, z \in \mathcal{M}, d(x, y) \leq d(x, z) + d(z, y)$ (triangle inequality). Consider a set of n voters $V = \{1, 2, \dots, n\}$ and a set of m candidates $C = \{a, b, \dots\}$; candidates will be typically represented with lowercase letters such as a, b, w, x , but it will be sometimes convenient to use numerical values as well. We assume that every voter $i \in V$ is associated with a point $v_i \in \mathcal{M}$, and every candidate $a \in C$ to a point $c_a \in \mathcal{M}$. Our goal is to select a candidate x who minimizes the *social cost*: $cost(x) = \sum_{i=1}^n d(v_i, c_x)$. This task would be trivial if we had access to the agents' distances from all the candidates. However, in the standard *metric distortion* framework every agent i provides only a *ranking* (a total order) σ_i over the points in C according to the *order* of i 's distances from the candidates. We assume that ties are broken arbitrarily, subject to transitivity, but we will not abuse the tie-breaking assumption.

In this work, we are considering a substantially more general setting wherein every agent provides a subset of σ_i . More precisely, we assume that agent i provides as input a set \mathcal{P}_i of ordered pairs of distinct candidates, such that $(a, b) \in \mathcal{P}_i \implies a \succ_i b$, where $a, b \in C$; it will always be assumed that \mathcal{P}_i corresponds to the *transitive closure* of the input. We will allow \mathcal{P}_i to be the empty set, in which case i does not provide any information to the mechanism; with a slight abuse of notation, we will let $\mathcal{P}_i \equiv \sigma_i$ when i provides the entire order of preferences. We will say that the input $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_n)$ is consistent with the metric d if $(a, b) \in \mathcal{P}_i \implies d(v_i, c_a) \leq d(v_i, c_b), \forall i \in V$, and this will be denoted with $d \triangleright \mathcal{P}$. We will represent with $\text{top}(i)$ and $\text{sec}(i)$ i 's first and second most preferred candidates, respectively. We may also sometimes use the notation $ab = \{i \in V : a \succ_i b\}$.

A *deterministic social choice rule* is a function that maps an *election* in the form of a 3-tuple $\mathcal{E} = (V, C, \mathcal{P})$ to a single candidate $a \in C$. We will measure the performance of f for a given input of preferences \mathcal{P} in terms of its *distortion*, namely, the worst-case approximation ratio it provides with respect to the social cost:

$$\text{distortion}(f; \mathcal{P}) = \sup \frac{\text{cost}(f(\mathcal{P}))}{\min_{a \in C} \text{cost}(a)}, \quad (5.1)$$

where the supremum is taken over all metrics such that $d \triangleright \mathcal{P}$. That is, once the mechanism selects a candidate (or a distribution over candidates if the social choice rule is *randomized*), an adversary can select any metric space subject to being consistent with the input preferences. Similarly, in Section 5.3 where we

study mechanisms that perform pairwise comparisons, the adversary can select any metric space consistent with the elicited comparisons. The distortion of a social choice rule f is defined as the supremum of $\text{distortion}(f; \mathcal{P})$ (5.1) taken over all possible input preferences \mathcal{P} (and possible values of n), under a fixed value for the number of candidates m . We point out the following:

Proposition 5.2.1. *Under any given preferences \mathcal{P} , there exists a metric space consistent with \mathcal{P} .*

This proposition follows immediately from Proposition 1 of Anshelevich et al. [18], which established the claim when $\mathcal{P} = \sigma$.

5.2.1 Instance-Optimal Voting

An important observation is that under any input preferences \mathcal{P} there exists a deterministic *instance-optimal* mechanism; this was noted by Goel et al. [132] (see also [52]) when $\mathcal{P} = \sigma$, but their mechanism directly applies to our more general setting. We briefly present their idea, as we will also employ this mechanism for our experiments.

The first ingredient is an optimization problem that allows to compare a pair of distinct candidates, subject to the set of preferences given to the mechanism. Specifically, for $a, b \in C$, with $a \neq b$, consider the following *linear program* METRIC-LP(a, b):

$$\begin{aligned}
& \text{maximize} && \sum_{i=1}^n x_{i,a} \\
& \text{subject to} && \sum_{i=1}^n x_{i,b} = 1; \\
& && x_{i,p} \leq x_{i,q}, \quad \forall (p, q) \in \mathcal{P}_i, \forall i \in V; \\
& && x_{i,i} = 0, \quad \forall i \in V \cup C; \\
& && x_{i,j} = x_{j,i}, \quad \forall i, j \in V \cup C; \\
& && x_{i,j} \leq x_{i,k} + x_{k,j}, \forall i, j, k \in V \cup C.
\end{aligned} \tag{5.2}$$

It should be pointed out that some of the constraints in METRIC-LP (5.2) are redundant, in the sense that they are implied by others, but we will not dwell on such optimizations here. We will represent with $\mathfrak{D}(a|b)$ the value of the linear program METRIC-LP(a, b); if it is unbounded, we let $\mathfrak{D}(a|b) = +\infty$. We also note that the linear program is always feasible by virtue of Theorem 5.2.1. In this context, the mechanism of Goel et al. [132] consists of the following steps:

- For any pair $a, b \in C$, with $a \neq b$, compute $\mathfrak{D}(a|b)$; also let $\mathfrak{D}(a|a) = 1$.
- Set $\mathfrak{D}(a) = \max_{b \in C} \mathfrak{D}(a|b)$.
- Return the candidate b with the minimum value $\mathfrak{D}(a)$ over all $a \in C$; ties are broken arbitrarily.

This mechanism will be referred to as MINIMAX-LP to distinguish from the *minimax* voting scheme of Simpson and Kramer [162]. The MINIMAX-LP rule essentially performs brute-force search over all possible metrics in order to identify the candidate who minimizes the distortion; nonetheless, it can be solved in $\text{poly}(n, m)$ time given that the METRIC-LP admits a strongly polynomial time algorithm; this follows because the *bit complexity* L —the number of bits required to represent it [147]—is small: $L = \mathcal{O}(\log(n+m))$. Moreover, it is easy to establish the following:

Theorem 5.2.2. *For any given preferences \mathcal{P} , the MINIMAX-LP rule is instance-optimal in terms of distortion.*

In particular, when $\mathcal{P} = \sigma$ we note that MINIMAX-LP always yields distortion at most 3 by virtue of the upper-bound of Gkatzelis et al. [131]. Nevertheless, in this work we will be mostly interested in providing upper bounds on the distortion of MINIMAX-LP under incomplete rankings.

5.3 Sequence of Pairwise Comparisons

In this section, we are considering voting rules that perform a sequence of pairwise comparisons between two candidates, with the result of each comparison being determined by the majority rule over the entire population of voters. To put it differently, consider the tournament graph $T = (C, E)$ where $(a, b) \in E$ if and only if candidate a pairwise-defeats candidate b ; that is, $|ab| \geq n/2$. (It will be tacitly assumed without any loss of generality that ties are broken arbitrarily so that T is indeed a tournament.) We are studying mechanisms which elicit edges from T , and we want to establish a trade-off between the number of elicited edges and the distortion of the mechanism. We commence with the following lower bound:

Proposition 5.3.1. *There are instances for which any deterministic mechanism which elicits (strictly) fewer than $m - 1$ edges from T has unbounded distortion.*

Sketch of Proof. Consider a family of tournaments \mathcal{T} as illustrated in Figure 5.1, with the set C^* containing a single candidate. Then, there are metric spaces for which all the voters are arbitrarily close to the candidate in C^* and arbitrarily far from any other candidate. Thus, any mechanism with bounded distortion has to identify the candidate in C^* . However, any pairwise comparison can eliminate at most one candidate from being in C^* . As a result, if $\widehat{T} = (C, \widehat{E})$ is the subgraph based on the elicited edges, there will be at least two distinct candidates which could lie in C^* for some tournament in \mathcal{T} consistent with \widehat{T} , leading to the desired conclusion. \square

In fact, the same limitation applies even if we allow randomization, either during the elicitation or the winner determination phase. On the other hand, we




Figure 5.1: A hard class of tournament graphs when fewer than $m - 1$ pairwise comparisons are elicited (see Theorem 5.3.1). C^* contains a single candidate whose is located arbitrarily closed to all the voters in the underlying metric space.

will show that $m - 1$ edges from T suffice to obtain near-optimal distortion. To this end, we will employ a powerful technical lemma due to Kempe, proved via an LP-duality argument.

Lemma 5.3.2 ([149]). *Let a_1, a_2, \dots, a_ℓ be a sequence of distinct candidates such that for every $i = 2, \dots, \ell$ at least half of the voters prefer candidate a_{i-1} over candidate a_i . Then, $\text{cost}(a_1) \leq (2\ell - 1)\text{cost}(a_\ell)$.*

Armed with this important lemma, we introduce the DOMINATIONROOT mechanism. DOMINATIONROOT determines a winning candidate with access only to a pairwise comparison oracle; namely, \mathfrak{D} takes as input two distinct candidates $a, b \in C$ and returns the *losing* candidate based on the voters' preferences (recall that in case of a tie the oracle returns an arbitrary candidate).

Mechanism 8: DOMINATIONROOT

Input: Set of candidates C , Pairwise comparison oracle \mathfrak{D} ;

Output: Winner $w \in C$;

1. Initialize $S := C$;
 2. Construct arbitrarily a set Π of $\lfloor S/2 \rfloor$ pairings from S ;
 3. For every $\{a, b\} \in \Pi$ remove $\mathfrak{D}(a, b)$ from S ;
 4. If $|S| = 1$ **return** $w \in S$; otherwise, continue from step 2;
-

We refer to Figure 5.2 for an illustration of DOMINATIONROOT. The analysis of this mechanism boils down to the following simple claims:

Claim 5.3.3. DOMINATIONROOT *elicits exactly $m - 1$ edges from T .*

Proof. The claim follows given that for every elicited edge we remove a candidate for the rest of the mechanism, until only a single candidate survives. \square




Figure 5.2: An implementation of DOMINATIONROOT for $m = 14$ candidates. Nodes correspond to candidates and edges to pairwise comparisons. We have highlighted with different colors pairwise comparisons that correspond to different rounds of the mechanism. Also, the “height” of every candidate indicates the order of elimination.

Claim 5.3.4. DOMINATIONROOT returns a candidate w which can reach every other node in T in paths of edge-length at most $\lceil \log m \rceil$ in T .

Proof. Consider the partition of candidates C_1, \dots, C_r such that C_i contains the candidates who were eliminated during the i -th round for $i \in \{1, 2, \dots, r-1\}$, and $C_r = \{w\}$. Observe that every candidate $a \in C_i$ (with $i \in \{1, 2, \dots, r-1\}$) was pairwise-defeated by some candidate in C_j for $j > i$; thus, the claim follows inductively since $r = \lceil \log m \rceil$. \square

We now arrive to one of our central results, establishing that only $m-1$ pairwise comparisons suffice to obtain near-optimal distortion:

Theorem 5.3.5. DOMINATIONROOT elicits only $m-1$ edges from T and guarantees distortion at most $2\lceil \log m \rceil + 1$.

Proof. The theorem follows directly from Theorem 5.3.3, Theorem 5.3.4, and Theorem 5.3.2. \square

This theorem, along with Theorem 5.3.1, imply a remarkable gap depending on whether the mechanism is able to elicit at least $m-1$ pairwise comparisons. We also provide a matching lower bound for the distortion of DOMINATIONROOT:


Proposition 5.3.6. There exist instances for which DOMINATIONROOT yields distortion at least $2\log m + 1$.

Proof. We will first show that the bound established in Theorem 5.3.2 is tight. Indeed, consider a set of ℓ candidates $\{1, 2, \dots, \ell\}$, for an even number ℓ , and two voters (the instance directly extends to an arbitrary even number of voters) positioned according to the pattern of Figure 5.3a. Then, the following profile of preferences is consistent with the underlying metric space:

- $1 \succ_1 3 \succ_1 2 \succ_1 5 \succ_1 \dots \succ_1 \ell - 2 \succ_1 \ell$;
- $2 \succ_2 1 \succ_2 4 \succ_2 3 \succ_2 \dots \succ_2 \ell \succ_2 \ell - 1$.

Now observe that—at least under some tie-breaking rule—candidate i pairwise-defeats candidate $i - 1$ for $i = 2, 3, \dots, \ell$. Moreover, it follows that $\text{cost}(i) = 2i - 1$, for all i , implying that $\text{cost}(\ell)/\text{cost}(1) = 2\ell - 1$, as desired.

Next, consider m candidates such that m is a power of 2. We first consider $\ell = \log m + 1$ candidates positioned according to our previous argument (Figure 5.3a); the rest of the candidates are located arbitrarily far from the voters. It is easy to see that there exists a sequence of pairings (Figure 5.3b) such that c_ℓ will be declared victorious, leading to a distortion of $2 \log m + 1$ by virtue of our previous argument. \square



(a) A metric embedding of voters and candidates. (b) A sequence of pairings such that c_ℓ emerges victorious. We have highlighted pairings that correspond to different rounds.

Figure 5.3: Proof of Theorem 5.3.6.

5.4 Distortion of Deterministic Rules under Incomplete Orders

In this section, we complement our previous results that assumed access to pairwise comparisons (Section 5.3) by characterizing the distortion of deterministic voting mechanisms under different classes of incomplete preferences. We commence this section with another useful lemma by [149].

Lemma 5.4.1 ([149]). *Consider three distinct candidates $w, y, x \in C$ such that $|wy| \geq \alpha n$ and $|yx| \geq \alpha n$, with $\alpha \in (0, 1]$. Then,*

$$\frac{\text{cost}(w)}{\text{cost}(x)} \leq \frac{2}{\alpha} + 1.$$

In particular, notice that if w is the winner in Copeland's rule, it follows that for any candidate x there exists some other candidate y such that w pairwise-defeats y and y pairwise-defeats x [182]; thus, applying Theorem 5.4.1 for $\alpha = 1/2$ implies that the winner in Copeland's rule has distortion upper-bounded by 5. This was initially established by Anshelevich et al. [18].

As a warm-up, we first employ this lemma to characterize the distortion when, for all pairs of candidates, at least a small fraction of voters has provided their pairwise preferences. We stress that all of our upper bounds are attainable by the MINIMAX-LP rule, but nonetheless our proofs are constructive in the sense that we provide (efficiently implementable) mechanisms that obtain the desired bounds.

Proposition 5.4.2. *Consider an election $\mathcal{E} = (V, C, \mathcal{P})$ such that for every pair of distinct candidates $a, b \in C$, it holds that $\sum_{i=1}^n \mathbb{1}\{(a, b) \in \mathcal{P}_i \vee (b, a) \in \mathcal{P}_i\} \geq \alpha \cdot n$, with $\alpha \in (0, 1]$. Then, there exists a voting rule that obtains distortion at most $4/\alpha + 1$.*

Proof. Consider a complete, weighted and directed graph $G = (C, E, w)$ such that

$$w_{a,b} = \frac{\sum_{i=1}^n \mathbb{1}\{(a, b) \in \mathcal{P}_i\}}{n}.$$

In words, $w_{a,b}$ represents the fraction of voters who *certainly* prefer a to b ; observe that if we had the complete rankings it would follow that $w_{a,b} + w_{b,a} = 1$, but here we can only say that $w_{a,b} + w_{b,a} \leq 1$. Moreover, by assumption, we know that $w_{a,b} + w_{b,a} \geq \alpha$, implying that $w_{a,b} \geq \alpha/2$ or $w_{b,a} \geq \alpha/2$. With that in mind, we construct from G an unweighted and directed graph $\widehat{G} = (C, \widehat{E})$ according to the following threshold rule: $(a, b) \in \widehat{E} \iff w_{a,b} \geq \alpha/2$. We have argued that our assumption implies that $(a, b) \in \widehat{E} \vee (b, a) \in \widehat{E}$. As a result, we can deduce that \widehat{G} contains as a subgraph a tournament; thus, there exists a *king vertex* w [199] so that every node $a \in C$ is reachable from w in at most 2 steps, and our claim follows directly from Theorem 5.4.1. \square

We remark that this upper bound is tight up to constant factors, at least for certain instances. Indeed, if we only have an α fraction of the votes in the presence of 2 candidates, it is easy to show an $\Omega(1/\alpha)$ lower bound for any mechanism, even if we allow randomization.

5.4.1 Missing Voters

Building on Theorem 5.4.2, consider an election $\mathcal{E} = (V, C, \mathcal{P})$ and a mechanism that has access to the votes of only a subset $V \setminus Q$ of voters, where $Q \subset V$ is the set

of *missing voters* such that $|Q| = \epsilon \cdot n$. If the mechanism performs well on $V \setminus Q$, can we characterize the distortion over the entire set of voters as ϵ increases? Observe that this setting is tantamount to $\mathcal{P}_i = \emptyset$ for all $i \in Q$. In the following theorem we provide a sharp bound:

Theorem 5.4.3. *Consider a mechanism with distortion at most ℓ w.r.t. an arbitrary subset with $(1 - \epsilon)$ fraction of all the voters, for some $\epsilon \in (0, 1)$. Then, the distortion of the mechanism w.r.t. the entire population is upper-bounded by*

$$\ell + \frac{\epsilon}{1 - \epsilon}(\ell + 1).$$

Proof. Consider a candidate $b \in C$ with distortion at most ℓ w.r.t. the agents in $V \setminus Q$. Moreover, consider some arbitrary candidate $a \in C$, and let $S_b = \sum_{i \in V \setminus Q} d(v_i, c_b)$, and $S_a = \sum_{i \in V \setminus Q} d(v_i, c_a)$; observe that, by assumption, $S_b/S_a \leq \ell$. Our analysis will distinguish between the following two cases:

Case I: $S_b \geq S_a > 0$.² Then, for all $i \in Q$ it follows that

$$S_b d(v_i, c_a) + S_a d(c_a, c_b) \geq S_a (d(v_i, c_a) + d(c_a, c_b)) \geq S_a d(v_i, c_b),$$

and hence,

$$S_a d(v_i, c_b) \leq S_a d(c_a, c_b) + S_b d(v_i, c_a) + d(c_a, c_b) d(v_i, c_a);$$

summing over all $i \in Q$ gives

$$\begin{aligned} S_a \sum_{i \in Q} d(v_i, c_b) &\leq |Q| S_a d(c_a, c_b) + S_b \sum_{i \in Q} d(v_i, c_a) + d(c_a, c_b) \sum_{i \in Q} d(v_i, c_a) \\ &\leq |Q| S_a d(c_a, c_b) + S_b \sum_{i \in Q} d(v_i, c_a) + |Q| d(c_a, c_b) \sum_{i \in Q} d(v_i, c_a). \end{aligned} \quad (5.3)$$

Moreover, observe that

$$(5.3) \iff \frac{S_b + \sum_{i \in Q} d(v_i, c_b)}{S_a + \sum_{i \in Q} d(v_i, c_a)} \leq \frac{S_b + |Q| d(c_a, c_b)}{S_a}. \quad (5.4)$$

Next, we have that $d(c_a, c_b) \leq d(v_i, c_a) + d(v_i, c_b), \forall i$; summing over all $i \in V \setminus Q$ implies that $(n - |Q|)d(c_a, c_b) \leq S_a + S_b \leq (\ell + 1)S_a$. Therefore, along with (5.4) we obtain that

$$\frac{\text{cost}(b)}{\text{cost}(a)} \leq \ell + \frac{|Q|}{n - |Q|}(\ell + 1) = \ell + \frac{\epsilon}{1 - \epsilon}(\ell + 1).$$

²The case where $S_a = 0$ can be trivially handled. Indeed, it implies that $S_b \leq \ell \times S_a = 0$, which in turn yields that $d(c_a, c_b) = 0$; thus, $\text{cost}(a) = \text{cost}(b)$.

Case II: $S_b < S_a$. In this case, we can simply observe that

$$\frac{\text{cost}(b)}{\text{cost}(a)} \leq \frac{S_b + \sum_{i \in Q} d(v_i, c_a) + |Q|d(c_a, c_b)}{S_a + \sum_{i \in Q} d(v_i, c_a)} \leq 1 + |Q| \frac{d(c_a, c_b)}{S_a}.$$

Thus, the proof follows given that $(n - |Q|)d(c_a, c_b) \leq S_a + S_b < 2S_a$. \square

A few remarks are in order. First, Borodin et al. [48] provided a similar result when the revealed votes contain complete preferences; it is plausible that their approach can be extended along the line of Theorem 5.4.3. In fact, if all the voters in the set $V \setminus Q$ had provided their entire rankings, we could derive a similar result via Theorem 5.4.2, but, unlike Theorem 5.4.2, Theorem 5.4.3 is parameterized in terms of the distortion with respect to the voters in $V \setminus Q$. Theorem 5.4.3 will be a useful tool in the sequel in order to establish bounds under k -top preferences (Theorem 5.4.7) and random sampling (Theorem 5.5.9), but we consider it to be of independent interest beyond those applications. We also point out that the derived bound in Theorem 5.4.3 is tight for certain instances. For example, consider an instance on the line with only two candidates a and b , so that every candidate receives half of the votes among the voters in $V \setminus Q$; assume without loss of generality that a is selected as the winning candidate, having distortion 3 w.r.t. the voters in $V \setminus Q$. However, we have to accept that $(1 - \epsilon)/2$ fraction of the voters could reside in the midpoint $(c_a + c_b)/2$, while the rest of the agents could all lie in c_b ; thus, the distortion of candidate a reads

$$\frac{\text{cost}(a)}{\text{cost}(b)} = \frac{\frac{1-\epsilon}{2} \frac{d(c_a, c_b)}{2} + \frac{1-\epsilon}{2} d(c_a, c_b) + \epsilon d(c_a, c_b)}{\frac{1-\epsilon}{2} \frac{d(c_a, c_b)}{2}} = 3 + 4 \frac{\epsilon}{1 - \epsilon},$$

which matches the bound of Theorem 5.4.3 when $\ell = 3$.

5.4.2 Top Preferences

In this subsection, we investigate how the distortion increases when every voter provides only her k -top preferences, for some parameter $k \in [m]$. It should be noted that the two extreme cases are by now well understood. Specifically, when $k = m$ the mechanism has access to the entire rankings, and we know that any deterministic mechanism has distortion at least 3, matching the upper bound established by Gkatzelis et al. [131]. On the other end of the spectrum, when $k = 1$, the plurality rule—which incidentally is the optimal deterministic mechanism when only the top preference is given—yields distortion at most $2m - 1$ [18]. Consequently, the question is to quantify the decay of distortion as we gradually increase k . We commence by reminding a lower bound given by Kempe [150]:

Proposition 5.4.4. *Any deterministic mechanism that elicits only the k -top preferences from every voter out of the m alternatives has distortion $\Omega(m/k)$.*

More precisely, the best lower bound is $2m/k$, ignoring some additive constant factors;

Proof. First of all, assume without any loss of generality that $k \mid (m-1)$,³ and let $n = (m-1)/k$ be the number of voters. For simplicity, let us enumerate the number of candidates as $C = \{1, 2, \dots, n \times k\} \cup \{x\}$. Now consider some preference profile \mathcal{P} in which the k -top preferences of voter i correspond to the set of candidates $\{(i-1)k+1, \dots, (i-1)k+k\}$ according to some arbitrary order; observe that all of these sets are pairwise disjoint.

Based on these preferences, the mechanism has to select some candidate. If x is selected, the lower bound follows trivially since x could actually be the last choice for every voter. Therefore, let us assume that candidate 1 was selected by the mechanism; this hypothesis is without loss of generality due to the symmetry of the input \mathcal{P} . The main observation is that the agents and the candidates could be located on the metric space of Figure 5.4. Indeed, it is easy to check that the induced metric space is consistent with the given preferences. As a result, it follows that

$$\frac{\text{cost}(1)}{\text{cost}(x)} = \frac{D + (n-1) \times (\delta + 2D)}{D + (n-1) \times \delta} = \frac{1 + (n-1) \times (\delta/D + 2)}{1 + (n-1) \times \delta/D}. \quad (5.5)$$

Thus, for $\delta/D \downarrow 0$, we obtain that $\text{cost}(1)/\text{cost}(x) \rightarrow 2n-1 = \Omega(m/k)$. \square

We should note that although in our worst-case example the number of voters n is smaller than the number of candidates m , which is not the canonical case, our argument directly extends whenever n is a multiple of $(m-1)/k$, allowing n to be arbitrarily large. Moreover, a similar construction shows an $\Omega(m/k)$ lower bound for α -decisive metrics, in the sense of Anshelevich and Postl [14], for any $\alpha \in [0, 1]$; indeed, it suffices to place the voters within the “cluster” of their k -most preferred candidates.

Theorem 5.4.5. *There exists a deterministic mechanism that elicits only the k -top preferences from every voter out of m candidates and has distortion at most $6m/k + 1$.*

Before we proceed with the proof, it is important to point out that having only the k -top preferences is not subsumed by our previous consideration in Theorem 5.4.2; e.g., even if $k = m-2$, there could be two candidates which lie on the last two positions of every voter’s list, and hence, it is impossible to know which one is mostly preferred among the voters.

