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

Metric distortion of voting rules with predictions

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

ΜΑΡΚΟΣ ΓΙΑΝΝΟΠΟΥΛΟΣ

Επιβλέπων : Δημήτριος Φωτάκης
Καθηγητής Ε.Μ.Π.

Αθήνα, Οκτώβριος 2024

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

Metric distortion of voting rules with predictions

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

ΜΑΡΚΟΣ ΓΙΑΝΝΟΠΟΥΛΟΣ

Επιβλέπων : Δημήτριος Φωτάκης
Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 30η Οκτωβρίου 2024.

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

Αθήνα, Οκτώβριος 2024

....

Μάρκος Γιαννόπουλος

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

Περίληψη

Σε αυτή τη διπλωματική εργασία, μελετάμε την απόδοση των κανόνων ψηφοφορίας μετρώντας τη *μετρική παραμόρφωσή τους* και ενσωματώνουμε σε αυτούς *προβλέψεις* προκειμένου να επιτύχουμε καλύτερα αποτελέσματα.

Έχουμε ένα σύνολο n *ψηφοφόρων* και ένα σύνολο m *υποψηφίων* οι οποίοι θεωρούμε ότι είναι τοποθετημένοι σε κάποιον μετρικό χώρο. Στόχος μας είναι να εκλέξουμε έναν υποψήφιο που ελαχιστοποιεί το *κοινωνικό κόστος*, δηλαδή το άθροισμα των αποστάσεων προς όλους τους ψηφοφόρους, έχοντας όμως πρόσβαση μόνο στις *κατατάξεις προτιμήσεων* των ψηφοφόρων και όχι στις πραγματικές τιμές των αποστάσεων. Η παραμόρφωση ενός κανόνα ψηφοφορίας μετρά τη χειρότερη δυνατή απόδοσή του σε σχέση με το βέλτιστο κοινωνικό κόστος. Παρουσιάζουμε αποτελέσματα από τη βιβλιογραφία, τα οποία παρέχουν φράγματα για τη μετρική παραμόρφωση αρκετών γνωστών κανόνων ψηφοφορίας, καθώς και έναν κανόνα ψηφοφορίας με βέλτιστη μετρική παραμόρφωση.

Στη συνέχεια, μελετάμε το νέο *μαθησιακά-ενισχυμένο πλαίσιο*, όπου οι αλγόριθμοι χρησιμοποιούν προβλέψεις προερχόμενες από μηχανική μάθηση με στόχο να βελτιώσουν την απόδοσή τους. Αυτοί οι αλγόριθμοι αξιολογούνται με τις παραμέτρους της *συνέπειας* (απόδοση στην περίπτωση που η πρόβλεψη είναι σωστή) και της *ευρωστίας* (απόδοση στην περίπτωση που η πρόβλεψη είναι αυθαίρετα λανθασμένη). Χρησιμοποιώντας το μαθησιακά-ενισχυμένο πλαίσιο, αποδεικνύουμε φράγματα για την συνέπεια και την ευρωστία κανόνων ψηφοφορίας. Εστιάζουμε στο πρόβλημα της εκλογής επιτροπών με k-μέλη στην πραγματική ευθεία, όπου εκλέγονται $k \geq 3$ υποψήφιοι και το κόστος κάθε ψηφοφόρου για μια επιτροπή ορίζεται ως η απόστασή του από το πλησιέστερο μέλος της επιτροπής. Προηγούμενες εργασίες δείχνουν ότι η παραμόρφωση για αυτό το πρόβλημα δεν είναι φραγμένη και ότι $O(k)$ ερωτήματα για ακριβείς αποστάσεις είναι ικανά και αναγκαία για μια γραμμική (ως προς το πλήθος των ψηφοφόρων) παραμόρφωση. Ο αλγόριθμός μας χρησιμοποιεί προβλέψεις σχετικά με τη βέλτιστη επιτροπή και επιτυγχάνει σταθερή συνέπεια και γραμμική ευρωστία με $O(k)$ ερωτήματα για ακριβείς αποστάσεις.

Λέξεις κλειδιά

Υπολογιστική Θεωρία Κοινωνικής Επιλογής, Εκλογή με έναν Νικητή, Εκλογή Επιτροπής, Κανόνες ψηφοφορίας, Μετρική Παραμόρφωση, Μαθησιακά-ενισχυμένοι Αλγόριθμοι.

Abstract

In this thesis, we study the performance of voting rules by measuring their *metric distortion* and incorporate *predictions* in order to achieve better results.

We have a set of n *voters* and a set of m *candidates*, which we consider located in a metric space. Our goal is to elect a candidate that minimizes the *social cost*, i.e. the sum of distances to all voters, when we only have access to the voters' *ordinal* preferences and not to the actual distances. The distortion of a voting rule measures its worst-case performance with respect to the optimal social cost. We present results from the literature that provide metric distortion bounds for several well-known voting rules, as well as a voting rule with optimal metric distortion.

Subsequently, we apply the new *learning-augmented framework*, where algorithms use machine-learned predictions in order to improve their performance. These algorithms are evaluated in terms of their *consistency* (performance in the case where the prediction is correct) and their *robustness* (performance in the case where prediction is arbitrarily wrong). Using the learning-augmented framework, we obtain bounds for the consistency and the robustness of voting rules. We focus on the problem of a k-committee election on the real line, where $k \geq 3$ candidates are elected and each voter's cost for a committee is defined as her distance to the nearest committee member. Previous work shows that the distortion for this setting is unbounded and that $O(k)$ distance queries are both sufficient and necessary for a linear (in the number of voters) distortion. Our algorithm uses predictions about the optimal committee and achieves constant consistency and linear robustness with $O(k)$ distance queries.

Key words

Computational Social Choice, Single-winner Election, Committee Election, Voting Rules, Metric Distortion, Learning-augmented Algorithms.

Ευχαριστίες

Αρχικά θα ήθελα να ευχαριστήσω τον καθηγητή κ. Δημήτρη Φωτάκη, για την έμπνευση που μου έδωσε για να ασχοληθώ με το συγκεκρικριμένο αντικείμενο μέσα από τα μαθήματα του, για την εμπιστοσύνη που μου έδειξε αναλαμβάνοντας την επίβλεψη αυτής της διπλωματικής εργασίας και για όλη τη βοήθεια και την υποστήριξη που μου παρείχε. Θα ήθελα να ευχαριστήσω επίσης τον Παναγιώτη Πατσιλινάκο για τη συνεργασία μας, τον χρόνο που μου αφιέρωσε και την πολύτιμη βοήθειά του σε όλη τη διάρκεια της εκπόνησης της εργασίας. Ευχαριστώ επίσης τους καθηγητές κ. Αριστείδη Παγουρτζή και κ. Ευάγγελο Μαρκάκη που συμμετείχαν στην τριμελή επιτροπή. Τέλος, ευχαριστώ τους γονείς μου, την αδελφή μου και τους φίλους μου για τη στήριξη τους την περίοδο των σπουδών μου και συνολικά στη ζωή μου.

Μάρκος Γιαννόπουλος,

Αθήνα, 30η Οκτωβρίου 2024

Περιεχόμενα

ΚΕΦΆΛΑΙΟ 1

Εκτενής Ελληνική Περίληψη

Η θεωρία της κοινωνικής επιλογής [25] μελετά τον τρόπο με τον οποίο οι ατομικές προτιμήσεις των μελών μιας ομάδας μπορούν να συνδυαστούν σε μια ενιαία συλλογική απόφαση. Αν και δεν είναι η μόνη εφαρμογή που παρουσιάζει ενδιαφέρον, θα χρησιμοποιήσουμε την ορολογία μιας εκλογής όπου οι συμμετέχοντες αναφέρονται ως *ψηφοφόροι*, οι πιθανές εναλλακτικές λύσεις αναφέρονται ως *υποψήφιοι* και η συνάρτηση που επιλέγει έναν υποψήφιο με βάση τις προτιμήσεις των ψηφοφόρων αναφέρεται ως *κανόνας ψηφοφορίας*. Υποθέτουμε ότι η επιθυμία ενός ψηφοφόρου να εκλεγεί κάποιος συγκεκριμένος υποψήφιος ποσοτικοποιείται από μια τιμή ωφέλειας. Τότε, ένας λογικός στόχος είναι να εκλεγεί εκείνος ο υποψήφιος που μεγιστοποιεί την *κοινωνική ωφέλεια* – δηλαδή το άθροισμα των τιμών ωφέλειας πάνω από όλους τους ψηφοφόρους. Η επίλυση αυτού του προβλήματος είναι τετριμμένη – αρκεί να υπολογίσουμε την κοινωνική ωφέλεια για όλους τους υποψήφιους και να εξαγάγουμε αυτόν που την μεγιστοποιεί. Ωστόσο, συνήθως δεν έχουμε πρόσβαση στις ακριβείς τιμές ωφέλειας, καθώς είναι δύσκολο ακόμη και για τους ψηφοφόρους να τις αξιολογήσουν για γνωσιακούς λόγους. Αντ' αυτού, απαιτούμε από τους ψηφοφόρους μόνο πληροφορίες στη μορφή της κατάταξης των υποψηφίων βάσει συγκρίσεων. Οι κανόνες ψηφοφορίας που χρησιμοποιούν μόνο πληροφορίες αυτού του τύπου δεν είναι σε θέση να εγγυηθούν την εύρεση μιας βέλτιστης λύσης για τον στόχο της μεγιστοποίησης της κοινωνικής ωφέλειας. Έτσι, όπως συμβαίνει και με τους προσεγγιστικούς αλγόριθμους [79] ή τους online αλγόριθμους όπου οι διαθέσιμες πληροφορίες είναι περιορισμένες [22], στοχος μας είναι να επιτύχουμε καλή προσέγγιση της βέλτιστης λύσης. Η *παραμόρφωση* ενός κανόνα ψηφοφορίας, ένας όρος ο οποίος εισήχθη από τους Procaccia και Rosenchein [72], είναι η μέγιστη δυνατή τιμή που μπορεί να λάβει ο λόγος της μέγιστης κοινωνικής ωφέλειας προς την κοινωνική ωφέλεια που επιτυγχάνει ο υποψήφιος που εκλέγεται από τον κανόνα ψηφοφορίας. Η παραμόρφωση έχει χρησιμοποιηθεί για την αξιολόγηση των επιδόσεων διαφόρων κανόνων ψηφοφορίας και έχουν δημιουργηθεί νέοι κανόνες προκειμένου να ελαχιστοποιηθεί όσο γίνεται περισσότερο η παραμόρφωση σε διάφορα περιβάλλοντα.

Στη γενική περίπτωση, υπάρχουν ισχυρά αρνητικά αποτελέσματα όσον αφορά την παραμόρφωση των κανόνων ψηφοφορίας. Γι' αυτόν τον λόγο, οι Anshelevich et al. [8], πρόσθεσαν την υπόθεση ότι οι ψηφοφόροι και οι υποψήφιοι βρίσκονται σε έναν μετρικό χώρο, προκειμένου να αποδείξουν πιο ουσιαστικά φράγματα για την παραμόρφωση. Σε αυτό το πλαίσιο, οι ψηφοφόροι προτιμούν τους υποψηφίους που βρίσκονται πιο κοντά τους σε σύγκριση με εκείνους που βρίσκονται πιο μακριά, και η απόσταση μεταξύ ενός ψηφοφόρου και ενός υποψηφίου μπορεί να θεωρηθεί ως ένα επαγόμενο κόστος. Αντί να μεγιστοποιούμε την κοινωνική ωφέλεια, στοχεύουμε στην εκλογή του υποψηφίου που ελαχιστοποιεί το κοινωνικό κόστος, το οποίο είναι το άθροισμα των αποστάσεων του υποψηφίου από όλους τους ψηφοφόρους. Στην παρούσα διπλωματική εργασία θα επικεντρωθούμε σε αυτή τη μετρική περίπτωση.

Οι περιορισμοί της πληροφορίας που παρουσιάζονται στο πρόβλημα της μετρικής παραμόρφωσης το καθιστούν κατάλληλο για την εφαρμογή του νέου *μαθησιακά ενισχυμένου πλαισίου*. Σε αυτό το πλαίσιο, οι αλγόριθμοι (στην προκειμένη περίπτωση οι κανόνες ψηφοφορίας) ενισχύονται με μια πρόβλεψη, η οποία χρησιμοποιείται προκειμένου να βελτιωθεί η απόδοσή τους. Η πρόβλεψη αυτή, η οποία μπορεί να λάβει διάφορες μορφές, μπορεί να προκύψει από ένα μοντέλο μηχανικής μάθησης, χρησιμοποιώντας σχετικά ιστορικά δεδομένα. Ο αλγόριθμος αξιολογείται ταυτόχρονα με βάση την απόδοσή του όταν η πρόβλεψη είναι ακριβής (ο δείκτης αυτός ονομάζεται *συνέπεια*), καθώς και την απόδοσή του όταν η πρόβλεψη μπορεί να είναι αυθαίρετα ανακριβής (ο δείκτης αυτός ονομάζεται *ευρωστία*). Στόχος μας σε αυτή τη διπλωματική εργασία είναι να βελτιώσουμε τα φράγματα για τη μετρική παραμόρφωση χρησιμοποιώντας μαθησιακά ενισχυμένους αλγόριθμους, κυρίως στην περίπτωση όπου έχουμε πολλαπλούς νικητές, όπου αντί για έναν μόνο υποψήφιο εκλέγουμε μια επιτροπή k υποψηφίων.

1.1 Κανόνες ψηφοφορίας και παραμόρφωση

Στην κλασική θεωρία κοινωνικής επιλογής, ένας κανόνας ψηφοφορίας λαμβάνει ως είσοδο από κάθε ψηφοφόρο μια γραμμική διάταξη των υποψηφίων και εξάγει έναν νικητή υποψήφιο. Το μοντέλο αυτό αντιστοιχεί στον τρόπο με τον οποίο οι άνθρωποι συνήθως εκφράζουν τις προτιμήσεις τους για τις εναλλακτικές λύσεις, κατατάσσοντάς τις και χωρίς να τους αποδίδουν κάποια ακριβή τιμή ωφέλειας. Λόγω της έλλειψης αριθμητικών τιμών που να μετρούν την ποιότητα του εκλεγμένου υποψηφίου, μια λογική προσέγγιση για την αξιολόγηση των κανόνων ψηφοφορίας είναι η *αξιωματική προσέγγιση*. Σε αυτήν την προσέγγιση, διατυπώνονται ορισμένες αξιωματικές ιδιότητες ή κριτήρια που πρέπει να ικανοποιούν οι κανόνες ψηφοφορίας. Στη συνέχεια μπορεί κανείς να επιλέξει έναν κανόνα που έχει μια επιθυμητή ιδιότητα έναντι ενός άλλου κανόνα που δεν την έχει. Κάποιες σημαντικές πρώιμες εργασίες σε αυτόν τον τομέα είναι αυτές των Arrow [14], May [65], Gibbard [46], Satterthwaite [75] και Young [83] (βλέπε επίσης το άρθρο επισκόπησης του Zwicker [84]).

Στην παρούσα διπλωματική εργασία θα αξιολογήσουμε τους κανόνες ψηφοφορίας ακολουθώντας την *ωφελιμιστική προσέγγιση* που χρησιμοποιείται στη θεωρία παιγνίων [80] και στον αλγοριθμικό σχεδιασμό μηχανισμών [70]. Ο ωφελιμισμός, που θεμελιώθηκε από τον Bentham, υποστηρίζει ότι η "ευτυχία" που αποκτά ένα συγκεκριμένο άτομο από μια συγκεκριμένη κατάσταση στον κόσμο μπορεί να ποσοτικοποιηθεί από μια *συνάρτηση ωφέλειας* και στοχεύει στη μεγιστοποίηση της συνολικής «ευτυχίας» του πληθυσμού μεγιστοποιώντας το άθροισμα των ατομικών ωφελειών, δηλαδή την κοινωνική ωφέλεια. Οι Boutilier κ.ά. [24] επισημαίνουν ότι, αν και δεν είναι όλα τα προβλήματα κοινωνικής επιλογής κατάλληλα για την ωφελιμιστική προσέγγιση (για παράδειγμα υπάρχουν περιπτώσεις όπου η διαπροσωπική σύγκριση των ωφελειών δεν είναι δυνατή), υπάρχουν πολλές πραγματικές καταστάσεις που ταιριάζουν στην ωφελιμιστική άποψη. Για παράδειγμα, στα συστήματα συστάσεων και σε πολλούς παρόμοιους τομείς από το σχεδιασμό μηχανισμών και το ηλεκτρονικό εμπόριο, οι υπολογιστικοί πράκτορες συνήθως αντιστοιχίζουν τιμές ωφέλειας στις διάφορες εναλλακτικές λύσεις αντί να διατάσσουν το σύνολο των υποψηφίων. Παρόλα αυτά, γίνεται η υπόθεση ότι οι τιμές ωφέλειας δεν είναι γνωστές και, όπως και στην κλασική θεωρία κοινωνικής επιλογής, οι κανόνες ψηφοφορίας έχουν πρόσβαση μόνο στις κατατάξεις των υποψηφίων. Αυτές οι κατατάξεις είναι συμβατές με τις τιμές ωφέλειας, πράγμα που σημαίνει ότι ένας υποψήφιος με υψηλότερη τιμή ωφέλειας κατατάσσεται υψηλότερα σε σύγκριση με έναν υποψήφιο με χαμηλότερη τιμή ωφέλειας. Ο περιορισμός αυτός δικαιολογείται από τους συμπεριφορικούς οικονομολόγους που έχουν δείξει ότι είναι γνωσιακά δύσκολο να αποδώσουμε ακριβείς τιμές ωφέλειας σε διάφορες εναλλακτικές λύσεις.

Ένας κανόνας ψηφοφορίας που χρησιμοποιεί τις κατατάξεις των υποψηφίων από τους ψηφοφόρους δεν είναι πάντα σε θέση να εντοπίσει έναν υποψήφιο που μεγιστοποιεί την κοινωνική ωφέλεια. Επομένως, μπορούμε να σκεφτούμε έναν κανόνα ψηφοφορίας ως έναν προσεγγιστικό αλγόριθμο που προσπαθεί να επιλέξει τον καλύτερο δυνατό υποψήφιο με βάση περιορισμένες πληροφορίες (τις κατατάξεις αντί για τις τιμές ωφέλειας). Αυτή η οπτική της ψηφοφορίας με κατατάξεις προτάθηκε από τους Procaccia και Rosenschein στο [72], όπου εισήγαγαν τον όρο παραμόρφωση για να αναφερθούν στην ποιότητα της προσέγγισης που παρέχει ένας κανόνας ψηφοφορίας. Η παραμόρφωση ενός κανόνα ψηφοφορίας είναι ο χειρότερος δυνατός λόγος της μέγιστης κοινωνικής ωφέλειας προς την κοινωνική ωφέλεια του υποψηφίου που εκλέγεται. Η έννοια της παραμόρφωσης παρέχει έναν ποσοτικό τρόπο σύγκρισης διαφόρων κανόνων ψηφοφορίας: χαμηλή παραμόρφωση είναι προφανώς ένα επιθυμητό χαρακτηριστικό. Μια πρόσφατη επισκόπηση των σημαντικότερων αποτελεσμάτων σχετικά με την παραμόρφωση μπορεί να βρεθεί στο [9].

Έστω V το σύνολο των ψηφοφόρων, $|V| = n$, και έστω C το σύνολο των υποψηφίων, $|C| = m$. Για αυθαίρετες τιμές ωφέλειας, η παραμόρφωση είναι απεριόριστη ακόμη και για μικρούς αριθμούς ψηφοφόρων και υποψηφίων. Ως εκ τούτου, οι Procaccia και Rosenchein [72] υπέθεσαν κανονικοποιημένες τιμές ωφέλειας όπου το άθροισμα των τιμών ωφέλειας για κάθε ψηφοφόρο είναι σταθερό, π.χ. 1. Ακόμη και με αυτή την υπόθεση, έδειξαν ότι οι δημοφιλείς κανόνες ψηφοφορίας Borda και Veto έχουν απεριόριστη παραμόρφωση [72]. Με την ίδια υπόθεση κανονικοποιημένων ωφελειών, οι Καραγιάννης και Procaccia [29] έδειξαν ότι ο απλός κανόνας Plurality, ο οποίος εκλέγει τον υποψήφιο που κατατάσσεται πρώτος από τους περισσότερους ψηφοφόρους, έχει παραμόρφωση $O(m^2)$. Οι Καραγιάννης et.al. [28] έδειξαν ότι η παραμόρφωση για τους ντετερμινιστικούς κανόνες ψηφοφορίας έχει κάτω φράγμα $\Omega(m^2)$, άρα ο κανόνας Plurality είναι βέλτιστος. Επιτρέποντας την τυχαιότητα στους κανόνες ψηφοφορίας μπορούμε να βελτιώσουμε περαιτέρω τα φράγματα για την παραμόρφωση, όπου η κοινωνική ωφέλεια του υποψηφίου που εκλέγεται στον ορισμό της παραμόρφωσης αντικαθίσταται από την *αναμενόμενη* κοινωνική ωφέλεια. Οι Boutilier et.al. [24], υποθέτοντας επίσης κανονικοποιημένες τιμές ωφέλειας, απέδειξαν ένα κάτω φράγμα της τάξης $\Omega(\sqrt{m})$ για όλους τους πιθανοτικούς κανόνες ψηφοφορίας και σχεδίασαν έναν πιθανοτικό κανόνα ψηφοφορίας με παραμόρφωση $O(\sqrt{m} \cdot \log^* m)$, που σχεδόν ταιριάζει με το κάτω φράγμα. Έχουν μελετηθεί και άλλες παραλλαγές της παραμόρφωσης, π.χ. η πολυπλοκότητα επικοινωνίας [63], [64] και η παραμόρφωση σε κατανεμημένα περιβάλλοντα [44].

1.2 Μετρική παραμόρφωση

Σε αυτή τη διπλωματική εργασία, θα επικεντρωθούμε στο πλαίσιο της *μετρικής παραμόρφωσης*, το οποίο εισήχθη από τους Anshelevich et al. [8] και υποθέτει ότι οι ψηφοφόροι και οι υποψήφιοι βρίσκονται σε έναν μετρικό χώρο. Είναι φυσιολογικό να υποθέσουμε ότι οι ψηφοφόροι προτιμούν υποψηφίους που βρίσκονται κοντά τους. Επομένως, σε αυτό το πλαίσιο, η επιθυμία ενός ψηφοφόρου να ελαχιστοποιήσει την απόστασή του (η οποία μπορεί να θεωρηθεί ως κόστος) από τον υποψήφιο που θα εκλεγεί αντιστοιχεί στην επιθυμία του να μεγιστοποιήσει την τιμή ωφέλειας στο ωφελιμιστικό πλαίσιο. Αυτή η προσέγγιση είναι παρόμοια με τα χωρικά μοντέλα ψηφοφορίας από την πολιτική επιστήμη [40], [13], [66], [35], [76]. Μπορούμε να θεωρήσουμε ότι η γνώμη ενός ψηφοφόρου για ένα θέμα ενδιαφέροντος μπορεί να αναπαρασταθεί με μια συντεταγμένη σε έναν Ευκλείδειο χώρο, π.χ. τον απλό μονοδιάστατο άξονα αριστερά–δεξιά. Ωστόσο, οι μετρικοί χώροι που θα θεωρήσουμε εδώ είναι πιο γενικοί και συνεπώς πιο ισχυροί.

Στο Κεφάλαιο 3 παρουσιάζουμε αποτελέσματα από το [8] για τη μετρική παραμόρφωση γνωστών κανόνων ψηφοφορίας. Ξεκινάμε με τον τυπικό ορισμό των εννοιών που θα χρησιμοποιήσουμε. Έστω V και C δύο πεπερασμένα σύνολα. Λέμε ότι το V είναι το σύνολο των *ψηφοφόρων* και το C είναι το σύνολο των *υποψηφίων*. Θέτουμε $n = |V|$ (το πλήθος των ψηφοφόρων) και $m = |C|$ (το πλήθος των υποψηφίων). Σε ό,τι ακολουθεί, συνήθως συμβολίζουμε τους ψηφοφόρους με u, v και τους υποψήφιους με c, x, y. Υποθέτουμε ότι τόσο το V όσο και το C είναι τοποθετημένα μέσα σε έναν μετρικό χώρο (X, d), οπότε η απόσταση $d(a, b)$ μεταξύ των a και b ορίζεται καλά για κάθε $a, b \in V \cup C$.

Το *κοινωνικό κόστος* ενός υποψηφίου $c \in C$ ως προς τη μετρική d είναι το άθροισμα

$$\mathrm{SC}(c, d) = \sum_{v \in V} d(v, c).$$

Γράφουμε $\mathrm{SC}(c)$ όταν η μετρική d είναι σαφής από τα συμφραζόμενα.

Μια τριάδα (V, C, d) όπως παραπάνω, ονομάζεται *παράδειγμα*. Η απόσταση $d(v, c)$ ανάμεσα σε έναν ψηφοφόρο v και έναν υποψήφιο c είναι ένα μέτρο του πόσο ο v προτιμάει τον c. Λέμε ότι ο $v \in V$ προτιμάει τον $c \in C$ έναντι του $c' \in C$ αν $d(v, c) \leq d(v, c')$. Σε αυτή την περίπτωση γράφουμε

$$c \succcurlyeq_v c'.$$

Κάθε δοσμένο παράδειγμα (V, C, d) επάγει κατατάξεις προτιμήσεων των ψηφοφόρων. Για κάθε ψηφοφόρο $v \in V$ έχουμε μια *κατάταξη προτιμήσεων* $\sigma_v : C \to \{1, \dots, m\}$ που ικανοποιεί το εξής: αν $\sigma_v(c) < \sigma_v(c')$ τότε $c \succcurlyeq_v c'$. Τότε λέμε ότι η d είναι *συμβατή* με την σ_v και γράφουμε $d \vartriangleright \sigma_v$.

Ένα *προφίλ προτιμήσεων* $\sigma := (\sigma_v)_{v \in V}$ είναι μια n-άδα κατατάξεων προτιμήσεων των ψηφοφόρων. Λέμε ότι η d είναι συμβατή με το προφίλ προτιμήσεων σ, και γράφουμε $d \vartriangleright \sigma$, αν $d \vartriangleright \sigma_v$ για κάθε $v \in V$.

Το πρόβλημα της μετρικής παραμόρφωσης περιγράφεται ως εξής: Ένας αλγόριθμος ALG λαμβάνει ως είσοδο κάποιο προφίλ προτιμήσεων σ το οποίο επάγεται από ένα παράδειγμα (V, C, d). Ο αλγόριθμος δεν έχει πρόσβαση στην υποκείμενη συνάρτηση αποστάσεων d. Ο στόχος είναι να αποδώσει ως έξοδο έναν υποψήφιο c^* ο οποίος ελαχιστοποιεί το κοινωνικό κόστος, δηλαδή να ικανοποιεί την

$$\mathrm{SC}(c^*) = \min_{c \in C} \mathrm{SC}(c).$$

Θα χρησιμοποιούμε επίσης τον όρο "*κανόνας ψηφοφορίας*" για κάθε αλγόριθμο γι" αυτό το πρόβλημα και τον όρο "*νικητής*" του κανόνα ψηφοφορίας για την έξοδο του αλγορίθμου. Η *παραμόρφωση* του ALG είναι η χειρότερη δυνατή πολλαπλασιαστική προσέγγιση που μπορεί να επιτύχει σε σχέση με το βέλτιστο, η οποία τυπικά ορίζεται ως εξής:

$$\mathrm{distortion}(\mathrm{ALG}) = \sup_{\sigma} \sup_{d : d \vartriangleright \sigma} \frac{\mathrm{SC}(\mathrm{ALG}(\sigma), d)}{\mathrm{SC}(c^*(d), d)}.$$

Ένα πρώτο βασικό ερώτημα που προκύπτει είναι πόσο καλή μπορεί να είναι η απόδοση *οποιουδήποτε* αλγορίθμου. Το πρώτο θεώρημα που παρουσιάζουμε προέρχεται από το [8] και δίνει ένα κάτω φράγμα για την επίδοση όλων των ντετερμινιστικών κανόνων ψηφοφορίας: κανένας τέτοιος κανόνας δεν μπορεί να έχει παραμόρφωση μικρότερη από 3.

Στη συνέχεια δίνουμε κάτω και άνω φράγματα από το [8] για την παραμόρφωση πολλών γνωστών κανόνων ψηφοφορίας, ξεκινώντας με κάποιους κλασσικούς κανόνες που ορίζονται από "διανύσματα βαθμολογίας θέσης". Ένας *κανόνας με βαθμολογίες θέσης* $f_{\vec{s}}$ για το C προσδιορίζεται από ένα διάνυσμα βαθμολογίας θέσης $\vec{s} = (s_1, s_2, \dots, s_m)$, όπου $s_i \in \mathbb{Q}$, $s_1 \geq s_2 \geq \cdots \geq s_m$ και $s_1 > s_m$. Όταν κάποιος ψηφοφόρος $v \in V$ κατατάσσει έναν υποψήφιο $c \in C$ στη θέση ℓ, τότε ο υποψήφιος c λαμβάνει $r_v(c) = s_\ell$ πόντους από τον v. Κατόπιν, η συνολική βαθμολογία του υποψηφίου c είναι το άθροισμα

$$\mathrm{score}_{\vec{s}}(c) = \sum_{v \in V} r_v(c).$$

Οι νικητές σύμφωνα με το διάνυσμα βαθμολογίας θέσης $\vec{s}$ είναι εκείνοι οι υποψήφιοι $c^* \in C$ για τους οποίους

$$\text{score}_{\vec{s}}(c^*) = \max_{c \in C} \text{score}_{\vec{s}}(c).$$

Μελετάμε τους ακόλουθους κανόνες με βαθμολογίες θέσης:

(i) Plurality (πλειοψηφία), όπου $\vec{s} = (1, 0, \ldots, 0)$.

(ii) Veto (βέτο), όπου $\vec{s} = (1, 1, \ldots, 1, 0)$.

(iii) Borda, όπου $\vec{s} = (m-1, m-2, \ldots, 1, 0)$.

(iv) Harmonic rule (αρμονικός κανόνας), όπου $\vec{s} = \left(1, \frac{1}{2}, \frac{1}{3}, \ldots, \frac{1}{m}\right)$.

(v) k-Approval (k-αποδοχή), για $1 \leq k < m$, όπου $\vec{s} = (1, 1, \ldots, 1\,[k], 0, \ldots, 0\,[m-k])$.

Αρχικά παρουσιάζουμε αποτελέσματα από το [8] που δίνουν φράγματα για γνωστούς κανόνες με βαθμολογίες θέσης. Για τον κανόνα k-Approval με $k > 1$ και για τον Veto, η παραμόρφωση δεν είναι φραγμένη. Για τον Plurality και τον Borda, η παραμόρφωση είναι $2m - 1$, όπου m είναι το πλήθος των υποψηφίων, με το άνω φράγμα να συμπίπτει με το κάτω φράγμα. Επιπλέον, δεν υπάρχει κανένας κανόνας με βαθμολογίες θέσης που η παραμόρφωσή του να είναι φραγμένη από μια σταθερά ανεξάρτητη του m. Για κάθε τιμή του m, η παραμόρφωση οποιουδήποτε κανόνα με βαθμολογίες θέσης για m υποψηφίους είναι τουλάχιστον $1 + 2\sqrt{\ln m} - 1$.

Είναι ενδιαφέρον το γεγονός ότι ούτε ο κανόνας Plurality ούτε ο κανόνας Borda επιτυγχάνουν την καλύτερη παραμόρφωση μέσα στην κλάση των κανόνων με βαθμολογίες θέσης. Ο Αρμονικός κανόνας των Boutilier et al. [24] έχει ασυμπτωτικά καλύτερη παραμόρφωση τόσο από τον κανόνα Plurality όσο και από τον κανόνα Borda. Όμως, η παραμόρφωση του Αρμονικού κανόνα είναι κι αυτή σχεδόν γραμμική, με ένα κάτω φράγμα $\Omega\left(\frac{m}{\sqrt{\ln m}}\right)$ και ένα άνω φράγμα $O\left(\frac{m}{\sqrt{\ln m}}\right)$.

Στο [8] αποδεικνύεται επίσης ότι, παρόλο που κάποιοι από τους κανόνες που χρησιμοποιούνται ευρέως έχουν μεγάλη παραμόρφωση, υπάρχουν σημαντικοί κανόνες ψηφοφορίας των οποίων η παραμόρφωση είναι φραγμένη' από μια μικρή σταθερά ή αυξάνει αργά καθώς αυξάνει το πλήθος των υποψηφίων. Ένα παράδειγμα δίνει ο γνωστός κανόνας Copeland. Το *σκορ Copeland* ενός υποψηφίου c είναι το πλήθος των υποψηφίων που νικάει ο c. Ο κανόνας Copeland [36] αποδίδει στην έξοδο όλους τους υποψηφίους που έχουν μέγιστο σκορ Copeland. Παρουσιάζουμε ένα αποτέλεσμα των Anshelevich et al. [8] οι οποίοι απέδειξαν ότι ο κανόνας Copeland έχει παραμόρφωση 5, με το άνω φράγμα να ταιριάζει με το κάτω φράγμα. Αυτό σημαίνει ότι, παρόλο που ο κανόνας Copeland δεν γνωρίζει τίποτα για τις τιμές της μετρικής πέρα από τις κατατάξεις προτιμήσεων που επάγονται από αυτές, και ίσως δεν μπορεί να βρει τον πραγματικό βέλτιστο υποψήφιο, εντούτοις επιλέγει πάντα κάποιον υποψήφιο του οποίου η ποιότητα απέχει μόνο ένα συντελεστή 5 από αυτήν του βέλτιστου υποψηφίου. Υπενθυμίζουμε ότι, από το γενικό κάτω φράγμα 3, κανένας ντετερμινιστικός κανόνας ψηφοφορίας δεν μπορεί να τα πάει πολύ καλύτερα από τον κανόνα Copeland ως προς το κοινωνικό κόστος.

Ένας άλλος δημοφιλής κανόνας ψηφοφορίας, γνωστός ως κανόνας της απλής μεταβιβάσιμης ψήφου (STV), μελετήθηκε στο [77] και αργότερα στο [6]. Ο κανόνας απλής μεταβιβάσιμης ψήφου (STV), που είναι ένας από τους ελάχιστους μη τετριμμένους κανόνες ψηφοφορίας που χρησιμοποιούνται σε πραγματικές εκλογές, είναι ένας επαναληπτικός κανόνας που ορίζεται με τον ακόλουθο τρόπο. Σε κάθε γύρο, ο υποψήφιος που κατατάσσεται πρώτος από τους λιγότερους ψηφοφόρους (αυτός με το χαμηλότερο σκορ πλειοψηφίας) αφαιρείται από το σύνολο των υποψηφίων και από τις κατατάξεις προτιμήσεων των ψηφοφόρων, πράγμα που σημαίνει ότι τα σκορ πλειοψηφίας πρέπει να υπολογιστούν εκ νέου. Μετά από $m - 1$ γύρους απομένει μόνο ένας υποψήφιος, ο οποίος είναι ο νικητής.

Οι Elkind και Skowron [77] απέδειξαν ένα άνω φράγμα $O(\ln m)$ και ένα κάτω φράγμα $\Omega(\sqrt{\ln m})$ για την παραμόρφωση του STV, κάτι που σημαίνει ότι ο STV δεν είναι εξίσου καλός με τον κανόνα Copeland.

Τα παραπάνω αποτελέσματα συνοψίζονται στον ακόλουθο Πίνακα.

Πίνακας 1. Άνω και κάτω φράγματα για την παραμόρφωση

Κανόνας ψηφοφορίας	**Άνω φράγμα**	**Κάτω φράγμα**
k-Approval, $k > 1$	∞	∞
Veto	∞	∞
Plurality	$2m - 1$	$2m - 1$
Borda	$2m - 1$	$2m - 1$
Harmonic Rule	$O\left(\frac{m}{\sqrt{\ln m}}\right)$	$\Omega\left(\frac{m}{\ln m}\right)$
Scoring Rules	∞	$1 + 2\sqrt{\ln m - 1}$
Copeland	5	5
STV	$O(\ln m)$	$\Omega(\sqrt{\ln m})$

Οι Anshelevich et al. [8] έκαναν την εικασία ότι ο κανόνας Ranked Pairs [78] επιτυγχάνει τη βέλτιστη παραμόρφωση 3, όμως αργότερα αποδείχθηκε ότι αυτό δεν ισχύει, από τους Goel et al. [51]. Στη συνέχεια, οι Munagala και Wang [69] παρουσίασαν έναν ντετερμινιστικό αλγόριθμο που επιτυγχάνει παραμόρφωση $2 + \sqrt{5} \approx 4.236$. Στην πραγματικότητα, αυτός ο αλγόριθμος δεν απαιτεί πρόσβαση στο πλήρες προφίλ προτιμήσεων των ψηφοφόρων. Χρησιμοποιεί μια περιορισμένη είσοδο, το σταθμισμένο γράφημα του τουρνουά, το οποίο περιέχει, για κάθε ζεύγος υποψηφίων, το πλήθος των ψηφοφόρων που προτιμούν τον έναν από τους δύο έναντι του άλλου. Με βάση αυτό, έκαναν την εικασία ότι εάν δίνονται οι πλήρεις καταστάσεις προτιμήσεων τότε η καλύτερη δυνατή παραμόρφωση θα πρέπει να είναι 3. Οι ίδιοι συγγραφείς περιέγραψαν μια ικανή συνθήκη που θα συνεπαγόταν τη βέλτιστη παραμόρφωση 3. Περίπου την ίδια εποχή, ο Kempe [56] απέδειξε ανεξάρτητα το ίδιο φράγμα $2 + \sqrt{5}$ χρησιμοποιώντας γραμμικό προγραμματισμό, έδωσε διάφορες εναλλακτικές διατυπώσεις της ικανής συνθήκης του [69] και απέδειξε ότι ο κανόνας Ranked Pairs έχει παραμόρφωση $\Theta(\sqrt{m})$.

Στην Ενότητα 3.5 παρουσιάζουμε ένα θεώρημα των Γκατζέλη, Halpern και Shah [47], οι οποίοι εισήγαγαν έναν ντετερμινιστικό αλγόριθμο, τον τον αλγόριθμο Plurality Matching, που ολοκληρώνει αυτή την κατεύθυνση έρευνας και εγγυάται τη βέλτιστη παραμόρφωση 3. Το κύριο τεχνικό συστατικό αυτής της εργασίας είναι το λήμμα "κατάταξης-ταιριάσματος" το οποίο εξασφαλίζει την ύπαρξη κλασματικών τέλειων ταιριασμάτων μέσα σε μια οικογένεια διμερών γραφημάτων με σταθμισμένες κορυφές που επάγεται από το πρόβλημα.

Μεταγενέστερες εργασίες προσέφεραν ακόμα καλύτερη κατανόηση του προβλήματος. Οι Kizilkaya και Kempe [58] πέτυχαν τη βέλτιστη παραμόρφωση με τον αλγόριθμο Plurality Veto, ο οποίος είναι πολύ απλούστερος σε σχέση με τον αλγόριθμο Plurality Matching, και στη συνέχεια οι ίδιοι συγγραφείς [59] πρότειναν μια βελτίωση του Plurality Veto που ονομάζεται Simultaneous Veto και επιλύει ζητήματα ισοπαλίας που εμφανίζονται στα [47] και [58] διατηρώντας τη βέλτιστη παραμόρφωση.