³If it is not the case that $k \mid (m-1)$, take k' to be the smallest number larger than k such that $k' \mid (m-1)$, and apply our argument for k' ; given that $k' < 2k$, we will establish again a lower bound of $\Omega(m/k)$, even though the mechanism had more information than the k -top preferences.




Figure 5.4: The metric space considered for the proof of Theorem 5.4.4, where $\delta/D \downarrow 0$ for some positive numbers δ and D . The distance between two points is simply the shortest path in the graph.

Proof of Theorem 5.4.5. Let \mathcal{L}_i be the set with the k -top preferences of voter i . For a candidate $a \in C$, we let

$$\mathcal{V}_a = \frac{\sum_{i=1}^n \mathbb{1}\{a \in \mathcal{L}_i\}}{n};$$

i.e., the fraction of voters for which a is among the k -top preferences. Notice that $\sum_{a \in C} \mathcal{V}_a = k$, and hence, by the pigeonhole principle there exists some candidate x such that $\mathcal{V}_x \geq k/m$. Similarly to Theorem 5.4.2 we consider the weighted, complete and directed graph $G = (C, E, w)$, so that

$$w_{a,b} = \frac{\sum_{i=1}^n \mathbb{1}\{(a, b) \in \mathcal{P}_i\}}{n}.$$

Moreover, based on G we construct the unweighted and directed graph $\widehat{G} = (C, \widehat{E})$, so that $(a, b) \in \widehat{E} \iff w_{a,b} \geq k/(3m)$; the constant $1/3$ in the threshold is selected as the largest number which makes the following argument work. In particular, we will show that \widehat{G} has a king vertex, and then the theorem will follow by virtue of Theorem 5.4.1.

Let $C' = \{a \in C : \exists b \in C \setminus \{a\}. (a, b) \notin \widehat{E} \wedge (b, a) \notin \widehat{E}\}$ and $C^* = C \setminus C'$. Observe that the *induced* subgraph on C^* contains as a subgraph a tournament, and as such, it has a king vertex $w \in C^*$ (we will argue very shortly that indeed $C^* \neq \emptyset$). As a result, if $C' = \emptyset$ the theorem follows.

In the contrary case, C' contains at least two (distinct) nodes; let $a, b \in C'$ be such that $(a, b) \notin \widehat{E} \wedge (b, a) \notin \widehat{E}$. An important observation is that $\mathcal{V}_a \leq 2k/(3m)$ and $\mathcal{V}_b \leq 2k/(3m)$. Indeed, for the sake of contradiction let us assume that $\mathcal{V}_a > 2k/(3m)$. Given that $(a, b) \notin \widehat{E}$ we can infer that b is preferred over a in at least

a $k/(3m)$ fraction of the voters; however, this would imply that $(b, a) \in \widehat{E}$, which is a contradiction. Similarly, we can show that $\mathcal{V}_b \leq 2k/(3m)$. Consequently, x cannot belong in the set C' , where recall that x is a candidate for which $\mathcal{V}_x \geq k/m$, verifying that $C^* \neq \emptyset$.

Next, it is easy to see that for all $a \in C'$, $(x, a) \in \widehat{E}$; this follows since $\mathcal{V}_a \leq 2k/(3m)$ for all $a \in C'$ while $\mathcal{V}_x \geq k/m$. As a result, if $x = w$ or if there exists the edge $(w, x) \in \widehat{E}$, then w can reach every node in at most 2 steps, and the theorem follows. Otherwise, it follows that there exists a path of length 2 from w to x since w is a king vertex in the induced subgraph on C^* and $x \in C^*$. We shall distinguish between two cases.

First, assume that for all $z \in C_1^*$, $(x, z) \in \widehat{E}$, where C_1^* is the subset of C^* which is reachable from w via a single edge. Then, given that we have assumed that $(w, x) \notin \widehat{E}$ and the induced graph on C^* is a tournament, it follows that $(x, w) \in \widehat{E}$, and subsequently x can reach every node in C in paths of length at most 2, as desired.

Finally, assume that there exists some $y \in C_1^*$ such that $(x, y) \notin \widehat{E}$. This implies that y is preferred over x in at least a $2k/(3m)$ fraction of the voters. If for every candidate $a \in C'$ it holds that $(w, a) \in \widehat{E}$ or $(z, a) \in \widehat{E}$ for some $z \in C_1^*$, we can conclude that w can reach every node in \widehat{G} in at most 2 steps, again reaching the desired conclusion. On the other hand, assume that there exists $b \in C'$ such that $(w, b) \notin \widehat{E}$ and $(z, b) \notin \widehat{E}$ for all $z \in C_1^*$. By the definition of the set C^* , we can infer that $(b, w) \in \widehat{E}$ and $(b, z) \in \widehat{E}$ for all $z \in C_1^*$. Moreover, we know that from all of the votes candidate y received, candidate b was below in at most a $k/(3m)$ fraction (over all the voters); otherwise, it would follow that $(y, b) \in \widehat{E}$. As a result, since y is preferred over x in at least a $2k/(3m)$ fraction of the voters, we can conclude that $(b, x) \in \widehat{E}$, in turn implying that b can reach every node in \widehat{G} in paths of length at most 2, concluding the proof. \square

From an algorithmic standpoint, although our proof of Theorem 5.4.5 is constructive, leading to an efficiently implementable voting rule, the established upper bound can also be subsequently attained by the instance-optimal MINIMAX-LP mechanism. Theorem 5.4.5 substantially improves over the previous best-known bound which was $12m/k$ [150, 149]. Nonetheless, there is still a gap between the aforementioned lower bound (Theorem 5.4.4). Before we conclude this section, we explain how one can further improve upon the bound obtained in Theorem 5.4.5.

Conjecture 5.4.6. *If we assume that every agent provides her k -top preferences for some $k \in [m]$, there is a candidate $a \in C$ and a subset $S \subseteq V$ such that*

- *There exists a perfect matching $M : S \rightarrow S$ in the integral domination graph of a (see Theorem 5.5.4 in the next section);*
- $|S| \geq n \times k/m$.




Figure 5.5: The anatomy of our proof for Theorem 5.4.5. The set of candidates is partitioned into a “good” set C^* and a “bad” set C' ; C^* has a king vertex w , and we can essentially apply the reasoning of Theorem 5.4.2. A key observation is that C' is always dominated by some node in C^* , namely x .

When $k = m$, this was shown to be true by [131]. On the other end of the spectrum, when $k = 1$ it is easy to verify that the plurality winner establishes this conjecture.

Proposition 5.4.7. *If Theorem 5.4.6 holds, then there exists a deterministic mechanism which elicits only the k -top preferences and yields distortion at most $4m/k - 1$.*

Proof. Let $a \in C$ be the candidate which satisfies Theorem 5.4.6. Then, it follows that a yields distortion at most 3 w.r.t. the voters in the set S [131]. As a result, Theorem 5.4.3 implies that the distortion of a is upper-bounded by

$$3 + 4 \frac{n - |S|}{|S|} \leq \frac{4m}{k} - 1.$$

□

5.5 Randomized Preference Elicitation and Sampling

Previously, we characterized the distortion when only a deterministically (and potentially adversarially) selected subset of voters has provided information to the mechanism. This raises the question of bounding the distortion when the mechanism elicits information from only a small *random sample* of voters. Here, a single

sample corresponds to the *entire* ranking of a voter. We stress that randomization is only allowed during the preference elicitation process; for any given random sample as input, the mechanism has to select a candidate *deterministically*. We commence this section with a simple lower bound, which essentially follows from a standard result due to Canetti et al. [59].

Proposition 5.5.1. *Any mechanism which yields distortion at most $3 + \epsilon$ with probability at least $1 - \delta$ requires $\Omega(\log(1/\delta)/\epsilon^2)$ samples, even for $m = 2$.*

Proof. Consider two candidates a, b , and assume that exactly $(1 - \epsilon)/2$ fraction of the voters prefer candidate a . It is easy to verify that a yields distortion strictly larger than $3 + \epsilon$; thus, any mechanism with distortion at most $3 + \epsilon$ has to return candidate b . However, the *winner determination* problem with *margin* ϵ requires $\Omega(\log(1/\delta)/\epsilon^2)$ samples to solve with probability $1 - \delta$ [59], concluding the proof. \square

5.5.1 Approximating Copeland

We begin by analyzing a *sampling approximation* of Copeland’s rule. Below we summarize the main result of this subsection.

Theorem 5.5.2. *For any sufficiently small $\epsilon > 0$ and $\delta > 0$ there exists a mechanism that takes a sample of size $c = \Theta(\log(m/\delta)/\epsilon^2)$ voters and yields at most $5 + \epsilon$ distortion with probability at least $1 - \delta$.*

We recall that when the entire input is available, Copeland yields distortion at most 5 [18]. It follows from Theorem 5.5.2 that $\tilde{\Theta}(m/\epsilon^2)$ bits of information in total from the voters suffice to yield $5 + \epsilon$ distortion with high probability, where the notation $\tilde{\Theta}(\cdot)$ suppresses poly-logarithmic factors. Before we proceed with the proof of Theorem 5.5.2, we state the following standard fact:

Lemma 5.5.3 (Chernoff-Hoeffding Bound). *Let $\{X_1, X_2, \dots, X_c\}$ be a set of i.i.d. random variables with $X_i \sim \text{Bern}(p)$ and $X_\mu = (X_1 + X_2 + \dots + X_c)/c$; then,*

$$\mathbb{P}(|X_\mu - p| \geq \epsilon) \leq 2e^{-2\epsilon^2 c}.$$

Proof of Theorem 5.5.2. Consider the complete, weighted and directed graph $G = (C, E, w)$ so that $w_{a,b} = |ab|/n$. We will show how to use the random sample in order to construct a graph $\hat{G} = (C, E, \hat{w})$ which approximately preserves the weights of G with high probability. In particular, consider some parameters $\epsilon \in (0, 1/2)$ and $\delta \in (0, 1)$, and take a sample S of size $|S| = c = \Theta(\log(m/\delta)/\epsilon^2)$ from the set of voters V ; for simplicity, we assume that the sampling occurs *with replacement* in order to guarantee independence, but the result holds even without replacement given that the dependencies are negligible; e.g., see the work of Kontorovich and Ramanan [154]. Now we let $\hat{w}_{a,b} = |\{i \in S : a \succ_i b\}|/c$. Theorem 5.5.3 implies

that $|\widehat{w}_{a,b} - w_{a,b}| < \epsilon$ with probability at least $1 - \delta/m^2$. Thus, the union bound implies that for all distinct pairs a, b we have approximately preserved the weights: $|\widehat{w}_{a,b} - w_{a,b}| < \epsilon$ with probability at least $1 - \delta$.

From \widehat{G} we construct the directed graph $T = (C, \widehat{E})$ so that $(a, b) \in \widehat{E} \iff \widehat{w}_{a,b} \geq 1/2$; if $\widehat{w}_{a,b} = \widehat{w}_{b,a}$ for some distinct candidates $a, b \in C$, we only retain one of the edges arbitrarily (this conundrum can be avoided by taking c to be odd). In this way, T will be a tournament, and as such, there exists a candidate w which can reach every node in T in at most 2 steps. Thus, for any $a \in C$ there exists some intermediate candidate $b \in C$ so that $|wb| \geq 1/2 - \epsilon$ and $|ba| \geq 1/2 - \epsilon$. As a result, Theorem 5.4.1 implies that the distortion of w is upper-bounded by $4/(1 - 2\epsilon) + 1 \leq 5 + 16\epsilon$, for any $\epsilon \leq 1/4$. Finally, rescaling ϵ by a constant factor concludes the proof. \square

5.5.2 Approximating Plurality Matching

In light of Theorem 5.5.1, the main question that arises is whether we can asymptotically reach the optimal distortion bound of 3. To this end, we will analyze a sampling approximation of PLURALITYMATCHING, a deterministic mechanism introduced by [131] that obtains the optimal distortion bound of 3. To keep the exposition reasonably self-contained, we recall some basic facts about PLURALITYMATCHING.

Definition 5.5.4 ([131], Definition 5). *For an election $\mathcal{E} = (V, C, \sigma)$ and a candidate $a \in C$, the integral domination graph of candidate a is the bipartite graph $G(a) = (V, V, E_a)$, where $(i, j) \in E_a$ if and only if $a \succeq_i \text{top}(j)$.*

Proposition 5.5.5 ([131], Corollary 1). *There exists a candidate $a \in C$ whose integral domination graph $G(a)$ admits a perfect matching.*

Before we proceed let us first introduce some notation. For this subsection it will be convenient to use numerical values in the set $\{1, 2, \dots, m\}$ to represent the candidates. We let $\Pi_j = \sum_{i \in V} \mathbb{1}\{\text{top}(i) = j\}$, i.e. the number of voters for which $j \in C$ is the top candidate. For candidate $j \in C$ we let $G(j)$ be the integral domination graph of j , and M_j be a maximum matching in $G(j)$. In the sequel, it will be useful to “decompose” M_j as follows. We consider the partition of V into $V_j^0, V_j^1, \dots, V_j^m$ such that $V_j^k = \{i \in V : M_j(i) = k\}$ for all $k \in [m]$, while V_j^0 represents the subset of voters which remained unmatched under M_j ; see Figure 5.6.

Moreover, consider a set $\mathcal{S} = S_j^0 \cup S_j^1 \cup \dots \cup S_j^m$ such that $S_j^k \subseteq V_j^k$ for all k ; we also let $c = |\mathcal{S}|$, and $\Pi'_j = c/n \times \Pi_j$. For now let us assume that $\Pi'_j \in \mathbb{N}$ for all j . We let $G^{\mathcal{S}}(j)$ represent the induced subgraph of $G(j)$ w.r.t. the subset $\mathcal{S} \subseteq V$ and the new plurality scores Π'_j . We start our analysis with the following observation:




Figure 5.6: An example of a matching decomposition in the integral domination graph of candidate 1.

Observation 5.5.6. *Assume that \mathcal{S} is such that $|S_j^k|/c = |V_j^k|/n$ for all k . Then, if $M_j^{\mathcal{S}}$ represents the maximum matching in $G^{\mathcal{S}}(j)$, it follows that $|M_j^{\mathcal{S}}|/c = |M_j|/n$.*

Sketch of Proof. First, it is clear that $|M_j^{\mathcal{S}}| \geq \sum_{k=1}^m |S_j^k| = c/n \sum_{k=1}^m |V_j^k| = c/n \times |M_j|$. Thus, it remains to show that $|M_j^{\mathcal{S}}| \leq c/n \times |M_j|$. Indeed, if we assume otherwise, we can infer via an exchange argument that M_j is not a maximum matching. \square

Let us denote with $\Phi_j = |M_j|/n$; roughly speaking, we know from the work of Gkatzelis et al. [131] that Φ_j is a good indicator for the “quality” of candidate j . Importantly, Theorem 5.5.6 tells us that we can determine Φ_j in a much smaller graph, if only we had a sampling-decomposition that satisfied the “proportionality” condition of the claim. Of course, determining explicitly such a decomposition makes little sense given that we do not know the sets V_j^k , but the main observation is that we can approximately satisfy the condition of Theorem 5.5.6 through sampling. It should be noted that we previously assumed that $\Pi_j' \in \mathbb{N}$, i.e., we ignored rounding errors. However, in the worst-case rounding errors can only induce an error of at most m/c in the value of Φ_j ; thus, we remark that our subsequent selection of c will be such that this error will be innocuous, in the sense that it will be subsumed by the “sampling error” (see Theorem 5.5.8). Before we proceed, recall that for $\mathbf{p}, \hat{\mathbf{p}} \in \Delta([k])$,

$$d_{\text{TV}}(\mathbf{p}, \widehat{\mathbf{p}}) \stackrel{\text{def}}{=} \sup_{S \subseteq [k]} |\mathbf{p}(S) - \widehat{\mathbf{p}}(S)| = \frac{1}{2} \|\mathbf{p} - \widehat{\mathbf{p}}\|_1,$$

where $\|\cdot\|_1$ represents the ℓ_1 norm. In this context, we will use the following standard fact:

Lemma 5.5.7 ([60]). *Consider a discrete distribution $\mathbf{p} \in \Delta([k])$ and let $\widehat{\mathbf{p}}$ be the empirical distribution derived from N independent samples. For any $\epsilon > 0$ and $\delta \in (0, 1)$, if $N = \Theta((k + \log(1/\delta))/\epsilon^2)$ it follows that $d_{\text{TV}}(\mathbf{p}, \widehat{\mathbf{p}}) \leq \epsilon$ with probability at least $1 - \delta$.*

As a result, if we draw a set \mathcal{S} with $|\mathcal{S}| = c = \Theta((m + \log(1/\delta))/\epsilon^2)$ samples (without replacement⁴) we can guarantee that

$$\sum_{k=0}^m \left| \frac{|S_j^k|}{c} - \frac{|V_j^k|}{n} \right| \leq 2\epsilon;$$

$$\sum_{k=1}^m \left| \frac{\widehat{\Pi}_k}{c} - \frac{\Pi_k}{n} \right| \leq 2\epsilon,$$

where S_j^k represents the subset of \mathcal{S} which intersects V_j^k , and $\widehat{\Pi}_k$ is the empirical plurality score of candidate k . Thus, the following lemma follows directly from Theorem 5.5.6 and Theorem 5.5.7.

Lemma 5.5.8. *Let $\widehat{\Phi}_j = |\widehat{M}_j|/c$, where \widehat{M}_j is the maximum matching in the graph $G^{\mathcal{S}}(j)$. Then, if $|\mathcal{S}| = \Theta((m + \log(1/\delta))/\epsilon^2)$ for some $\epsilon, \delta \in (0, 1)$, it follows that $(1 - \epsilon)\Phi_j \leq \widehat{\Phi}_j \leq (1 + \epsilon)\Phi_j$ with probability at least $1 - \delta$.*

Theorem 5.5.9. *For any sufficiently small $\epsilon > 0$ and $\delta > 0$ there exists a mechanism that takes a sample of size $\Theta((m + \log(m/\delta))/\epsilon^2)$ voters and yields distortion at most $3 + \epsilon$ with probability at least $1 - \delta$.*

Proof. Fix some $\epsilon \in (0, 1/4)$ and $\delta \in (0, 1)$. If we draw $\Theta((m + \log(m/\delta))/\epsilon^2)$ samples, Theorem 5.5.8 along with the union bound imply that $(1 - \epsilon)\Phi_j \leq \widehat{\Phi}_j \leq (1 + \epsilon)\Phi_j$ for all $j \in [m]$, with probability at least $1 - \delta$, where $\widehat{\Phi}_j$ is defined as in Theorem 5.5.8. Conditioned on the success of this event, let $w = \arg \max_{j \in C} \widehat{\Phi}_j$. Theorem 5.5.5 implies that there exists some candidate x for which $\Phi_x = 1$; hence, we know that $\widehat{\Phi}_w \geq \widehat{\Phi}_x \geq 1 - \epsilon$, in turn implying that $\Phi_w \geq (1 - \epsilon)/(1 + \epsilon) \geq 1 - 2\epsilon$ (Theorem 5.5.8). As a result, it follows that there exists a subset of voters V' for which there exists a perfect matching in the integral domination graph $G(w)$, with $|V'| \geq n(1 - 4\epsilon)$. Thus, it follows that for the subset of voters in V' candidate w yields distortion at most 3 [131], and Theorem 5.4.3 leads to the desired conclusion. \square

⁴Although the samples are not independent since we are not replacing them, observe that the induced bias is negligible for n substantially larger than m .

5.A Plurality Matching vs Minimax-LP

In this section, we compare the PLURALITYMATCHING mechanism of Gkatzelis et al. [131] with the instance-optimal mechanism, namely MINIMAX-LP; here, we tacitly posit that $\mathcal{P} = \sigma$, i.e. all the agents provide their entire rankings to the mechanism.

5.A.1 Instance Optimality

The first question that arises is how far could the distortion of PLURALITYMATCHING be with respect to the instance-optimal candidate. To address this question, we commence with the following proposition:

Lemma 5.A.1 (Lemma 6, [131]). *For any election $\mathcal{E} = (V, C, \sigma)$, a candidate $a \in C$ can be selected by PLURALITYMATCHING only if $\text{plu}(a) \geq \text{veto}(a)$.*

With this lemma at hand, we consider an instance with $n = m$ voters $V = \{1, 2, \dots, m\}$ and a set of m candidates $C = \{a, \dots\}$. We assume that for every voter $i \in [n - 1]$, $\text{sec}(i) = a$, while the (single) top-preferences of all the voters $i \in [n - 1]$ are assumed to be pairwise-distinct. Finally, the last voter places candidate a in her last place, while her preferences are otherwise arbitrary. An example with four candidates $\{a, b, e, f\}$ corresponds to the following input:

- $b \succ_1 a \succ_1 e \succ_1 f$;
- $f \succ_2 a \succ_2 b \succ_2 e$;
- $e \succ_3 a \succ_3 f \succ_3 b$;
- $b \succ_4 f \succ_4 e \succ_4 a$.

In general, observe that for any candidate $b \in C \setminus \{a\}$, it follows that $|ab| = n - 2$. Further, we will use the following standard lemma:

Lemma 5.A.2. *Consider two (distinct) candidates $a, b \in C$ such that $|ab| \geq \alpha n > 0$. Then,*

$$\frac{\text{cost}(a)}{\text{cost}(b)} \leq \frac{2}{\alpha} - 1.$$

This implies that the distortion of candidate a is $1 + \mathcal{O}(1/m)$. However, given that $\text{plu}(a) = 0 < 1 = \text{veto}(a)$, we know from Theorem 5.A.1 that a cannot be selected by PLURALITYMATCHING. We will show that every other candidate yields distortion close to 3. In particular, consider the metric space illustrated in Figure 5.7. It is easy to verify that the induced metric space is consistent with the given preferences. But, it follows that $\text{cost}(a) = m$, while $\text{cost}(b) = 2 + 3(m - 2) = 3m - 4$ for any $b \neq a$, implying that $\text{cost}(b)/\text{cost}(a) = 3 - \mathcal{O}(1/m)$. As a result, we have arrived at the following conclusion:

Proposition 5.A.3. *For any sufficiently small $\epsilon > 0$ and $m = \mathcal{O}(1/\epsilon)$ there is a profile of preferences σ such that MINIMAX-LP yields distortion at most $1 + \epsilon$, while PLURALITYMATCHING yields distortion at least $3 - \epsilon$.*




Figure 5.7: A metric space embedded on an *unweighted* and *undirected* graph; this example corresponds to $m = n = 4$, but the pattern should be already clear.

5.A.2 Decisive Metrics

Moreover, it is natural to compare these mechanisms in more refined metrics. Specifically, we espouse the α -decisiveness assumption of Anshelevich and Postl [14], according to which $d(v_i, c_p) \leq \alpha \cdot d(v_i, c_q)$, where $p = \text{top}(i)$ and $q = \text{sec}(i)$, and $\alpha \in [0, 1]$ some parameter; notice that the general case corresponds to $\alpha = 1$, while for $\alpha = 0$ every voter also serves as a candidate. The first observation is that this particular refinement can be addressed by simply incorporating some additional constraints in the linear program. More precisely, for a pair of distinct candidates a, b this leads to the following linear program $\text{METRIC}^\alpha\text{-LP}(a, b)$:

$$\begin{aligned}
 & \text{maximize} && \sum_{i=1}^n x_{i,a} \\
 & \text{subject to} && \sum_{i=1}^n x_{i,b} = 1; \\
 & && x_{i,\text{top}(i)} \leq \alpha \cdot x_{i,\text{sec}(i)}, \forall i \in V; \\
 & && x_{i,p} \leq x_{i,q}, \quad \forall (p, q) \in \mathcal{P}_i, \forall i \in V; \\
 & && x_{i,i} = 0, \quad \forall i \in V \cup C; \\
 & && x_{i,j} = x_{j,i}, \quad \forall i, j \in V \cup C; \\
 & && x_{i,j} \leq x_{i,k} + x_{k,j}, \quad \forall i, j, k \in V \cup C.
 \end{aligned} \tag{5.6}$$

Here we have assumed that every agent i provides her most preferred candidate $\text{top}(i)$, as well as her second most preferred candidate $\text{sec}(i)$. Having solved the $\text{METRIC}^\alpha\text{-LP}(a, b)$ (5.6) for every distinct pair of candidates a, b , we simply

select the candidate who minimizes the maximum cost obtained over all other candidates; this mechanism shall be referred to as the $\text{MINIMAX}^\alpha\text{-LP}$. Similarly to Theorem 5.2.2, we can establish the following:

Proposition 5.A.4. *For any given preferences \mathcal{P} and any $\alpha \in [0, 1]$ the $\text{MINIMAX}^\alpha\text{-LP}$ rule is instance-optimal in terms of distortion under α -decisive metrics.*

We should point out that for α -decisive metrics PLURALITYMATCHING always yields a candidate with distortion $2 + \alpha$. Moreover, [131] showed a lower bound of $2 + \alpha - 2(1 - \alpha)/m'$ for deterministic mechanisms, where $m' = 2\lfloor m/2 \rfloor$; thus, they showed that their mechanism obtains the optimal distortion only when $m \rightarrow \infty$ or when $\alpha = 1$, leaving a substantial gap.