Μετά από την οριστική απάντηση στο πρόβλημα της παραμόρφωσης για την κατηγορία των ντετερμινιστικών αλγορίθμων, ένα ανοιχτό πρόβλημα είναι να επιτευχθούν επίσης τα βέλτιστα φράγματα παραμόρφωσης για *τυχαιοποιημένους* αλγορίθμους. Οι Anshelevich και Postl [11] έδωσαν το κάτω φράγμα 2 και απέδειξαν ότι ο αλγόριθμος Random Dictatorship έχει παραμόρφωση μικρότερη από 3, όμως το φράγμα τους τείνει στα 3 καθώς αυξάνεται το πλήθος των ψηφοφόρων. Στη συνέχεια, οι Fain et al. [41] και ο Kempe [57] παρουσίασαν άλλους αλγορίθμους με παραμόρφωση που τείνει στα 3, αυτή τη φορά καθώς αυξάνεται το πλήθος των υποψηφίων. Το καλύτερο γνωστό κάτω φράγμα 2.1126 δόθηκε πρόσφατα από τους Charikar και Ramakrishnan [32], και οι Charikar et al. [33] πέτυχαν το άνω φράγμα 2.753 το οποίο ξεφεύγει για πρώτη φορά από την τιμή 3. Το ερώτημα αν μπορεί να μειωθεί το εναπομείναν χάσμα ανάμεσα στο κάτω και άνω φράγμα παραμένει ανοικτό.

Το άρθρο [49] προσπαθεί να συνδυάσει φράγματα για τη μετρική παραμόρφωση με φράγματα για την παραμόρφωση στο ωφελιμιστικό περιβάλλον (όπου οι προτιμήσεις των ψηφοφόρων είναι αυθαίρετες κανονικοποιημένες αποτιμήσεις). Στο [1] μελετάται μια παραλλαγή του προβλήματος, όπου είναι διαθέσιμες περισσότερες πληροφορίες πέρα από τις συνήθεις κατατάξεις προτιμήσεων, και στο [57] μπορεί κανείς να βρει μια μελέτη των συμβιβασμών ανάμεσα στης επιτεύξιμη παραμόρφωση και την πολυπλοκότητα επικοινωνίας των ντετερμινιστικών αλγορίθμων κοινωνικής επιλογής. Στο [10] μελετάται η μετρική παραμόρφωση σε κατανεμημένα περιβάλλοντα και στο [43] εξετάζεται η παραμόρφωση των *φιλαλήθων* (truthful) μηχανισμών. Έχουν επίσης μελετηθεί διάφορα άλλα πλαίσια, όπου οι διαθέσιμες πληροφορίες είναι πιο περιορισμένες από αυτό των πλήρων κατατάξεων προτιμήσεων (βλ. [11], [52], [41], [57], [23], [7]).

1.3 Μαθησιακά ενισχυμένοι αλγόριθμοι

Πρόσφατες κατευθύνσεις της έρευνας εστιάζουν όλο και περισσότερο σε ένα νέο πλαίσιο ανάλυσης, το οποίο ονομάζεται *αλγόριθμοι με προβλέψεις* ή *μαθησιακά ενισχυμένοι αλγόριθμοι*. Αυτή η προσέγγιση αξιοποιεί την καθοδήγηση των προβλέψεων με μηχανική μάθηση, για να επιτύχει βελτιωμένα φράγματα για την απόδοση των αλγορίθμων και να ξεπεράσει τους περιορισμούς της παραδοσιακής ανάλυσης της χειρότερης περίπτωσης. Η μορφή των προβλέψεων μπορεί να ποικίλλει ανάλογα με τον τύπο των δεδομένων που διαθέτει ο σχεδιαστής και το πρόβλημα που πρέπει να επιλυθεί. Με την ενσωμάτωση των προβλέψεων, ο αλγόριθμος μπορεί να βελτιώσει την απόδοσή του χρησιμοποιώντας τις ως οδηγό.

Είναι σημαντικό να σημειώσουμε ότι η ποιότητα αυτών των προβλέψεων είναι άγνωστη και όχι αξιόπιστη. Ο αλγόριθμος αξιολογείται με δύο κριτήρια: την απόδοσή του όταν η πρόβλεψη είναι ακριβής (τη λεγόμενη *συνέπεια*) καθώς και την απόδοσή του όταν η πρόβλεψη μπορεί να είναι αυθαίρετα ανακριβής (τη λεγόμενη *ευρωστία*). Στην ακραία περίπτωση όπου ακολουθείται πάντα η πρόβλεψη, το αποτέλεσμα θα ήταν ικανοποιητικό όταν η πρόβλεψη είναι ακριβής (η συνέπεια είναι καλή), αλλά το αποτέλεσμα θα ήταν κακό όταν η πρόβλεψη είναι ανακριβής (η ευρωστία είναι κακή). Στην άλλη ακραία περίπτωση όπου κάποιος αγνοεί την πρόβλεψη, η συνέπεια είναι πολύ αδύναμη. Γενικά, κάθε μαθησιακά ενισχυμένος αλγόριθμος παρέχει έναν συμβιβασμό μεταξύ ευρωστίας και συνέπειας, και ο στόχος είναι να εντοπιστεί το σύνορο Pareto μεταξύ αυτών των δύο μέτρων.

Τα τελευταία χρόνια, το πλαίσιο της μαθησιακής ενίσχυσης έχει χρησιμοποιηθεί εκτενώς στο σχεδιασμό και την ανάλυση αλγορίθμων για να αντιμετωπιστούν οι περιορισμοί των υπερβολικά απαισιόδοξων φραγμάτων που δίνει η μελέτη της χειρότερης περίπτωσης. Αρκετές πρόσφατες εργασίες έχουν εφαρμόσει αυτό το πλαίσιο σε κλασικές περιοχές αλγοριθμικών προβλημάτων, όπως η διαδικτυακή σελιδοποίηση [62], ο χρονοπρογραμματισμός [74], τα προβλήματα γραμματέων [37], [12], τα προβλήματα βελτιστοποίησης με κάλυψη [18], και τα προβλήματα σακιδίου [54], καθώς και διάφορα προβλήματα της θεωρίας γραφημάτων [15]. Οι πρώιμες συνεισφορές σε αυτόν τον τομέα περιγράφονται στο άρθρο επισκόπησης [67], και το [60] προσφέρει μια ενημερωμένη συλλογή σχετικών εργασιών. Η προσέγγιση της μαθησιακής ενίσχυσης έχει επίσης εφαρμοστεί σε προβλήματα κοινωνικής επιλογής, όπου γίνονται προβλέψεις σχετικά με μια ομάδα ψηφοφόρων και τις προτιμήσεις τους, με στόχο τη βελτιστοποίηση του κοινωνικού κόστους ή των συναρτήσεων ευημερίας. Σε αυτή την κατεύθυνση εντάσσεται η έρευνα πάνω σε προβλήματα κατανομής πόρων όπου οι προτιμήσεις των πρακτόρων αποκαλύπτονται σταδιακά σε ένα διαδικτυακό περιβάλλον (βλ. [19] και [20]), καθώς και σενάρια όπου οι πράκτορες είναι στρατηγικοί και οι προτιμήσεις τους είναι ιδιωτικές (βλ. [2], [82], [48], [16], [55], [71], [61], [27] και [17]).

Η μελέτη του προβλήματος της μετρικής παραμόρφωσης σε αυτό το πλαίσιο είναι εύλογη λόγω της περιορισμένης πληροφορίας στην οποία βασίζεται. Εκτός από τις κατατάξεις των ψηφοφόρων, είναι λογικό να υποθέσουμε ότι μπορούμε να έχουμε πρόσβαση σε ιστορικά δεδομένα σχετικά με τις αποφάσεις των ψηφοφόρων σε συναφή θέματα, τα οποία θα μπορούσαν να συσχετιστούν με τις προτιμήσεις τους για το συγκεκριμένο θέμα. Αυτές οι πρόσθετες πληροφορίες μπορούν να δώσουν καλύτερη διαίσθηση για το πιθανό αποτέλεσμα.

Στο άρθρο [21], οι Berger et al. εγκαινιάζουν την ανάλυση του προβλήματος της μετρικής παραμόρφωσης χρησιμοποιώντας το πλαίσιο μαθησιακής ενίσχυσης και χαρακτηρίζουν το σύνορο Pareto της ευρωστίας-συνέπειας γι' αυτό το πρόβλημα. Εξετάζουν αλγορίθμους που λαμβάνουν ως είσοδο ένα ζεύγος (σ, p), όπου σ είναι ένα προφίλ προτιμήσεων και $p \in C$ είναι μια *πρόβλεψη* σχετικά με τον βέλτιστο υποψήφιο $c^*(d)$. Ο στόχος είναι να αξιολογηθεί η απόδοση ενός αλγορίθμου μέσω της *συνέπειας* και της *ευρωστίας* του. Η συνέπεια του ALG ορίζεται ως η παραμόρφωση που εγγυάται ο ALG όταν η παρεχόμενη πρόβλεψη είναι ακριβής, δηλαδή $p = c^*(d)$. Πιο τυπικά,

$$\text{consistency}(\text{ALG}) = \sup_{\sigma} \sup_{d:d \,\triangleright\, \sigma} \frac{\text{SC}(\text{ALG}(\sigma, c^*(d)), d)}{\text{SC}(c^*(d), d)}.$$

Η ευρωστία της ALG ορίζεται ως η παραμόρφωση που εγγυάται η ALG με μια αυθαίρετη πρόβλεψη, ανεξάρτητα από το πόσο ακριβής μπορεί να είναι αυτή η πρόβλεψη. Πιο τυπικά,

$$\text{robustness}(\text{ALG}) = \sup_{\sigma} \sup_{p \in C} \sup_{d:d \,\triangleright\, \sigma} \frac{\text{SC}(\text{ALG}(\sigma, p), d)}{\text{SC}(c^*(d), d)}.$$

Στο Κεφάλαιο 4 μελετάμε τη δουλειά των Berger, Feldman, Gkatzelis και Tan, οι οποίοι εισήγαγαν μια οικογένεια αλγορίθμων, με παραμέτρους $0 \leq \delta < 1$, γνωστή ως BoostedSV$_\delta$, και απέδειξαν φράγματα για τη συνέπεια και την ευρωστία τους συναρτήσει του δ. Πρόκειται για μια μαθησιακά ενισχυμένη τροποποίηση του αλγορίθμου SimultaneousVeto που προτάθηκε στο [59]. Ο SimultaneousVeto αρχικά αποδίδει σε κάθε υποψήφιο $c \in C$ μια βαθμολογία ίση με τον αριθμό των ψηφοφόρων που κατέταξαν τον c στην κορυφή (το σκορ πλειοψηφίας του) και στη συνέχεια αφήνει τους ψηφοφόρους συνεχώς και ταυτόχρονα να μειώνουν τη βαθμολογία του λιγότερο προτιμώμενου υποψηφίου τους μεταξύ αυτών που εξακολουθούν να έχουν θετική βαθμολογία. Τέλος, επιλέγει ως νικητή έναν από τους υποψήφιους που η βαθμολογία τους φτάνει τελευταία στο 0. Ο αλγόριθμος BoostedSV$_\delta$ βελτιώνει τον SimultaneousVeto με την "ενίσχυση" της αρχικής βαθμολογίας του υποψηφίου $p \in C$ που προβλέπεται να είναι βέλτιστος. Το μέγεθος αυτής της ενίσχυσης είναι μια προσεκτικά επιλεγμένη αύξουσα συνάρτηση του δ, και στη συνέχεια, ενισχύεται επίσης κατάλληλα ο ρυθμός με τον οποίο όλοι οι ψηφοφόροι μειώνουν τις βαθμολογίες. Όσο υψηλότερη είναι η τιμή του δ, κάτι που αντιστοιχεί σε υψηλότερη εμπιστοσύνη του σχεδιαστή στην ποιότητα της πρόβλεψης, τόσο μεγαλύτερο είναι το μέγεθος αυτής της ενίσχυσης. Αποδεικνύεται ότι ο αλγόριθμος BoostedSV$_\delta$ επιτυγχάνει $\frac{3-\delta}{1+\delta}$-συνέπεια και $\frac{3+\delta+13\delta^2-\delta^3}{(1+\delta)(1-\delta)^2}$-ευρωστία. Πιο πρόσφατα, οι ίδιοι συγγραφείς παρουσίασαν στο [21] μια δεύτερη οικογένεια αλγορίθμων, που ονομάζεται LA$_\delta$, και απέδειξαν ότι για κάθε δ $[0, 1)$, ο αλγόριθμος LA$_\delta$

επιτυγχάνει $\frac{3-\delta}{1+\delta}$ συνέπεια και $\frac{3+\delta}{1-\delta}$ αξιοπιστία. Επιπλέον, αυτός είναι ο βέλτιστος συμβιβασμός. Δηλαδή, κανένας ντετερμινιστικός αλγόριθμος που είναι $\frac{3-\delta}{1+\delta}$-συνεπής δεν μπορεί να έχει ευρωστία αυστηρά μικρότερη από $\frac{3+\delta}{1-\delta}$, ακόμη και όταν θεωρούμε τη συνήθη μετρική στην ευθεία και δύο μόνο υποψήφιους.

Η συνέπεια και η ευρωστία αποτυπώνουν δύο ακραίες καταστάσεις όσον αφορά την ακρίβεια της πρόβλεψης, δηλαδή την πλήρη ακρίβεια και την πλήρη αυθαιρεσία, αντίστοιχα. Μια πιο εκλεπτυσμένη ανάλυση, η οποία δίνει φράγματα για την παραμόρφωση ως συνάρτηση του επιπέδου ακρίβειας της πρόβλεψης, μπορεί να επιτευχθεί αν ορίσουμε το *σφάλμα πρόβλεψης*. Η παράμετρος αυτή είναι η απόσταση ανάμεσα στον προβλεπόμενο και τον βέλτιστο υποψήφιο, κανονικοποιημένη με τη βέλτιστη μέση απόσταση $\mathrm{SC}(c^*)/n$. Δηλαδή,

$$\eta := \frac{n \cdot d(p, c^*)}{(\mathrm{SC}(c^*)}.$$

Στο [21] αποδεικνύεται ότι για κάθε $\delta \in [0, 1)$ και κάθε σφάλμα πρόβλεψης η, η παραμόρφωση του LA_δ είναι το πολύ

$$\min\left\{\frac{3-\delta+2\delta\eta}{1+\delta}, \frac{3+\delta}{1-\delta}\right\},$$

ενσωματώνοντας έτσι και αυτή την παράμετρο στην ανάλυση του προβλήματος.

Σημειώνουμε ότι αυτό το φράγμα ανακτά την εγγύηση συνέπειας για $\eta = 0$, και στη συνέχεια αυξάνει γραμμικά με το η μέχρι να φτάσει στο φράγμα ευρωστίας, όπου σταθεροποιείται. Για παράδειγμα, αν υποθέσουμε ότι $\eta \leq 2$ τότε η παραμόρφωση είναι το πολύ 3, για οποιαδήποτε τιμή του δ.

1.4 Παραμόρφωση των κανόνων εκλογής επιτροπής

Στις προηγούμενες ενότητες εξετάσαμε το πρόβλημα μιας εκλογής μοναδικό νικητή και της μετρικής παραμόρφωσής της. Ένα άλλο ενδιαφέρον πρόβλημα που δεν έχει μελετηθεί τόσο πολύ είναι η μετρική παραμόρφωση των εκλογών με πολλαπλούς νικητές, όπου στόχος μας είναι να εκλέξουμε μια επιτροπή από $k \geq 2$ μέλη (μέσα από $m \geq k+1$ υποψήφιους) με βάση τις κατατάξεις προτιμήσεων που δίνονται από n ψηφοφόρους. Όπως και στην περίπτωση του μοναδικού νικητή, οι ψηφοφόροι και οι υποψήφιοι αντιστοιχίζονται με θέσεις σε κάποιον μετρικό χώρο και οι προτιμήσεις των ψηφοφόρων προσδιορίζονται από την απόστασή τους από τις θέσεις των υποψηφίων. Όμως, σε αυτή την περίπτωση υπάρχουν διάφοροι τρόποι για να ορίσουμε το κόστος μιας επιτροπής, γεγονός που οδηγεί σε διαφορετικούς τύπους εκλογών επιτροπής που ικανοποιούν διαφορετικές ιδιότητες [39], [42].

Οι Goel et. al. [50] και Chen et. al. [34] θεώρησαν ότι το κόστος μιας επιτροπής για έναν ψηφοφόρο είναι το άθροισμα των αποστάσεών του από όλα τα μέλη της επιτροπής. Αποδείχθηκε στην [50] ότι εφαρμόζοντας επανειλημμένα k φορές έναν κανόνα εκλογής με μοναδικό νικητή που έχει παραμόρφωση α, επιτυγχάνουμε παραμόρφωση το πολύ ίση με α για το κόστος της εκλογής με πολλαπλούς νικητές. Επομένως, η βέλτιστη παραμόρφωση 3 που εξασφαλίζεται στο [47] για την περίπτωση του μοναδικού νικητή είναι εφικτή και σε αυτό το περιβάλλον. Ωστόσο, αυτή η επιλογή κόστους τείνει να ευνοεί τις "ομοιογενείς" πλειοψηφίες.

Σε αυτή τη διπλωματική εργασία μελετάμε το πλαίσιο όπου το κόστος κάθε ψηφοφόρου για μια επιτροπή ορίζεται ως η απόστασή του από το πλησιέστερο μέλος. Αυτό το πλαίσιο έχει ως κίνητρο τους κανόνες των Chamberlin–Courant [31] και Monroe [68], οι οποίοι στοχεύουν στην εκλογή μιας ποικιλόμορφης επιτροπής που αντιπροσωπεύει καλύτερα το σύνολο του πληθυσμού των ψηφοφόρων. Χρησιμοποιήθηκε επίσης από τους Καραγιάννη et al. [28] για την ανάλυση της *ωφελιμιστικής* παραμόρφωσης των κανόνων με πολλαπλούς νικητές. Στο Κεφάλαιο 5 παρουσιάζουμε αποτελέσματα των Φωτάκη, Gourvès και Πατσιλινάκου [45] οι οποίοι εστιάζουν στην απλούστερη περίπτωση των *γραμμικών προτιμήσεων*, όπου οι ψηφοφόροι και οι υποψήφιοι είναι τοποθετημένοι στην πραγματική ευθεία. Τα αποτελέσματα των Καραγιάννη et. al. στην εργασία [30], όπου μελετάται ένα πιο γενικό πλαίσιο (το κόστος ενός ψηφοφόρου ορίζεται ως η απόστασή του από το q-οστό πλησιέστερο μέλος της επιτροπής) δείχνουν ότι για $q = 1$ η παραμόρφωση είναι μη φραγμένη για κάθε $k \geq 3$ ακόμη και στη γραμμική περίπτωση.

Παρά το αποτέλεσμα αυτό, στην εργασία [45] αποδεικνύεται ότι μπορούμε να χρησιμοποιήσουμε έναν περιορισμένο αριθμό ερωτημάτων προσδιορισμού απόστασης (η ιδέα αυτή προέρχεται από τα [3], [5], [4]) για να επιτύχουμε φραγμένη ή και σταθερή παραμόρφωση σε εκλογές επιτροπές με k-μέλη με γραμμικές προτιμήσεις, για $k \geq 3$. Συγκεκριμένα, επιτυγχάνεται φραγμένη παραμόρφωση με $O(k)$ ερωτήματα και σταθερή παραμόρφωση με $O(k \log n)$ ερωτήματα. Μια παρόμοια προσέγγιση σε αυτό το πρόβλημα (δηλαδή, βασισμένη σε περιορισμένο αριθμό ερωτημάτων προσδιορισμού απόστασης) σε γενικούς μετρικούς χώρους εμφανίζεται στο [73] και στο [26].

1.5 Συνεισφορά

Ο στόχος αυτής της διπλωματικής εργασίας είναι να μελετήσει τους κανόνες ψηφοφορίας για εκλογές με έναν νικητή και εκλογές επιτροπών χρησιμοποιώντας το μαθησιακά ενισχυμένο πλαίσιο, έχοντας ως αφετηρία την επιτυχημένη εφαρμογή αυτού του πλαισίου στην [21]. Τα αποτελέσματά μας παρουσιάζονται στο Κεφάλαιο 6.

Η πρώτη μας συμβολή είναι μια συζήτηση για τη συνέπεια και την ευρωστία κατάλληλα ορισμένων ενισχυμένων εκδοχών των αλγορίθμων Plurality και Borda. Στην περίπτωση του Plurality, για κάθε $\delta \in [0, 1)$ ορίζουμε έναν αλγόριθμο $\text{BoostedPlurality}_\delta$ ο οποίος χρησιμοποιεί μια πρόβλεψη p για τον βέλτιστο υποψήφιο και εξαρτάται από την παράμετρο εμπιστοσύνης δ. Επιλέγει είτε τον υποψήφιο με το μεγαλύτερο σκορ πλειοψηφίας ή τον p αν το σκορ πλειοψηφίας του είναι αρκετά υψηλό. Αποδεικνύουμε ότι ο αλγόριθμος $\text{BoostedPlurality}_\delta$ έχει ευρωστία

$$\text{robustness}(\text{BoostedPlurality}_\delta) \leq \frac{2m}{1-\delta} - 1$$

και συνέπεια

$$\text{consistency}(\text{BoostedPlurality}_\delta) \geq 2m - 1 - 2\delta.$$

Αν $\delta = 0$ τότε το κάτω φράγμα για τη συνέπεια ταιριάζει με την παραμόρφωση του απλού αλγορίθμου Plurality. Όταν $\delta \to 1$, το κάτω φράγμα για τη συνέπεια τείνει στο $2m - 3$, που σημαίνει ότι η πρόβλεψη δεν εξασφαλίζει σχεδόν καμία βελτίωση στην παραμόρφωση του αλγορίθμου Plurality. Τα αποτελέσματα είναι παρόμοια για την μαθησιακά-ενισχυμένη εκδοχή του κανόνα Borda, $\text{BoostedBorda}_\delta$. Δίνουμε επίσης φράγματα για την παραμόρφωση που επιτυγχάνεται από τους $\text{BoostedPlurality}_\delta$ και $\text{BoostedBorda}_\delta$ σε περιπτώσεις όπου η πρόβλεψη p έχει ένα δεδομένο σφάλμα η. Αυτά τα αποτελέσματα υποδεικνύουν ότι οι προβλέψεις δεν είναι χρήσιμες όταν η μόνη διαθέσιμη πληροφορία είναι το σκορ πλειοψηφίας ή το σκορ Borda των υποψηφίων.

Στη συνέχεια, εστιάζουμε στο πρόβλημα της εκλογής επιτροπής. Υποθέτουμε ότι οι αλγόριθμοί μας έχουν πρόσβαση σε μια πρόβλεψη $\mathcal{P} = \{p_1, \ldots, p_k\} \subseteq C$ για τη βέλτιστη επιτροπή. Ο πρώτος αλγόριθμος που εξετάζουμε, είναι μια μαθησιακά ενισχυμένη έκδοση του Greedy αλγορίθμου από το [45], με παράμετρο το $\delta \in [0, 1)$. Σε κάθε επανάληψη εκλέγει στην επιτροπή είτε τον πιο απομακρυσμένο μέχρι στιγμής υποψήφιο ή τον τον πιο απομακρυσμένο μέχρι στιγμής υποψήφιο *που έχει προταθεί*. Χρησιμοποιεί $\Theta(k)$ ερωτήσεις για πραγματικές αποστάσεις, αλλά για κάθε $\delta \in [0, 1)$ η συνέπειά του εξακολουθεί να είναι $\Omega(n)$.

Στη συνέχεια, παρουσιάζουμε έναν αλγόριθμο ο οποίος, χρησιμοποιώντας μια πρόβλεψη $\mathcal{P}$ και τον Greedy αλγόριθμο, υπολογίζει ένα **καλό** αντιπροσωπευτικό σύνολο υποψηφίων και εκλέγει τη βέλτιστη k-επιτροπή στο περιορισμένο παράδειγμα που επάγεται από αυτό το σύνολο. Ο αλγόριθμός μας επιτυγχάνει σταθερή συνέπεια και γραμμική ευρωστία με $O(k)$ ερωτήσεις για πραγματικές αποστάσεις. Συγκεκριμένα, επιτυγχάνει συνέπεια το πολύ 3 και ευρωστία το πολύ $10n + 1$.

CHAPTER 2

Introduction

The field of social choice theory [25] studies how the individual preferences of people from a group can be combined into a single collective decision. Although this is not the only application, we will use the terminology of an election where the participants are referred to as *voters*, the possible alternatives are referred to as *candidates* and the function that chooses a candidate based on voters' preferences is referred to as a *voting rule*. We assume that the desire a voter has for a certain candidate to be elected is quantified by a *cardinal* utility value. So, a reasonable objective is to elect the candidate that maximizes the (utilitarian) *social welfare*; that is the sum of the utilities over all voters. Solving this problem is trivial; we just have to compute the social welfare for all candidates and output the one who achieves the maximum. However, usually we do not have access to the exact cardinal values since it is difficult even for the voters to evaluate them due to cognitive reasons. Instead, we require from the voters only *ordinal* information in the form of rankings of the candidates based on comparisons. Voting rules that use only ordinal information are unable to guarantee finding an optimal solution for the cardinal objective of maximizing the social welfare. Thus, similarly to approximation algorithms [79] or online algorithms where the available information is limited [22], we aim to approximate the optimal solution. The *distortion* of a voting rule, which was introduced by Procaccia and Rosenchein [72], is the worst-case approximation ratio between the maximum social welfare and the social welfare of the candidate elected by the voting rule. Distortion has been used to measure the performance of various voting rules and new rules have been created in order to minimize distortion as much as possible in different settings.

In the general case, there are strong impossibility results regarding the distortion of voting rules. Therefore, Anshelevich et al. [8] added the assumption that voters and candidates are located in a metric space in order to prove more meaningful distortion bounds. In this setting, voters prefer candidates that are closer to them compared to those that are farther and the distance between a voter and a candidate can be viewed as an induced cost. Instead of maximizing the social welfare, we want to elect the candidate that minimizes the social cost, which is the sum of distances to all voters. In this thesis we will focus on the metric case.

The information limitations that are present in the metric distortion problem make it suitable for an application of the new *learning-augmented framework*. In this framework, algorithms (in this case voting rules) are enhanced with a *prediction*, which is used in order to improve their performance. This prediction, which can take different forms, can be obtained from a machine learning model, using relevant historical data. The algorithm is simultaneously evaluated based on its performance when the prediction is accurate (known as its *consistency*) as well as its performance when the prediction can be arbitrarily inaccurate (known as its *robustness*). Our goal in this thesis is to improve metric distortion bounds using learning-augmented algorithms, mainly in the multi-winner case, where instead of a single candidate we elect a committee of k members.

2.1 Voting rules and distortion

In classical social choice theory, a voting rule takes as an input from each voter their linear ordering of the candidates and outputs a winning candidate. This model corresponds to the way humans usually express their preferences over alternatives by ranking them instead of assigning a precise utility value to them. Due to the lack of numerical values measuring the quality of the elected candidate, one reasonable approach to evaluating voting rules is the *axiomatic approach*. In this approach, one formulates certain axiomatic properties or criteria that voting rules should satisfy. Then, one can choose a rule that has a desirable property over another rule that does not. Some important early papers in this area are those of Arrow [14], May [65], Gibbard [46], Satterthwaite [75] and Young [83] (see also the survey by Zwicker [84]).

In this thesis we will evaluate voting rules following the *utilitarian approach* which is used in game theory [80] and algorithmic mechanism design [70]. Utilitarianism, founded by Bentham, argues that the 'happiness' a certain person obtains from a certain state of the world can be quantified with a *utility function* and aims to maximize the overall 'happiness' of the population by maximizing the sum of individual utilities, i.e. the social welfare. Boutilier et al. [24] point out that although not all social choice problems are suitable to the utilitarian approach (for example there are cases where interpersonal comparison of utilities is not possible), there are many real-life settings that fit the utilitarian view. For example, in recommender systems and many similar domains from mechanism design and e-commerce, the computational agents typically assign real-valued utilities to alternatives rather than have ordinal preferences over the set of candidates. Still, it is assumed that utilities are latent and similarly to classical social choice, voting rules have access only to ordinal rankings of candidates. These rankings are aligned with the cardinal utilities, meaning that a candidate with higher utility ranks higher compared to one with lower utility. This limitation is justified from behavorial economists who have shown that it is cognitively difficult to assign precise values to alternatives.

Any given voting rule that uses ranked ballots is not always able to identify a candidate that maximizes the social welfare. Therefore, we may think of a voting rule as an approximation algorithm that tries to choose the best possible candidate based on limited information (ordinal preferences instead of utilities). This perspective on voting with ranked ballots was proposed by Procaccia and Rosenschein in [72] who introduced the term distortion to refer to the quality of approximation provided by a voting rule. The distortion of a voting rule is the worst-case ratio of the maximum social welfare over the social welfare of the candidate that is elected. The notion of distortion provides a quantitative way to compare various voting rules: low distortion is obviously a desirable feature. A recent survey on the most important results on distortion can be found in [9].

Let V be the set of voters, $|V| = n$, and let C be the set of candidates, $|C| = m$. For arbitrary utility values, distortion is unbounded even for small numbers of voters and candidates. Therefore, Procaccia and Rosenchein [72] assumed normalized utilities where the utility sum for each voter is fixed, e.g. equal to 1. Even with this assumption, they showed that the popular voting rules Borda and Veto have unbounded distortion [72]. Under the same normalized utilities assumption, Caragiannis and Procaccia [29] showed that the simple Plurality rule, which elects the candidate ranked first by the largest number of voters, has distortion $O(m^2)$. Caragiannis et al. [28] showed that distortion for deterministic voting rules has a lower bound of $\Omega(m^2)$, hence Plurality is optimal. By allowing randomization in voting rules we can further improve distortion bounds, where the social welfare of the candidate elected in the definition of distortion is replaced by the *expected* social welfare. Boutilier et al. [24], also assuming normalized utilities, proved a lower bound of $\Omega(\sqrt{m})$ for all randomized voting rules and designed a randomized voting rule with distortion $O(\sqrt{m} \cdot \log^* m)$, nearly matching the lower bound. Other variants of distortion have been studied, e.g. communication complexity [63], [64] and distortion in distributed settings [44].

2.2 Metric distortion

In this thesis, we will focus on the *metric distortion* framework, introduced by Anshelevich et al. [8], which assumes that voters and candidates are located in a metric space. It is natural to assume that voters prefer candidates in their proximity. Therefore, in this setting, a voter's desire to minimize her distance (which can be viewed as a cost) from the elected candidate corresponds to her desire to maximize utility in the utilitarian setting. This approach is similar to spatial models of voting from political science [40], [13], [66], [35], [76]. We can think that a voter's opinion on an issue of interest can be represented with a coordinate in a Euclidean space, e.g. the simple one-dimensional left-right axis. However, the metric spaces considered here are more general and thus more powerful.

In Chapter 3 we present results from [8] about the metric distortion of common voting rules. We start with the formal definition of the notions that will be used. Let V and C be two finite sets. We say that V is the set of *voters* and C is the set of *candidates*. We set $n = |V|$ (the number of voters) and $m = |C|$ (the number of candidates). In what follows, we usually denote voters by u, v and candidates by c, x, y. We assume that both V and C are located in a metric space (X, d), so that the distance $d(a, b)$ between a and b is defined for all $a, b \in V \cup C$.

The *social cost* of a candidate $c \in C$ with respect to the metric d is the sum

$$\mathrm{SC}(c, d) = \sum_{v \in V} d(v, c).$$

We write $\mathrm{SC}(c)$ when the metric d is clear from the context.

A triplet (V, C, d) as above is called an *instance*. The distance $d(v, c)$ between a voter v and a candidate c is a measure of how much v prefers c. We say that $v \in V$ prefers $c \in C$ over $c' \in C$ if $d(v, c) \leq d(v, c')$, Then we write

$$c \succcurlyeq_v c'.$$

Any given instance (V, C, d) induces preference rankings for each voter. For each voter $v \in V$ we have a *preference ranking* $\sigma_v : C \to \{1, \dots, m\}$ such that $\sigma_v(c) < \sigma_v(c')$ implies $c \succcurlyeq_v c'$ and then we say that d is *aligned* with σ_v and we write $d \ \triangleright \ \sigma_v$.

A *preference profile* $\sigma := (\sigma_v)_{v \in V}$ is an n-tuple of preference rankings for each voter. We say that d is aligned with the preference profile σ, and we write $d \ \triangleright \ \sigma$, if $d \ \triangleright \ \sigma_v$ for all $v \in V$.

The metric distortion problem may be described as follows: An algorithm ALG receives as input a preference profile σ which is induced by an instance (V, C, d). The algorithm does not have access to the underlying distance function d. The goal is to output a candidate c^* that minimizes the social cost, i.e.

$$\mathrm{SC}(c^*) = \min_{c \in C} \mathrm{SC}(c).$$

An algorithm for this problem will also be referred to as a *voting rule* and its output as the *winner* of the voting rule. The *distortion* of ALG is the worst-case multiplicative approximation it achieves to the optimum, i.e.

$$\mathrm{distortion}(\mathrm{ALG}) = \sup_{\sigma} \sup_{d : d \, \triangleright \, \sigma} \frac{\mathrm{SC}(\mathrm{ALG}(\sigma), d)}{\mathrm{SC}(c^*(d), d)}.$$

A first main question to address is how well can *any* voting rule perform. We present a first theorem from [8] that provides a lower bound on the performance of all deterministic voting rules: no such rule can have distortion better than 3.

Then, we provide lower and bounds from [8] for the distortion of many well-known voting rules, starting with some classical positional scoring rules. A *positional scoring rule* $f_{\vec{s}}$ is determined by a scoring vector $\vec{s} = (s_1, s_2, \dots, s_m)$, where $s_i \in \mathbb{Q}$, $s_1 \geq s_2 \geq \dots \geq s_m$ and $s_1 > s_m$. Whenever a voter $v \in V$ ranks a candidate $c \in C$ in position ℓ, then the candidate c receives $r_v(c) = s_\ell$ points from v. Then, the total score of the candidate c is the sum

$$\text{score}_{\vec{s}}(c) = \sum_{v \in V} r_v(c).$$

The winners according to the scoring vector $\vec{s}$ are the candidates $c^* \in C$ for which

$$\text{score}_{\vec{s}}(c^*) = \max_{c \in C} \text{score}_{\vec{s}}(c).$$

We discuss the following positional scoring rules:

(i) Plurality, where $\vec{s} = (1, 0, \dots, 0)$.

(ii) Veto, where $\vec{s} = (1, 1, \dots, 1, 0)$.

(iii) Borda, where $\vec{s} = (m-1, m-2, \dots, 1, 0)$.

(iv) Harmonic rule, where $\vec{s} = \left(1, \frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{m}\right)$.

(v) k-Approval $(1 \leq k < m)$, where $\vec{s} = (1, 1, \dots, 1\,[k], 0, \dots, 0\,[m-k])$.

First we provide bounds from [8] for common positional scoring rules. For k-Approval with $k > 1$ and for Veto, the distortion is unbounded. For Plurality and Borda, the distortion is $2m - 1$, where m is the number of candidates, with the upper bound matching the lower bound. Moreover, there is no positional scoring rule whose distortion is bounded by a constant independent of m. For each value of m, the distortion of every positional scoring rule for m candidates is at least $1 + 2\sqrt{\ln m - 1}$.

It is interesting to note that neither Plurality nor Borda achieve the best possible distortion in the class of positional scoring rules. The Harmonic rule of Boutilier et al. [24] has better distortion than either Plurality or Borda. However, the distortion of the Harmonic rule is still almost linear in m, with a lower bound of $\Omega\left(\frac{m}{\ln m}\right)$ and an upper bound of $O\left(\frac{m}{\sqrt{\ln m}}\right)$.

Subsequently, it is shown in [8] that, while some commonly used rules have high distortion, there are important voting rules for which distortion is bounded by a small constant or grows slowly with the number of candidates. An example is given by the known Copeland rule. The *Copeland score* of a candidate c is the number of candidates that c defeats. The Copeland rule [36] outputs the candidate that has maximum Copeland score. Anshelevich et al. [8] proved that the Copeland rule has distortion 5, with the upper bound matching the lower bound. This means that, although the Copeland rule knows nothing about the metric costs other than the ordinal preferences induced by them, and cannot possibly find the true optimal alternative, it nevertheless always selects a candidate whose quality is only a factor of 5 away from optimal. Recall that, by the general lower bound of 3, no deterministic voting rule can do much better than Copeland for the social cost.

Another popular voting rule, known as single transferable vote (STV), was studied in [77] and later in [6]. Single Transferable Vote (STV), which is used in real-life elections, is an iterative rule which is defined in the following way. In each round, the candidate that is ranked first by the fewest voters (the one with the lowest Plurality score) is removed from the set of candidates and from the rankings of the voters, which means that the Plurality scores have to be computed again. After $m-1$ rounds there is only one candidate left, and this is the winner. Elkind and Skowron [77] proved an upper bound of $O(\ln m)$ and a lower bound of $\Omega(\sqrt{\ln m})$ for the distortion of STV , which means that STV is not as good as the Copeland rule.

The results above are summarized in the following Table.

Table 1. Upper and lower bounds for the distortion

Rule	**Upper bound**	**Lower bound**
k-Approval, $k > 1$	∞	∞
Veto	∞	∞
Plurality	$2m-1$	$2m-1$
Borda	$2m-1$	$2m-1$
Harmonic Rule	$O\left(\frac{m}{\sqrt{\ln m}}\right)$	$\Omega\left(\frac{m}{\ln m}\right)$
Scoring Rules	∞	$1+2\sqrt{\ln m - 1}$
Copeland	5	5
STV	$O(\ln m)$	$\Omega(\sqrt{\ln m})$

Anshelevich et al. [8] conjectured that the Ranked Pairs rule [78] can achieve the optimal distortion of 3 but this was later disproved by Goel et al. [51]. Subsequently, Munagala and Wang [69] presented a deterministic algorithm that achieves a distortion of $2+\sqrt{5} \approx 4.236$. In fact, this algorithm does not require access to the full ordinal preferences of the voters; it makes use of a limited input, the weighted tournament graph, which contains, for every pair of candidates, the number of voters that prefer one over another. Based on this, they conjectured that if full ordinal preferences are provided then the best possible distortion should be 3. The same authors described a sufficient condition that would imply the optimal distortion 3. Around the same time, Kempe [56] obtained independently the same bound of $2+\sqrt{5}$ using a linear programming duality framework, provided several alternative formulations of the sufficient condition of [69] and proved that Ranked Pairs has distortion $\Theta(\sqrt{m})$.

In Section 3.5 we present a theorem of Gkatzelis, Halpern and Shah [47], who introduced a deterministic algorithm that concludes this line of research and guarantees the optimal distortion of 3. The main technical ingredient in this work is the “ranking-matching lemma” which establishes the existence of fractional perfect matchings within a family of vertex-weighted bipartite graphs induced by an instance of the problem.

Subsequent works have provided an even better understanding of the problem. Kizilkaya and Kempe [58] achieved the optimal distortion with the Plurality Veto algorithm, which is much simpler when compared to Plurality Matching, and afterwards the same authors [59] proposed a refinement of Plurality Veto called Simultaneous Veto that resolves the arbitrary tie-breaking issues in [47] and [58] while maintaining optimal distortion.

After the final answer to the distortion problem for the class of deterministic algorithms, an open problem is to also achieve optimal distortion bounds for *randomized* algorithms. Anshelevich and Postl [11] provided a lower bound of 2 and proved that the Random Dictatorship algorithm has distortion lower than 3, but it tends to 3 as the number of voters grows. Subsequently, Fain et. al. [41] and Kempe [57] provided other algorithms with distortion that tends to 3, yet as the number of candidates grows.Recently, the best known lower bound of 2.1126 was shown by Charikar and Ramakrishnan [32] and an upper bound of 2.753 was achieved by Charikar et. al. [33], breaking away from 3 for the first time. Closing the remaining gap is still an open question.

The article [49] tries to combine metric distortion bounds with distortion bounds for the utilitarian setting (where the voters' preferences are arbitrary normalized valuations). In [1] a variant of the problem is studied, where more information besides the usual preference rankings is available, and in [57] one can find a study of the trade-offs between the achievable distortion and the communication complexity of deterministic social choice algorithms. In [10] metric distortion is studied in distributed settings and in [43] the distortion of *truthful* mechanisms is examined. Different settings where the information available is more limited than the full preference rankings have also been studied [11], [52], [41], [57], [23], [7].