We will show that $\text{MINIMAX}^\alpha\text{-LP}$ can substantially outperform PLURALITYMATCHING even for α -decisive metrics with α close to 0. Specifically, consider an election with 3 candidates and 2 voters⁵ with the following preferences: $\sigma_1 = a \succ b \succ e$, and $\sigma_2 = e \succ b \succ a$. For this election, b could be returned by PLURALITYMATCHING [131]. However, we claim that b yields distortion $2 + \alpha$, while a and e have distortion $1 + 2\alpha$. Indeed, we will show that candidate a has always distortion upper-bounded by $1 + 2\alpha$ (by symmetry, the same holds for e), while for candidate b there exists a metric space for which b yields distortion $2 + \alpha$. Specifically, we have that $d(c_a, c_b) \leq d(v_1, c_a) + d(v_1, c_b) \leq (1 + \alpha)d(v_1, c_b)$; thus, we obtain that

$$d(v_1, c_a) \leq \alpha d(v_1, c_b), \quad (5.7)$$

$$d(v_2, c_a) \leq d(v_2, c_b) + d(c_a, c_b) \leq (1 + \alpha)d(v_1, c_b) + d(v_2, c_b). \quad (5.8)$$

Summing these inequalities yields that $\text{cost}(a) \leq (1 + 2\alpha)d(v_1, c_b) + d(v_2, c_b) \leq (1 + 2\alpha)\text{cost}(b)$. Similarly, we can prove that $\text{cost}(a) \leq (1 + 2\alpha)\text{cost}(e)$. On the other hand, for candidate b , Gkatzelis et al. [131] considered the metric space of Figure 5.8.

Naturally, the distance between a pair of nodes is the corresponding shortest path in the graph. Thus, for this instance it follows that $\text{cost}(e) = 1$, while $\text{cost}(b) = 2 + \alpha$, implying that the distortion of b is $2 + \alpha$. Thus, for $\alpha \rightarrow 0$ PLURALITYMATCHING loses a factor of 2 with respect to the optimal candidate, which would be identified by the $\text{MINIMAX}^\alpha\text{-LP}$ rule by virtue of Theorem 5.A.4.

Proposition 5.A.5. *There exists a preference profile σ such that $\text{MINIMAX}^\alpha\text{-LP}$ yields distortion $1 + 2\alpha$, while PLURALITYMATCHING yields distortion at least $2 + \alpha$ under α -decisive metrics.*

Nonetheless we should point out that PLURALITYMATCHING does *not* require knowing the value of parameter α , unlike the instance-optimal mechanism.

⁵This example is taken from [131].




Figure 5.8: A metric space embedded on a graph.

Chapter 6

Dimensionality and Coordination in Voting: The Distortion of STV

The framework of distortion under *metric preferences* was first introduced a few years ago by Anshelevich et al. [18] (see also [20]). Specifically, they observed a lower bound of 3 for any deterministic mechanism, while they also showed—among others—that Copeland’s method, a very popular voting system, always incurs distortion at most 5, with the bound being tight for certain instances. This threshold was subsequently improved by Munagala and Wang [183], introducing a novel (deterministic) mechanism with distortion $2 + \sqrt{5}$, while the same bound was independently obtained by Kempe [149] through an approach based on LP duality. The lower bound of 3 was only recently matched by PLURALITYMATCHING, a mechanism introduced by Gkatzelis, Halpern, and Shah [131]. In Section 6.5 we investigate the performance of this mechanism under certain refinements and extensions, leveraging an important property established in [131] regarding the existence of a *perfect fractional matching* on a certain bipartite graph.

All of the aforementioned results apply under arbitrary metric spaces. Several special cases have also attracted attention in the literature. For one-dimensional spaces, Feldman et al. [104] establish several improved bounds, while a comprehensive characterization in a *distributed* setting was recently given by Filos-Ratsikas and Voudouris in [109]. Another notable refinement germane to our considerations in Section 6.5 was studied by Anshelevich and Postl [14] in the form of α -*decisiveness*, imposing that voters support their top choices by a non-negligible margin. This condition has led to several refined upper and lower bounds; cf. see [131]. The interested reader is referred to the concise survey of Anshelevich et al. [22] for detailed accounts on the rapidly growing literature on the subject. Moreover, for related research beyond the framework of distortion we refer to [128], and references therein.

The model we introduce in Section 6.4 is related to the seminal work of Branzei, Caragiannis, Morgenstern, and Procaccia [55] (see also the extensive follow-up work, such as [187]), viewing voting from the standpoint of *price of anarchy* (PoA). In particular, the authors study the discrepancy between the plurality scores under truthfulness, and under worst-case limit points of *best-response* dynamics. Instead, we argue that the utilitarian performance of a voting rule—in terms of distortion—offers a very compelling alternative to study this discrepancy, similarly to the original formulation of PoA in the context of routing games [155], while going beyond best-response dynamics is very much in line with the modern approach in the context of learning in games [69]. Finally, we stress that Question 2 has already received extensive attention in the literature (cf. see [56, 225] and references therein), but it was not addressed within the framework of (metric) distortion.

6.1 Contributions

Our first contribution is to relate the distortion of STV to the dimensionality of the underlying metric space. Specifically, our first insight is to employ the following fundamental concept from metric geometry:¹

Definition 6.1.1 (Doubling Dimension). *The doubling constant of a metric space $(\mathcal{M}, \text{dist})$ is the least integer $\lambda \geq 1$ such that for all $x \in \mathcal{M}$ and for all $r > 0$, every ball $\mathcal{B}(x, 2r)$ can be covered by the union of at most λ balls of the form $\mathcal{B}(s, r)$, where $s \in \mathcal{M}$; that is, there exists a subset $\mathcal{S} \subseteq \mathcal{M}$ with $|\mathcal{S}| \leq \lambda$ such that*

$$\mathcal{B}(x, 2r) \subseteq \bigcup_{s \in \mathcal{S}} \mathcal{B}(s, r). \quad (6.1)$$

*The doubling dimension is then defined as $\dim(\mathcal{M}) := \log_2 \lambda$.*²

This concept generalizes the standard notion of dimension since $\dim(\mathbb{R}^d) = \Theta(d)$ when \mathbb{R}^d is endowed with the ℓ_p norm. Moreover, it is clear that for a finite metric space $(\mathcal{M}, \text{dist})$, $\dim(\mathcal{M}) \leq \log_2 |\mathcal{M}|$; for example, this is essentially tight for the high-dimensional metric of Figure 2.1. The concept of doubling dimension was introduced by Larman [159] and Assouad [26], and was first used in algorithm design by Clarkson [75] in the context of the nearest neighbors problem. Nevertheless, we are not aware of any prior characterization that leverages the doubling dimension in the realm of voting theory. In the sequel, it will be assumed that $(\mathcal{M}, \text{dist}(\cdot, \cdot))$ stands for the metric space induced by the set of candidates and voters. In this context, our first main contribution is the following theorem:

¹To keep the exposition reasonably smooth, the formal definition of standard notation is deferred to the preliminaries in Section 6.2.

²To avoid trivialities it will be assumed that $\lambda \geq 2$.

Theorem 6.1.2. *If d is the doubling dimension of \mathcal{M} , then the distortion of STV is $O(d \log \log m)$.*

For *doubling metrics*³ this theorem already implies an exponential improvement in the distortion over the $\Omega(\sqrt{\log m})$ lower bound for general metrics. Moreover, it addresses as a special case Question 1:

Corollary 6.1.3. *The distortion of STV under low-dimensional Euclidean spaces is $O(\log \log m)$.*

To the best of our knowledge, this is the first result that relates the performance of any voting rule to the “intrinsic dimensionality” of the underlying metric space. It also corroborates the experimental findings of Elkind et al. [88] regarding the superiority of STV on the 2-dimensional Euclidean plane. More broadly, we suspect that our characterization applies for a wide range of iterative voting rules, to which STV serves as a canonical example. We should note that the $O(\log \log m)$ factor appears to be an artifact of our analysis. Indeed, we put forward the following conjecture:

Conjecture 6.1.4. *If d is the doubling dimension of \mathcal{M} , then the distortion of STV is $O(d)$.*


Verifying this conjecture in light of our result might be of small practical importance, but nonetheless we believe that it can be established by extending our techniques. In fact, for one-dimensional spaces we actually confirm this conjecture, proving that the distortion of STV on the line is $O(1)$ in Theorem 6.3.1. It should be noted, however, that the underlying phenomenon is inherently different once we turn our attention to higher-dimensional spaces. In addition, to complement our positive results we refine the lower bound of Skowron and Elkind [214], showing an $\Omega(\sqrt{d})$ lower bound, where d represents the doubling dimension of the submetric induced by the set of candidates \mathcal{M}_C . Thus, it should be noted that there are still small gaps left to be bridged in future research.

Other Notions of Dimension. An important advantage of the doubling dimension is that it essentially subsumes other commonly-used notions of dimension. Most notably, Karger and Ruhl [146] have introduced a concept of dimension based on the *growth rate* of a (finite) metric space, and it is known ([139, Proposition 1.2]) that the doubling dimension can only be a factor of 4 larger than the growth rate of Karger and Ruhl. Moreover, a similar statement applies for the *local density* of an unweighted graph, another natural notion of volume that has been employed in the analysis of a graph’s bandwidth [100].

High-Level Intuition. In this paragraph we briefly attempt to explain why the distortion of STV depends on the “covering dimension” of the underlying metric

³A doubling metric refers to a metric space with doubling dimension upper-bounded by some universal constant.

space. First, we have to describe the technique developed by Skowron and Elkind [214]. Specifically, their method for deriving an upper bound for the distortion of an iterative voting rule consists of letting a substantial fraction of agents reside within close proximity to the optimal candidate, and then analyze how the support of these agents propagates throughout the evolution of the iterative process. More precisely, the overall distance covered immediately implies an upper bound on the distortion (see Theorem 6.2.2). The important observation is that the underlying dimension drastically affects this phenomenon. In particular, when a large fraction of agents lies in a low-dimensional ball supporting many different candidates, we can infer that their (currently) second most-preferred alternatives ought to be “close”—for most of the agents—by a covering argument (and the triangle inequality). This directly circumscribes the propagation of the support, as hinted in Figure 6.1b, juxtaposed to the phenomenon in high dimensions in Figure 6.1a. We stress that we shall make use of this basic skeleton developed by Skowron and Elkind [214]. We should also remark that we prove the $O(\log m)$ bound under general metrics through a simpler analysis (see Theorem 6.3.3), which incidentally reveals a very clean recursive structure; this argument will be directly invoked for the proof of our main theorem.



(a) The propagation of the support in high dimensions.

(b) The propagation of the support in low dimensions.

Figure 6.1: The impact of the underlying dimension on STV.

The next theme of our work is motivated by the performance of STV, and in particular offers a preliminary answer to Question 2. Specifically, to formally address such questions we first propose a natural iterative model: In each day every agent has to select a *single* candidate, and at the end of the round agents are informed about the (plurality) scores of all the candidates (cf., see [49]). This process is repeated for sufficiently many days, and it is assumed that the candidate who

enjoyed the largest amount of support in the ultimate day will eventually prevail. Observe that in this scenario truthful engagement appears to be very unrealistic since agents would endeavor to adapt their support based on the popularity of each candidate; for example, it would make little sense to squander one’s vote (at least towards the last stages) to an unpopular candidate. More broadly, there is an interesting nexus between distortion and stability, as we elaborate in Section 6.4, emphasizing on a connection with the notion of *core* in cooperative game theory (Theorem 6.4.2).

In this context, STV already suggests a particularly natural strategic engagement, improving exponentially over the outcome of the truthful dynamics. Yet, it yields super-constant distortion due to the greedy aspect of the induced dynamics. We address this issue by designing a simple and decentralized exploration/exploitation scheme:

Theorem 6.1.5. *There exist simple, deterministic and distributed dynamics that converge to a candidate with $O(1)$ distortion.*

We elaborate on the proposed dynamics, as well as on all the aforementioned issues in Section 6.4.

The final contribution of our work concerns refinements and extensions of prior results under metric preferences, providing new insights along two main lines. First, we study preference aggregation under ordinal information when agents and candidates are located in *ultra-metric spaces*, which is a strengthening of the standard metric assumption. This setting is mostly motivated by the fundamental *bottleneck* variant in facility location games, wherein the cost of a path between an agent and a server corresponds to the largest weight among the edges in the path, instead of their sum [125]; it should be noted that ultra-metrics also commonly arise in branches of mathematics such as metric geometry [1], and p -adic analysis [197]. In this context, our main observation is that the PLURALITYMATCHING mechanism of Gkatzelis, Halpern, and Shah [131] always obtains distortion 2, which incidentally is the provable lower bound for any deterministic mechanism. It is particularly interesting that the optimal mechanism under metric spaces retains its optimality under an important refinement, illustrating the robustness of PLURALITYMATCHING.

We also study the performance of PLURALITYMATCHING under distance functions that satisfy a ρ -relaxed triangle inequality. This consideration is directly driven by the fact that many well-studied and commonly-arising distances are only approximately metrics (most notably, the *squared Euclidean* distance is a 2-approximate metric), but we believe that there is another concrete reason. Most research in the realm of distortion has thus far been divided between the metric case and the *unit-sum* case, with these two lines of research being largely disconnected. Studying approximate metrics serves as an attempt to bridge this gap. In this context, we prove a lower bound of $\rho^2 + \rho + 1$, while PLURALITYMATCHING incurs distortion at most $2\rho^2 + \rho$, thus leaving a small gap for future research. No-

tion that for the special case $\rho := 1$ this recovers the result of Gkatzelis, Halpern, and Shah [131].

6.2 Preliminaries

A *metric space* is a pair $(\mathcal{M}, \text{dist}(\cdot, \cdot))$, where $\text{dist} : \mathcal{M} \times \mathcal{M} \mapsto \mathbb{R}$ is a *metric* on \mathcal{M} , i.e., (i) $\forall x, y \in \mathcal{M}, \text{dist}(x, y) = 0 \iff x = y$ (identity of indiscernibles), (ii) $\forall x, y \in \mathcal{M}, \text{dist}(x, y) = \text{dist}(y, x)$ (symmetry), and (iii) $\forall x, y, z \in \mathcal{M}, \text{dist}(x, y) \leq \text{dist}(x, z) + \text{dist}(z, y)$ (triangle inequality). Now consider a set of n voters $V = \{1, 2, \dots, n\}$, and a set of m candidates C ; we will reference candidates with lowercase letters such as a, b, w, x . Voters and candidates are associated with points in a finite metric space $(\mathcal{M}, \text{dist})$, while it is assumed that \mathcal{M} is the (finite) set induced by the set of voters and candidates. The goal is to select a candidate x who minimizes the *social cost*: $\text{SC}(x) = \sum_{i=1}^n \text{dist}(i, x)$. This task would be trivial if we had access to the agents' distances from all the candidates. However, in the *metric distortion* framework every agent i provides only a *ranking* (a total order) σ_i over the points in C according to the *order* of i 's distances from the candidates, with ties broken arbitrarily. We also define $\sigma := (\sigma_1, \dots, \sigma_n)$, while we will sometimes use $\text{top}(i)$ to represent i 's most preferred alternative.

A *deterministic social choice rule* is a function that maps an *election* in the form of a 3-tuple $\mathcal{E} = (V, C, \sigma)$ to a single candidate $a \in C$. We will measure the performance of f for a given input of preferences σ in terms of its *distortion*; namely, the worst-case approximation ratio it provides with respect to the social cost:

$$\text{distortion}(f; \sigma) = \sup \frac{\text{SC}(f(\sigma))}{\min_{a \in C} \text{SC}(a)}, \quad (6.2)$$

where the supremum is taken over all metrics consistent with the voting profile. The distortion of a social choice rule f is the maximum of $\text{distortion}(f; \sigma)$ over all possible input preferences σ . To put it differently, once the mechanism selects a candidate (or a distribution over candidates if the social choice rule is *randomized*) an adversary can select any metric space subject to being consistent with the input preferences. These definitions naturally apply for refinements and extensions studied in the present work.

We define the *open ball* on the metric space $(\mathcal{M}, \text{dist})$ with center $x \in \mathcal{M}$ and radius $r > 0$ as $\mathcal{B}(x, r) := \{z \in \mathcal{M} : \text{dist}(z, x) < r\}$. An alternative definition for the doubling dimension considers the diameter of subsets, instead of the radius of balls; that is, the doubling constant is the smallest value of λ such that every subset of \mathcal{M} can be covered by at most λ subsets of (at most) half the diameter. According to this definition, for any submetric $\mathcal{X} \subseteq \mathcal{M}$ it follows that $\text{dim}(\mathcal{X}) \leq \text{dim}(\mathcal{M})$. Nonetheless, it will be convenient to work with the initial notion (Theorem 6.1.1) since switching between the two definitions can only affect the dimension by at

most a factor of 2 (see [139]). The following standard covering lemma will be useful for the analysis of STV in doubling metrics.

Lemma 6.2.1. *Consider a metric space $(\mathcal{M}, \text{dist})$ with doubling constant $\lambda \geq 1$. Then, for any $x \in \mathcal{M}$ and $r > 0$, the ball $\mathcal{B}(x, r)$ can be covered by at most $\lambda^{\lceil \log(r/\epsilon) \rceil}$ balls of radius at most ϵ .*

Proof. If we apply Theorem 6.1.1 successively we can conclude that any ball of radius r can be covered by at most λ^i balls of radius $r/2^i$. Thus, taking $i := \lceil \log(r/\epsilon) \rceil$ leads to the desired conclusion. \square

It should be noted that (when unspecified) the $\log(\cdot)$ will always be implied to the base 2. We conclude this section with a useful lemma observed by Skowron and Elkind [214], which will be used for analyzing iterative voting rules.

Lemma 6.2.2 ([214]). *Consider two distinct candidates $a, b \in C$. If $r := \text{dist}(a, b)/h$ for some parameter $h > 0$, and at most γn agents reside in $\mathcal{B}(a, r)$ for some $\gamma \in [0, 1)$, then*

$$\frac{\text{SC}(b)}{\text{SC}(a)} \leq 1 + \frac{h}{1 - \gamma}. \quad (6.3)$$

Proof. The triangle inequality implies that

$$\begin{aligned} \frac{\text{SC}(b)}{\text{SC}(a)} &= \frac{\sum_{i \in V} \text{dist}(i, b)}{\sum_{i \in V} \text{dist}(i, a)} \leq \frac{\sum_{i \in V} (\text{dist}(i, a) + \text{dist}(a, b))}{\sum_{i \in V} \text{dist}(i, a)} \\ &= 1 + n \frac{\text{dist}(a, b)}{\sum_{i \in V} \text{dist}(i, a)} \\ &\leq 1 + \frac{\text{dist}(a, b)}{(1 - \gamma)r} \\ &= 1 + \frac{h}{1 - \gamma}. \end{aligned}$$

\square

6.3 STV in Doubling Metrics

6.3.1 STV on the Line

As a warm-up, we will analyze the performance of STV on the line. In particular, the purpose of this subsection is to establish the following result:

Theorem 6.3.1. *The distortion of STV on the line is at most 15.*

Before we proceed with the proof of this theorem a few remarks are in order. First of all, we did not pursue optimizing the constant in the theorem, although this might be an interesting avenue for future research. It should also be noted that Theorem 6.3.1 already implies a stark separation between STV and PLURALITY, as the latter is known to admit a one-dimensional $\Omega(m)$ lower bound (recall Table 2.1).

Proof of Theorem 6.3.1. Let $w \in C$ be the winner of STV under some (fixed) sequence of eliminations, and $x \in C$ be the candidate who minimizes the social cost. In the sequel it will be assumed that $\text{dist}(x, w) > 0$; in the contrary case the theorem follows trivially. Moreover, we let $r := d(x, w)/7$, and we consider a sequence of balls $\{\mathcal{B}_i\}_{i=1}^4$ so that every ball \mathcal{B}_i has center at x and radius $(2i-1) \times r$, for $i = 1, 2, 3, 4$. We will show that at most half of the voters could reside in \mathcal{B}_1 .

For the sake of contradiction, let us assume that at least a $\gamma > 1/2$ fraction of the voters are in \mathcal{B}_1 ; that is, $\sum_{i \in V} \mathbb{1}\{i \in \mathcal{B}_1\} = \gamma n > n/2$. First, we will argue that at the time the last candidate in \mathcal{B}_i gets eliminated there is always a candidate located in $\mathcal{B}_{i+1} \setminus \mathcal{B}_i$. Indeed, in the contrary case we can deduce that the last candidate to be eliminated from \mathcal{B}_i would receive the support of all the voters in \mathcal{B}_1 , which is a contradiction since by construction all the candidates in \mathcal{B}_4 have to be eliminated (this follows given that $w \notin \mathcal{B}_4$). Now consider the stage of STV just before the last candidate from \mathcal{B}_1 was eliminated; observe that this is well-defined as \mathcal{B}_1 initially contains at least one candidate, namely $x \in C$. Let us denote with $a^\ell, a^r \in C$ the leftmost and the rightmost (respectively) nearest *active* candidates from \mathcal{B}_1 —which are not in \mathcal{B}_1 . We shall distinguish between two cases:

Case I. $a^\ell \in \mathcal{B}_2 \setminus \mathcal{B}_1$ and $a^r \in \mathcal{B}_2 \setminus \mathcal{B}_1$. Following the elimination of the last candidate from \mathcal{B}_1 every voter in \mathcal{B}_1 will support either a^ℓ or a^r . Thus, (by the pigeonhole principle) we can conclude that one of these two candidates accumulates at least $n/4$ supporters in the round following the elimination of the last candidate from \mathcal{B}_1 . Let us assume without any loss of generality that this candidate is a^ℓ . It is important to point out that the support of a candidate can only grow throughout the execution of STV, until elimination. Now consider the stage just before a^ℓ gets eliminated. We previously argued that there exists a candidate $y \in \mathcal{B}_4 \setminus \mathcal{B}_3$ that has to remain active after the elimination of a^ℓ . This implies that y received at least $n/4$ votes at this particular stage—just before the elimination of a^ℓ . However, observe that none of the supporters of y can derive from \mathcal{B}_1 since for every voter in \mathcal{B}_1 candidate a^ℓ is (strictly) superior to y . Moreover, the eventual winner w should also have at least $n/4$ supporters at this stage in order to qualify for the next round, but the supporters of w are certainly not from \mathcal{B}_1 , and are also different from the supporters of y . This follows since both are active at this stage and $y \neq w$, implied by the fact that $d(x, y) < 7r \leq d(x, w)$. As a result, we have concluded that the total number of voters is strictly more than n , which is an obvious contradiction.

Case II. Only one of the candidates a^ℓ and a^r resides in $\mathcal{B}_2 \setminus \mathcal{B}_1$. Notice that our previous argument implies that at least one of the two should be in $\mathcal{B}_2 \setminus \mathcal{B}_1$, and hence, there is indeed no other case to consider. Let us assume without any

loss of generality that the candidate in $\mathcal{B}_2 \setminus \mathcal{B}_1$ is a^ℓ . Observe that just before the last candidate from \mathcal{B}_1 gets eliminated every voter in \mathcal{B}_1 either supports the candidate in \mathcal{B}_1 or candidate a^ℓ . Moreover, we have argued that there are at least 5 remaining candidates. Thus, we can infer that the last candidate from \mathcal{B}_1 received at most 20% of the votes at the round of elimination, in turn implying that a^ℓ enjoyed the support of at least 30% of the voters. Consequently, we can easily reach a contradiction similarly to the previously considered case.




Figure 6.2: An illustration of our argument for Theorem 6.3.1.

As a result, we have established that $\gamma \leq 1/2$, where recall that γ represents the fraction of agents in \mathcal{B}_1 , and the theorem follows by Theorem 6.2.2. \square

6.3.2 Main Result

Moving on to the main result of this section, we will prove the following theorem:

Theorem 6.3.2. *If d is the doubling dimension of \mathcal{M} , then the distortion of STV is $O(d \log \log m)$.*

Before we proceed with the proof of this theorem, let us first present an analysis for general metric spaces. Our argument will uncover the same bound $O(\log m)$ in terms of distortion, as in [214], but it is considerably simpler, and it will also be used in the proof of Theorem 6.3.2.

Theorem 6.3.3. *The distortion of STV is $O(\log m)$.*

Proof. Let $w \in \mathcal{C}$ be the winner of STV under some sequence of eliminations, and $x \in \mathcal{C}$ be the candidate who minimizes the social cost. Moreover, let $r := \text{dist}(x, w)/(4\mathcal{H}_m + 2)$, where \mathcal{H}_m denotes the m -th harmonic number. If V_1 represents the subset of voters in $\mathcal{B}(x, r)$ and $\gamma := |V_1|/n$, we will show that $\gamma \leq 1/2$. Then, our claim will follow from Theorem 6.2.2.