2.3 Learning-augmented algorithms

Recent research has increasingly focused on a novel analytical framework known as *algorithms with predictions* or *learning-augmented algorithms*. This approach leverages machine-learned predictions to refine performance bounds and overcome the limitations of traditional worst-case analysis. The form of the predictions may vary depending on the type of data available to the designer and the problem to be solved. By incorporating predictions, the algorithm can improve its performance by using them as a guide.

However, the accuracy of these predictions is uncertain and cannot be fully trusted. The algorithm is evaluated based on two criteria: its performance when the prediction is correct (referred to as *consistency*) and its performance when the prediction is highly inaccurate (referred to as *robustness*). In an extreme scenario where the prediction is followed blindly, good predictions will lead to excellent outcomes (high consistency), while inaccurate predictions will result in poor performance (low robustness). On the other hand, completely ignoring the prediction weakens the consistency. Typically, each learning-augmented algorithm strikes a balance between robustness and consistency, with the goal of identifying the optimal trade-off between the two.

In recent years, the learning-augmented framework has been extensively used in the design and analysis of algorithms to address the limitations of overly pessimistic worst-case bounds. Several recent works have applied this framework to classical algorithmic challenges, including online paging [62], scheduling [74], secretary problems [37], [12], optimization problems with covering [18], and knapsack constraints [54], as well as various graph problems [15]. Early contributions to this field are discussed in the survey [67], and [60] offers an up-to-date compilation of relevant papers. The learning-augmented approach has also been applied to social choice problems, where predictions are made about a group of agents and their preferences, with the objective of optimizing social cost or welfare functions. This includes research on resource allocation problems where agents' preferences are revealed incrementally in an online setting (see [19] and [20]), as well as scenarios where agents are strategic and their preferences are private (see [2], [82], [48], [16], [55], [71], [61], [27] and [17]).

Considering the metric distortion problem under this framework is reasonable because of the information limitations it contains. In addition to voter rankings, it is logical to assume that we may have access to historical data on voters' decisions in related matters, which could be correlated with their preferences for the issue at hand. This additional information can provide insights into the likely outcome.

In the article [21], Berger et al. initiate the analysis of the metric distortion problem using the learning-augmented framework and characterize the robustness-consistency Pareto frontier for this problem. They consider algorithms that receive as input a pair (σ, p), where σ is a preference profile and $p \in C$ is a *prediction* about the optimal candidate $c^*(d)$. The goal is to evaluate the performance of an algorithm through its *consistency* and its *robustness*. The consistency of ALG is defined as the distortion that ALG guarantees when the provided prediction is accurate, i.e. $p = c^*(d)$. The robustness of ALG is defined as the distortion that ALG guarantees with an arbitrary prediction, independently of how accurate this prediction may be.

In Chapter 4 we study the work of Berger, Feldman, Gkatzelis and Tan, who introduced a family of algorithms, which is parameterized by $0 \leq \delta < 1$, known as BoostedSV$_\delta$, and obtained consistency and robustness bounds in terms of δ. The family BoostedSV$_\delta$ represents a learning-augmented adaptation of the SimultaneousVeto algorithm proposed in [59]. The SimultaneousVeto algorithm starts by assigning each candidate $c \in C$ a score equal to the number of voters ranking c first (its plurality score). Voters then continuously and simultaneously reduce the score of their least preferred candidate among those with remaining positive scores. The candidate whose score reaches zero last is selected as the winner. The BoostedSV$_\delta$ algorithm enhances SimultaneousVeto by boosting the initial score of the candidate $p \in C$ predicted to be optimal. The size of this boost is a carefully calibrated increasing function of δ, which also adjusts the rate at which voters reduce scores. As δ increases—indicating greater confidence in the prediction—the size of the boost grows. It is proved that the algorithm BoostedSV$_\delta$ achieves $\frac{3-\delta}{1+\delta}$-consistency and $\frac{3+\delta+13\delta^2-\delta^3}{(1+\delta)(1-\delta)^2}$-robustness. More recently, the authors presented in [21] a second family of algorithms, called LA$_\delta$, and proved that for any $\delta \in [0, 1)$, the algorithm LA$_\delta$ achieves $\frac{3-\delta}{1+\delta}$-consistency and $\frac{3+\delta}{1-\delta}$-robustness. Moreover, they showed that this is the optimal trade-off. Namely, no deterministic algorithm that is $\frac{3-\delta}{1+\delta}$-consistent can be strictly better than $\frac{3+\delta}{1-\delta}$-robust, even for the line metric and just two candidates.

Consistency and robustness represent two extreme cases of prediction accuracy: perfect accuracy and total unpredictability. A more detailed analysis, which establishes bounds on the distortion based on the prediction's accuracy level, can be conducted by introducing the concept of *prediction error*. This error is defined as the distance between the predicted and optimal candidates, normalized by the optimal average distance, given by $\mathrm{SC}(c^*)/n$. Specifically, we define the prediction error as

$$\eta := \frac{n \cdot d(p, c^*)}{\mathrm{SC}(c^*)}.$$

It is shown in [21] that for any $\delta \in [0, 1)$ and prediction error η, the distortion of LA$_\delta$ is at most

$$\min\left\{\frac{3 - \delta + 2\delta\eta}{1 + \delta}, \frac{3 + \delta}{1 - \delta}\right\},$$

thus incorporating this additional parameter into the analysis of the problem.

It is interesting to note that the above bound recovers the consistency guarantee when $\eta = 0$. As η increases, the bound grows linearly with the prediction error until it reaches the robustness threshold, at which point it flattens.

2.4 Distortion of committee election

In the previous sections we considered the problem of a single-winner election and its metric distortion. Another interesting problem that has not been studied as much is the metric distortion of multi-winner elections, where we aim to elect a committee of $k \geq 2$ members (out of $m \geq k+1$ candidates) based on ordinal preferences provided by n voters. Similarly to the single-winner case, the voters and candidates are associated with locations in a metric space and the voters' cardinal preferences correspond to their distances from the candidates' locations. However, in this case there are multiple ways to define the cost of a committee, leading to different types of committee elections that satisfy different properties [39], [42].

Goel et al. [50] and Chen et al. [34] considered the cost of a committee for a voter to be the sum of her distances to all committee members. It was proved in [50] that by repeatedly applying k times a single-winner rule with distortion α, a distortion of at most α for the multi-winner cost is achieved. Therefore, the optimal distortion of 3 [47] for the single-winner case is also possible in this setting. However, this cost selection tends to favor "homogeneous" majorities.

Here we consider the setting where the cost of each voter for a committee is defined as her distance to the nearest member. This setting is motivated by the rules of Chamberlin and Courant [31] and Monroe [68], which aim at electing a diverse committee that best represents the entirety of the voters' population. It was also used by Caragiannis et al. [28] for the analysis of the *utilitarian* distortion of multi-winner rules. In Chapter 5, we present results of Fotakis, Gourvès and Patsilinakos [45] who focus on the simplest case of *linear preferences*, where the voters and candidates are embedded in the real line. The results of Caragiannis et al. [30], where a more general setting is studied (the cost of a voter is defined as her distance to the q-th nearest member of the committee) imply that for $q = 1$ the distortion is $\Theta(n)$ for $k = 2$ and unbounded for all $k \geq 3$ even in the linear case.

Despite this result, it is shown in [45] that one can use a restricted amount of cardinal distance queries (inspired by [3], [5], [4]) to achieve bounded or even constant distortion in k-committee election with linear preferences, for $k \geq 3$. Specifically, bounded distortion is achieved with $O(k)$ queries and constant distortion is achieved with $O(k \log n)$ queries. A similar approach to the related k-median and k-center problems in general metric spaces, i.e. the availability of limited cardinal queries in addition to the ordinal rankings, appears in [73] and in [26].

2.5 Contribution

The goal of this thesis is to study voting rules for single-winner and committee elections under the learning-augmented framework, inspired by the successful application of the framework in [21]. Our results are presented in Chapter 6.

Our first contribution is a discussion of consistency and robustness bounds for suitably defined boosted versions of the Plurality and Borda rules. In the case of the plurality rule, for any $\delta \in [0, 1)$ we define the algorithm $\text{BoostedPlurality}_\delta$ that uses a prediction p for the optimal candidate and depends on the confidence parameter δ. It elects either the candidate with the highest plurality score or p if his plurality score is high enough. We show that $\text{BoostedPlurality}_\delta$ has robustness

$$\text{robustness}(\text{BoostedPlurality}_\delta) \leq \frac{2m}{1-\delta} - 1$$

and consistency

$$\text{consistency}(\text{BoostedPlurality}_\delta) \geq 2m - 1 - 2\delta.$$

If $\delta = 0$ then the lower bound for the consistency matches the distortion of the simple plurality algorithm. When $\delta \to 1$ the lower bound for the consistency tends to $2m - 3$, meaning that the prediction does not give nearly any improvement in the distortion of the plurality algorithm. The results are similar for the learning-augmented version of the Borda rule, $\text{BoostedBorda}_\delta$. We also provide bounds on the distortion achieved by $\text{BoostedPlurality}_\delta$ and $\text{BoostedBorda}_\delta$ on instances where the prediction p has a given error η. These results indicate that predictions are not useful when the only information available is the plurality score or the Borda score of the candidates.

Subsequently, we focus on the committee election problem. We assume that our algorithms have access to a prediction $\mathcal{P} = \{p_1, \ldots, p_k\} \subseteq C$ for the optimal committee. The first algorithm that we examine, is a learning-augmented version of the Greedy algorithm from [45], parameterized by $\delta \in [0, 1)$. At each iteration it elects in the committee either the most distant candidate thus far or the most distant *predicted* candidate thus far. It uses $\Theta(k)$ distance queries, but for all $\delta \in [0, 1)$ its consistency is still $\Omega(n)$.

Afterwards, we introduce an algorithm which, using a prediction $\mathcal{P}$ and the Greedy algorithm, computes a **good** representative set of candidates and elects the optimal k-committee in the restricted instance induced by this set. Our algorithm achieves constant consistency and linear robustness with $O(k)$ distance queries. Specifically, it achieves a consistency of at most 3 and a robustness of at most $10n + 1$.

CHAPTER 3

Metric distortion

In this chapter we discuss voting under metric preferences. Both voters and candidates are associated with points in a metric space, and each voter prefers candidates that are closer to her to ones that are further away. The goal is to select a candidate that minimizes the social cost, i.e. the sum of distances to the voters. A measure of the quality of a voting rule is its distortion, defined as the worst-case ratio between the performance of a candidate selected by the rule and that of an optimal candidate, i.e. a candidate with minimal social cost. Thus, distortion measures how good a voting rule is at approximating a candidate with minimum social cost, while using only ordinal preference information. The underlying costs can be arbitrary, implicit, and unknown; our only assumption is that they form a metric space. We describe a number of results of Anshelevich, Bhardwaj, Elkind and Postl from [8]. The first one is a lower bound 3 on the distortion of any deterministic voting rule. Then, an analysis of the distortion of positional scoring rules which shows that the distortion cannot be bounded above by a constant, and for several popular rules in this family distortion is linear in the number of candidates or even unbounded independently of the number of candidates. On the other hand, rules that select from the so-called uncovered set achieve a small constant-factor approximation to the optimal candidate; this class of rules includes the well-known Copeland rule for which it is shown that distortion is bounded by a factor of 5. For Single Transferable Vote (STV) the main results are an upper bound of $O(\ln m)$, where m is the number of candidates, as well as a lower bound of $\Omega(\sqrt{\ln m})$.

It was conjectured in [8] that the optimal deterministic algorithm has distortion 3. This conjecture was confirmed in [47] by Gkatzelis, Halpern and Shah who provided a polynomial-time algorithm that achieves distortion 3, matching the known lower bound. We present the proof which is based on a novel lemma, the ranking-matching lemma, about matching voters to candidates. This lemma induces a family of novel algorithms, and we will see that a special algorithm in this family achieves distortion 3.

3.1 Definitions and preliminaries

Let V and C be two finite sets. We say that V is the set of *voters* and C is the set of *candidates*. We set $n = |V|$ (the number of voters) and $m = |C|$ (the number of candidates). In what follows, we usually denote voters by u, v and candidates by c, x, y.

Recall the definition of a metric. If X is a non-empty set then a metric on X is a function $d : X \times X \to \mathbb{R}^+$ with the following properties:

(i) $d(a,b) \geq 0$ for all $a,b \in X$, and $d(a,b) = 0$ if and only if $a = b$.

(ii) $d(b,a) = d(a,b)$ for all $a,b \in X$.

(iii) $d(a,c) \leq d(a,b) + d(b,c)$ for all $a,b,c \in X$.

Then, we say that (X,d) is a metric space.

We assume that both V and C are located in a metric space (X,d), so that the distance $d(a,b)$ between a and b is defined for all $a,b \in V \cup C$.

Definition 3.1.1 (Social cost)**.** The *social cost* of a candidate $c \in C$ with respect to the metric d is the sum

$$\mathrm{SC}(c,d) = \sum_{v \in V} d(v,c).$$

We write $\mathrm{SC}(c)$ when the metric d is clear from the context.

Definition 3.1.2 (Preference ranking)**.** A triplet (V,C,d) as above is called an *instance*. The distance $d(v,c)$ between a voter v and a candidate c is a measure of how much v prefers c. We say that $v \in V$ prefers $c \in C$ over $c' \in C$ if $d(v,c) \leq d(v,c')$, Then we write

$$c \succcurlyeq_v c'.$$

Any given instance (V,C,d) induces preference rankings to the voters. For each voter $v \in V$ we have a *preference ranking* $\sigma_v : C \to \{1,\dots,m\}$ such that $\sigma_v(c) < \sigma_v(c')$ implies $c \succcurlyeq_v c'$ and then we say that d is *aligned* with σ_v and we write $d \ \triangleright\ \sigma_v$. Note that σ_v is not always determined uniquely, since it may happen that $d(v,c) = d(v,c')$ for some $c \neq c'$ in C.

Definition 3.1.3 (Preference profile)**.** A *preference profile* $\sigma := (\sigma_v)_{v \in V}$ is an n-tuple of preference rankings to the voters. We say that d is aligned with the preference profile σ, and we write $d \ \triangleright\ \sigma$, if $d \ \triangleright\ \sigma_v$ for all $v \in V$.

We are now ready to describe the **metric distortion problem**: An algorithm ALG receives as input a preference profile σ which is induced by an instance (V,C,d). The algorithm does not have access to the underlying distance function d. The goal is to output a candidate $c^*(d)$ that minimizes the social cost, i.e.

$$\mathrm{SC}(c^*(d)) = \min_{c \in C} \mathrm{SC}(c,d).$$

An algorithm for this problem will also be referred to as a *voting rule* and its output as the *winner* of the voting rule. The *distortion* of ALG is the worst-case multiplicative approximation it achieves to the optimum, i.e.

$$\mathrm{distortion}(\mathrm{ALG}) = \sup_{\sigma} \sup_{d : d \,\triangleright\, \sigma} \frac{\mathrm{SC}(\mathrm{ALG}(\sigma),d)}{\mathrm{SC}(c^*(d),d)}.$$

An important first question that we can pose is how well can *any* algorithm perform. The next theorem from [8] shows that one cannot hope to approximate the optimal candidate within a factor better than 3.

Theorem 3.1.4 (Anshelevich-Bhardwaj-Elkind-Postl)**.** *Any deterministic algorithm has worst-case distortion at least* 3 *for the social cost.*

Proof. We examine the case where there are only two candidates, x and w. Suppose that half of the voters prefer x to w and half of the voters prefer w to x. This means that there exists a partition of V into two sets V_1, V_2 such that $|V_1| = |V_2| = \frac{n}{2}$ and the preference ranking σ is defined as follows:

- For $v \in V_1$, $\sigma_v(x) = 1$ and $\sigma_v(w) = 2$.
- For $v \in V_2$, $\sigma_v(x) = 2$ and $\sigma_v(w) = 1$.

Now, consider an algorithm ALG and without loss of generality assume that w is its output. We define a metric d aligned to σ as follows: All $n/2$ voters v who prefer x to w satisfy $d(v,x) = 0$ and $d(v,w) = 2$. On the other hand, all $n/2$ voters v who prefer w to x satisfy $d(v,x) = 1+\epsilon$ and $d(v,w) = 1-\epsilon$ for some small $\epsilon > 0$ (they are approximately halfway between x and w).

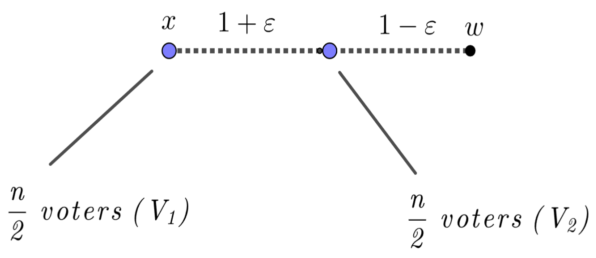

Figure 3.1: Example for Theorem 3.1.4.

We compute

$$\mathrm{SC}(x,d) = \sum_{v \in V} d(v,x) = (1+\epsilon)\,\frac{n}{2}$$

and

$$\mathrm{SC}(w,d) = \sum_{v \in V} d(v,w) = 2\,\frac{n}{2} + (1-\epsilon)\,\frac{n}{2}.$$

This shows that

$$\mathrm{distortion(ALG)} \geq \frac{2n + (1-\epsilon)n}{(1+\epsilon)n} = \frac{3-\epsilon}{1+\epsilon},$$

which tends to 3 as $\epsilon \to 0$. □

Remark: In the example of the proof of Theorem 3.1.4 we allowed $d(x,v) = 0$ while $x \neq v$, that is, d was not a metric but a pseudometric. This is not an essential transgression, since we could have replaced 0 by any sufficiently small quantity without changing the essence of the example. In the sequel, we will often make use of this convention.

It will be useful to fix a notation for the following sets of voters:

(i) $xy = \{v \in V : x \succcurlyeq_v y\}$ is the set of voters v which prefer x over y.

(ii) $xyz = \{v \in V : x \succcurlyeq_v y \succcurlyeq_v z\}$ is the set of voters v which prefer x over y and y over z.

We shall make frequent use of the next two elementary lemmas that appear in [8].

Lemma 3.1.5. *Let $c, x, y, z \in C$. Then, the following hold:*

(i) *For all $v \in cx$ we have that $d(x,c) \leq 2d(v,x)$.*

(ii) *For all $v \in cy$ we have that $2d(v,x) \geq d(x,c) - d(x,y)$.*

(iii) *For all $v \in V$ we have that $d(v,c) \leq d(v,x) + \min_{c \succcurlyeq_v z} d(x,z)$.*

Proof. (i) Let $v \in cx$, Then, $c \succcurlyeq_v x$, which means that $d(v,c) \leq d(v,x)$. By the triangle inequality we get

$$d(x,c) \leq d(x,v) + d(v,c) = d(v,x) + d(v,c) \leq 2d(v,x).$$

(ii) Let $v \in cy$. By the triangle inequality we have that

$$d(x,c) \leq d(v,c) + d(v,x).$$

Since $c \succcurlyeq_v y$, we also have $d(v,c) \leq d(v,y)$, and hence

$$d(x,c) \leq d(v,y) + d(v,x) \leq d(v,x) + d(x,y) + d(v,x) = d(x,y) + 2d(v,x).$$

Therefore, $2d(v,x) \geq d(x,c) - d(x,y)$.