For the sake of contradiction, let us assume that $\gamma > 1/2$. We will establish that no voter in V_1 will support candidate w at any stage of STV, which is an obvious contradiction since $\gamma > 1/2$ and w was assumed to be the winner. In particular, let $\bar{D}^{(t)}$ be defined as follows:

$$\bar{D}^{(t)} := \frac{1}{\gamma n} \sum_{i \in V_1} \text{dist}(x, \text{top}(i; t)), \quad (6.4)$$

where $\text{top}(i; t)$ represents the most preferred (active) candidate for voter i after round $t = 1, \dots, m - 1$. The quantity $\overline{D}^{(0)}$ is also defined as in Equation (6.4), assuming that $\text{top}(i; 0) := \text{top}(i)$. Thus, observe that the triangle inequality yields that

$$\overline{D}^{(0)} \leq \frac{1}{\gamma n} \sum_{i \in V_1} (\text{dist}(x, i) + \text{dist}(i, \text{top}(i))) < 2r. \quad (6.5)$$

Moreover, we claim that if a voter i supports a candidate a at round t , then $\text{dist}(x, a) \leq \overline{D}^{(t-1)} + 2r$. Indeed, we will show the following: If two voters i, j in $\mathcal{B}(x, r)$ support two candidates a, b respectively, then it follows that $\text{dist}(x, a) \leq \text{dist}(x, b) + 2r$. In particular, successive applications of the triangle inequality yield that $\text{dist}(i, a) \leq \text{dist}(i, b) \leq \text{dist}(x, i) + \text{dist}(x, b) \leq r + \text{dist}(x, b)$, while $\text{dist}(i, a) \geq \text{dist}(x, a) - \text{dist}(x, i) \geq \text{dist}(x, a) - r$, in turn implying that

$$\text{dist}(x, a) \leq \text{dist}(x, b) + 2r. \quad (6.6)$$

By symmetry, it also follows that $\text{dist}(x, b) \leq \text{dist}(x, a) + 2r$. Next, we will (inductively) establish that the quantity $\overline{D}^{(m-2)}$ is strictly less than $4\mathcal{H}_m \times r$. Indeed, first note that under the invariance $\overline{D}^{(t)} < 4\mathcal{H}_m \times r$ the voters in $\mathcal{B}(x, r)$ support at least two distinct candidates; otherwise, the unique supported candidate a would prevail since $\gamma > 1/2$, which in turn is a contradiction given that $\text{dist}(x, w) > \overline{D}^{(t)} + 2r \implies a \neq w$. Next, observe that at round $t = 1, \dots, m - 1$ at most $n/(m - t + 1)$ agents alter their support. This follows since there are exactly $m - t + 1$ candidates, while STV eliminates the one who enjoys the least amount of support. Moreover, all the agents who recast their support will end up coalescing with a candidate whose distance from x increases by at most a $2r$ additive factor—compared to the previously supported alternative; this follows by the bound of Equation (6.6), and the fact that there is indeed another candidate supported by voters in $\mathcal{B}(x, r)$ in round t (as we previously argued). As a result, we have shown the following recursive structure:

$$\overline{D}^{(t)} \leq \overline{D}^{(t-1)} + \frac{1}{\gamma} \frac{2r}{m - t + 1} \leq \overline{D}^{(t-1)} + \frac{4r}{m - t + 1}, \quad (6.7)$$

for all $t = 1, \dots, m - 2$. This verifies the assertion that $\overline{D}^{(m-2)} < 4\mathcal{H}_m \times r$. Thus, in the ultimate round of STV more than half the voters support a candidate a for which $\text{dist}(a, x) < 2r + 4r \times \mathcal{H}_m$, which is a contradiction since $\text{dist}(x, w) = 2r \times (2\mathcal{H}_m + 1)$. \square

Next, we provide the proof of Theorem 6.3.2. In particular, the main technical challenge of the analysis lies in maintaining the appropriate invariance during STV. We address this with a simple trick, essentially identifying a subset of the domain with a sufficient degree of regularity. We should also note that the second part of the proof makes use of the technique devised by Skowron and Elkind [214].

Proof of Theorem 6.3.2. As before, let $w \in C$ be the winner of STV under some sequence of eliminations, and $x \in C$ be the candidate who minimizes the social cost. Moreover, let $r := \text{dist}(x, w)/(4h + 7)$, where h is defined as $h := 1 + \lceil \log_2(6\lambda^{\log \mathcal{H}_m + 1}) \rceil = \Theta(d \log \log m)$. If γ represents the fraction of the voters in $\mathcal{B}(x, r)$, we will establish that $\gamma \leq 2/3$.

For the sake of contradiction, let us assume that $\gamma > 2/3$. Our argument will characterize the propagation of the support of the voters in $\mathcal{B}(x, r)$. In particular, we proceed in the following two phases:

Phase I. Our high-level strategy is to essentially employ the argument in the proof of Theorem 6.3.3, but not for the entire set of voters in $\mathcal{B}(x, r)$. Instead, we will establish the existence of a set with a helpful invariance, which still contains most of the voters. More precisely, we first consider a covering $\{\mathcal{B}(z_j, r_j)\}_{j=1}^\mu$ of the ball $\mathcal{B}(x, r)$, where the radius of every ball is at most $\epsilon \times r$ for some parameter $\epsilon \in (0, 1)$. The balls that do not contain any voter may be discarded for the following argument. We let $\mathcal{S}^{(0)}$ be the union of these balls. We know from Theorem 6.2.1 that $\mu = \mu(\epsilon; \lambda) \leq \lambda^{\log(1/\epsilon)+1}$. For Phase I we assume that more than M candidates remain active in STV, where $M := 6\mu$, while $\epsilon := 1/\mathcal{H}_m$.

Let us consider a round $t = 1, \dots, m - M$ of STV. In particular, let $a \in C$ be the candidate who is eliminated at round t . Observe that if a is not supported by any voter residing in $\mathcal{B}(x, r)$, the support of these agents remains invariant under round t . Thus, let us focus on the contrary case. Specifically, if there exists a ball in the covering which contains exclusively supporters of candidate a , we shall remove every such ball from the current covering, updating analogously the set $\mathcal{S}^{(t)}$. Given that we are at round t , we can infer that the number of such supporters is at most $n/(m - t + 1) < n/M$. Thus, since we can only remove μ balls from the initial covering, it follows that the set $\mathcal{S} := \mathcal{S}^{(t)}$ with $t = m - M$ contains strictly more than $2n/3 - n\mu/M = n/2$ voters.

Next, we will argue about the propagation of the support for the voters in \mathcal{S} during the first $m - M$ rounds of STV. By construction of the set \mathcal{S} , we have guaranteed the following invariance: Whenever a candidate a supported by voters in \mathcal{S} gets eliminated, every supporter of a from \mathcal{S} lies within a ball of radius at most ϵ with agents championing a different candidate. Now, let us define $\bar{D}^{(t)}$ as follows:

$$\bar{D}^{(t)} := \frac{1}{\gamma' n} \sum_{i \in \mathcal{S}} \text{dist}(i, \text{top}(i; t)), \quad (6.8)$$

where γ' represents the fraction of the voters residing in \mathcal{S} . Note that $\bar{D}^{(t)}$ is defined slightly differently than in the proof of Theorem 6.3.3. Consider two voters i, j supporting two candidates a, b respectively. We will show that $\text{dist}(i, b) \leq \text{dist}(i, a) + 2 \text{dist}(i, j)$, and similarly, $\text{dist}(j, a) \leq \text{dist}(j, b) + 2 \text{dist}(i, j)$. Indeed, successive applications of the triangle inequality imply that

$$\begin{aligned}
\text{dist}(i, b) &\leq \text{dist}(i, j) + \text{dist}(j, b) \\
&\leq \text{dist}(i, j) + \text{dist}(j, a) \\
&\leq \text{dist}(i, j) + \text{dist}(j, i) + \text{dist}(i, a) \\
&= \text{dist}(i, a) + 2 \text{dist}(i, j).
\end{aligned}$$

Thus, if the voters i and j happen to reside within a ball of radius at most ϵ , we can infer that $\text{dist}(i, b) \leq \text{dist}(i, a) + 4\epsilon$. As a result, by the recursive argument of Theorem 6.3.3 we can conclude that

$$\overline{D}^{(t)} \leq \overline{D}^{(t-1)} + \frac{1}{\gamma'} \frac{4\epsilon r}{m-t+1} \leq \overline{D}^{(t-1)} + \frac{8\epsilon r}{m-t+1}, \quad (6.9)$$

in turn implying that

$$\overline{D}^{(m-M)} \leq 8(\epsilon r) \mathcal{H}_m. \quad (6.10)$$

Consequently, we have essentially shown that the propagation of the support is “decelerated” by a factor of ϵ . In particular, for $\epsilon = 1/\mathcal{H}_m$ this implies that during the first phase the agents in \mathcal{S} support candidates within $O(1) \times r$ distance from candidate x .

Phase II. At the beginning of the second phase there are M remaining candidates. Let us denote with $\mathcal{B}_j := \mathcal{B}(x, (2j-1) \times r)$. In this phase we will argue about the entire set of voters in $\mathcal{B}(x, r)$. Let $m_1 \leq M$ be the number of candidates supported by voters in $\mathcal{B}(x, r)$ at the beginning of the second phase. Our previous argument implies that every such candidate will reside in \mathcal{B}_7 ; this follows directly by applying the triangle inequality. Let us denote with m_j the number of candidates residing outside \mathcal{B}_{4+2j} for $j \geq 2$ at the round the last candidate from \mathcal{B}_{3+2j} gets eliminated.

By the pigeonhole principle, we can infer that there exists a candidate a in \mathcal{B}_7 who enjoys the support of at least $\gamma n/m_1$ voters. Moreover, observe that the triangle inequality implies that no voter will support a candidate outside \mathcal{B}_8 as long as candidate a remains active. Thus, at the round a gets eliminated we can deduce that $(1-\gamma)n/m_2 \geq \gamma n/m_1 \iff m_2 \leq m_1 \times (1-\gamma)/\gamma$, where we used that the number of candidates in every subset can only decrease during STV. Inductively, we can infer that

$$m_h \leq \left(\frac{1-\gamma}{\gamma}\right)^{h-1} m_1 < \left(\frac{1}{2}\right)^{h-1} M \leq 1, \quad (6.11)$$

for $h = \lceil \log_2 M \rceil + 1$, where we used that $\gamma > 2/3$. This implies that the winner of STV should lie within \mathcal{B}_{4+2h} , i.e. $\text{dist}(x, w)/r < 4h + 7$, which is a contradiction since $\text{dist}(x, w) = (4h + 7) \times r$. Thus, the theorem follows directly from Theorem 6.2.2. \square

6.3.3 The Lower Bound

In this subsection we refine the $\Omega(\sqrt{\log m})$ lower bound of Skowron and Elkind [214, Theorem 4] based on the doubling dimension of the submetric induced by the set of candidates \mathcal{M}_C . In particular, we will establish the following theorem:

Theorem 6.3.4 (Lower Bound for STV). *For any $\lambda \geq 2$ there exists a metric space induced by the set of candidates $(\mathcal{M}_C, \text{dist})$, with $d = \Theta(\log \lambda)$ being the doubling dimension of \mathcal{M}_C , and a voting profile such that the distortion of STV is $\Omega(\sqrt{d})$.*

For the proof, we consider first a tree \mathcal{T} with λ number of leaves; it will be assumed that λ is such that $\lambda = a_i$ for some $i \in \mathbb{N}$, where $\{a_i\}_{i \in \mathbb{N}}$ is a sequence such that $a_1 = 2$ and $a_{i+1} = 2(a_i + 1)$ for $i \geq 1$. Notice that if this is not the case we can always select the maximal λ' smaller than λ that satisfies this property; given that $\lambda' = \Theta(\lambda)$ this would not affect the conclusion (up to constant factors). Then, the next (or first) *layer* will consist of nodes which are parents of leaves, and in particular, every node in layer 1 will be parent to exactly 2 (mutually distinct) leaves, and we will say that the *branching factor* is $b_1 := 2$. This construction is continued iteratively until we reach the root, with the branching factor of layer $i > 1$ satisfying $b_{i+1} = 2(b_i + 1)$; the first two layers of this construction are illustrated in Figure 6.3. Observe that by construction the branching factor increases exponentially fast. Moreover, the number of nodes in the i -th layer is $m_i = m_{i-1}/b_i$, with $m_0 := \lambda$. Now let h be the height of the induced tree. We can infer that

$$m_h = \frac{m_{h-1}}{b_h} = \frac{\lambda}{\prod_{i=1}^h b_i} \geq \frac{\lambda}{\prod_{i=1}^h 4^i} = \frac{\lambda}{4^{\sum_{i=1}^h i}} = \frac{\lambda}{2^{h(h+1)}}, \quad (6.12)$$

where we used that $b_{i+1} \leq 4b_i$. Thus, since $m_h = 1$, it follows that $h = \Omega(\sqrt{\log_2 \lambda})$. Finally, we incorporate a node which is connected via edges to all the leaves. Then, the distance between two nodes is defined as the length of the shortest path in the induced *unweighted* graph. The metric space we introduced will be henceforth represented as $(\mathcal{M}_C, \text{dist})$.

Claim 6.3.5. *The doubling dimension of the metric space $(\mathcal{M}_C, \text{dist})$ is $\Theta(\log \lambda)$.*

Proof. It is easy to see that the doubling constant of the metric space $(\mathcal{M}_C, \text{dist})$ is at least λ . Thus, the claim follows since $|\mathcal{M}_C| \leq 2\lambda$ and $\dim(\mathcal{M}_C) \leq \log_2 |\mathcal{M}_C|$. \square

The Voting Instance. We assume that voters and candidates are mapped to points on the metric space $(\mathcal{M}_C, \text{dist})$. In particular, for every point $x \in \mathcal{T}$ we assign a (distinct) candidate, while every remaining candidate will be allocated to the point connected to all the leaves. In this way, $(\mathcal{M}_C, \text{dist})$ is indeed the metric space induced by the set of candidates. We will let $x \in C$ be a candidate located




Figure 6.3: A 2-layer instance of the tree \mathcal{T} employed for the lower bound in Theorem 6.3.4.

to the point connected to all the leaves, and $w \in C$ be the candidate at the root of \mathcal{T} . Moreover, for a layer i of \mathcal{T} we place ν_i number of voters at each point of the layer, such that $\nu_0 = 1$ and $\nu_{i+1} = (b_i + 1)\nu_i$ for all $i \geq 0$ (here we tacitly assume that $b_0 = 0$); no voters are collocated with candidate x .

Claim 6.3.6. *There exists an elimination sequence such that $w \in C$ is the winner of STV.*

Proof. Given that $\nu_{i+1} = (b_i + 1)\nu_i$ we can inductively infer that there exists an elimination sequence such that every candidate in layer i is eliminated before any candidate in the layer $i + 1$, while every candidate collocated with x will be eliminated before any candidate in \mathcal{T} . This concludes the proof. \square

Claim 6.3.7. $SC(w)/SC(x) = \Omega(h)$.

Proof. Let n_i be the total number of agents residing in the i -th layer of the tree \mathcal{T} . We will show that $n_{i+1} = n_i/2$, for all $i = 0, 1, \dots, h-1$. Indeed, for $i \geq 0$ it follows that $n_{i+1} = \nu_{i+1}m_{i+1} = (b_i + 1)\nu_i m_{i+1} = n_i(b_i + 1)/b_{i+1} = n_i/2$, where we used that $m_i = m_{i-1}/b_i$ and $b_{i+1} = 2(b_i + 1)$. As a result, it follows that $SC(w) \geq n_0 \times h$, while it is easy to show that $SC(x) \leq 4n_0$, concluding the proof. \square

Proof of Theorem 6.3.4. The lower bound follows by directly applying and combining Theorem 6.3.5, Theorem 6.3.6, and Theorem 6.3.7. \square

Notice that this theorem implies as a special case the $\Omega(\sqrt{\log m})$ lower bound for general metrics, which only applies when the metric $(\mathcal{M}_C, \text{dist})$ is *near-uniform*.

Remark 6.3.8. *It is not difficult to show that the distortion of STV is always $O(\Delta)$, where Δ represents the aspect ratio of \mathcal{M}_C —the ratio between the largest (pairwise) distance to the smallest distance in \mathcal{M}_C . In fact, this bound is tight—up to constant factors—for certain instances, as implied by the construction in Theorem 6.3.4.*

6.4 Coordination Dynamics

In this section we explore the degree to which natural and distributed learning dynamics can converge to social choices with near-optimal distortion. We should point out that there is a concrete connection between such considerations and the results of the previous section, which will be revealed in detail very shortly. First, let us commence with the following observation:

Observation 6.4.1. *Consider a voting instance under a metric space so that some candidate $a \in C$ has distortion at least \mathfrak{D} .⁴ Then, there exists a candidate $x \neq a$ and subset $W \subseteq V$ such that*

1. *Every agent in W strictly prefers x to a ;*
2. *$|W|/n \geq 1 - 2/(\mathfrak{D} + 1)$.*

This statement essentially tells us that candidates with large distortion are inherently unstable, in the sense that there will exist a large “coalition” of voters that strictly prefer a different outcome. Interestingly, this observation implies a connection between (metric) distortion and the notion of *core* in cooperative game theory. To be more precise, we will say that a set of coalitions \mathcal{W} is α -*large*, with $\alpha \in [0, 1]$, if it contains every coalition $W \subseteq V$ such that $|W|/n \geq \alpha$; a candidate a is said to be in the core if there does *not* exist a coalition $W \in \mathcal{W}$ such that every agent in W (strictly) prefers a different alternative.⁵ In this context, the following proposition follows directly from Theorem 6.4.1:

Proposition 6.4.2. *Consider a voting instance under a metric space so that some candidate $a \in C$ has distortion at least \mathfrak{D} . Then, candidate a cannot be in the core with respect to an α -large set of coalitions, as long as $\alpha \leq 1 - 2/(\mathfrak{D} + 1)$.*

As a result, it is interesting to study the strategic behavior and the potential coordination dynamics that may arise in the face of an inefficient voting system.

6.4.1 The Model

We consider the following abstract model: For some given voting system, agents are called upon to cast their votes for a series of T days or rounds, where T is sufficiently large. After the end of each day, voters are informed about the results of the round, and the winner is determined based on the results of the *ultimate* day. This is essentially an iterative implementation of a given voting rule, in place of the *one-shot* execution typically considered, and it is introduced to take into account external information typically accumulated before the actual voting (e.g. through polls). For concreteness, we will assume that the voting rule employed

⁴That is, $\text{SC}(a)/\min_{x \in C} \text{SC}(x) \geq \mathfrak{D}$.

⁵Considering only “large” coalitions is standard in the literature; cf. [53].

in each day is simply the PLURALITY mechanism, not least due to its popularity both in theory and in practice.

Before we describe and analyze natural dynamics in this model, let us first note that if all the voters engage truthfully throughout this game, the victor will coincide with the plurality winner, and as we know there are instances for which this candidate may have $\Omega(m)$ distortion. As a result, Theorem 6.4.1 implies that there will be a large coalition with a $1 - \Theta(1/m)$ fraction of the voters that strictly prefer a different outcome. Indeed, the lower bound of PLURALITY is built upon $m - 1$ clusters of voters formed arbitrarily close, while a different extreme party with roughly the same plurality score could eventually prevail. However, the access to additional information renders this scenario rather unrealistic given that we expect some type of adaptation or coordination mechanism from the agents.

6.4.2 A Greedy Approach

Let us denote with $n_a^{(t)}$ the plurality score of candidate a at round $t \in [T]$. A particularly natural approach for an agent to engage in this scenario consists of maintaining a time-varying parameter $\theta^{(t)}$, which will essentially serve as the “temperature”. Then, at some round $t > 1$ agent i will support the candidate b for which $b \succeq_i a$ for all $a, b \in C^{(t)}$, where $C^{(t)} := \{a \in C : n_a^{(t-1)} \geq \theta^{(t)}\}$.⁶ That is, agents only consider candidates who exceeded some level of support during the previous day. Then, the temperature parameter is updated accordingly, for example with some small constant increment $\theta^{(t+1)} := \theta^{(t)} + \epsilon$, for some $\epsilon > 0$. In this context, observe that for a sufficiently small ϵ these dynamics will converge to an STV winner (based on the parallel universe model). This implies that the greedy tactics already offer an exponential improvement—in terms of the utilitarian efficiency—compared to the truthful dynamics. Nevertheless, the lower bound for STV (Theorem 6.3.4) suggests that we have to design a more careful adaptive rule in order to attain $O(1)$ distortion.

6.4.3 Exploration/Exploitation

The inefficiency of the previous approach—and subsequently of STV—stems from the greedy nature of the iterative process: Agents may choose to dismiss candidates prematurely. For example, this becomes immediately apparent by inspecting the elimination pattern in the lower bound of Theorem 6.3.4. In light of this, the remedy we propose—and what arguably occurs in many practical scenarios—is an exploration phase. In particular, voters initially do not possess any information about the preferences of the rest of the population. Thus, they may attempt to explore several alternatives in order to evaluate the viability of

⁶The definition of the set $C^{(t)}$ for $t > 1$ is subject to $|C^{(t)}| \geq |C^{(t-1)}| - 1$, i.e. agents never disregard more than 1 candidate in the course of a single round; in the contrary case the guarantee we state for the dynamics does not hold due to some pathological instances.

each candidate; while doing so, agents will endeavor to somehow indicate or favor their own preferences. After the exploration phase, agents will leverage the information they have learnt to adapt their support. More concretely, we will consider the following dynamics:

1. Exploration phase: In each round $t \in [m]$ every agent i maintains a list $\mathcal{L}_i^{(t)}$, initialized as $\mathcal{L}_i^{(1)} := \emptyset$. If $C_i^{(t)} := C \setminus \mathcal{L}_i^{(t)}$, then at round t an agent i shall vote for the candidate $a \in C_i^{(t)}$ such that $a \succeq_i b$ for all $b \in C_i^{(t)}$. Then, agent i updates her list accordingly: $\mathcal{L}_i^{(t+1)} := \mathcal{L}_i^{(t)} \cup \{a\}$;
2. Exploitation phase: Every agent supports the first candidate⁷ within her list that managed to accumulate—over all prior rounds—at least $n/2$ votes.

In a sense, voters try to balance between voting for their most-preferred candidates and having an impact on the final result. We shall refer to this iterative process as COORDINATION dynamics.

Theorem 6.4.3. *COORDINATION dynamics lead to a candidate with distortion at most 11.*

Proof. Let $w \in C$ be the winner under COORDINATION dynamics, and $x \in C$ be the candidate who minimizes the social cost. If $r := \text{dist}(x, w)/5$, we consider the sequence of balls $\{\mathcal{B}_i\}_{i=1}^3$ such that $\mathcal{B}_i := \mathcal{B}(x, (2i-1)r)$ for $i = 1, 2, 3$. If γ is the fraction of the voters in \mathcal{B}_1 , we will argue that $\gamma \leq 1/2$.

For the sake of contradiction, let us assume that $\gamma > 1/2$. Let t be the first round for which a voter in \mathcal{B}_1 supports a candidate outside \mathcal{B}_3 . Then, it follows by the triangle inequality that the list of this voter just after round $t-1$ included all the candidates in \mathcal{B}_2 . This in turn implies that by round $t-1$ every agent in \mathcal{B}_1 had already voted for all candidates in \mathcal{B}_1 . Given that $\gamma > 1/2$, we can conclude that no agent from \mathcal{B}_1 voted for w during the exploitation phase.

Now let us consider the first round for which some candidate $a \in C$ accumulated at least $n/2$ votes, which clearly happens during the exploration phase. Then, at the exact same round at least $n/2$ agents have a in their list; this follows since agents vote for different candidates during the exploration phase, and a candidate is always included in the list once voted for. As a result, our tie-breaking assumption implies that there will be a candidate with the support of at least $n/2$ agents during the exploitation phase. But our previous argument shows that this candidate cannot be w , which is an obvious contradiction. As a result, we have shown that $\gamma \leq 1/2$, and the theorem follows by Theorem 6.2.2. \square

Before we conclude this section, let us briefly mention some intriguing open problems related to our results. Specifically, we have attempted to argue that

⁷For simplicity, it is assumed that in case multiple such agents exist we posit some arbitrary but *common* among all agents tie-breaking mechanism.

candidates with small distortion may arise through natural learning rules. This was motivated in part by Theorem 6.4.1, which implies the instability of outcomes with large distortion. However, the converse of this statement is not quite true: Although there always exists a candidate with distortion at most 3 [131], there might be a subset with at least half of the voters that strictly prefer a different outcome (a.k.a. Condorcet’s paradox). Still, there might be an appropriate notion of stability which ensures that near-optimal candidates are in some sense stable. In spirit, this is very much pertinent to the main result of Gkatzelis et al. [131] concerning the existence of an *undominated* candidate, leading to the following question:

Question 3. *Are there deterministic and distributed learning rules which converge to a candidate with distortion 3?*

6.5 Intrinsic Robustness of Plurality Matching

Gkatzelis et al. [131] introduced the PLURALITYMATCHING (deterministic) mechanism, and they showed that it always incurs distortion at most 3 under metric preferences. Nonetheless, it is natural to ask how it performs in more refined, as well as in more general spaces. It should be noted that Gkatzelis et al. [131] established the robustness of PLURALITYMATCHING under different objective functions (measuring the social cost); namely, they showed that the same distortion bound can be achieved for the more stringent *fairness ratio* of Goel et al. [132]. In this section we extend the robustness of PLURALITYMATCHING along two regimes.

6.5.1 Ultra-Metrics

First, we study the power of PLURALITYMATCHING under *ultra-metric spaces*; in particular, recall the following definition:

Definition 6.5.1. *An ultra-metric on a set \mathcal{M} is a function $\text{dist} : \mathcal{M} \times \mathcal{M} \mapsto \mathbb{R}$ such that $\forall x, y, z$,*

1. $\text{dist}(x, y) = 0$ if and only if $x = y$ (*identity of indiscernibles*);
2. $\text{dist}(x, y) = \text{dist}(y, x)$ (*symmetry*);
3. $\text{dist}(x, z) \leq \max\{\text{dist}(x, y), \text{dist}(y, z)\}$ (*ultra-metric inequality*).

Notice that these axioms also imply that $\text{dist}(x, y) \geq 0, \forall x, y \in \mathcal{M}$. We will say that an ultra-metric space is an ordered pair $(\mathcal{M}, \text{dist})$ consisting of a set \mathcal{M} along with an ultra-metric dist on \mathcal{M} . Naturally, every ultra-metric is also a metric since $\max\{\text{dist}(x, y), \text{dist}(y, z)\} \leq \text{dist}(x, y) + \text{dist}(y, z)$, but the converse is

not necessarily true. Perhaps the simplest conceivable ultra-metric is the *discrete metric*, which is defined on a set \mathcal{M} as follows:

$$\text{dist}(x, y) = \begin{cases} 1 & \text{if } x \neq y; \\ 0 & \text{if } x = y, \end{cases}$$

where $x, y \in \mathcal{M}$. As we explained in our introduction, we study this setting mostly driven by the fundamental bottleneck variant in facility location games. Specifically, if the cost of a path corresponds to the maximum-weight edge in the path (posit non-negative weights), and the distance between two nodes in the graph is the minimum-cost path among all possible paths, it is well-known that these (so-called *minimax*) distances satisfy the ultra-metric inequality of Theorem 6.5.1. In Table 6.1 we summarize some lower bounds for well-studied mechanisms; they mostly follow directly from the techniques of Anshelevich et al. [18], and thus we omit their proof.

Mechanism	Lower Bound
Plurality & Borda	m
k -top	$\Omega(m/k)$
Approval & Veto	n
Any Deterministic	2

Table 6.1: Lower bounds for standard mechanisms under ultra-metric spaces. We use k -top to represent any *deterministic* mechanism which elicits only the k -top preferences.

In particular, we first prove a lower bound of 2 for any deterministic mechanism under ultra-metrics:

Proposition 6.5.2. *There exists a voting profile for which the distortion of any deterministic mechanism under an ultra-metric space is at least 2.*

Proof. Consider a voting instance with two candidates a, b and n voters, such that the votes between the two candidates are split equally. Assume without any loss of generality that the mechanism eventually selects candidate b . We will present an ultra-metric space for which the social cost of a is half than the social cost of b . Specifically, consider an unweighted path graph with 3 nodes endowed with the minimax distance. We assume that candidate a resides in the leftmost node of the graph along with all of the voters who supported a ; on the other hand, candidate b resides in the rightmost node of the graph, while all of her supporters lie in the intermediate node (see Figure 6.4). Then, it follows that $\text{SC}(a) = n/2$ whereas $\text{SC}(b) = n$, as desired. \square




Figure 6.4: An example of our construction for the proof of Theorem 6.5.2 for $n = 2$ voters; the argument is similar to the one for metric spaces, but observe that in this case $\text{dist}(1, b) = 1 \neq 2$ since we have considered minimax distances.

Importantly, we will show that PLURALITYMATCHING always matches this lower bound. To keep the exposition reasonable self-contained we shall first recall some basic ingredients developed in [131].

Definition 6.5.3 ([131], Definition 5). *For an election $\mathcal{E} = (V, C, \sigma)$ and a candidate $a \in C$, the integral domination graph of candidate a is the bipartite graph $G^\mathcal{E}(a) = (V, V, E_a)$, where $(i, j) \in E_a$ if and only if $a \succeq_i \text{top}(j)$.*

Proposition 6.5.4 ([131], Corollary 1). *There exists a candidate $a \in C$ whose integral domination graph $G^\mathcal{E}(a)$ admits a perfect matching.*

We should also note that a candidate whose integral domination graph admits a perfect matching can be identified in strongly polynomial time. In particular, PLURALITYMATCHING always returns such a candidate. These ingredients suffice in order to establish the following:

Theorem 6.5.5. PLURALITYMATCHING *returns a candidate with distortion at most 2 under any ultra-metric space.*

Proof of Theorem 6.5.5. Let $a \in C$ be a candidate whose integral domination graph $G^\mathcal{E}(a)$ admits a perfect matching $M : V \mapsto V$ (recall Theorem 6.5.4), such

that $a \succeq_i \text{top}(M(i))$ for all $i \in V$. Then, it follows that

$$\begin{aligned}
\text{SC}(a) &= \sum_{i \in V} \text{dist}(i, a) \\
&\leq \sum_{i \in V} \text{dist}(i, \text{top}(M(i))) && (a \succeq_i \text{top}(M(i)), \forall i \in V) \\
&\leq \sum_{i \in V} \text{dist}(i, b) + \text{dist}(b, \text{top}(M(i))) && (\text{triangle inequality}) \\
&= \text{SC}(b) + \sum_{i \in V} \text{dist}(b, \text{top}(i)) && (M \text{ is a perfect matching}) \\
&\leq \text{SC}(b) + \sum_{i \in V} \max\{\text{dist}(i, b), \text{dist}(i, \text{top}(i))\} && (\text{ultra-metric inequality}) \\
&= \text{SC}(b) + \sum_{i \in V} \text{dist}(i, b) && (\text{top}(i) \succeq_i b) \\
&= 2 \text{SC}(b).
\end{aligned}$$

Given that the choice of b was arbitrary the theorem follows. \square

6.5.2 Approximate Metrics

Next, we study the distortion of deterministic mechanisms when the distances *approximately* satisfy the triangle inequality, as formalized in the following definition:

Definition 6.5.6. *For some parameter $\rho \geq 1$, a ρ -approximate metric on a set \mathcal{M} is a function $\text{dist} : \mathcal{M} \times \mathcal{M} \mapsto \mathbb{R}$ such that $\forall x, y, z$,*

1. $\text{dist}(x, y) = 0$ if and only if $x = y$ (*identity of indiscernibles*);
2. $\text{dist}(x, y) = \text{dist}(y, x)$ (*symmetry*);
3. $\text{dist}(x, z) \leq \rho(\text{dist}(x, y) + \text{dist}(y, z))$ (ρ -*relaxed triangle inequality*).

Again we point out that these axioms imply that $\text{dist}(x, y) \geq 0, \forall x, y \in \mathcal{M}$. We commence with the following lower bound:

Proposition 6.5.7. *There exists a voting profile for which the distortion of every deterministic mechanism under a ρ -approximate metric space is at least $\rho^2 + \rho + 1$.*

Proof of Theorem 6.5.7. As usual, consider an instance with $2n$ voters and 2 candidates $a, b \in C$, such that every candidate obtains exactly half of the votes. Let us assume without any loss of generality that the social choice rule selects candidate b . Now consider a ρ -approximate metric $\text{dist}(\cdot, \cdot)$ on the set $\mathcal{M} = \{x, y, z, \omega\}$ defined as follows:

$$\begin{pmatrix} x & y & z & \omega \\ 0 & 1 & 2\rho & \epsilon \\ 1 & 0 & 1 & \rho + \rho\epsilon \\ 2\rho & 1 & 0 & \rho^2 + \rho + \rho\epsilon \\ \epsilon & \rho + \rho\epsilon & \rho^2 + \rho + \rho\epsilon & 0 \end{pmatrix} \begin{matrix} x \\ y \\ z \\ \omega \end{matrix};$$

here we assume that $\epsilon \in (0, 1)$. It is a simple exercise to verify that $\text{dist}(\cdot, \cdot)$ indeed satisfies the axioms of a ρ -approximate metric. We also assume that $a := x$ and $b := z$; the n supporters of candidate a are located on ω , while the n supporters of candidate b on y . Then, it follows that the agents' locations are consistent with their preferences, while the distortion of candidate b reads

$$\frac{\text{SC}(b)}{\text{SC}(a)} = \frac{nd(z, \omega) + nd(y, z)}{nd(x, \omega) + nd(x, y)} = \frac{\rho^2 + \rho + \rho\epsilon + 1}{\epsilon + 1}. \quad (6.13)$$

Taking the supremum of this ratio over $\epsilon \in (0, 1)$ concludes the proof. \square

To provide some intuition let us assume that dist corresponds to the squared Euclidean distance.⁸ If we consider the usual voting scenario wherein the votes are splitted equally among two candidates (see the proof of Theorem 6.5.7 and Figure 6.5), then it follows that the distortion of candidate b reads

$$\frac{\text{SC}(b)}{\text{SC}(a)} = \frac{n \times 1^2 + n \times (2 + \delta)^2}{n \times 1^2 + n \times \delta^2} = \frac{\delta^2 + 4\delta + 5}{\delta^2 + 1}. \quad (6.14)$$




Figure 6.5: The lower bound for squared Euclidean distances.

Interestingly, this ratio increases as δ goes from 0 to a sufficiently small constant, implying a notable qualitative difference compared to the standard metric case. In particular, for the squared Euclidean distance it follows that the distortion is lower-bounded by $\sup_{\delta} \frac{\delta^2 + 4\delta + 5}{\delta^2 + 1} = (4 + 2\sqrt{2}) / (4 - 2\sqrt{2}) \cong 5.8284$. For general approximate metrics we can employ the techniques of Gkatzelis et al. [131] to show the following:

Theorem 6.5.8. *PLURALITYMATCHING returns a candidate with distortion at most $2\rho^2 + \rho$ under ρ -approximate metrics.*

⁸Note that Young's inequality implies that $\text{dist}(\cdot, \cdot)$ is a 2-approximate metric.

As a result, this theorem leaves a gap between the upper bound derived for PLURALITYMATCHING and the lower bound of Theorem 6.5.7 when $\rho > 1$. Nonetheless, it should be noted that for ρ -approximate metrics there exists an *instance-optimal* and computationally efficient mechanism based on linear programming (cf. [132]).

Chapter 7

Node-Max-Cut and the Complexity of Equilibrium in Linear Weighted Congestion Games

Existence and efficient computation of (exact or approximate) equilibria for weighted congestion games have received significant research attention. We briefly discuss here some of the most relevant previous work. There has been significant research interest in the convergence rate of best response dynamics for weighted congestion games (see e.g., [73, 62, 93, 99, 134]). Gairing et al. [126] presented a polynomial algorithm for computing a PNE for load balancing games on restricted parallel links. Caragiannis et al. [64] established existence and presented efficient algorithms for computing approximate PNE in weighted CGs with polynomial latencies (see also [106, 129]).

Bhalgat et al. [41] presented an efficient algorithm for computing a $(3 + \varepsilon)$ -approximate equilibrium in MAX-CUT games, for any $\varepsilon > 0$. The approximation guarantee was improved to $2 + \varepsilon$ in [64]. We highlight that the notion of approximate equilibrium in cut games is much stronger than the notion of approximate local optimum of MAX-CUT, since the former requires that no vertex can significantly improve the total weight of its incidence edges that cross the cut (as e.g., in [41, 64]), while the latter simply requires that the total weight of the cut cannot be significantly improved (as e.g., in [64]).

Johnson et al. [143] introduced the complexity class PLS and proved that CIRCUIT-FLIP is PLS-complete. Subsequently, Schäffer and Yannakakis [208] proved that MAX-CUT is PLS-complete. From a technical viewpoint, our work is close to previous work by Elsässer and Tscheuschner [89] and Gairing and Savani [127], where they show that LOCAL-MAX-CUT in graphs of maximum degree five [89] and computing a PNE for hedonic games [127] are PLS-complete, and by

Ackermann et al. [2], where they reduce LOCAL-MAX-CUT to computing a PNE in network congestion games.

7.1 Contributions

In this work we show that equilibrium computation in linear weighted CGs is significantly harder than for standard CGs, in the sense that it is PLS-complete either for very restricted strategy spaces, namely when player strategies are paths on a series-parallel network with a single origin and destination, or for very restricted latency functions, namely when resource costs are equal to the congestion. Our main step towards proving the latter result is to show that computing a local optimum of NODE-MAX-CUT, a natural and interesting restriction of MAX-CUT where the weight of each edge is the product of the weights of its endpoints, is PLS-complete.

For the complexity of equilibrium computation we will mostly focus on the first part here. More specifically, using a tight reduction from LOCAL-MAX-CUT, we first show, in Section 7.3, that equilibrium computation for linear weighted CGs on series-parallel networks with a single origin and destination is PLS-complete (Theorem 7.3.1). The reduction results in games where both the player weights and the latency slopes are exponential. Our result reveals a remarkable gap between weighted and standard CGs regarding the complexity of equilibrium computation, since for standard CGs on series-parallel networks with a single origin and destination, a PNE can be computed by a simple greedy algorithm [119].

Aiming at a deeper understanding of how different player weights affect the complexity of equilibrium computation in CGs, we show that computing a PNE of weighted network CGs with asymmetric player strategies and identity latency functions is PLS-complete. Again the gap to standard CGs is remarkable, since for standard CGs with identity latency functions, any better response dynamics converges to a PNE in polynomial time. NODE-MAX-CUT plays a role similar to that of threshold games in [2, Sec. 4] in the constructed reduction.

NODE-MAX-CUT is a natural restriction of MAX-CUT and settling the complexity of its local optima computation may be of independent interest, both conceptually and technically. NODE-MAX-CUT coincides with the restriction of MAX-CUT shown (weakly) NP-complete on complete graphs in the seminal paper of Karp [148], while a significant generalization of NODE-MAX-CUT with polynomial weights was shown P-complete in [208].

As a complement to this, we show that a $(1 + \varepsilon)$ -approximate equilibrium for NODE-MAX-CUT, where no vertex can switch sides and increase the weight of its neighbors across the cut by a factor larger than $1 + \varepsilon$, can be computed in time exponential in the number of different weights (see Theorem 7.4.1 for a precise statement). Thus, we can efficiently compute a $(1 + \varepsilon)$ -approximate equilibrium for NODE-MAX-CUT, for any $\varepsilon > 0$, if the number of different vertex weights is constant. Since similar results are not known for MAX-CUT, we believe that

Theorem 7.4.1 may indicate that approximate equilibrium computation for NODE-MAX-CUT may not be as hard as for MAX-CUT. An interesting direction for further research is to investigate (i) the quality of efficiently computable approximate equilibria for NODE-MAX-CUT; and (ii) the smoothed complexity of its local optima.

7.2 Basic Definitions and Notation

Polynomial-Time Local Search (PLS). A *polynomial-time local search* (PLS) problem L [143, Sec. 2] is specified by a (polynomially recognizable) set of instances I_L , a set $S_L(x)$ of feasible solutions for each instance $x \in I_L$, with $|s| = O(\text{poly}(|x|))$ for every solution $s \in S_L(x)$, an objective function $f_L(s, x)$ that maps each solution $s \in S_L(x)$ to its value in instance x , and a *neighborhood* $N_L(s, x) \subseteq S_L(x)$ of feasible solutions for each $s \in S_L(x)$. Moreover, there are three polynomial-time algorithms that for any given instance $x \in I_L$: (i) the first generates an initial solution $s_0 \in S_L(x)$; (ii) the second determines whether a given s is a feasible solution and (if $s \in S_L(x)$) computes its objective value $f_L(s, x)$; and (iii) the third returns either that s is *locally optimal* or a feasible solution $s' \in N_L(s, x)$ with better objective value than s . If L is a maximization (resp. minimization) problem, a solution s is locally optimal if for all $s' \in N_L(s, x)$, $f_L(s, x) \geq f_L(s', x)$ (resp. $f_L(s, x) \leq f_L(s', x)$). If s is not locally optimal, the third algorithm returns a solution $s' \in N_L(s, x)$ with $f(s, x) < f(s', x)$ (resp. $f(s, x) > f(s', x)$). The complexity class PLS consists of all polynomial-time local search problems. By abusing the terminology, we always refer to polynomial-time local search problem simply as local search problems.

PLS Reductions and Completeness. A local search problem L is *PLS-reducible* to a local search problem L' , if there are polynomial-time algorithms ϕ_1 and ϕ_2 such that (i) ϕ_1 maps any instance $x \in I_L$ of L to an instance $\phi_1(x) \in I_{L'}$ of L' ; (ii) ϕ_2 maps any (solution s' of instance $\phi_1(x)$, instance x) pair, with $s' \in S_{L'}(\phi_1(x))$, to a solution $s \in S_L(x)$; and (iii) for every instance $x \in I_L$, if s' is locally optimal for $\phi_1(x)$, then $\phi_2(s', x)$ is locally optimal for x .

By definition, if a local search problem L is PLS-reducible to a local search problem L' , a polynomial-time algorithm that computes a local optimum of L' implies a polynomial time algorithm that computes a local optimum of L . Moreover, a PLS-reduction is transitive. As usual, a local search problem Q is *PLS-complete*, if $Q \in \text{PLS}$ and any local search problem $L \in \text{PLS}$ is PLS-reducible to Q .

Max-Cut and Node-Max-Cut. An instance of MAX-CUT consists of an undirected edge-weighted graph $G(V, E)$, where V is the set of vertices and E is the set of edges. Each edge e is associated with a positive weight w_e . A cut of G is a vertex partition $(S, V \setminus S)$, with $\emptyset \neq S \neq V$. We usually identify a cut with one of its sides (e.g., S). We denote $\delta(S) = \{\{u, v\} \in E : u \in S \wedge v \notin S\}$ the set of edges that cross the cut S . The weight (or the value) of a cut S , denoted $w(S)$,

is $w(S) = \sum_{e \in \delta(S)} w_e$. In MAX-CUT, the goal is to compute an optimal cut S^* of maximum value $w(S^*)$.

In NODE-MAX-CUT, each vertex v is associated with a positive weight w_v and the weight of each edge $e = \{u, v\}$ is $w_e = w_u w_v$, i.e. equal to the product of the weights of e 's endpoints. Again the goal is to compute a cut S^* of maximum value $w(S^*)$. As optimization problems, both MAX-CUT and NODE-MAX-CUT are NP-complete [148].

In this work, we study MAX-CUT and NODE-MAX-CUT as local search problems under the FLIP neighborhood. Then, they are referred to as LOCAL-MAX-CUT and LOCAL-NODE-MAX-CUT. The neighborhood $N(S)$ of a cut $(S, V \setminus S)$ consists of all cuts $(S', V \setminus S')$ where S and S' differ by a single vertex. Namely, the cut S' is obtained from S by moving a vertex from one side of the cut to the other. A cut S is locally optimal if for all $S' \in N(S)$, $w(S) \geq w(S')$. In LOCAL-MAX-CUT (resp. LOCAL-NODE-MAX-CUT), given an edge-weighted (resp. vertex-weighted) graph, the goal is to compute a locally optimal cut. Clearly, both MAX-CUT and NODE-MAX-CUT belong to PLS. In the following, we abuse the terminology and refer to LOCAL-MAX-CUT and LOCAL-NODE-MAX-CUT as MAX-CUT and NODE-MAX-CUT, for brevity, unless we need to distinguish between the optimization and the local search problem.

Weighted Congestion Games. A *weighted congestion game* \mathcal{G} consists of n players, where each player i is associated with a positive weight w_i , a set of resources E , where each resource e is associated with a non-decreasing latency function $\ell_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, and a non-empty strategy set $\Sigma_i \subseteq 2^E$ for each player i . A game is linear if $\ell_e(x) = a_e x + b_e$, for some $a_e, b_e \geq 0$, for all $e \in E$. The identity latency function is $\ell(x) = x$. The player strategies are *symmetric*, if all players share the same strategy set Σ , and *asymmetric*, otherwise.

We focus on *network* weighted congestion games, where the resources E correspond to the edges of an underlying network $G(V, E)$ and the player strategies are paths on G . A network game is *single-commodity*, if G has an origin o and a destination d and the player strategies are all (simple) $o - d$ paths. A network game is *multi-commodity*, if G has an origin o_i and a destination d_i for each player i , and i 's strategy set Σ_i consists of all (simple) $o_i - d_i$ paths. A single-commodity network $G(V, E)$ is *series-parallel*, if it either consists of a single edge (o, d) or can be obtained from two series-parallel networks composed either in series or in parallel (see e.g., [218] for details on composition and recognition of series-parallel networks).

A configuration $\vec{s} = (s_1, \dots, s_n)$ consists of a strategy $s_i \in \Sigma_i$ for each player i . The congestion s_e of resource e in configuration \vec{s} is $s_e = \sum_{i: e \in s_i} w_i$. The cost of resource e in \vec{s} is $\ell_e(s_e)$. The *individual cost* (or *cost*) $c_i(\vec{s})$ of player i in configuration \vec{s} is the total cost for the resources in her strategy s_i , i.e., $c_i(\vec{s}) = \sum_{e \in s_i} \ell_e(s_e)$. A configuration \vec{s} is a *pure Nash equilibrium* (equilibrium or PNE, for brevity), if for every player i and every strategy $s' \in \Sigma_i$, $c_i(\vec{s}) \leq c_i(\vec{s}_{-i}, s')$ (where (\vec{s}_{-i}, s') denotes the configuration obtained from \vec{s} by replacing s_i with s').

Namely, no player can improve her cost by unilaterally switching her strategy.

Equilibrium Computation and Local Search. [118] shows that for linear weighted congestion games, with latencies $\ell_e(x) = a_e x + b_e$, $\Phi(\vec{s}) = \sum_{e \in E} (a_e s_e^2 + b_e s_e) + \sum_i w_i \sum_{e \in s_i} (a_e w_i + b_e)$ changes by $2w_i(c_i(\vec{s}) - c_i(\vec{s}_{-i}, s'))$, when a player i switches from strategy s_i to strategy s' in \vec{s} . Hence, Φ is a weighted potential function, whose local optimal (wrt. single player deviations) correspond to PNE of the underlying game. Hence, equilibrium computation for linear weighted congestion games is in PLS. Specifically, configurations corresponds to solutions, the neighborhood $N(\vec{s})$ of a configuration \vec{s} consists of all configurations (\vec{s}_{-i}, s') with $s' \in \Sigma_i$, for some player i , and local optimality is defined wrt. the potential function Φ .

Max-Cut and Node-Max-Cut as Games. LOCAL-MAX-CUT and LOCAL-NODE-MAX-CUT can be cast as cut games, where players correspond to vertices of $G(V, E)$, strategies $\Sigma = \{0, 1\}$ are symmetric, and configurations $\vec{s} \in \{0, 1\}^{|V|}$ correspond to cuts, e.g., $S(\vec{s}) = \{v \in V : s_v = 0\}$. Each player v aims to maximize $w_v(\vec{s}) = \sum_{e=\{u,v\} \in E: s_u \neq s_v} w_e$, that is the total weight of her incident edges that cross the cut. For NODE-MAX-CUT, this becomes $w_v(\vec{s}) = \sum_{u:\{u,v\} \in E \wedge s_u \neq s_v} w_u$, i.e., v aims to maximize the total weight of her neighbors across the cut. A cut \vec{s} is a PNE if for all players v , $w_v(\vec{s}) \geq w_v(\vec{s}_{-i}, 1 - s_v)$. Equilibrium computation for cut games is equivalent to local optimum computation, and thus, is in PLS.

A cut \vec{s} is a $(1 + \varepsilon)$ -approximate equilibrium, for some $\varepsilon > 0$, if for all players v , $(1 + \varepsilon)w_v(\vec{s}) \geq w_v(\vec{s}_{-i}, 1 - s_v)$. Note that the notion of $(1 + \varepsilon)$ -approximate equilibrium is stronger than the notion of $(1 + \varepsilon)$ -approximate local optimum, i.e., a cut S such that for all $S' \in N(S)$, $(1 + \varepsilon)w(S) \geq w(S')$ (see also the discussion in [64]).

7.3 Hardness of Computing Equilibria in Weighted Congestion Games on Series-Parallel Networks

Theorem 7.3.1. *Computing a pure Nash equilibrium in weighted congestion games on single-commodity series-parallel networks with linear latency functions is PLS-complete.*

Proof sketch. Membership in PLS follows from the potential function argument of [118]. To show hardness, we present a reduction from MAX-CUT.