(iii) Let $v \in V$ and $c, x \in C$. If $z \in C$ satisfies $c \succcurlyeq_v z$ then

$$d(v,c) \leq d(v,z) \leq d(v,x) + d(x,z).$$

Since this is true for all z satisfying $c \succcurlyeq_v z$, we get $d(v,c) \leq d(v,x) + \min_{c \succcurlyeq_v z} d(x,z)$. □

Lemma 3.1.6 (Anshelevich-Bhardwaj-Elkind-Postl)**.** *For every pair of candidates c, x we have*

$$\frac{\mathrm{SC}(c)}{\mathrm{SC}(x)} \leq \frac{2n}{|cx|} - 1.$$

Proof. Let $c \in C$. Note that if $v \in cx$ then $d(v,c) \leq d(v,x)$, while if $v \in xc$ we have that $d(v,c) \leq d(v,x) + d(x,c)$ by the triangle inequality. It follows that

$$\begin{aligned}
\mathrm{SC}(c) &= \sum_{v \in V} d(v,c) = \sum_{v \in cx} d(v,c) + \sum_{v \in xc} d(v,c) \\
&\leq \sum_{v \in cx} d(v,x) + \sum_{v \in xc} (d(v,x) + d(x,c)) \\
&= \sum_{v \in V} d(v,x) + \sum_{v \in xc} d(x,c) = \mathrm{SC}(x) + |xc| \cdot d(x,c).
\end{aligned}$$

Therefore,

$$\frac{\mathrm{SC}(c)}{\mathrm{SC}(x)} \leq 1 + \frac{|xc| \cdot d(x,c)}{\mathrm{SC}(x)} = 1 + \frac{(n - |cx|) \cdot d(x,c)}{\mathrm{SC}(x)}. \tag{3.1}$$

On the other hand, Lemma 3.1.5 (i) shows that $d(x,c) \leq 2d(v,x)$ for all $v \in cx$, and hence

$$\mathrm{SC}(x) = \sum_{v \in V} d(v,x) \geq \sum_{v \in cx} d(v,x) \geq \frac{|cx| \cdot d(x,c)}{2}.$$

Combining the above we obtain

$$\frac{\mathrm{SC}(c)}{\mathrm{SC}(x)} \leq 1 + \frac{(n - |cx|) \cdot d(x,c)}{\mathrm{SC}(x)} \leq 1 + \frac{2(n - |cx|)}{|cx|} = \frac{2n}{|cx|} - 1.$$

This proves the lemma. □

3.2 Distortion of scoring rules

Definition 3.2.1 (Positional scoring rule)**.** A *positional scoring rule* $f_{\vec{s}}$ for C is determined by a **scoring vector** $\vec{s} = (s_1, s_2, \ldots, s_m)$, where $s_i \in \mathbb{Q}$, $s_1 \geq s_2 \geq \cdots \geq s_m$ and $s_1 > s_m$. Whenever a voter $v \in V$ ranks a candidate $c \in C$ in position ℓ, i.e. $\sigma_v(c) = \ell$, then the candidate c receives $r_v(c) = s_\ell$ points from v. Then, the total score of the candidate c is the sum

$$\text{score}_{\vec{s}}(c) = \sum_{v \in V} r_v(c).$$

The winners according to the scoring vector $\vec{s}$ are the candidates $c^* \in C$ for which

$$\text{score}_{\vec{s}}(c^*) = \max_{c \in C} \text{score}_{\vec{s}}(c).$$

Some well known voting rules are in fact positional scoring rules. Well-known examples are the following:

(i) Plurality, where $\vec{s} = (1, 0, \ldots, 0)$.

(ii) Veto, where $\vec{s} = (1, 1, \ldots, 1, 0)$.

(iii) Borda, where $\vec{s} = (m-1, m-2, \ldots, 1, 0)$.

(iv) Harmonic rule, where $\vec{s} = \left(1, \frac{1}{2}, \frac{1}{3}, \ldots, \frac{1}{m}\right)$.

(v) k-Approval $(1 \leq k < m)$, where $\vec{s} = (1, 1, \ldots, 1[k], 0, \ldots, 0[m-k])$.

Our aim in this section is to provide lower and upper bounds for the worst-case distortion of these voting rules. All the results that are presented in this section are from the work of Anshelevich, Bhardwaj, Elkind, Postl and Skowron [8].

Lower bounds for the distortion

As we will see, the distortion of such rules cannot be bounded by a constant. We will see that the distortion grows with m, that is, given any $C > 0$ we may find m_0 such that if we consider any positional scoring rule with $m \geq m_0$ candidates then the worst-case distortion of the rule exceeds C. Moreover, for several families of scoring rules the distortion grows linearly with the number of candidates or is infinite.

The next proposition from [8] shows that the distortion of Veto and k-Approval with $k > 1$ is infinite for every value of m.

Proposition 3.2.2. *Consider a positional scoring rule with a score vector $\vec{s} = (s_1, \ldots, s_m)$ such that $s_1 = s_2$. Then, the distortion of the rule is infinite.*

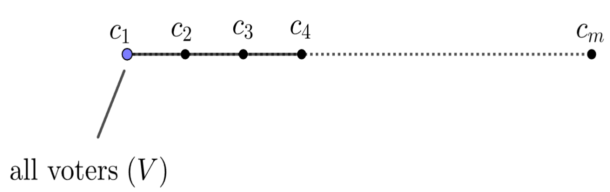

Figure 3.2: Example for Proposition 3.2.2.

Proof. Let $C = \{c_1, \ldots, c_m\}$ and $V = \{v_1, \ldots, v_n\}$. Assume that c_i is located at the point $x = i$ and all voters are located at $x = 1$ on the real axis $\mathbb{R}$ equipped with the usual Euclidean metric. Then, all voters provide the same ranking $c_1 \succ c_2 \succ \cdots \succ c_m$. Since $s_1 = s_2$ both c_1 and c_2 achieve the same score $ns_1 = ns_2$, and hence they are among the election winners. On the other hand, the sum of the distances from the voters to c_1 is 0, while the sum of distances from the voters to c_2 is n. Therefore, the distortion of the rule is $+\infty$. □

The next theorem covers the more interesting complementary case where $s_1 > s_2$.

Theorem 3.2.3 (Anshelevich-Bhardwaj-Elkind-Postl-Skowron)**.** *Let $\vec{s} = (s_1, \ldots, s_m)$ be a score vector with $s_1 > s_2$. There exists a profile σ on m candidates such that*

$$\text{distortion}(f_{\vec{s}}, \sigma) \geq 1 + 2\sqrt{\ln m - 1}.$$

If $\vec{s}$ is the scoring vector of the Plurality or Borda rule then there exists a profile σ on m candidates such that

$$\text{distortion}(f_{\vec{s}}, \sigma) \geq 2m - 1.$$

Finally, if $\vec{s}$ is the scoring vector of the Harmonic rule then there exists a profile σ on m candidates such that

$$\text{distortion}(f_{\vec{s}}, \sigma) = \Omega\left(\frac{m}{\ln m}\right).$$

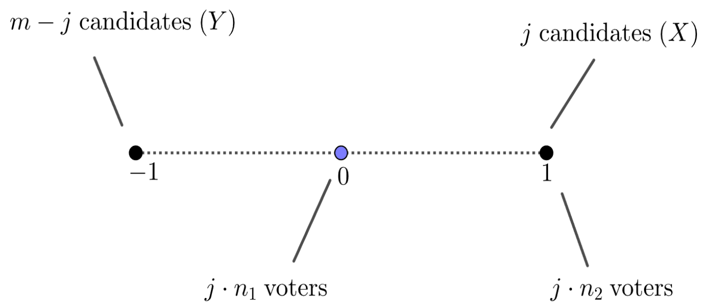

Figure 3.3: Example for Theorem 3.2.3.

Proof. We first note that the output of a scoring rule $\vec{s}$ does not change if the scoring rule "normalized" in the sense that $\vec{s}$ is replaced by $\lambda \cdot (s_1 - t, s_2 - t, \ldots, s_m - t)$ where t and λ are constants with $\lambda > 0$. Since $s_1 > s_2 \geq s_m$, taking $t = s_m$ and $\lambda = \frac{1}{s_1 - s_m}$, we may assume that $s_1 = 1$ and $s_m = 0$.

We fix j with $0 < j < m$ and n_1, n_2 such that the number of voters is $n = j \cdot n_1 + j \cdot n_2$. The numbers j, n_1, n_2 will be determined later. We consider voters and candidates placed on the real axis $\mathbb{R}$, equipped with the Euclidean metric, as follows. We partition the set C of candidates into two sets $X = \{x_1, \ldots, x_j\}$ and $Y = \{y_1, \ldots, y_{m-j}\}$. Let the j candidates of the set X be placed at the point 1 and the $m - j$ candidates of the set Y be placed at the point -1. The voters are placed as follows: There are $j \cdot n_1$ voters at 0 and $j \cdot n_2$ voters at 1. This means that the social cost of any candidate $x_i \in X$ is

$$\text{SC}(x_i, d) = j \cdot n_1$$

while the social cost of any candidate $y_i \in Y$ is

$$\mathrm{SC}(y_i, d) = j \cdot n_1 + 2j \cdot n_2.$$

Therefore, it is optimal to choose any candidate from the set X and, in case an algorithm ALG chooses a candidate from Y then its distortion will be

$$\text{distortion}(\text{ALG}) \geq \frac{j \cdot n_1 + 2j \cdot n_2}{j \cdot n_1} = 1 + 2\frac{n_2}{n_1}. \tag{3.2}$$

Thus, we will try to choose $\frac{n_2}{n_1}$ as large as possible.

We shall now describe a profile σ with $d \ \triangleright \ \sigma$. We assume the following: all voters situated at 0 prefer candidates from the set Y to candidates from the set X. All voters rank the candidates from Y in the same way: $y_1 \succ y_2 \succ \cdots \succ y_{m-j}$. For each $i = 1, \ldots, j$ there are n_1 voters at 0 who rank the candidates as follows:

$$y_1 \succ y_2 \succ \cdots \succ y_{m-j} \succ x_1 \succ x_{i+1} \succ \cdots \succ x_j \succ x_1 \succ \cdots \succ x_{i-1}$$

and n_2 voters at 1 who rank the candidates as:

$$x_i \succ \cdots \succ x_j \succ x_1 \succ \cdots \succ x_{i-1} \succ y_1 \succ \cdots \succ y_{m-j}.$$

So, the score of y_1 is larger than the score of every other $y \in Y$ according to this profile, and equal to

$$\text{score}_{\vec{s}}(y_1) = j \cdot n_1 + j \cdot n_2 s_{i+1}$$

(recall that $s_1 = 1$) while the scores of the candidates from X are all equal to

$$\text{score}_{\vec{s}}(x_i) = n_1(s_{m-j+1} + \cdots + s_m) + n_2(s_1 + \cdots + s_j).$$

We need to have

$$j \cdot n_1 + j \cdot n_2 s_{j+1} \geq n_1(s_{m-j+1} + \cdots s_m) + n_2(s_1 + \cdots + s_j)$$

in order for the algorithm to output y_1. Dividing with n_1 we get the condition

$$j\left(1 + \frac{n_2}{n_1}s_{j+1}\right) \geq (s_{m-j+1} + \cdots s_m) + \frac{n_2}{n_1}(s_1 + \cdots + s_j)$$

or, equivalently

$$\frac{n_2}{n_1}[(s_1 + \cdots + s_j) - js_{j+1}] \leq j - (s_{m-j+1} + \cdots + s_m),$$

or

$$\frac{n_2}{n_1} \leq \frac{j - (s_{m-j+1} + \cdots + s_m)}{(s_1 + \cdots + s_j) - j \cdot s_{j+1}}.$$

So, our goal is to choose j so that

$$T_j = \frac{j - (s_{m-j+1} + \cdots + s_m)}{(s_1 + \cdots + s_j) - j \cdot s_{j+1}}$$

becomes as large as possible and then we can directly choose integers n_1, n_2 such that $\frac{n_2}{n_1} = T_j$.

We shall first choose j for each positional scoring rule separately:

Borda rule: According to our normalization assumption, $\vec{s} = \left(1, \frac{m-2}{m-1}, \frac{m-3}{m-1}, \dots, 0\right)$. Choosing $j = 1$ we have

$$T_1 = \frac{1 - s_m}{1 - s_2} = \frac{1}{1 - \frac{m-2}{m-1}} = m - 1.$$

Then we can choose $n_2 = m-1, n_1 = 1$ and applying (3.2) we get the lower bound distortion(Borda) $\geq 2m - 1$.

Plurality rule: We have $\vec{s} = (1, 0, \dots, 0)$ so we choose $j = m - 1$ and see that

$$T_{m-1} = \frac{m - 1 - (s_2 + \cdots + s_m)}{(s_1 + \cdots + s_{m-1}) - (m-1)s_m} = \frac{m-1}{1} = m - 1.$$

Then, applying (3.2) we get the lower bound distortion(Plurality) $\geq 2m - 1$.

Harmonic rule: We choose $j = m - 1$ and check that

$$T_{m-1} = \frac{m - 1 - (s_2 + \cdots + s_m)}{(s_1 + \cdots + s_{m-1}) - (m-1)s_m} \geq \frac{m - 1 - \ln(m-1) - 1}{\ln m}.$$

Then, applying (3.2) we get the lower bound distortion(Harmonic) $= \Omega\left(\frac{m}{\ln m}\right)$.

We now turn to the general case. We need to show that for every score vector $\vec{s}$ with $1 = s_1 > s_2$ and $s_m = 0$ we may choose j $(1 \leq j \leq m-1)$ such that $T_j \geq \sqrt{\ln m - 1}$. We set $\lambda = 1/\sqrt{\ln m - 1}$ and consider two cases:

$s_2 \geq 1 - \lambda$. Then we choose $j = 1$ and we have

$$T_1 = \frac{1 - s_m}{s_1 - s_2} = \frac{1}{1 - s_2} \geq 1/\lambda = \sqrt{\ln m - 1}.$$

$s_2 < 1 - \lambda$. Then, for each $j = 1, \dots, m-1$ we have that

$$\frac{s_{m-j+1} + \cdots + s_m}{j} < s_2 < 1 - \lambda. \tag{3.3}$$

We shall show that there exists $j \in \{1, \dots, m-1\}$ such that

$$\frac{s_1 + \cdots + s_j}{j} - s_{j+1} < \lambda^2. \tag{3.4}$$

Then, for this j, we will have

$$T_j = \frac{1 - \frac{s_{m-j+1} + \cdots + s_m}{j}}{\frac{s_1 + \cdots + s_j}{j} - s_{j+1}} > \frac{\lambda}{\lambda^2} = \frac{1}{\lambda} = \sqrt{\ln m - 1},$$

where the inequality follows from (3.3) and (3.4).

So, suppose that there is no $j \in \{1, \dots, m-1\}$ such that (3.4) holds and we will get a contradiction. For each $j = 1, \dots, m$ we set $f_j = \frac{s_1 + \cdots + s_j}{j}$, so we suppose that for all $j \in \{1, \dots, m-1\}$ we have $f_j - s_{j+1} \geq \lambda^2$ or, equivalently $s_{j+1} < f_j - \lambda^2$. Then,

$$s_1 + \cdots + s_{j+1} = (s_1 + \cdots + s_j) + s_{j+1} \leq j f_j + f_j - \lambda^2 = (j+1) f_j - \lambda^2,$$

and hence $f_{j+1} \leq f_j - \frac{\lambda^2}{j+1}$ for all $j = 1, \dots, m-1$. By induction, we get

$$f_m \leq f_1 - \lambda^2 \left(\frac{1}{2} + \cdots + \frac{1}{m}\right)$$

and using the inequality $\frac{1}{2} + \cdots + \frac{1}{m} > \ln m - 1$ we conclude that

$$f_m < 1 - \lambda^2(\ln m - 1) = 0,$$

which is a contradiction since $f_m > 0$. □

Upper bounds for the distortion

In this subsection we provide upper bounds for the distortion of Plurality and Borda rule, which match the lower bounds given by Theorem 3.2.3, and an upper bound for the Harmonic rule, which is close to the corresponding lower bound.

Theorem 3.2.4 (Anshelevich-Bhardwaj-Elkind-Postl-Skowron). *The distortion of Plurality and Borda rule on m candidates is bounded by $2m - 1$.*

Proof. Let d be a metric which is aligned with the preference profile σ, and let x be an optimal candidate with respect to d. We shall show that if w is a winning candidate then both for the Plurality and for the Borda rule we have $|wx| \geq \frac{n}{m}$. Then, the result follows from Lemma 3.1.6.

The claim is obvious for the Plurality rule, because at least $\frac{n}{m}$ voters have w as their first preference. To see this, note that if $\text{plu}(c)$ is the number of voters who have c as their first preference then

$$n = \sum_{x \in C} \text{plu}(c) \leq |C| \cdot \text{plu}(w) = m \cdot \text{plu}(w).$$

For the Borda rule we argue as follows: If $v \in V$ we set $\delta_v = \sigma_v(x) - \sigma_v(w)$. Since w is a Borda winner, we have that $\sum_{v \in V} \delta_v \geq 0$. Note that if $v \in xw$ then $\delta_v \leq -1$, while if $v \in wx$ then $\delta_v \leq m-1$. Assume that $|wx| < \frac{n}{m}$. Then, $|xw| > n - \frac{n}{m}$, and hence

$$\sum_{v \in V} \delta_v \leq (-1)|xw| + (m-1)|wx| < -\frac{n(m-1)}{m} + \frac{n(m-1)}{m} = 0.$$

We have arrived at a contradiction, therefore $|wx| \geq \frac{n}{m}$. □

The upper bound for the Harmonic rule is sublinear in m.

Theorem 3.2.5 (Anshelevich-Bhardwaj-Elkind-Postl-Skowron). *The distortion of the Harmonic rule on m candidates is $O\left(\frac{m}{\sqrt{\ln m}}\right)$.*

Proof. For every $t > 0$ we define $s_t = 1/t$. We also set $F_j = 1 + \frac{1}{2} + \cdots + \frac{1}{j}$ for every positive integer j. Let d be a metric which is aligned with the preference profile σ, let x be an optimal candidate for σ, let w be a winner under the Harmonic rule, and set $\delta = d(w, x)$.

We consider the set

$$M = \left\{ v \in V : d(v, x) < \frac{\delta}{6} \right\}$$

and denote the cardinality of M by m. Then,

$$\begin{aligned} \frac{\sum_{v \in V} d(v, w)}{\sum_{v \in V} d(v, x)} &\leq \frac{\sum_{v \in V}(d(v, x) + d(x, w))}{\sum_{v \in V} d(v, x)} = 1 + \frac{n\delta}{\sum_{v \in V} d(v, x)} \\ &\leq 1 + \frac{n\delta}{\sum_{v \in V \setminus M} d(v, x)} \leq 1 + \frac{n\delta}{(n-m)\delta/6} = 1 + \frac{6n}{n-m}. \end{aligned}$$

For the proof of the theorem it remains to show that $n-m=n\,\Omega\left(\frac{\sqrt{\ln m}}{m}\right)$. Consider the set

$$P=\left\{y\in C: d(y,x)<\frac{\delta}{3}\right\}$$

and denote the cardinality of P by p. For each $v\in M$, let s_v denote the score that x receives from v. Then,

$$s_v=s_{\sigma_i(x)}=\frac{1}{\sigma_i(x)}.$$

Let also

$$\zeta_v=F_{\sigma_i(x)}=s_1+\cdots+s_{\sigma_i(x)}.$$

If we fix a voter $v\in M$ then, for every $y\in P$ we have that

$$d(v,y)\leq d(v,x)+d(x,y)<\frac{\delta}{6}+\frac{\delta}{3}=\frac{\delta}{2}$$

and

$$d(v,w)\geq d(x,w)-d(v,x)>\delta-\frac{\delta}{6}=\frac{5\delta}{6},$$

which means that v prefers every candidate from P over w. Next, we observe that if $z\in C\setminus P$ then

$$d(v,z)\geq d(z,x)-d(v,x)>\frac{\delta}{3}-\frac{\delta}{6}=\frac{\delta}{6}>d(v,x),$$

which means that v prefers x to every candidate from $C\setminus P$. So, the preference order of v is of the form

$$v:\ a_{i_1}\succ\cdots\succ a_{i_\ell}\succ x\succ\cdots\succ w\succ\cdots,$$

where $\{a_{i_1},\ldots,a_{i_\ell}\}$ is a subset of $P\setminus\{x\}$. It follows that the total score of w is at most $ms_{j+1}+(n-m)$, while the total score of x is at least $\sum\limits_{v\in M}s_v$.