Let $H(V, A)$ be an instance of LOCAL-MAX-CUT with n vertices and m edges. Based on H , we construct a weighted congestion game on a single-commodity series-parallel network G with $3n$ players, where for every $i \in [n]$, there are three players with weight $w_i = 16^i$. Network G is a parallel composition of two identical copies of a simpler series-parallel network. We refer to these copies as G_1 and G_2 . Each of G_1 and G_2 is a series composition of m simple series-parallel networks F_{ij} ,




Figure 7.1: The series-parallel network F_{ij} that corresponds to edge $\{i, j\} \in A$.




Figure 7.2: An example of the network G constructed in the proof of Theorem 7.3.1 for graph $H(V, A)$, with $V = \{1, 2, 3, 4\}$ and $A = \{\{1, 2\}, \{1, 3\}, \{1, 4\}, \{2, 4\}\}$. G is a parallel composition of two parts, each consisting of the smaller networks F_{12} , F_{13} , F_{14} and F_{24} (see also Figure 7.1) connected in series.

each corresponding to an edge $\{i, j\} \in A$. Network F_{ij} is depicted in Figure 7.1, where D is assumed to be a constant chosen (polynomially) large enough. An example of the entire network G is shown in Figure 7.2.

In each of G_1 and G_2 , there is a unique path that contains all edges with latency functions $\ell_i(x) = Dx/4^i$, for each $i \in [n]$. We refer to these paths as p_i^u for G_1 and p_i^l for G_2 . In addition to the edges with latency $\ell_i(x)$, p_i^u and p_i^l include all edges with latencies $\ell_{ij}(x) = \frac{w_{ij}x}{w_i w_j} = \frac{w_{ij}x}{16^{i+j}}$, which correspond to the edges incident to vertex i in H .

Due to the choice of the player weights and the latency slopes, a player with weight w_i must choose either p_i^u or p_i^l in any PNE. We can prove this claim by induction on the player weights. The players with weight $w_n = 16^n$ have a dominant strategy to choose either p_n^u or p_n^l , since the slope of $\ell_n(x)$ is significantly smaller than the slope of any other latency $\ell_i(x)$. In fact, the slope of ℓ_n is so small that even if all other $3n - 1$ players choose one of p_n^u or p_n^l , a player with weight w_n would prefer either p_n^u or p_n^l over all other paths. Therefore, we can assume that each of p_n^u and p_n^l are used by at least one player with weight w_n in any PNE,

which would increase their latency so much that no player with smaller weight would prefer them any more. The inductive argument applies the same reasoning for players with weights w_{n-1} , who should choose either p_{n-1}^u or p_{n-1}^l in any PNE, and subsequently, for players with weights w_{n-2}, \dots, w_1 . Hence, we conclude that for all $i \in [n]$, each of p_i^u and p_i^l is used by at least one player with weight w_i .

Moreover, we note that two players with different weights, say w_i and w_j , go through the same edge with latency $\ell_{ij}(x) = \frac{w_{ij}x}{w_i w_j}$ in G only if the corresponding edge $\{i, j\}$ is present in H . The correctness of the reduction follows the fact that a player with weight w_i aims to minimize her cost through edges with latencies ℓ_{ij} in G in the same way that in the MAX-CUT instance, we want to minimize the weight of the edges incident to a vertex i and do not cross the cut. Formally, we next show that a cut S is locally optimal for the MAX-CUT instance if and only if the configuration where for every $k \in S$, two players with weight w_k use p_k^u and for every $k \notin S$, two players with weight w_k use p_k^l is a PNE of the weighted congestion game on G .

Assume an equilibrium configuration and consider a player a of weight w_k that uses p_k^u together with another player of weight w_k (if this is not the case, vertex k is not included in S and we apply the symmetric argument for p_k^l). By the equilibrium condition, the cost of player a on p_k^u is at most her cost on p_k^l , which implies that

$$\sum_{k=1}^m \frac{2D16^k}{4^k} + \sum_{j:\{k,j\} \in A} \frac{w_{kj}(2 \cdot 16^k + x_j^u 16^j)}{16^{k+j}} \leq \sum_{k=1}^m \frac{2D16^k}{4^k} + \sum_{j:\{k,j\} \in A} \frac{w_{kj}(2 \cdot 16^k + x_j^l 16^j)}{16^{k+j}},$$

where x_j^u (resp. x_j^l) is either 1 or 2 (resp. 2 or 1) depending on whether, for each vertex j connected to vertex k in H , one or two players (of weight w_j) use p_j^u . Simplifying the inequality above, we obtain that:

$$\sum_{j:\{k,j\} \in A} w_{kj}(x_j^u - 1) \leq \sum_{j:\{k,j\} \in A} w_{kj}(x_j^l - 1) \quad (7.1)$$

Let $S = \{i \in V : x_i^u = 2\}$. By hypothesis, $k \in S$ and the left-hand side of (7.1) corresponds to the total weight of the edges in H that are incident to k and do not cross the cut S . Similarly, the right-hand side of (7.1) corresponds to the total weights of the edges in H that are incident to k and cross the cut S . Therefore, (7.1) implies that we cannot increase the value of the cut S by moving vertex k from S to $V \setminus S$. Since this or its symmetric condition holds for any vertex k of H , the cut $(S, V \setminus S)$ is locally optimal. To conclude the proof, we argue along the same lines that any locally optimal cut of H corresponds to a PNE in the weighted congestion game on G . \square

7.4 Computing Approximate Equilibria for Node-Max-Cut

We complement our PLS-completeness proof for NODE-MAX-CUT, with an efficient algorithm computing $(1 + \varepsilon)$ -approximate equilibria for NODE-MAX-CUT, when the number of different vertex weights is a constant. We note that similar results are not known (and a similar approach fails) for MAX-CUT. Investigating if stronger approximation guarantees are possible for efficiently computable approximate equilibria for NODE-MAX-CUT is beyond the scope of this work and an intriguing direction for further research.

Given a vertex-weighted graph $G(V, E)$ with n vertices and m edges, our algorithm, called BRIDGEGAPS, computes a $(1 + \varepsilon)^3$ -approximate equilibrium for a NODE-MAX-CUT, for any $\varepsilon > 0$, in $(m/\varepsilon)(n/\varepsilon)^{O(D_\varepsilon)}$ time, where D_ε is the number of different vertex weights in G , when the weights are rounded down to powers of $1 + \varepsilon$. We next sketch the algorithm and the proof of Theorem 7.4.1.

For simplicity, we assume that n/ε is an integer and that vertices are indexed in nondecreasing order of weight, i.e., $w_1 \leq w_2 \leq \dots \leq w_n$. BRIDGEGAPS first rounds down vertex weights to the closest power of $(1 + \varepsilon)$. Namely, each weight w_i is replaced by weight $w'_i = (1 + \varepsilon)^{\lfloor \log_{1+\varepsilon} w_i \rfloor}$. Clearly, an $(1 + \varepsilon)^2$ -approximate equilibrium for the new instance G' is an $(1 + \varepsilon)^3$ -approximate equilibrium for the original instance G . The number of different weights D_ε , used in the analysis, is defined wrt. the new instance G' .

Then, BRIDGEGAPS partitions the vertices of G' into groups g_1, g_2, \dots , so that the vertex weights in each group increase with the index of the group and the ratio of the maximum weight in group g_j to the minimum weight in group g_{j+1} is no less than n/ε . This can be performed by going through the vertices, in nondecreasing order of their weights, and assign vertex $i + 1$ to the same group as vertex i , if $w'_{i+1}/w'_i \leq n/\varepsilon$. Otherwise, vertex $i + 1$ starts a new group. The idea is that for an $(1 + \varepsilon)^2$ -approximate equilibrium in G' , we only need to enforce the $(1 + \varepsilon)$ -approximate equilibrium condition for each vertex i only for i 's neighbors in the highest-indexed group (that includes some neighbor of i). To see this, let g_j be the highest-indexed group that includes some neighbor of i and let ℓ be the lowest indexed neighbor of i in g_j . Then, the total weight of i 's neighbors in groups g_1, \dots, g_{j-1} is less than $\varepsilon w'_\ell$. This holds because i has at most $n - 2$ neighbors in these groups and by definition, $w'_q \leq (\varepsilon/n)w'_\ell$, for any i 's neighbor q in groups g_1, \dots, g_{j-1} . Therefore, we can ignore all neighbors of i in groups g_1, \dots, g_{j-1} , at the expense of one more $1 + \varepsilon$ factor in the approximate equilibrium condition.

Since for every vertex i , we need to enforce its (approximate) equilibrium condition only for i 's neighbors in a single group, we can scale down vertex weights in the same group uniformly (i.e., dividing all the weights in each group by the same factor), as long as we maintain the key property in the definition of groups (i.e., that the ratio of the maximum weight in group g_j to the minimum weight in group g_{j+1} is no less than n/ε). Hence, we uniformly scale down the weights

in each group so that (i) the minimum weight in group g_1 becomes 1; and (ii) for each $j \geq 2$, the ratio of the maximum weight in group g_{j-1} to the minimum weight in group g_j becomes exactly n/ε . This results in a new instance G'' where the minimum weight is 1 and the maximum weight is $(n/\varepsilon)^{D_\varepsilon}$. Therefore, a $(1+\varepsilon)$ -approximate equilibrium in G'' can be computed, in a standard way, after at most $(m\varepsilon)(n/\varepsilon)^{2D_\varepsilon}$ ε -best response moves.

Putting everything together and using $\varepsilon' = \varepsilon/7$, so that $(1+\varepsilon')^3 \leq 1+\varepsilon$, for all $\varepsilon \in (0, 1]$, we obtain the following. We note that the running time of BRIDGEGAPS is polynomial, if $D_\varepsilon = O(1)$ (and quasipolynomial if $D_\varepsilon = \text{poly}(\log n)$).

Theorem 7.4.1. *For any vertex-weighted graph G with n vertices and m edges and any $\varepsilon > 0$, BRIDGEGAPS computes a $(1+\varepsilon)$ -approximate pure Nash equilibrium for NODE-MAX-CUT on G in $(m/\varepsilon)(n/\varepsilon)^{O(D_\varepsilon)}$ time, where D_ε denotes the number of different vertex weights in G , after rounding them down to the nearest power of $1 + \varepsilon$.*

7.A The Proof of Theorem 7.3.1

We will reduce from the PLS-complete problem MAX-CUT and given an instance of MAX-CUT we will construct a network weighted network Congestion Game for which the Nash equilibria will correspond to maximal solutions of MAX-CUT and vice versa. First we give the construction and then we prove the theorem. For the formal PLS-reduction, which needs functions ϕ_1 and ϕ_2 , ϕ_1 returns the (polynomially) constructed instance described below and ϕ_2 will be revealed later in the proof.

Let $H(V, E)$ be an edge-weighted graph of a MAX-CUT instance and let $n = |V|$ and $m = |E|$. In the constructed network weighted CG instance there will be $3n$ players which will share n different weights inside the set $\{16^i : i \in [n]\}$ so that for every $i \in [n]$ there are exactly 3 players having weight $w_i = 16^i$. All players share a common origin-destination pair $o - d$ and choose $o - d$ paths on a series-parallel graph G . Graph G is a parallel composition of two identical copies of a series-parallel graph. Call these copies G_1 and G_2 . In turn, each of G_1 and G_2 is a series composition of m different series-parallel graphs, each of which corresponds to the m edges of H . For every $\{i, j\} \in E$ let F_{ij} be the series-parallel graph that corresponds to $\{i, j\}$. Next we describe the construction of F_{ij} , also shown in Fig. 7.1.

F_{ij} has 3 vertices, namely o_{ij}, v_{ij} and d_{ij} and $n + 1$ edges. For any $k \in [n]$ other than i, j there is an $o_{ij} - d_{ij}$ edge with latency function $\ell_k(x) = \frac{Dx}{4^k}$, where D serves as a big constant to be defined later. There are also two $o_{ij} - v_{ij}$ edges, one with latency function $\ell_i(x) = \frac{Dx}{4^i}$ and one with latency function $\ell_j(x) = \frac{Dx}{4^j}$. Last, there is a $v_{ij} - d_{ij}$ edge with latency function $\ell_{ij}(x) = \frac{w_{ij}x}{w_i w_j}$, where w_{ij} is the weight of edge $\{i, j\} \in E$ and w_i and w_j are the weights of players i and j , respectively, as described earlier. Note that in every F_{ij} and for any $k \in [n]$ the

latency function $\ell_k(x) = \frac{Dx}{k}$ appears in exactly one edge. With F_{ij} defined, an example of the structure of such a network G is given in Fig. 7.2.

Observe that in each of G_1 and G_2 there is a unique path that contains all the edges with latency functions $\ell_i(x)$, for $i \in [n]$, and call these paths p_i^u and p_i^l for the upper (G_1) and lower (G_2) copy respectively. Note that each of p_i^u and p_i^l in addition to those edges, contains some edges with latency function of the form $\frac{w_{ij}x}{w_i w_j}$. These edges for path p_i^u or p_i^l is in one to one correspondence to the edges of vertex i in H and this is crucial for the proof.

We go on to prove the correspondence of Nash equilibria in G to maximal cuts in H , i.e., solutions of MAX-CUT. We will first show that at a Nash equilibrium, a player of weight w_i chooses either p_i^u or p_i^l . Additionally, we prove that p_i^u and p_i^l will have at least one player (of weight w_i). This already provides a good structure of a Nash equilibrium and players of different weights, say w_i and w_j , may go through the same edge in G (the edge with latency function $w_{ij}x/w_i w_j$) only if $\{i, j\} \in E$. The correctness of the reduction lies in the fact that players in G try to minimize their costs incurred by these type of edges in the same way one wants to minimize the sum of the weights of the edges in each side of the cut when solving MAX-CUT.

To begin with, we will prove that at equilibrium any player of weight w_i chooses either p_i^u or p_i^l and at least one such player chooses each of p_i^u and p_i^l . For that, we will need the following proposition as a building block, which will also reveal a suitable value for D .

Proposition 7.A.1. *For some $i, j \in [n]$ consider F_{ij} (Fig. 7.1) and assume that for all $k \in [n]$, there are either one, two or three players of weight w_k that have to choose an $o_{ij} - d_{ij}$ path. At equilibrium, all players of weight w_k (for any $k \in [n]$) will go through the path that contains a edge with latency function $\ell_k(x)$.*

Proof. The proof is by induction on the different weights starting from bigger weights. For any $k \in [n]$ call e_k the edge of F_{ij} with latency function $\ell_k(x)$ and call e_{ij} the edge with latency function $\frac{w_{ij}x}{w_i w_j}$. For some $k \in [n]$ assume that for all $l > k$ all players of weight w_l have chosen the path containing e_l and lets prove that this is the case for players of weight w_k as well. Since D is going to be big enough, for the moment ignore edge e_{ij} and assume that in F_{ij} there are only n parallel paths each consisting of a single edge.

Let the players be at equilibrium and consider any player, say player K , of weight w_k . The cost she computes on e_k is upper bounded by the cost of e_k if all players with weight up to w_k are on e_k , since by induction players with weight $> w_k$ are not on e_k at equilibrium. This cost is upper bounded by $c^k = \frac{D(3 \sum_{l=1}^k 16^l)}{4^k} = \frac{3D \frac{16^{k+1}-1}{16-1}}{4^k}$.

For any edge e_l for $l < k$, the cost that K computes is lower bounded by $c^{<} = \frac{D16^k}{4^{k-1}}$ since she must include herself in the load of e_l and the edge with the

smallest slope in its latency function is e_{k-1} . But then $c^k < c^<$, since

$$c^k < c^< \Leftrightarrow \frac{3D \frac{16^{k+1}-1}{16-1}}{4^k} < \frac{D16^k}{4^{k-1}} \Leftrightarrow 48 \cdot 16^k - 3 < 60 \cdot 16^k$$

Thus, at equilibrium players of weight w_k cannot be on any of the e_l 's for all $l < k$. On the other hand, the cost that K computes for e_l for $l > k$ is at least $c_l^> = \frac{D(16^l+16^k)}{4^l}$, since by induction e_l is already chosen by at least one player of weight w_l . But then $c^k < c_l^>$ since

$$\begin{aligned} \frac{3D \frac{16^{k+1}-1}{16-1}}{4^k} &< \frac{D(16^l + 16^k)}{4^l} \Leftrightarrow \\ 48 \cdot 16^k - 3 &< 15 \frac{16^l + 16^k}{4^{l-k}} \Leftrightarrow \\ 48 \cdot 4^{l-k} 16^k &< 15 \cdot 16^l = 15 \cdot 4^{l-k} 4^{l-k} 16^k. \end{aligned}$$

Thus, at equilibrium players of weight w_k cannot be on any of the e_l 's for all $l > k$.

This completes the induction for the simplified case where we ignored the existence of e_{ij} , but lets go on to include it and define D so that the same analysis goes through. By the above, $c^< - c^k = \frac{D16^k}{4^{k-1}} - \frac{3D \frac{16^{k+1}-1}{16-1}}{4^k} > D$ and also for any $l > k$ it is $c_l^> - c^k = \frac{D(16^l+16^k)}{4^l} - \frac{3D \frac{16^{k+1}-1}{16-1}}{4^k} > D$ (this difference is minimized for $l = k + 1$). On the other hand the maximum cost that edge e_{ij} may have is bounded above by $c^{ij} = \frac{w_{ij} 3 \sum_{l=1}^n 16^l}{w_i w_j}$, as e_{ij} can be chosen by at most all of the players and note that $c^{ij} \leq 16^{n+1} \max_{q,r \in [n]} w_{qr}$. Thus, one can choose a big value for D , namely $D = 16^{n+1} \max_{q,r \in [n]} w_{qr}$, so that even if a player with weight w_k has to add the cost of e_{ij} when computing her path cost, it still is $c^{ij} + c^k < c^<$ (since $c^< - c^k > D \geq c_{ij}$) and for all $l > k$: $c^{ij} + c^k < c_l^>$ (since $c_l^> - c^k > D \geq c^{ij}$), implying that at equilibrium all players of weight w_k may only choose the path through e_k . \square

Other than revealing a value for D , the proof of Proposition 7.A.1 reveals a crucial property: a player of weight w_k in F_{ij} strictly prefers the path containing e_k to the path containing e_l for any $l < k$, independent to whether players of weight $> w_k$ are present in the game or not. With this in mind we go back to prove that at equilibrium any player of weight w_i chooses either p_i^u or p_i^l and at least one such player chooses each of p_i^u and p_i^l . The proof is by induction, starting from bigger weights.

Assume that by the inductive hypothesis for every $i > k$, players with weights w_i have chosen paths p_i^u or p_i^l and at least one such player chooses each of p_i^u and p_i^l . Consider a player of weight w_k , and, wlog, let her have chosen an $o-d$ path through G_1 . Since at least one player for every bigger weight is by induction already in the paths of G_1 (each in her corresponding p_i^u), Proposition 7.A.1 and

the remark after its proof give that in each of the $F_{i,j}$'s the player of weight w_k has chosen the subpath of p_k^u , and this may happen only if her chosen path is p_k^u . It remains to show that there is another player of weight w_k that goes through G_2 , which, with an argument similar to the previous one, is equivalent to this player choosing path p_k^l .

To reach a contradiction, let p_k^u be chosen by all three players of weight w_k , which leaves p_k^l empty. Since all players of bigger weights are by induction settled in paths completely disjoint to p_k^l , the load on this path if we include a player of weight w_k is upper bounded by the sum of all players of weight $< w_k$ plus w_k , i.e., $16^k + 3 \sum_{t=1}^{k-1} 16^t = 16^k + 3 \frac{16^k - 1}{16 - 1}$, which is less than the lower bound on the load of p_k^u , i.e., $3 \cdot 16^k$ (since p_k^u carries 3 players of weight 16^k). This already is a contradiction to the equilibrium property, since p_k^u and p_k^l share the exact same latency functions on their edges which, given the above inequality on the loads, makes p_k^u more costly than p_k^l for a player of weight w_k . To summarize, we have the following.

Proposition 7.A.2. *At equilibrium, for every $i \in [n]$ a player of weight w_i chooses either p_i^u or p_i^l . Additionally, each of p_i^u and p_i^l have been chosen by at least one player (of weight w_i).*

Finally, we prove that every equilibrium of the constructed instance corresponds to a maximal solution of MAX-CUT and vice versa. Given a maximal solution S of MAX-CUT we will show that the configuration Q that for every $k \in S$ routes 2 players through p_k^u and 1 player through p_k^l and for every $k \in V \setminus S$ routes 1 player through p_k^u and 2 players through p_k^l is an equilibrium. Conversely, given an equilibrium Q the cut $S = \{k \in V : 2 \text{ players have chosen } p_k^u \text{ at } Q\}$ is a maximal solution of MAX-CUT.

Assume that we are at equilibrium and consider a player of weight w_k that has chosen p_k^u and wlog p_k^u is chosen by two players (of weight w_k). By the equilibrium conditions the cost she computes for p_k^u is at most the cost she computes for p_k^l , which, given Proposition 7.A.2, implies

$$\sum_{i=1}^m \frac{2D16^k}{4^k} + \sum_{\{k,j\} \in E} \frac{w_{kj}(2 \cdot 16^k + x_j^u 16^j)}{16^k 16^j} \leq \sum_{i=1}^m \frac{2D16^k}{4^k} + \sum_{\{k,j\} \in E} \frac{w_{kj}(2 \cdot 16^k + x_j^l 16^j)}{16^k 16^j}$$

where x_j^u (resp. x_j^l) is either 1 or 2 (resp. 2 or 1) depending whether, for any $j : \{k,j\} \in E$, one or two players (of weight w_j) respectively have chosen path p_j^u . By canceling out terms, the above implies

$$\sum_{\{k,j\} \in E} w_{kj} x_j^u \leq \sum_{\{k,j\} \in E} w_{kj} x_j^l \Leftrightarrow \sum_{\{k,j\} \in E} w_{kj} (x_j^u - 1) \leq \sum_{\{k,j\} \in E} w_{kj} (x_j^l - 1) \quad (7.2)$$

Define $S = \{i \in V : x_i^u = 2\}$. By our assumption it is $k \in S$ and the left side of (7.2), i.e., $\sum_{\{k,j\} \in E} w_{kj} (x_j^u - 1)$, is the sum of the weights of the edges

of H with one of its vertices being k and the other belonging in S . Similarly, the right side of (7.2), i.e., $\sum_{\{k,j\} \in E} w_{kj}(x_j^l - 1)$ is the sum of the weights of the edges with one of its vertices being k and the other belonging in $V \setminus S$. But then (7.2) directly implies that for the (neighboring) cut S' where k goes from S to $V \setminus S$ it holds $w(S) \geq w(S')$. Since k was arbitrary (given the symmetry of the problem), this holds for every $k \in [n]$ and thus for every $S' \in N(S)$ it is $w(S) \geq w(S')$ proving one direction of the claim. Observing that the argument works backwards we complete the proof. For the formal proof, to define function ϕ_2 , given the constructed instance and one of its solutions, say s' , ϕ_2 returns solution $s = \{k \in V : 2 \text{ players have chosen } p_k^u \text{ at } s'\}$. \square

7.B Missing Technical Details from the Analysis of BridgeGaps

In this section, we present an algorithm that computes approximate equilibria for NODE-MAX-CUT. Let $G(V, E)$ be vertex-weighted graph with n vertices and m edges, and consider any $\varepsilon > 0$. The algorithm, called BRIDGEGAPS and formally presented in Algorithm 9, returns a $(1 + \varepsilon)^3$ -approximate equilibrium (Lemma 7.B.2) for G in time $O(\frac{m}{\varepsilon} \lceil \frac{n}{\varepsilon} \rceil^{2D_\varepsilon})$ (Lemma 7.B.1), where D_ε is the number of different rounded weights, i.e., the weights produced by rounding down each of the original weights to its closest power of $(1 + \varepsilon)$. To get a $(1 + \varepsilon)$ -approximate equilibrium, for $\varepsilon < 1$, it suffices to run the algorithm with $\varepsilon' = \frac{\varepsilon}{7}$.

Description of the Algorithm. BRIDGEGAPS first creates an instance G' with weights rounded down to their closest power of $(1 + \varepsilon)$, i.e., weight w_i is replaced by weight $w'_i = (1 + \varepsilon)^{\lfloor \log_{1+\varepsilon} w_i \rfloor}$ in G' , and then computes a $(1 + \varepsilon)^2$ -approximate equilibrium for G' . Observe that any $(1 + \varepsilon)^2$ -approximate equilibrium for G' is a $(1 + \varepsilon)^3$ -approximate equilibrium for G , since

$$\sum_{j \in V_i: s_i = s_j} w_j \leq (1 + \varepsilon) \sum_{j \in V_i: s_i = s_j} w'_j \leq (1 + \varepsilon)^3 \sum_{j \in V_i: s_i \neq s_j} w'_j \leq (1 + \varepsilon)^3 \sum_{j \in V_i: s_i \neq s_j} w_j,$$

where V_i denotes the set of vertices that share an edge with vertex i , with the first and third inequalities following from the rounding and the second one following from the equilibrium condition for G' .