Note also that the total score that a voter $v\in M$ gives to the candidates in P is at least ζ_v (because $x\in P$) and hence, by the pigeonhole principle, there exists some candidate in P whose total score is at least $\frac{1}{j}\sum\limits_{v\in M}\zeta_v$. Since w is a winner under the Harmonic rule, we see that

$$ms_{j+1}+(n-m)\geq\sum_{v\in M}s_v \tag{3.5}$$

and

$$ms_{j+1}+(n-m)\geq\frac{1}{j}\sum_{v\in M}\zeta_v. \tag{3.6}$$

If we set $\xi=\frac{1}{m}\sum\limits_{v\in M}\sigma_v(x)$ then by the harmonic-arithmetic mean inequality we get

$$\frac{m}{\sum_{v\in M}s_v}=\frac{m}{\sum_{v\in M}\frac{1}{\sigma_v(x)}}\leq\frac{\sum_{v\in M}\sigma_v(x)}{m}=\xi,$$

and (3.5) gives

$$ms_{j+1}+(n-m)\geq\frac{m}{\xi}=ms_{\xi}. \tag{3.7}$$

On the other hand, since $\sigma_v(x) \leq j$ for every $v \in M$ and $\frac{F_q}{q}$ is a decreasing function of q, we get $F_{\sigma_v(x)} \geq \sigma_v(x)\frac{F_j}{j}$ for all $v \in M$, which implies that

$$\frac{\sum_{v\in M}\zeta_v}{m\xi} = \frac{\sum_{v\in M}\zeta_v}{\sum_{v\in M}\sigma_v(x)} = \frac{\sum_{v\in M}F_{\sigma_v(x)}}{\sum_{v\in M}\sigma_v(x)} \geq \frac{F_j}{j}.$$

Then, (3.6) gives

$$ms_{j+1} + (n-m) \geq \frac{m\xi\, F_j}{j^2}. \tag{3.8}$$

Now, if $\xi \leq j/\sqrt{\ln j}$ then $s_\xi \geq \sqrt{\ln j}/j$, and if $\xi > j/\sqrt{\ln j}$ then $F_j > \ln j$, and hence $\xi F_j/j^2 > \sqrt{\ln j}/j$. Taking into account (3.7) and (3.8) we see that $\frac{m}{j+1} + (n-m) \geq m\frac{\sqrt{\ln j}}{j}$, or equivalently

$$n \geq m\left(1 + \frac{\sqrt{\ln j}}{j} - \frac{1}{j+1}\right).$$

Since $j \leq m$, we conclude that $m \leq \dfrac{n}{1+\frac{\sqrt{\ln m}}{m} - \frac{1}{m+1}}$, and this finally gives

$$n - m \geq n\frac{\frac{\sqrt{\ln m}}{m} - \frac{1}{m+1}}{1+\frac{\sqrt{\ln m}}{m} - \frac{1}{m+1}} = n\cdot\Omega\left(\frac{\sqrt{\ln m}}{m}\right),$$

which proves the theorem. □

3.3 Distortion of the Copeland rule

In this section we show that there are simple voting rules whose distortion with respect to the social cost is bounded by an absolute constant. We start with a few definitions.

Definition 3.3.1 (uncovered set)**.** Let (V, C, d) be an instance and let σ be an induced preference profile. If $x, y \in C$ then we always have $|xy| + |yx| = n$. We say that x *universally defeats* y if $|xy| > \frac{n}{2}$ and that x *weakly universally defeats* y if $|xy| \geq \frac{n}{2}$. A candidate c is called *Condorcet winner* (respectively, *weak Condorcet winner*) if c universally defeats (respectively, weakly universally defeats) all other candidates.

We say that a candidate $w \in C$ belongs to the *uncovered set* of σ if for every $x \in C$ we have that w weakly universally defeats x or there is a candidate y such that w weakly universally defeats y and y weakly universally defeats x. The uncovered set of σ is denoted by UC(σ).

Definition 3.3.2 (Copeland rule)**.** Let (V, C, d) be an instance and let σ be the induced preference profile. The Copeland score of a candidate c is the number of candidates that c universally defeats. The Copeland rule outputs all candidates that have maximal Copeland score.

One can check that the uncovered set of any preference profile is always non-empty and that the output of the Copeland rule is always a subset of the uncovered set. Indeed, suppose that w is a Copeland winner with Copeland score k, i.e. w universally defeats the candidates $y_1, \ldots, y_k$. We shall show that $w \in$ UC(σ). If not, there must be a candidate x such that none of the candidates $w, y_1, \ldots, y_k$ weakly universally defeats x. Then, we have

$$|wx| < \frac{n}{2}, |y_1x| < \frac{n}{2}, \ldots, |y_kx| < \frac{n}{2}$$

and hence

$$|xw| > \frac{n}{2}, |xy_1| > \frac{n}{2}, \dots, |xy_k| > \frac{n}{2}$$

which means that x universally defeats $w, y_1, \dots, y_k$. It follows that x has Copeland score at least $k+1$, which is higher than that of w, and we arrive at a contradiction.

Our aim in this section is to show that the distortion of Copeland rule is less than or equal to 5.

Theorem 3.3.3 (Anshelevich-Bhardwaj-Elkind-Postl-Skowron)**.** *Let σ be a preference profile and let x be an optimal candidate for σ with respect to the metric d. Then, for every $w \in \mathrm{UC}(\sigma)$ we have that*

$$\mathrm{SC}(w) \leq 5\,\mathrm{SC}(x).$$

For the proof of Theorem 3.3.3 we need the next two lemmas.

Lemma 3.3.4. *Let $u = (u_1, \dots, u_m) \in \mathbb{R}^m$ with $u_1 \geq u_2 \geq \dots \geq u_m \geq 0$ and let $\alpha, \beta \in \mathbb{R}^m$ such that $\sum_{i=1}^{k} \alpha_i \geq \sum_{i=1}^{k} \beta_i$ for all $k \in [m]$. Then, $\sum_{i=1}^{m} \alpha_i u_i \geq \sum_{i=1}^{m} \beta_i u_i$.*

Proof. For each $i \in [m]$ we define $\gamma_i = \alpha_i - \beta_i$. From Abel's summation formula we know that

$$\sum_{i=1}^{m} t_i s_i = t_m \sum_{i=1}^{m} s_i - \sum_{i=1}^{m-1} \left(\sum_{j=1}^{k} s_j \right) (t_{i+1} - t_i)$$

for any $t_1, \dots, t_m, s_1, \dots, s_m \in \mathbb{R}$. Using this identity with $s_i = \gamma_i$ and $t_i = u_i$, we get

$$\begin{aligned} \sum_{i=1}^{m} \alpha_i u_i - \sum_{i=1}^{m} \beta_i u_i = \sum_{i=1}^{m} \gamma_i u_i &= u_m \sum_{i=1}^{m} \gamma_i - \sum_{i=1}^{m-1} \left(\sum_{j=1}^{k} \gamma_j \right) (u_{i+1} - u_i) \\ &= u_m \sum_{i=1}^{m} \gamma_i + \sum_{i=1}^{m-1} \left(\sum_{j=1}^{k} \gamma_j \right) (u_i - u_{i+1}). \end{aligned} \tag{3.9}$$

The assumption of the lemma implies that $\sum_{j=1}^{k} \gamma_j \geq 0$ for all $k \in [m]$ and $u_i - u_{i+1} \geq 0$ for all $i \in [m-1]$. Therefore, the sum in (3.9) is non-negative, and the lemma follows. □

The next lemma gives a lower bound for the cost of a given candidate x in terms of a second candidate w, which is useful when x is an optimal candidate and w is a winning candidate.

Lemma 3.3.5. *Let σ be a preference profile and let x, w be a pair of candidates. If*

$$\sum_{v \in V} d(v, x) \geq \frac{1}{\gamma} \sum_{v \in xw} \min_{w \succcurlyeq_v z} d(x, z) \tag{3.10}$$

for some γ, then

$$\mathrm{SC}(w) \leq (1 + \gamma)\mathrm{SC}(x).$$

Proof. From Lemma 3.1.5 (iii) we know that if $w, x \in C$ then for all $v \in V$ we have that $d(v, w) \leq d(v, x) + \min_{w \succcurlyeq_v z} d(x, z)$. Using this inequality together with the fact that $d(v, w) \leq d(v, x)$ for all

$v \in wx$, we obtain the following upper bound for the cost of the candidate w:

$$\begin{aligned}\frac{\mathrm{SC}(w)}{\mathrm{SC}(x)} &= \frac{\sum_{v\in V} d(v,w)}{\sum_{v\in V} d(v,x)} = \frac{\sum_{v\in wx} d(v,w) + \sum_{v\in xw} d(v,w)}{\sum_{v\in V} d(v,x)} \\ &\le \frac{\sum_{v\in V} d(v,x) + \sum_{v\in xw} \min_{w\succcurlyeq_v z} d(x,z)}{\sum_{v\in V} d(v,x)} = 1 + \frac{\sum_{v\in xw} \min_{w\succcurlyeq_v z} d(x,z)}{\sum_{v\in V} d(v,x)}.\end{aligned}$$

Then, taking into account (3.10) we obtain the assertion of the lemma. □

Proof of Theorem 3.3.3. We fix a preference profile σ, an optimal candidate x for σ and a candidate $w \in \mathrm{UC}(\sigma)$.

If w weakly universally defeats x then $|wx| \ge \frac{n}{2}$ and Lemma 3.1.6 shows that

$$\frac{\mathrm{SC}(w)}{\mathrm{SC}(x)} \le \frac{2n}{|wx|} - 1 \le 4 - 1 = 3$$

or, equivalently

$$\mathrm{SC}(w) \le 3\,\mathrm{SC}(x).$$

Suppose that w does not weakly universally defeat x. Since $w \in \mathrm{UC}(\sigma)$, there exists a candidate y such that w weakly universally defeats y and y weakly universally defeats x. Therefore, we have $|wy| \ge n/2 \ge |yw|$ and $|yx| \ge n/2 \ge |xy|$. We distinguish two cases:

(i) First, assume that $d(x,y) \ge d(x,w)$. There are at least $n/2$ voters that prefer y over x. Lemma 3.1.5 shows that each of these voters contributes a term $\ge \frac{1}{2}d(x,y) \ge \frac{1}{2}d(x,w)$ to the social cost $\mathrm{SC}(x)$ of x. Then,

$$\begin{aligned}\sum_{v\in V} d(v,x) \ge \sum_{v\in yx} d(v,x) \ge \frac{1}{2}\sum_{v\in yx} d(x,y) = \frac{1}{2}|yx|\, d(x,y) \ge \frac{n}{4}d(x,y) \\ \ge \frac{n}{4}d(x,w) \ge \frac{1}{4}|xw|\, d(x,w) \ge \frac{1}{4}\sum_{v\in xw} \min_{w\succcurlyeq_v z} d(x,z).\end{aligned}$$

Applying Lemma 3.3.5 with $\gamma = 4$ we obtain an upper bound of 5 for the distortion.

(ii) Next, assume that $d(x,y) < d(x,w)$. Note that in the previous case, in order to give a lower bound for $\sum_{v\in V} d(v,x)$ we took into account only voters $v \in yx$. This time, we consider the sets wx, xwy and yxw. Observe that, by the choice of x, y and w, we have that

$$\begin{aligned}|wx| + |xwy| &= |ywx| + |wyx| + |wxy| + |xwy| \\ &\ge |wy| \ge |yw| \ge |yxw| + |xyw| \ge \frac{1}{2}|yxw| + \frac{1}{2}|xyw|\end{aligned} \tag{3.11}$$

and

$$\begin{aligned}|wx| + |yxw| &= |ywx| + |wyx| + |wxy| + |yxw| \\ &\ge |yx| \ge \frac{n}{2} \ge \frac{1}{2}(|yxw| + |xyw| + |xwy|).\end{aligned} \tag{3.12}$$

We apply Lemma 3.3.4 with

$$\alpha_1 = |wx| + |xwy|, \quad \alpha_2 = |yxw|, \quad \alpha_3 = -|xwy|$$
$$\beta_1 = \frac{1}{2}|yxw| + \frac{1}{2}|xyw|, \quad \beta_2 = \frac{1}{2}|xwy|, \quad \beta_3 = 0$$
$$u_1 = d(x,w), \quad u_2 = u_3 = d(x,y).$$

From (3.11) we see that $\alpha_1 \geq \beta_1$. Also, from (3.12) we see that $\alpha_1 + \alpha_2 + \alpha_3 \geq \beta_1 + \beta_2 + \beta_3$, and since $\alpha_3 \leq 0$ and $\beta_3 = 0$, this implies that $\alpha_1 + \alpha_2 \geq \beta_1 + \beta_2$. The assumption also shows that $u_1 \geq u_2 = u_3$. Therefore, we may apply Lemma 3.3.4 to write

$$\begin{aligned} \alpha_1 u_1 + \alpha_2 u_2 + \alpha_3 u_3 &= (|wx| + |xwy|)\, d(x,w) + (|yxw| - |xwy|)\, d(x,y) \\ &\geq \beta_1 u_1 + \beta_2 u_2 + \beta_3 u_3 = \frac{1}{2}(|yxw| + |xyw|)\, d(x,w) + \frac{1}{2}|xwy|\, d(x,y). \end{aligned} \tag{3.13}$$

Now, we apply Lemma 3.1.5 to write

$$\begin{aligned} \sum_{v \in V} d(v,x) &\geq \sum_{v \in wx} d(v,x) + \sum_{v \in xwy} d(v,x) + \sum_{v \in yxw} d(v,x) \\ &\geq \frac{1}{2}|wx|\, d(x,w) + |xwy| \left(\frac{d(x,w) - d(x,y)}{2} \right) + \frac{1}{2}|yxw|\, d(x,y) \\ &= \frac{1}{2}(|wx| + |xwy|)\, d(x,w) + \frac{1}{2}(|yxw| - |xwy|)\, d(x,y) \\ &\geq \frac{1}{4}(|yxw| + |xwy|)\, d(x,w) + \frac{1}{4}|xwy|\, d(x,y) \\ &\geq \frac{1}{4} \sum_{v \in xw} \min_{w \succcurlyeq_v z} d(x,z), \end{aligned}$$

where the first inequality comes from Lemma 3.1.5, the second inequality follows from (3.13) and the last inequality is a consequence of $xw = yxw \cup xyw \cup xwy$, combined with the facts that $w \succcurlyeq_v w$ for all $v \in V$ and $w \succcurlyeq_v y$ for all $v \in xwy$.

Having obtained this estimate for $\sum_{v \in V} d(v,x)$, we may apply Lemma 3.3.5 with $\gamma = 4$ to obtain again an upper bound of 5 for the distortion. ☐

Theorem 3.3.3 shows that the distortion of the Copeland rule is at most 5. The next result shows that, in fact, this upper bound is tight.

Theorem 3.3.6 (Anshelevich-Bhardwaj-Elkind-Postl-Skowron)**.** *The distortion of the Copeland rule with respect to the social cost is at most* 5*. Moreover, for every* $\epsilon > 0$ *there exists a preference profile* σ *such that the Copeland rule on* σ *with respect to the social cost has distortion at least* $5 - \epsilon$*.*

Proof. The upper bound is a consequence of Theorem 3.3.3 after we recall that the output of the Copeland rule is a subset of the uncovered set.

In order to prove that this upper bound is tight, we consider a preference profile σ with three candidates w, x, y, where $\frac{n}{2} - 1$ voters choose the rank $y \succ x \succ w$, $\frac{n}{2} - 1$ voters choose the rank $x \succ w \succ y$, and the remaining two voters choose the rank $w \succ y \succ x$. This implies that w universally defeats y, y universally defeats x, and x universally defeats w. Then, the Copeland score of all the candidates is equal to 1, and in particular w is a Copeland winner.

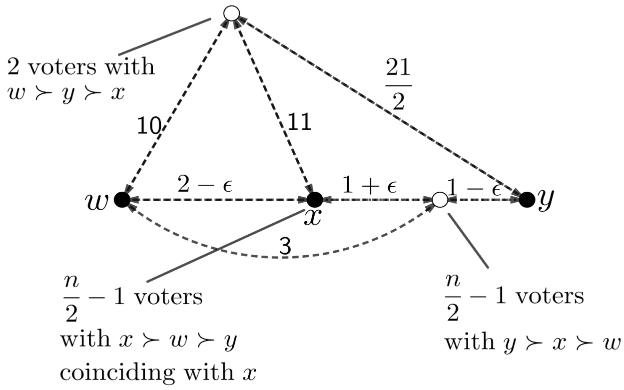

Figure 3.4: The distortion of any rule that outputs w is arbitrarily close to 5 as $\epsilon \to 0$ and $n \to \infty$.

We consider the metric in Figure 3.4. It is easy to check that the triangle inequality is satisfied. Observe that

$$\frac{\sum_{v \in V} d(v, w)}{\sum_{v \in V} d(v, x)} = \frac{\left(\frac{n}{2} - 1\right) \cdot (2 - \epsilon) + \left(\frac{n}{2} - 1\right) \cdot 3 + 20}{\left(\frac{n}{2} - 1\right) \cdot (1 + \epsilon) + 22}.$$

Note that if we let $n \to \infty$ and $\epsilon \to 0$, we obtain instances for which the distortion of the Copeland rule is arbitrarily close to 5. □

3.4 Distortion of single transferable vote

Single transferable vote (STV) is an iterative rule which is defined in the following way. In each round, the candidate that is ranked first by the fewest voters (the one with the lowest Plurality score) is removed from the set of candidates and from the rankings of the voters, which means that the Plurality scores have to be computed again. After $m - 1$ rounds there is only one candidate left, and this is the winner.

It was proved in [8] that the distortion of STV grows logarithmically with the number m of candidates. On the other hand, the authors also provide a non-constant lower bound, which means that STV is not as good as the Copeland rule. First, we explain the upper bound.

Theorem 3.4.1 (Anshelevich-Bhardwaj-Elkind-Postl-Skowron)**.** *The distortion of single transferable vote is* $O(\ln m)$.

Proof. We fix a preference profile σ, an optimal candidate x for σ and a winning candidate w under STV. We set $d = d(x, w)$, we fix a constant $\gamma \in \left(\frac{2}{3}, 1\right)$ and define

$$p = 2\left\lceil \log_{\frac{\gamma}{1-\gamma}} \left(m \cdot \frac{2\gamma - 1}{3\gamma - 2}\right)\right\rceil + 1.$$

Note that p is an odd integer and $p = O(\ln m)$. We also set

$$r = \frac{d}{2p}.$$

For every $i = 1, \ldots, p + 1$ we consider the ball $\mathcal{B}_i = B(x, (2i - 1)r)$. Note that

$$(2p - 1)r = \frac{2p - 1}{2p} d < d = d(x, w) < \frac{2p + 1}{2p} d = (2p + 1)r$$

which means that $w \in \mathcal{B}_{p+1} \setminus \mathcal{B}_p$.

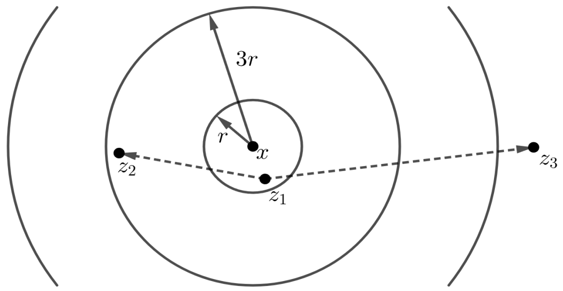

Figure 3.5: The sequence of balls used in the proof of Theorem 3.4.1.

We shall show that $\mathcal{B}_1$ contains at most γn voters. This will allow us to obtain the claimed upper bound for the distortion. We just have to note that

$$\frac{\sum_{v\in V} d(v,w)}{\sum_{v\in V} d(v,x)} \le \frac{\sum_{v\in V}(d(v,x)+d(x,w))}{\sum_{v\in V} d(v,x)} = 1 + \frac{nd}{\sum_{v\in V} d(v,x)}$$
$$\le 1 + \frac{nd}{\sum_{v\notin\mathcal{B}_1} d(v,x)} \le 1 + \frac{nd}{n(1-\gamma)r} = 1 + \frac{2p}{1-\gamma}$$

and then recall that γ is a fixed constant and $p = O(\ln m)$.