To compute a $(1 + \varepsilon)^2$ -approximate equilibrium, BRIDGEGAPS first sorts the vertices in increasing weight order and note that, wlog, we may assume that $w'_1 = 1$, as we may simply divide all weights by w'_1 . Then, it groups the vertices so that the fraction of the weights of consecutive vertices in the same group is bounded above by $\lceil n/\varepsilon \rceil$, i.e., for any i , vertices i and $i + 1$ belong in the same group if and only if $\frac{w'_{i+1}}{w'_i} \leq \lceil \frac{n}{\varepsilon} \rceil$. This way, groups g_j are formed on which we assume an increasing order, i.e., for any j , the vertices in g_j have smaller weights than those in g_{j+1} .

The next step is to bring the groups closer together using the following process which will generate weights w''_i . For all j , all the weights of vertices on heavier

groups, i.e., groups g_{j+1}, g_{j+2}, \dots , are divided by $d_j = \frac{1}{\lceil n/\varepsilon \rceil} \frac{w_{j+1}^{min}}{w_j^{max}}$ so that $\frac{w_{j+1}^{min}/d_j}{w_j^{max}} = \lceil \frac{n}{\varepsilon} \rceil$, where w_{j+1}^{min} is the smallest weight in g_{j+1} and w_j^{max} is the biggest weight in g_j . For vertex i , let the resulting weight be w''_i , i.e., $w''_i = \frac{w'_i}{\prod_{j \in I_i} d_j}$, where I_i contains the indexes of groups below i 's group, and keep the increasing order on the vertex weights. Observe that by the above process for any i : $\frac{w''_{i+1}}{w''_i} \leq \lceil \frac{n}{\varepsilon} \rceil$, either because i and $i+1$ are in the same group or because the groups are brought closer together. Additionally if i and $i+1$ belong in different groups then $\frac{w''_{i+1}}{w''_i} = \lceil \frac{n}{\varepsilon} \rceil$, implying that for vertices i, i' in different groups with $w''_{i'} > w''_i$ it is $\frac{w''_{i'}}{w''_i} \geq \lceil \frac{n}{\varepsilon} \rceil$. Thus, if we let D_ε be the number of different weights in G' , i.e., $D_\varepsilon = |\{w'_i : i \text{ vertex of } G'\}|$, then the maximum weight w''_n is $w''_n = \frac{w''_n}{w''_{n-1}} \frac{w''_{n-1}}{w''_{n-2}} \dots \frac{w''_2}{w''_1} \leq \lceil \frac{n}{\varepsilon} \rceil^{D_\varepsilon}$

In a last step, using the w'' weights, BRIDGEGAPS starts from an arbitrary configuration (a 0-1 vector) and lets the vertices play ε -best response moves, i.e., as long as there is an index i of the vector violating the $(1 + \varepsilon)$ -approximate equilibrium condition, BRIDGEGAPS flips its bit. When there is no such index BRIDGEGAPS ends and returns the resulting configuration.

Lemma 7.B.1. *For any $\varepsilon > 0$, BRIDGEGAPS terminates in time $O(\frac{m}{\varepsilon} \lceil \frac{n}{\varepsilon} \rceil^{2D_\varepsilon})$*

Proof. We are going to show the claimed bound for the last step of BRIDGEGAPS since all previous steps can be (naively) implemented to end in $O(n^2)$ time.

The proof relies on a potential function argument. For any $\vec{s} \in \{0, 1\}^n$, let

$$\Phi(\vec{s}) = \frac{1}{2} \sum_{i \in V} \sum_{j \in V_i: s_i = s_j} w''_i w''_j.$$

Since for the maximum weight w''_n it is $w''_n \leq \lceil \frac{n}{\varepsilon} \rceil^{D_\varepsilon}$, it follows that $\Phi(\vec{s}) \leq m \lceil \frac{n}{\varepsilon} \rceil^{2D_\varepsilon}$. On the other hand whenever an ε -best response move is made by BRIDGEGAPS producing \vec{s}' from some \vec{s} , Φ decreases by at least ε , i.e., $\Phi(\vec{s}) - \Phi(\vec{s}') \geq \varepsilon$. This is because, if i is the index flipping bit from \vec{s} to \vec{s}' , then by the violation of the $(1 + \varepsilon)$ -equilibrium condition

$$w''_i \sum_{j \in V_i: s_i = s_j} w''_j \geq w''_i (1 + \varepsilon) \sum_{j \in V_i: s_i \neq s_j} w''_j \Rightarrow \sum_{j \in V_i: s_i = s_j} w''_i w''_j - \sum_{j \in V_i: s'_i = s'_j} w''_i w''_j \geq \varepsilon,$$

since $\sum_{j \in V_i: s'_i = s'_j} w''_i w''_j \geq 1$, and

$$\Phi(\vec{s}) - \Phi(\vec{s}') = \sum_{j \in V_i: s_i = s_j} w''_i w''_j - \sum_{j \in V_i: s'_i = s'_j} w''_i w''_j \geq \varepsilon.$$

Consequently, the last step of the algorithm will do at most $\frac{m \lceil \frac{n}{\varepsilon} \rceil^{2D_\varepsilon}}{\varepsilon}$ ε -best response moves. \square

Mechanism 9: BRIDGEGAPS, computing $(1 + \varepsilon)^3$ -approximate equilibria

Input: A NODE-MAX-CUT instance $G(V, E)$ with n vertices and weights

$\{w_i\}_{i \in [n]}$ sorted increasingly with $w_1 = 1$, and an $\varepsilon > 0$.

Output: A vector $\vec{s} \in \{0, 1\}^n$ partitioning the vertices in two sets.

```

1 for  $i \in [n]$  do  $w_i := (1 + \varepsilon)^{\lfloor \log_{1+\varepsilon} w_i \rfloor}$ 
2 groups := 1;
  insert  $w_1$  into  $g_{groups}$ ; ; // Assign the weights into groups
   $\{g_j\}_{j \in [groups]}$ 
  for  $i \in \{2, \dots, n\}$  do
    if  $\frac{w_i}{w_{i-1}} > \lceil \frac{n}{\varepsilon} \rceil$  then groups++;
    insert  $w_i$  into  $g_{groups}$ ;
3 for  $j \in \{2, \dots, groups\}$  do
   $w_j^{min} :=$  minimum weight of group  $g_j$ ; ; // Bring the groups  $\lceil \frac{n}{\varepsilon} \rceil$ 
  close
   $w_{j-1}^{max} :=$  maximum weight of group  $g_{j-1}$ ;
   $d_j = \frac{1}{\lceil n/\varepsilon \rceil} \frac{w_{j+1}^{min}}{w_j^{max}}$ 
  for  $w_i \in g_j \cup \dots \cup g_{group}$  do  $w_i := w_i/d_j$ ;
4  $\vec{s} :=$  an arbitrary  $\{0, 1\}^n$  vector;
5 For all  $i$ , let  $V_i = \{j : \{i, j\} \in E\}$  be the neighborhood of  $i$  in  $G$ ;
  while  $\exists i : \sum_{j \in V_i: s_i = s_j} w_j > (1 + \varepsilon) \sum_{j \in V_i: s_i \neq s_j} w_j$  do
     $s_i := 1 - s_i$ ; ; // Moves towards equilibrium
6 return  $\vec{s}$ .
```

Lemma 7.B.2. For any vertex-weight graph G and any $\varepsilon > 0$, BRIDGEGAPS returns a $(1 + \varepsilon)^3$ -approximate equilibrium for NODE-MAX-CUT in G .

Proof. Clearly, BRIDGEGAPS terminates with a vector \vec{s} that is a $(1 + \varepsilon)$ -approximate equilibrium for the instance with the w'' weights. It suffices to show that \vec{s} is a $(1 + \varepsilon)^2$ -approximate equilibrium for G' , i.e., the instance with the w' weights, since this will directly imply that \vec{s} is a $(1 + \varepsilon)^3$ -approximate equilibrium for G , as already discussed at the beginning of the description of the algorithm.

Consider any index i and let V_i^h be the neighbors of i that belong in the heaviest group among the neighbors of i . By the $(1 + \varepsilon)$ -approximate equilibrium condition it is

$$\sum_{j \in V_i \setminus V_i^h: s_i = s_j} w_j'' + \sum_{j \in V_i^h: s_i = s_j} w_j'' \leq (1 + \varepsilon) \left(\sum_{j \in V_i \setminus V_i^h: s_i \neq s_j} w_j'' + \sum_{j \in V_i^h: s_i \neq s_j} w_j'' \right). \quad (7.3)$$

Recalling that for every j , $w_j'' = \frac{w_j'}{\prod_{k \in I_j} d_k}$, where I_j contains the indexes of groups

below j 's group, and letting $D = \prod_{k \in I_j} d_k$, for a $j \in V_i^h$, gives

$$\sum_{j \in V_i \setminus V_i^h : s_i = s_j} w'_j + \sum_{j \in V_i^h : s_i = s_j} w'_j \leq D \left(\sum_{j \in V_i \setminus V_i^h : s_i = s_j} w''_j + \sum_{j \in V_i^h : s_i = s_j} w''_j \right) \quad (7.4)$$

On the other hand for any j and j' , if j' belongs in a group lighter than j then by construction $\frac{w''_j}{w''_{j'}} \geq \frac{n}{\varepsilon}$ (recall the way the groups were brought closer), which gives $nw''_{j'} \leq \varepsilon w''_j$, yielding

$$\sum_{j \in V_i \setminus V_i^h : s_i \neq s_j} w''_j + \sum_{j \in V_i^h : s_i \neq s_j} w''_j \leq (1 + \varepsilon) \sum_{j \in V_i^h : s_i \neq s_j} w''_j \quad (7.5)$$

Using equations (7.4), (7.3) and (7.5), in this order, and that $D \sum_{j \in V_i^h : s_i \neq s_j} w''_j = \sum_{j \in V_i^h : s_i \neq s_j} w'_j \leq \sum_{j \in V_i : s_i \neq s_j} w'_j$, we get

$$\sum_{j \in V_i : s_i = s_j} w'_j \leq (1 + \varepsilon)^2 \sum_{j \in V_i : s_i \neq s_j} w'_j$$

as needed. □

Remark 7.B.3. *We observe the following trade off: we can get a $(1 + \varepsilon)^2$ -approximate equilibrium if we skip the rounding step at the beginning of the algorithm but then the number of different weights D_ε and thus the running time of the algorithm may increase. Also, if Δ is the maximum degree among the vertices of G , then replacing $\frac{n}{\varepsilon}$ with $\frac{\Delta}{\varepsilon}$ in the algorithm and following a similar analysis gives $O\left(\frac{m}{\varepsilon} \lceil \frac{\Delta}{\varepsilon} \rceil^{2D_\varepsilon}\right)$ running time.*

Bibliography

- [1] Ittai Abraham, Yair Bartal, and Ofer Neiman. Advances in metric embedding theory. *Advances in Mathematics*, 228(6):3026–3126, 2011.
- [2] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6):25:1–25:22, 2008.
- [3] P.K. Agarwal, H.-C. Chang, K. Munagala, E. Taylor, and E. Welzl. Clustering under perturbation stability in near-linear time. In *Proc. of the 40th IARCS Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2020)*, volume 182 of *LIPICs*, pages 8:1–8:16, 2020.
- [4] Adrian K. Agogino and Kagan Tumer. Handling communication restrictions and team formation in congestion games. *Auton. Agents Multi Agent Syst.*, 13(1):97–115, 2006.
- [5] N. Alon, M. Feldman, A.D. Procaccia, and M. Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2010.
- [6] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. Peeking behind the ordinal curtain: Improving distortion via cardinal queries. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 1782–1789. AAAI Press, 2020.
- [7] Ioannis Anagnostides, Dimitris Fotakis, and Panagiotis Patsilinakos. Asymptotically optimal communication in simple mechanisms. In *SAGT*, volume 12283 of *Lecture Notes in Computer Science*, pages 17–31. Springer, 2020.
- [8] Ioannis Anagnostides, Dimitris Fotakis, and Panagiotis Patsilinakos. Dimensionality and coordination in voting: The distortion of STV. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI, 2022*, pages 4776–4784. AAAI Press, 2022.

- [9] Ioannis Anagnostides, Dimitris Fotakis, and Panagiotis Patsilinakos. Metric-distortion bounds under limited information. *J. Artif. Intell. Res.*, 74:1449–1483, 2022. doi: 10.1613/jair.1.13338.
- [10] Ioannis Anagnostides, Dimitris Fotakis, and Panagiotis Patsilinakos. Dimensionality and coordination in voting: The distortion of STV. In *AAAI*, pages 4776–4784. AAAI Press, 2022.
- [11] Ioannis Anagnostides, Dimitris Fotakis, and Panagiotis Patsilinakos. Metric-distortion bounds under limited information. *J. Artif. Intell. Res.*, 74:1449–1483, 2022.
- [12] Omer Angel, Sébastien Bubeck, Yuval Peres, and Fan Wei. Local Max-Cut in Smoothed Polynomial Time. In *Proc. of the 49th ACM SIGACT Symposium on Theory of Computing (STOC 2017)*, pages 429–437, 2017.
- [13] H. Angelidakis, K. Makarychev, and Y. Makarychev. Algorithms for stable and perturbation-resilient problems. In *Proc. of the 49th ACM Symposium on Theory of Computing (STOC 2017)*, pages 438–451, 2017.
- [14] Elliot Anshelevich and John Postl. Randomized social choice functions under metric preferences. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*, pages 46–59. IJCAI/AAAI Press, 2016.
- [15] Elliot Anshelevich and Shreyas Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 390–396. AAAI Press, 2016.
- [16] Elliot Anshelevich and Shreyas Sekar. Truthful mechanisms for matching and clustering in an ordinal world. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics - 12th International Conference, WINE 2016*, volume 10123 of *Lecture Notes in Computer Science*, pages 265–278. Springer, 2016. doi: 10.1007/978-3-662-54110-4_19.
- [17] Elliot Anshelevich and Wennan Zhu. Ordinal approximation for social choice, matching, and facility location problems given candidate positions. In George Christodoulou and Tobias Harks, editors, *Web and Internet Economics - 14th International Conference, WINE 2018*, volume 11316 of *Lecture Notes in Computer Science*, pages 3–20. Springer, 2018. doi: 10.1007/978-3-030-04612-5_1.

- [18] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. Approximating optimal social choice under metric preferences. In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, 2015*, pages 777–783. AAAI Press, 2015.
- [19] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artif. Intell.*, 264:27–51, 2018.
- [20] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artif. Intell.*, 264:27–51, 2018.
- [21] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. Distortion in social choice problems: The first 15 years and beyond. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021*, pages 4294–4301. ijcai.org, 2021.
- [22] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. Distortion in social choice problems: The first 15 years and beyond. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021*, pages 4294–4301. ijcai.org, 2021. doi: 10.24963/ijcai.2021/589.
- [23] A. Archer and R. Kleinberg. Truthful germs are contagious: A local-to-global characterization of truthfulness. In *Proc. of the 9th ACM Conference on Electronic Commerce (EC '08)*, pages 21–30, 2008.
- [24] Kenneth Arrow. *Advances in the Spatial Theory of Voting*. Cambridge University Press, 1990.
- [25] David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)*, pages 1027–1035. SIAM, 2007.
- [26] Patrice Assouad. Plongements lipschitziens dans \setminus^n . *Bulletin de la Société Mathématique de France*, 111:429–448, 1983.
- [27] V. Auletta, R. De Prisco, P. Penna, and G. Persiano. The power of verification for one-parameter agents. *Journal of Computer and System Sciences*, 75:190–211, 2009.
- [28] Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, December 2004.

- [29] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1*. MIT Press, 2006.
- [30] Marco Avella-Medina and Victor-Emmanuel Brunel. Differentially private sub-gaussian location estimators, 2019.
- [31] P. Awasthi, A. Blum, and O. Sheffet. Center-based clustering under perturbation stability. *Inf. Process. Lett.*, 112(1-2):49–54, 2012.
- [32] M. Babaioff, S. Dobzinski, and S. Oren. Combinatorial auctions with endowment effect. In *Proc. of the 2018 ACM Conference on Economics and Computation (EC 2018)*, pages 73–90, 2018.
- [33] M.-F. Balcan, N. Haghtalab, and C. White. k -Center Clustering Under Perturbation Resilience. In *Proc. of the 43rd International Colloquium on Automata, Languages and Programming (ICALP 2016)*, volume 55 of *LIPICs*, pages 68:1–68:14, 2016.
- [34] Maria-Florina Balcan and Yingyu Liang. Clustering under perturbation resilience. *SIAM Journal on Computing*, 45(1):102–155, 2016.
- [35] Maria-Florina Balcan, Avrim Blum, and Anupam Gupta. Clustering under approximation stability. *Journal of the ACM*, 60(2), 2013.
- [36] Gleb Beliakov, Tomasa Calvo, and Simon James. *Aggregation of Preferences in Recommender Systems*, pages 705–734. Springer US, 2011. doi: 10.1007/978-0-387-85820-3_22.
- [37] Gerdus Benadè, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Preference elicitation for participatory budgeting. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 376–382. AAAI Press, 2017.
- [38] Gerdus Benadè, Ariel D. Procaccia, and Mingda Qiao. Low-distortion social welfare functions. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, pages 1788–1795. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33011788.
- [39] Salem Benferhat, Sylvain Lagrue, and Odile Papini. Reasoning with partially ordered information in a possibilistic logic framework. *Fuzzy Sets and Systems*, 144(1):25–41, 2004. doi: <https://doi.org/10.1016/j.fss.2003.10.012>. Possibilistic Logic and Related Issues.
- [40] Matthias Bentert and Piotr Skowron. Comparing election methods where

- each voter ranks only few candidates. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 2218–2225. AAAI Press, 2020.
- [41] Anand Bhalgat, Tanmoy Chakraborty, and Sanjeev Khanna. Approximating pure Nash equilibrium in cut, party affiliation, and satisfiability games. In *Proc. of the 11th ACM Conference on Electronic Commerce (EC 2010)*, pages 73–82, 2010.
- [42] Umang Bhaskar, Varsha Dani, and Abheek Ghosh. Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 925–932. AAAI Press, 2018.
- [43] Y. Bilu and N. Linial. Are Stable Instances Easy? In *Proc. of the 1st Symposium on Innovations in Computer Science (ICS 2010)*, pages 332–341. Tsinghua University Press, 2010.
- [44] Yonatan Bilu, Amit Daniely, Nati Linial, and Michael E. Saks. On the practically interesting instances of MAXCUT. In Natacha Portier and Thomas Wilke, editors, *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013)*, volume 20 of *LIPIcs*, pages 526–537. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013.
- [45] Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948. ISSN 00223808, 1537534X.
- [46] Liad Blumrosen and Michal Feldman. Implementation with a bounded action space. In Joan Feigenbaum, John C.-I. Chuang, and David M. Pennock, editors, *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006)*, pages 62–71. ACM, 2006.
- [47] Liad Blumrosen, Noam Nisan, and Ilya Segal. Auctions with severely bounded communication. *J. Artif. Intell. Res.*, 28:233–266, 2007.
- [48] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Primarily about primaries. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, pages 1804–1811. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33011804.
- [49] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Primarily about primaries. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, pages 1804–1811. AAAI Press, 2019.

- [50] Allan Borodin, Daniel Halpern, Mohamad Latifian, and Nisarg Shah. Distortion in voting with top-t preferences. In *IJCAI '22*, 2022.
- [51] Allan Borodin, Daniel Halpern, Mohamad Latifian, and Nisarg Shah. Distortion in voting with top-t preferences. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022*, pages 116–122. ijcai.org, 2022.
- [52] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artif. Intell.*, 227:190–213, 2015. doi: 10.1016/j.artint.2015.06.003.
- [53] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, USA, 1st edition, 2016.
- [54] Simina Brânzei and Ariel D. Procaccia. Verifiably truthful mechanisms. In Tim Roughgarden, editor, *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 297–306. ACM, 2015.
- [55] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In Marie desJardins and Michael L. Littman, editors, *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*. AAAI Press, 2013.
- [56] Markus Brill and Vincent Conitzer. Strategic voting and strategic candidacy. In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 819–826. AAAI Press, 2015.
- [57] Victor-Emmanuel Brunel and Marco Avella-Medina. Propose, test, release: Differentially private estimation with high probability, 2020.
- [58] Ran Canetti, Guy Even, and Oded Goldreich. Lower bounds for sampling algorithms for estimating the average. *Information Processing Letters*, 53(1):17 – 25, 1995.
- [59] Ran Canetti, Guy Even, and Oded Goldreich. Lower bounds for sampling algorithms for estimating the average. *Inf. Process. Lett.*, 53(1):17–25, 1995. doi: 10.1016/0020-0190(94)00171-T.
- [60] Clément L. Canonne. A short note on learning discrete distributions, 2020.

- [61] I. Caragiannis, E. Elkind, M. Szegedy, and L. Yu. Mechanism design: from partial to probabilistic verification. In *Proc. of the 13th ACM Conference on Electronic Commerce (EC '12)*, pages 266–283, 2012.
- [62] Ioannis Caragiannis and Angelo Fanelli. On Approximate Pure Nash Equilibria in Weighted Congestion Games with Polynomial Latencies. In *Proc. of the 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132, pages 133:1–133:12, 2019.
- [63] Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. Efficient Computation of Approximate Pure Nash Equilibria in Congestion Games. In *Proc. of the IEEE 52nd Symposium on Foundations of Computer Science, (FOCS 2011)*, pages 532–541, 2011.
- [64] Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. Approximate Pure Nash Equilibria in Weighted Congestion Games: Existence, Efficient Computation, and Structure. *ACM Transactions on Economics and Computation*, 3(1):2:1–2:32, 2015.
- [65] Ioannis Caragiannis, Aris Filos-Ratsikas, Søren Kristoffer Stiel Frederiksen, Kristoffer Arnsfelt Hansen, and Zihan Tan. Truthful facility assignment with resource augmentation: An exact analysis of serial dictatorship. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics - 12th International Conference, WINE 2016*, volume 10123 of *Lecture Notes in Computer Science*, pages 236–250. Springer, 2016. doi: 10.1007/978-3-662-54110-4_17.
- [66] Ioannis Caragiannis, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. Subset selection via implicit utilitarian voting. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 151–157. IJCAI/AAAI Press, 2016.
- [67] Ioannis Caragiannis, Aris Filos-Ratsikas, Swaprava Nath, and Alexandros A. Voudouris. Truthful mechanisms for ownership transfer with expert advice. *CoRR*, abs/1802.01308, 2018.
- [68] G. Carroll. When are local incentive constraints sufficient? *Econometrica*, 80(2):661–686, 2012.
- [69] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [70] Shuwei Chen, Jun Liu, Hui Wang, and Juan Carlos Augusto. Ordering based decision making – a survey. *Information Fusion*, 14(4):521–531, 2013. doi:

<https://doi.org/10.1016/j.inffus.2012.10.005>.

- [71] Xi Chen, Chenghao Guo, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Michalis Yannakakis, and Xinzhi Zhang. Smoothed complexity of local max-cut and binary max-csp. *CoRR*, abs/1911.10381, 2019. URL <http://arxiv.org/abs/1911.10381>.
- [72] Zhihuai Chen, Ken C.K. Fong, Minming Li, Kai Wang, Hongning Yuan, and Yong Zhang. Facility location games with optional preference. *Theoretical Computer Science*, 847:185–197, 2020.
- [73] Steve Chien and Alistair Sinclair. Convergence to approximate Nash equilibria in congestion games. *Games and Economic Behavior*, 71(2):315–327, 2011.
- [74] Sungdai Cho and James W. Endersby. Issues, the spatial theory of voting, and british general elections: A comparison of proximity and directional models. *Public Choice*, 114(3/4):275–293, 2003.
- [75] Kenneth L. Clarkson. Nearest neighbor queries in metric spaces. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, 1997*, pages 609–617. ACM, 1997.
- [76] Vincent Conitzer. Eliciting single-peaked preferences using comparison queries. *J. Artif. Int. Res.*, 35(1):161–191, 2009.
- [77] 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*, pages 392–397. AAAI Press / The MIT Press, 2002.
- [78] Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce, EC '05*, page 78–87, New York, NY, USA, 2005. Association for Computing Machinery.
- [79] Vincent Conitzer, Matthew Rognlie, and Lirong Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI'09*, page 109–115. Morgan Kaufmann Publishers Inc., 2009.
- [80] Esther David, Alex Rogers, Jeremy Schiff, Sarit Kraus, and N. R. Jennings. Optimal design of english auctions with discrete bid levels. In *ACM Confer-*

- ence on Electronic Commerce (EC'05), pages 98–107, 2005.
- [81] Herbert A. David. *Order Statistics*, pages 1039–1040. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [82] 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*, page 1421–1430, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.
- [83] Swapnil Dhamal and Y. Narahari. Scalable preference aggregation in social networks. In *HCOMP*, 2013.
- [84] Ning Ding and Fangzhen Lin. Voting with partial information: What questions to ask? In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13*, page 1237–1238, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems.
- [85] E. Dokow, M. Feldman, R. Meir, and I. Nehama. Mechanism design on discrete lines and cycles. In *Proc. of the 13th ACM Conference on Electronic Commerce (EC '12)*, pages 423–440, 2012.
- [86] P. Düetting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs. In *Proc. of the 58th Symposium on Foundations of Computer Science (FOCS 2017)*, pages 540–551, 2017.
- [87] Cynthia Dwork. Differential privacy: A survey of results. In Manindra Agrawal, Ding-Zhu Du, Zhenhua Duan, and Angsheng Li, editors, *Theory and Applications of Models of Computation, 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings*, volume 4978 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2008.
- [88] Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. What do multiwinner voting rules do? an experiment over the two-dimensional euclidean domain. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017*, pages 494–501. AAAI Press, 2017.
- [89] Robert Elsässer and Tobias Tscheuschner. Settling the Complexity of Local Max-Cut (Almost) Completely. In *Proc. of the 38th International Col-*

- loquium on Automata, Languages and Programming (ICALP 2011)*, pages 171–182, 2011.
- [90] James M. Enelow and Melvin J. Hinich. *The Spatial Theory of Voting*. Cambridge Books. Cambridge University Press, 1984.
- [91] B. Escoffier, L. Gourvès, N.K. Thang, F. Pascual, and O. Spanjaard. Strategy-proof mechanisms for Facility Location games with many facilities. In *Proc. of the 2nd International Conference on Algorithmic Decision Theory (ADT '11)*, volume 6992 of *LNAI*, pages 67–81, 2011.
- [92] Michael Etscheid and Heiko Röglin. Smoothed analysis of local search for the maximum-cut problem. *ACM Transactions on Algorithms*, 13(2):25:1–25:12, 2017.
- [93] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence Time to Nash Equilibria. In *Proc. of the 30th International Colloquium on Automata, Languages and Programming (ICALP '03)*, volume 2719 of *LNCS*, pages 502–513. Springer-Verlag, 2003.
- [94] T. Ezra, M. Feldman, and O. Friedler. A general framework for endowment effects in combinatorial markets. In *Proc. of the 2020 ACM Conference on Economics and Computation (EC 2020)*, 2020.
- [95] A. Fabrikant, C. Papadimitriou, and K. Talwar. The Complexity of Pure Nash Equilibria. In *Proc. of the 36th ACM Symposium on Theory of Computing (STOC '04)*, pages 604–612, 2004.
- [96] Ronald Fadel and Ilya Segal. The communication cost of selfishness. *Journal of Economic Theory*, 144(5):1895 – 1920, 2009. ISSN 0022-0531.
- [97] Brandon Fain, Ashish Goel, Kamesh Munagala, and Sukolsak Sakshuwong. Sequential deliberation for social choice. In *Web and Internet Economics - 13th International Conference, WINE 2017*, volume 10660 of *Lecture Notes in Computer Science*, pages 177–190. Springer, 2017. doi: 10.1007/978-3-319-71924-5_13.
- [98] Brandon Fain, Ashish Goel, Kamesh Munagala, and Nina Prabhu. Random dictators with a random referee: Constant sample complexity mechanisms for social choice. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019*, pages 1893–1900. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33011893.

- [99] Angelo Fanelli and Luca Moscardelli. On best response dynamics in weighted congestion games with polynomial delays. In *Proc. of the 5th Workshop on Internet and Network Economics (WINE 2009)*, pages 55–66, 2009.
- [100] Uriel Feige. Approximating the bandwidth via volume respecting embeddings. *Journal of Computer and System Sciences*, 60(3):510–539, 2000.
- [101] Itai Feigenbaum, Minming Li, Jay Sethuraman, Fangzhou Wang, and Shaokun Zou. Strategic facility location problems with linear single-dipped and single-peaked preferences. *Autonomous Agents and Multi-Agent Systems*, 34(2):49, 2020.
- [102] M. Feldman and Y. Wilf. Randomized strategyproof mechanisms for Facility Location and the mini-sum-of-squares objective. *CoRR abs*, 1108.1762, 2011.
- [103] M. Feldman, N. Gravin, and B. Lucier. Combinatorial Auctions via Posted Prices. In *Proc. of the 26th ACM-SIAM Symposium on Discrete Algorithms*, pages 123–135, 2014.
- [104] 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*, pages 269–286. ACM, 2016. doi: 10.1145/2940716.2940725.
- [105] Michal Feldman, Amos Fiat, and Iddan Golomb. On voting and facility location. In Vincent Conitzer, Dirk Bergemann, and Yiling Chen, editors, *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16*, pages 269–286. ACM, 2016.
- [106] Matthias Feldotto, Martin Gairing, Grammateia Kotsialou, and Alexander Skopalik. Computing approximate pure nash equilibria in shapley value weighted congestion games. In *Proc. of the 13th International Conference on Web and Internet Economics (WINE 2017)*, pages 191–204, 2017.
- [107] Giannis Fikioris and Dimitris Fotakis. Mechanism design for perturbation stable combinatorial auctions. In *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT 2020)*, volume 12283 of *Lecture Notes in Computer Science*, pages 47–63. Springer, 2020.
- [108] Alina Filimonov and Reshef Meir. Strategyproof facility location mechanisms on discrete trees. *CoRR*, abs/2102.02610, 2021. URL <https://arxiv.org/abs/2102.02610>.
- [109] Aris Filos-Ratsikas and Alexandros A. Voudouris. Approximate mecha-

- nism design for distributed facility location. In *Algorithmic Game Theory - 14th International Symposium, SAGT 2021*, volume 12885 of *Lecture Notes in Computer Science*, pages 49–63. Springer, 2021. doi: 10.1007/978-3-030-85947-3_4.
- [110] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In Ron Lavi, editor, *Algorithmic Game Theory - 7th International Symposium, SAGT 2014*, volume 8768 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2014. doi: 10.1007/978-3-662-44803-8_1.
- [111] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. *Theoretical Computer Science*, 472:90–103, 2013.
- [112] D. Fotakis and C. Tzamos. On the power of deterministic mechanisms for facility location games. *ACM Transactions on Economics and Computation*, 2(4):15:1–15:37, 2014.
- [113] Dimitris Fotakis. A selective tour through congestion games. In *Algorithms, Probability, Networks, and Games: Scientific Papers and Essays Dedicated to Paul G. Spirakis on the Occasion of His 60th Birthday*, pages 223–241, 2015.
- [114] Dimitris Fotakis and Panagiotis Patsilinakos. Strategyproof facility location in perturbation stable instances. In *WINE*, volume 13112 of *Lecture Notes in Computer Science*, pages 95–112. Springer, 2021.
- [115] Dimitris Fotakis and Christos Tzamos. On the power of deterministic mechanisms for facility location games. *ACM Trans. Econ. Comput.*, 2(4), October 2014.
- [116] Dimitris Fotakis and Christos Tzamos. Strategyproof facility location for concave cost functions. *Algorithmica*, 76(1):143–167, 2016.
- [117] Dimitris Fotakis and Emmanouil Zampetakis. Truthfulness flooded domains and the power of verification for mechanism design. *ACM Transactions on Economics and Computation*, 3(4):20:1–20:29, 2015.
- [118] Dimitris Fotakis, Spyros C. Kontogiannis, and Paul G. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2-3):226–239, 2005.
- [119] Dimitris Fotakis, Spyros C. Kontogiannis, and Paul G. Spirakis. Symmetry in Network Congestion Games: Pure Equilibria and Anarchy Cost. In *Proc.*

of the 3rd Workshop on Approximation and Online Algorithms, Revised Papers (WAOA 2005), pages 161–175, 2005.

- [120] Dimitris Fotakis, Spyros C. Koutogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul G. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. *Theoretical Computer Science*, 410(36): 3305–3326, 2009.
- [121] Dimitris Fotakis, Christos Tzamos, and Manolis Zampetakis. Mechanism design with selective verification. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC 2016)*, pages 771–788. ACM, 2016.
- [122] Dimitris Fotakis, Anthimos Vardis Kandiros, Thanasis Lianeas, Nikos Mouzakis, Panagiotis Patsilinakos, and Stratis Skoulakis. Node-max-cut and the complexity of equilibrium in linear weighted congestion games. In *ICALP*, volume 168 of *LIPICs*, pages 50:1–50:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [123] Dimitris Fotakis, Alkis Kalavasis, and Konstantinos Stavropoulos. Aggregating incomplete and noisy rankings. In Arindam Banerjee and Kenji Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 2278–2286. PMLR, 2021.
- [124] Johannes Fürnkranz and Eyke Hüllermeier. *Preference Learning*. Springer-Verlag, Berlin, Heidelberg, 1st edition, 2010.
- [125] Harold N Gabow and Robert E Tarjan. Algorithms for two bottleneck optimization problems. *Journal of Algorithms*, 9(3):411–417, 1988. doi: [https://doi.org/10.1016/0196-6774\(88\)90031-4](https://doi.org/10.1016/0196-6774(88)90031-4).
- [126] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. In *Proc. of the 36th ACM Symposium on Theory of Computing (STOC '04)*, pages 613–622, 2004.
- [127] Martin Gairing and Rahul Savani. Computing Stable Outcomes in Hedonic Games. In *Proc. of the 3rd Symposium on Algorithmic Game Theory (SAGT 2010)*, pages 174–185, 2010.
- [128] Alex Gershkov, Benny Moldovanu, and Xianwen Shi. Voting on multiple issues: What to put on the ballot? *Theoretical Economics*, 14(2):555–596, 2019.

- [129] Yiannis Giannakopoulos, Georgy Noarov, and Andreas S. Schulz. An improved algorithm for computing approximate equilibria in weighted congestion games. *CoRR*, abs/1810.12806, 2018.
- [130] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- [131] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. Resolving the optimal metric distortion conjecture. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, pages 1427–1438. IEEE, 2020. doi: 10.1109/FOCS46700.2020.00134.
- [132] Ashish Goel, Anilesh K. Krishnaswamy, and Kamesh Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17*, page 287–304. Association for Computing Machinery, 2017. ISBN 9781450345279. doi: 10.1145/3033274.3085138.
- [133] Sumit Goel and Wade Hann-Caruthers. Coordinate-wise median: Not bad, not bad, pretty good. *CoRR*, abs/2007.00903, 2020. URL <https://arxiv.org/abs/2007.00903>.
- [134] Paul W. Goldberg. Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game. In *Proc. of the 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC 2004)*, pages 131–140, 2004.
- [135] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.
- [136] Iddan Golomb and Christos Tzamos. Truthful facility location with additive errors. *CoRR*, abs/1701.00529, 2017. URL <http://arxiv.org/abs/1701.00529>.
- [137] J. Green and J. Laffont. Partially verifiable information and mechanism design. *Review of Economic Studies*, 53(3):447–456, 1986.
- [138] Stephen Gross, Elliot Anshelevich, and Lirong Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, 2017*, pages 544–550. AAAI Press, 2017.
- [139] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded ge-

- ometries, fractals, and low-distortion embeddings. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003*, pages 534–543. IEEE Computer Society, 2003.
- [140] Tobias Harks and Max Klimm. On the Existence of Pure Nash Equilibria in Weighted Congestion Games. In *Proc. of the 37th International Colloquium on Automata, Languages and Programming (ICALP 2010)*, pages 79–89, 2010.
- [141] Tobias Harks, Max Klimm, and Rolf H. Möhring. Characterizing the existence of potential functions in weighted congestion games. *Theory Computing Systems*, 49(1):46–70, 2011.
- [142] Rajesh Jayaram and David P. Woodruff. Perfect lp sampling in a data stream. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018*, pages 544–555. IEEE Computer Society, 2018.
- [143] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.
- [144] John H. Kagel and Dan Levin. Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders. *The Economic Journal*, 103(419):868–879, 1993. ISSN 0013-0133.
- [145] John H. Kagel, Ronald M. Harstad, and Dan Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55(6):1275–1304, 1987. ISSN 00129682, 14680262.
- [146] David R. Karger and Matthias Ruhl. Finding nearest neighbors in growth-restricted metrics. In John H. Reif, editor, *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002*, pages 741–750. ACM, 2002.
- [147] Narendra Karmarkar. A new polynomial-time algorithm for linear programming. *Comb.*, 4(4):373–396, 1984. doi: 10.1007/BF02579150.
- [148] Richard M. Karp. Reducibility among combinatorial problems. In *Proc. of Symposium on the Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103, 1972.
- [149] David Kempe. An analysis framework for metric voting based on LP dual-

- ity. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 2079–2086. AAAI Press, 2020.
- [150] David Kempe. Communication, distortion, and randomness in metric voting. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 2087–2094. AAAI Press, 2020.
- [151] David Kempe. Communication, distortion, and randomness in metric voting. 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, New York, NY, USA, February 7-12, 2020*, pages 2087–2094. AAAI Press, 2020.
- [152] Fatih Erdem Kizilkaya and David Kempe. Plurality veto: A simple voting rule achieving optimal metric distortion. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 349–355. ijcai.org, 2022.
- [153] Pieter Kleer and Guido Schäfer. Potential Function Minimizers of Combinatorial Congestion Games: Efficiency and Computation. In *Proc. of the 2017 ACM Conference on Economics and Computation (EC 2017)*, pages 223–240, 2017.
- [154] Leonid (Aryeh) Kontorovich and Kavita Ramanan. Concentration inequalities for dependent random variables via the martingale method. *The Annals of Probability*, 36(6):2126 – 2158, 2008. doi: 10.1214/07-AOP384.
- [155] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In Christoph Meinel and Sophie Tison, editors, *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer, 1999.
- [156] Maria Kyropoulou, Carmine Ventre, and Xiaomeng Zhang. Mechanism design for constrained heterogeneous facility location. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT 2019)*, volume 11801 of *Lecture Notes in Computer Science*, pages 63–76. Springer, 2019.
- [157] Jérôme Lang, Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Winner determination in sequential majority voting. In Manuela M. Veloso, editor, *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1372–1377, 2007.

- [158] Valentino Larcinese. The Instrumental Voter Goes to the News-Agent: Demand for Information, Election Closeness, and the Media. UFAE and IAE Working Papers 579.03, Unitat de Fonaments de l'Anàlisi Econòmica (UAB) and Institut d'Anàlisi Econòmica (CSIC), March 2003.
- [159] D. G. Larman. A New Theory of Dimension. *Proceedings of the London Mathematical Society*, s3-17(1):178–192, 01 1967.
- [160] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces: isoperimetry and processes*. Ergebnisse der Mathematik und ihrer Grenzgebiete A Series of Modern Surveys in Mathematics. Springer, Berlin, 1991.
- [161] Gaëlle Legrand and Nicolas Nicoloyannis. Feature selection method using preferences aggregation. In Petra Perner and Atsushi Imiya, editors, *Machine Learning and Data Mining in Pattern Recognition*, pages 203–217, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [162] Jonathan Levin and Barry Nalebuff. An introduction to vote-counting schemes. *Journal of Economic Perspectives*, 9(1):3–26, March 1995. doi: 10.1257/jep.9.1.3.
- [163] Minming Li, Pinyan Lu, Yuhao Yao, and Jialin Zhang. Strategyproof mechanism for two heterogeneous facilities with constant approximation ratio. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*, pages 238–245, 2020.
- [164] Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, November 2017.
- [165] Arend Lijphart. Democratization and constitutional choices in czechoslovakia, hungary and poland 1989-91. *Journal of Theoretical Politics*, 4(2):207–223, 1992.
- [166] P. Lu, Y. Wang, and Y. Zhou. Tighter bounds for facility games. In *Proc. of the 5th Workshop on Internet and Network Economics (WINE '09)*, volume 5929 of *LNCS*, pages 137–148, 2009.
- [167] P. Lu, X. Sun, Y. Wang, and Z.A. Zhu. Asymptotically Optimal Strategy-Proof Mechanisms for Two-Facility Games. In *Proc. of the 11th ACM Conference on Electronic Commerce (EC '10)*, pages 315–324, 2010.
- [168] Tyler Lu and Craig Boutilier. Vote elicitation with probabilistic preference models: Empirical estimation and cost tradeoffs. In Ronen I. Brafman, Fred S. Roberts, and Alexis Tsoukiàs, editors, *Algorithmic Decision Theory*,

- pages 135–149, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [169] Tyler Lu and Craig Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One*, IJCAI'11, page 287–293. AAAI Press, 2011.
- [170] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., 1996.
- [171] Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, and David P. Woodruff. Efficient and thrifty voting by any means necessary. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019*, pages 7178–7189, 2019.
- [172] Debmalya Mandal, Nisarg Shah, and David P. Woodruff. Optimal communication-distortion tradeoff in voting. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 795–813. Association for Computing Machinery, 2020. doi: 10.1145/3391403.3399510.
- [173] Jiří Matoušek. *Embedding Finite Metric Spaces into Normed Spaces*, pages 355–400. Springer New York, New York, NY, 2002.
- [174] Lili Mei, Minming Li, Deshi Ye, and Guochuan Zhang. Facility location games with distinct desires. *Discrete Applied Mathematics*, 264:148–160, 2019.
- [175] Reshef Meir. Strategyproof facility location for three agents on a circle. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT 2019)*, volume 11801 of *Lecture Notes in Computer Science*, pages 18–33. Springer, 2019.
- [176] Wil Michiels, Emile Aarts, and Jan Korst. *Theoretical Aspects of Local Search*. EATCS Monographs in Theoretical Computer Science. Springer, 2007.
- [177] E. Miyagawa. Locating libraries on a street. *Social Choice and Welfare*, 18: 527–541, 2001.
- [178] D. Monderer and L. Shapley. Potential Games. *Games and Economic Behavior*, 14:124–143, 1996.

- [179] Morteza Monemizadeh and David P. Woodruff. 1-pass relative-error l_p -sampling with applications. In Moses Charikar, editor, *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010*, pages 1143–1160. SIAM, 2010.
- [180] Dilip Mookherjee and Masatoshi Tsumagari. Mechanism design with communication constraints. *Journal of Political Economy*, 122(5):1094–1129, 2014. ISSN 00223808, 1537534X.
- [181] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980. ISSN 00485829, 15737101.
- [182] H. Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.
- [183] 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*, pages 245–262. ACM, 2019. doi: 10.1145/3328526.3329550.
- [184] J.F. Nash. Equilibrium points in n -person games. *Proc. of the National Academy of Sciences*, 36:48–49, 1950.
- [185] J.F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
- [186] Kobbi Nissim, Rann Smorodinsky, and Moshe Tennenholtz. Approximately optimal mechanism design via differential privacy. In *Innovations in Theoretical Computer Science (ITCS 2012)*, pages 203–213. ACM, 2012.
- [187] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 993–999. AAAI Press, 2015.
- [188] 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*, page 309–316. AAAI Press, 2013. ISBN 9781577356332.
- [189] Panagiota N. Panagopoulou and Paul G. Spirakis. Algorithms for pure Nash equilibria in weighted congestion games. *ACM Journal of Experimental Algorithmics*, 11, 2006.

- [190] David C. Parkes, Lyle H. Ungar, and Dean P. Foster. Accounting for cognitive costs in on-line auction design. In Pablo Noriega and Carles Sierra, editors, *Agent Mediated Electronic Commerce*, pages 25–40, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg. ISBN 978-3-540-48835-4.
- [191] Svatopluk Poljak. Integer Linear Programs and Local Search for Max-Cut. *SIAM Journal on Computing*, 21(3):450–465, 1995.
- [192] A.D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proc. of the 10th ACM Conference on Electronic Commerce (EC '09)*, pages 177–186, 2009.
- [193] Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In Matthias Klusch, Michael Rovatsos, and Terry R. Payne, editors, *Cooperative Information Agents X*, pages 317–331, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [194] Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In Matthias Klusch, Michael Rovatsos, and Terry R. Payne, editors, *Cooperative Information Agents X, 10th International Workshop, CIA 2006*, volume 4149 of *Lecture Notes in Computer Science*, pages 317–331. Springer, 2006. doi: 10.1007/11839354_23.
- [195] Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. *ACM Trans. Econ. Comput.*, 1(4), December 2013. ISSN 2167-8375.
- [196] Ariel D. Procaccia, Aviv Zohar, Yoni Peleg, and Jeffrey S. Rosenschein. The learnability of voting rules. *Artif. Intell.*, 173(12-13):1133–1149, 2009. doi: 10.1016/j.artint.2009.03.003.
- [197] R. A. Rankin. George bachman, introduction to p-adic numbers and valuation theory (academic press, new york, 1964), 173 pp., 24s. 6d. *Proceedings of the Edinburgh Mathematical Society*, 15(1):83–84, 1966. doi: 10.1017/S0013091500013237.
- [198] Stefan Reichelstein. Incentive compatibility and informational requirements. *Journal of Economic Theory*, 34(1):32 – 51, 1984. ISSN 0022-0531.
- [199] K.B. Reid. Every vertex a king. *Discrete Mathematics*, 38(1):93–98, 1982. doi: [https://doi.org/10.1016/0012-365X\(82\)90173-X](https://doi.org/10.1016/0012-365X(82)90173-X).
- [200] Juan M. Restrepo, Rosalyn C. Rael, and James M. Hyman. Modeling the influence of polls on elections: A population dynamics approach. *Public*

- Choice*, 140(3/4):395–420, 2009.
- [201] R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [202] T. Roughgarden. Beyond Worst-Case Analysis. *Communications of the ACM*, 62(3):88–96, 2019.
- [203] T. Roughgarden. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2020.
- [204] T. Roughgarden and I. Talgam-Cohen. Approximately Optimal Mechanism Design. *CoRR*, 2018. URL <http://arxiv.org/abs/1812.11896>.
- [205] Tim Roughgarden. Lecture 6: Perturbation-stable clustering. *CS264: Beyond Worst-Case Analysis*, 2017. URL <http://timroughgarden.org/w17/1/16.pdf>.
- [206] Donald G. Saari and Vincent R. Merlin. The copeland method: I.: Relationships and the dictionary. *Economic Theory*, 8(1):51–76, 1996.
- [207] Mark Allen Satterthwaite. Strategy-proofness and arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- [208] Alejandro A. Schäffer and Mihalis Yannakakis. Simple Local Search Problems That are Hard to Solve. *SIAM Journal on Computing*, 20(1):56–87, 1991.
- [209] J. Schummer and R.V. Vohra. Mechanism design without money. *Algorithmic Game Theory*, 10:243–299, 2007.
- [210] James Schummer and Rakesh V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405 – 428, 2002.
- [211] Ilya Segal. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory*, 136(1):341 – 378, 2007. ISSN 0022-0531.
- [212] Paolo Serafino and Carmine Ventre. Heterogeneous facility location without money. *Theoretical Computer Science*, 636:27–46, 2016.
- [213] A. Skopalik and B. Vöcking. Inapproximability of Pure Nash Equilibria. In *Proc. of the 40th ACM Symposium on Theory of Computing (STOC ’08)*, pages 355–364, 2008.

- [214] Piotr Krzysztof Skowron and Edith Elkind. Social choice under metric preferences: Scoring rules and STV. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 706–712. AAAI Press, 2017.
- [215] H. Smith, Tamara Dinev, and Heng Xu. Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35:989–1015, 2011.
- [216] H. Jeff Smith, Tamara Dinev, and Heng Xu. Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4):989–1015, 2011.
- [217] Xin Sui, Craig Boutilier, and Tuomas Sandholm. Analysis and optimization of multi-dimensional percentile mechanisms. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, pages 367–374. IJCAI/AAAI, 2013.
- [218] J. Valdez, R.E. Tarjan, and E.L. Lawler. The Recognition of Series-Parallel Digraphs. *SIAM J. Comput.*, 11(2):298–313, 1982.
- [219] Timothy Van Zandt. Communication complexity and mechanism design. *Journal of the European Economic Association*, 5(2/3):543–553, 2007. ISSN 15424766, 15424774.
- [220] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961. ISSN 00221082, 15406261.
- [221] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- [222] Maksims Volkovs and Richard S. Zemel. New learning methods for supervised and unsupervised preference aggregation. *J. Mach. Learn. Res.*, 15(1): 1135–1176, 2014.
- [223] Maksims Volkovs, Hugo Larochelle, and Richard S. Zemel. Learning to rank by aggregating expert preferences. In Xue-wen Chen, Guy Lebanon, Haixun Wang, and Mohammed J. Zaki, editors, *21st ACM International Conference on Information and Knowledge Management, CIKM’12, 2012*, pages 843–851. ACM, 2012. doi: 10.1145/2396761.2396868.
- [224] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- [225] Michael Zuckerman, Piotr Faliszewski, Vincent Conitzer, and Jeffrey S.

Rosenschein. An NTU cooperative game theoretic view of manipulating elections. In Ning Chen, Edith Elkind, and Elias Koutsoupias, editors, *Internet and Network Economics - 7th International Workshop, WINE 2011*, volume 7090 of *Lecture Notes in Computer Science*, pages 363–374. Springer, 2011.