It remains to prove that $|\mathcal{B}_1 \cap C| \le \gamma n$. Assume the contrary. We consider an elimination process for which w is the last remaining candidate. We say that $y \in C$ is *supported* by $v \in V$ (equivalently, that v *supports* y) at some stage of the process if y is the closest candidate to v at the given stage. For every $1 \le i \le p$ denote by z_{i-1} the last candidate in $\mathcal{B}_{i-1}$ that is removed and by s_i the number of candidates in $\mathcal{B}_i \setminus \mathcal{B}_{i-1}$ just before z_{i-1} is being removed.

Let $i \le p-2$. If $v \in \mathcal{B}_1$ and $y \notin \mathcal{B}_{i+1}$ then

$$d(v,z_i) \le d(v,x) + d(x,z_i) \le r + (2i-1)r = (2i+1)r - r < d(y,x) - d(v,x) \le d(v,y).$$

This shows that, just before z_i is being removed, every $v \in \mathcal{B}_1$ supports some candidate in $\mathcal{B}_{i+1}$. By the pigeonhole principle, we may find a candidate $y \in \mathcal{B}_{i+1}$ that is supported by more than $\frac{\gamma n}{s_{i+1}+1}$ voters in $\mathcal{B}_1$. So, at the stage where y is removed, all the remaining candidates in $\mathcal{B}_{i+3} \setminus \mathcal{B}_{i+2}$ are supported by more than $\frac{\gamma n}{s_{i+1}+1}$ voters. Since every voter in $\mathcal{B}_1$ prefers y to every $y' \in \mathcal{B}_{i+3} \setminus \mathcal{B}_{i+2}$, none of these voters is in $\mathcal{B}_1$. To see this, observe that

$$d(v,y) \le d(v,x) + d(x,y) \le 2(i+1)r < d(y',x) - d(v,x) \le d(v,y')$$

for all $v \in \mathcal{B}_1$.

It follows that $s_{i+3}\dfrac{\gamma n}{s_{i+1}+1} < n(1-\gamma)$, or equivalently,

$$s_{i+1} > \frac{\gamma}{1-\gamma}s_{i+3} - 1.$$

We set $\xi = \frac{\gamma}{1-\gamma}$. Since $\frac{2}{3} < \gamma < 1$, we see that $\xi > 1$ and $\frac{1}{1-\xi} < 0$. Note that $s_{p+1} \ge 1$ because

$w \in \mathcal{B}_{p+1} \setminus \mathcal{B}_p$. We write

$$\begin{aligned} s_1 &> \xi s_3 - 1 \geq \xi^2 s_5 - \xi - 1 \geq \cdots \geq \xi^{\frac{p-1}{2}} - \xi^{\frac{p-1}{2}-1} - \cdots - 1 \\ &= \xi^{\frac{p-1}{2}} - \frac{1-\xi^{\frac{p-1}{2}}}{1-\xi} = \xi^{\frac{p-1}{2}}\left(1 + \frac{1}{1-\xi}\right) - \frac{1}{1-\xi} \geq \xi^{\frac{p-1}{2}} \cdot \frac{2-\xi}{1-\xi} \\ &= \left(\frac{\gamma}{1-\gamma}\right)^{\frac{p-1}{2}} \frac{3\gamma-2}{2\gamma-1} \geq m \frac{2\gamma-1}{3\gamma-2}\frac{3\gamma-2}{2\gamma-1} = m. \end{aligned}$$

Since $|C| = m$, this is a contradiction, and the proof is complete. □

We pass now to the lower bound.

Theorem 3.4.2 (Anshelevich-Bhardwaj-Elkind-Postl-Skowron)**.** *The maximum distortion of single transferable vote over all profiles with m candidates is $\Omega(\sqrt{\ln m})$.*

Proof. Given a positive integer p, we construct a perfectly balanced tree of height p and connect all its leaves to an additional node. We agree that all leaves belong to the first layer and, for any $i > 1$, layer i consists of the parents of the nodes at level $i-1$. For any $2 \leq i \leq p$, each node at level i has $r_i = 2^i + 2^{i-2} - 2$ children. We denote by t_i the number of nodes at level i. The length of each edge of the graph is equal to 1, and the distance of two nodes is the length of the shortest path between them.

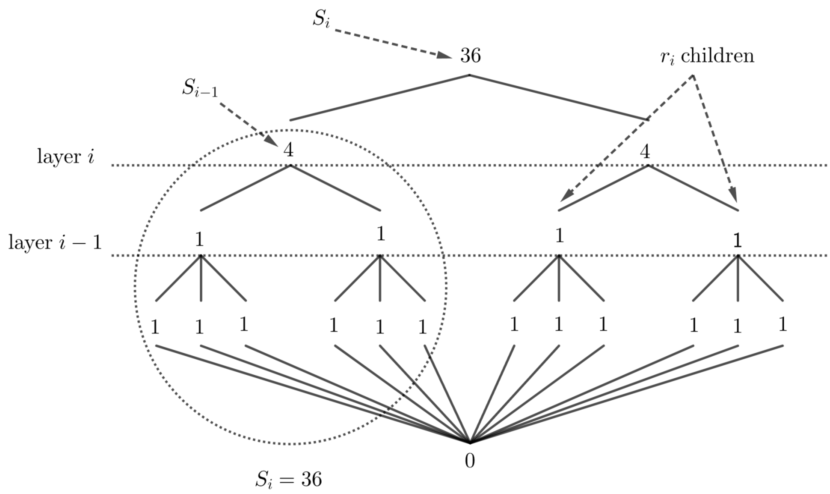

Figure 3.6: The metric space used in the proof of Theorem 3.4.2 for $p = 4$. The number that appears on every node is the number of voters in that node.

We place one candidate on each node, including the node connected to the leaves. Since $r_i = 2^i + 2^{i-2} - 2$ for all $2 \leq i \leq p$, we get $r_i \leq 2^{i+1}$ for all $1 \leq i \leq p$, and the total number of leaves is

$$t_1 \leq 2^{p+1} \cdot 2^p \cdots 2 = 2^{\frac{(p+1)(p+2)}{2}} \leq 2^{(p+1)^2}.$$

Since the degree of each internal node is at least 2, we see that the number m of candidates is at most $2 \cdot 2^{(p+1)^2}$, and this shows that

$$p \geq \sqrt{\log_2 m} - 2.$$

Next, we place the voters on the graph. We place one voter on each leaf. Let S_i be the total number of voters in a subtree rooted in a level-i node. We have $S_1 = 1$ and for each $2 \leq i \leq p$ we compute S_{i-1} and place S_{i-1} voters on each node of layer i. This means that $S_i = S_{i-1}(r_i + 1)$. There are exactly t_1 voters in the bottom layer and there are $t_i S_{i-1}$ voters at level i. Moreover, $t_i = t_{i+1} r_{i+1}$, and hence

$$\frac{t_{i+1} S_i}{t_i S_{i-1}} = \frac{S_i}{r_{i+1} S_{i-1}} = \frac{r_i + 1}{r_{i+1}} = \frac{2^i + 2^{i-2} - 1}{2^{i+1} + 2^{i-1} - 2} = \frac{1}{2}.$$

This means that the number of voters at layer $i+1$ is half the number of voters at layer i, and since layer 1 has t_1 voters, the number of voters in layer i is equal to $t_1/2^{i+1}$.

Let y_0 be the candidate that is placed on the node connected to the leaves. Then, STV removes y_0 first, because there is no voter that ranks him first and every other candidate is ranked first by at least one voter. In the next step, STV can remove all the candidates that are placed on the leaves, one by one: this is because, initially, each such candidate is ranked first by exactly one voter, and no such candidate gains new votes when other leaf candidates are removed. Suppose that STV has removed all candidates lying in layers $1, \ldots, i-1$, and all other candidates have not been removed. Then, a candidate in layer i is ranked first among the remaining candidates by the S_i voters in the respective subtree, and each candidate in layer j, $j > i$, is ranked first by the $S_{i+1} \geq S_i$ voters that are located in the same node as that candidate. Therefore, STV can remove the candidates in layer i one by one. So, we conclude that the root of the tree can be selected as the winner. Since there is a voter in each leaf, the total distance d_{stv} of the voters to the root is at least pt_1.

On the other hand, the total distance d_{bot} of the voters to the candidate in the node that connects all leaves satisfies

$$d_{\text{bot}} = t_1 + 2 \cdot \frac{t_1}{2} + \cdots + p \cdot \frac{t_1}{2^{p-1}} = t_1 \sum_{i=1}^{p} \frac{i}{2^{i-1}} = 4t_1 \left(1 - \frac{p+1}{2^p} + \frac{p}{2^{p+1}}\right) \leq 4t_1.$$

This implies that

$$\frac{d_{\text{st}}}{d_{\text{bot}}} \geq \frac{p}{4} \geq \frac{\sqrt{\log_2 m} - 2}{4},$$

and the theorem follows. □

Anagnostides, Fotakis and Patsilinakos [6] studied the distortion of STV with respect to the *dimensionality* of the undelying metric space. A crucial notion in their work is the *doubling dimension* of a metric space. The *doubling constant* of a metric space (X, d) is the least integer $N \geq 1$ such that for any $x \in X$ and any $r > 0$ the ball $B(x, 2r)$ can be covered by the union of at most N balls of radius r, i.e. there exists a subset $S \subseteq X$ with cardinality $|S| \leq N$ such that

$$B(x, 2r) \subseteq \bigcup_{y \in S} B(y, r).$$

Then, the doubling dimension of X is defined as

$$\dim(X) = k := \log_2 N.$$

A standard volumetric argument shows that the doubling dimension of the Euclidean space $\mathbb{R}^d$ is of the order of (its usual dimension) d. One can also check that if (X, d) is a finite metric space then $\dim(X) \leq \log_2 |X|$. The main result in [6] is the next theorem.

Theorem 3.4.3 (Anagnostides-Fotakis-Patsilinakos)**.** *The maximum distortion of single transferable* *vote over all profiles with m candidates located in a metric space with doubling dimension k is $O(k \ln\ln m)$.*

This establishes an upper bound for the distortion which is much better than the general lower bound $\Omega(\sqrt{\ln m})$ when the doubling dimension of the underlying metric space is small. For example, the distortion of STV under low-dimensional Euclidean spaces is $O(\ln\ln m)$.

It is actually conjectured in [6] that the $\ln\ln m$ factor can be removed and one can have that if the doubling dimension of (X, d) is equal to k then the distortion of STV is $O(k)$. In [6] this conjecture is verified in the case $k = 1$ of one-dimensional metric spaces.

3.5 Plurality Matching

In this section we present a theorem of Gkatzelis, Halpern and Shah [47], who introduced a deterministic algorithm that guarantees the optimal distortion of 3. In what follows, instead of an instance (V, C, d) and the induced preference profile σ, we consider an *election* $\mathcal{E} = (V, C, \sigma)$, i.e. we do not have full access to the metric d but we know that d is aligned to σ. We write $\Delta(V)$ and $\Delta(C)$ for the set of probability distributions over V and C respectively, that is, vectors of non-negative weights that add up to 1. For a given preference profile σ we denote by $\text{top}(v)$ the candidate ranked first in σ_v, and we set $\text{plu}(a) = |\{v \in V : \text{top}(v) = a\}|$; i.e. $\text{plu}(a)$ is the plurality score of the candidate a. For every $D \subseteq C$ we also set $\text{plu}(D) = \sum_{a\in D} \text{plu}(a)$.

We shall obtain a refinement of our distortion bounds using the notion of α-decisiveness, introduced by Anshelevich and Postl in [11]. Given $\alpha \in [0, 1]$, a voter v is called α-decisive if

$$d(v, \text{top}(v)) \leq \alpha \cdot d(v, c)$$

for all $c \in C$, $c \neq \text{top}(v)$, i.e. the distance of v from her top choice is at most α times its distance from any other candidate. We say that the metric space is α-decisive if all voters $v \in V$ are α-decisive. When we study the distortion of a voting rule in the framework of α-decisive metric spaces, we are interested in the worst case over only those metric spaces that are α-decisive and satisfy $d \rhd \sigma$.

We shall define a deterministic voting rule that has distortion at most 3. The key step is a lemma about matching voters to candidates. Before stating the lemma we need to introduce some terminology.

Definition 3.5.1. Let $\mathcal{E} = (V, C, \sigma)$ be an election, and consider two normalized weight vectors $p \in \Delta(V)$ and $q \in \Delta(C)$. The (p, q)-*domination graph* of the candidate a is the vertex-weighted bipartite graph $G^{\mathcal{E}}_{p,q}(a) = (V, C, E_a, p, q)$, where $(v, c) \in E_a$ if and only if $a \succcurlyeq_v c$. The vertex $v \in V$ has weight p_v and the vertex $c \in C$ has weight q_c. When $\mathcal{E}$ is clear from the context, we set $G_{p,q}(a) := G^{\mathcal{E}}_{p,q}(a)$.

Definition 3.5.2. We say that the (p, q)-domination graph $G_{p,q}(a)$ admits a *fractional perfect matching* if there exists a weight function $w : E_a \to \mathbb{R}^+$ such that the total weight of edges incident on each vertex equals the weight of the vertex. This means that for each $v \in V$ we have that $\sum\limits_{\{c\in C:(v,c)\in E_a\}} w(v, c) = p_v$ and for each $c \in C$ we have that $\sum\limits_{\{v\in V:(v,c)\in E_a\}} w(v, c) = q_c$.

For all $S \subseteq C$ we set $p(S) = \sum_{v\in S} p_v$ and for all $D \subseteq C$ we set $q(D) = \sum_{c\in D} q_c$. If $\mathcal{E} = (V, C, \sigma)$ is an election, for any $a \in C$ and $S \subseteq V$ we say that a candidate c is weakly defeated by a in S if there exists $v \in S$ such that $a \succcurlyeq_v c$ and we define the set

$$D_a(S) := D^{\mathcal{E}}_a(S) = \{c \in C : \ c \text{ is weakly defeated by } a \text{ in } S\}.$$

Note that $D_a(S)$ is precisely the set of neighbors of S in the graph $G_{p.q}(a)$, for any p and q. Using a generalization of Hall's condition we can show that a fractional perfect matching in $G_{p.q}(a)$ exists if and only if the set of neighbors of S has at least as much weight as S itself.

Lemma 3.5.3. *Let $\mathcal{E} = (V, C, \sigma)$ be an election, and let $p \in \Delta(V)$, $q \in \Delta(C)$ and $a \in C$. Then, $G_{p,q}(a)$ admits a fractional perfect matching if and only if $q(D_a(S)) \geq p(S)$ for all $S \subseteq V$. Moreover, we can check whether $G_{p,q}(a)$ admits a fractional perfect matching in strongly polynomial time.*

Proof. Consider an election $\mathcal{E} = (V, C, \sigma)$ and the weight vectors $p \in \Delta(V)$ and $q \in \Delta(C)$. For a given candidate $a \in C$ we shall first show that $G_{p,q}(a)$ admits a fractional perfect matching if and only if $q(D_a(S)) \geq p(S)$ for all $S \subseteq V$.

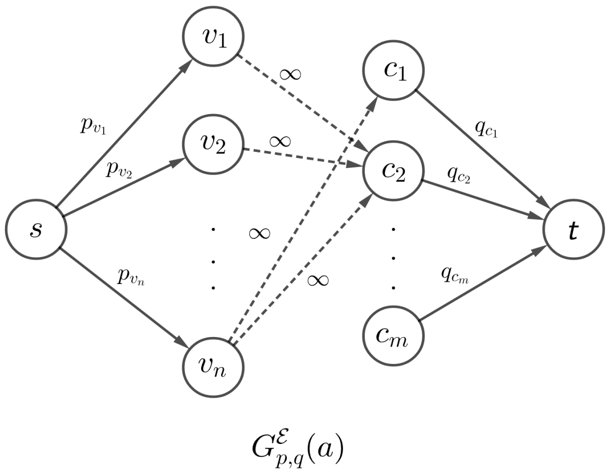

Figure 3.7: The flow network used in the proof of Lemma 3.5.3. The edge from source node s to each voter v has capacity p_v, the edge from each candidate c to target t has capacity q_c and edges from voters to candidates are as in $G^{\mathcal{E}}_{p,q}(a)$ with infinite capacity.

For the proof we use the max-flow min-cut theorem. We create a network N as in Figure 3.7: We convert all edges of $G_{p,q}(a)$ to directed edges from voters to candidates with infinite capacity, we add a source node s with an outgoing edge of capacity p_v to each voter v, and a target node t with an incoming edge of capacity q_c from each candidate c.

This network admits a flow of value 1 if and only if $G_{p,q}(a)$ admits a fractional perfect matching. To see this, note that if N has a flow of value 1, then s has an outgoing flow of 1 and t has an incoming flow of 1. Since $\sum\limits_{v \in V} p_v = \sum\limits_{c \in C} q_c = 1$, this implies that the node for each voter v must have a flow of p_v passing through her, and the node of each candidate c must have a flow of q_c passing through him. Then, we can define a fractional perfect matching in $G_{p,q}(a)$ by using the flow on each edge (v, c) as the weight of the edge. Conversely, if $G_{p,q}(a)$ admits a fractional perfect matching w, we can define a flow as follows: each edge (s, v) carries a flow of p_v, each edge (v, c) carries a flow of $w(v, c)$, and each edge (c, t) carries a flow of q_c. We easily check that this a flow of value 1.

Since p and q are normalized weight vectors, the network N has a flow of value 1 if and only if its maximum flow is 1. The max-flow min-cut theorem shows that, equivalently, the minimum cut capacity of N is equal to 1. Since $(\{s\}, V \cup (C \cup \{t\}))$ is a cut of capacity 1, we have that the minimum cut capacity will be 1 if and only if every cut has capacity equal to 1.

We will show that this last condition is equivalent to the fact that $q(D_a(S)) \geq p(S)$ for all $S \subseteq V$. Assume first that every cut has capacity at least 1. Let $S \subseteq V$ and consider the cut $(\{s\} \cup S \cup$

$D_a(S), (V \setminus S) \cup (C \setminus D_a(S)) \cup \{t\})$. Since $D_a(S)$ is the set of neighbors of S, there are no edges of infinite capacity in this cut. The only edges that are part of the cut are the ones that go from s to nodes in $V \setminus S$ and the ones that go from $D_a(S)$ to t. It follows that the capacity of this cut is $p(V \setminus S) + q(D_a(S))$. Since this sum is ≥ 1 and $p(S) = 1 - p(V \setminus S)$, we see that $q(D_a(S)) \geq p(S)$.

Conversely, assume that $q(D_a(S)) \geq p(S)$ for all $S \subseteq V$. Let $(\{s\} \cup A, B \cup \{t\})$ be a cut of N. If this cut contains an edge of infinite capacity, then it clearly has capacity at least 1. If not, then all the neighbors of $S = A \cap V$ are also in A, which shows that $D_a(S) \subseteq A$. Then, this cut has capacity at least $p(V \setminus S) + q(D_a(S)) = 1 - p(S) + q(D_a(S)) \geq 1$.

We have proved that $G_{p,q}(a)$ admits a fractional perfect matching if and only if the network that we constructed has max-flow 1 and that this holds if and only if $q(D_a(S)) \geq p(S)$ for all $S \subseteq V$. Moreover, since the maximum flow value can be computed in strongly polynomial time, we see that the existence of a fractional perfect matching in $G_{p,q}(a)$ can be checked in strongly polynomial time. □

Lemma 3.5.3 shows that if $\mathcal{E} = (V, C, \sigma)$ is an election, $p \in \Delta(V)$, $q \in \Delta(C)$ and there exists a candidate $a \in C$ such that $G^{\mathcal{E}}_{p,q}(a)$ admits a fractional perfect matching then in any subset S of V we have that a weakly defeats a set of candidates with total weight at least equal to that of S, and hence a is a very good choice. In fact, we will see that if we choose the weights p and q appropriately then this also implies low distortion. To make this plan work, our first goal is to show that such a candidate a always exists in any election $\mathcal{E}$ and for any choice of weight vectors p and q. This is indeed established by the next crucial lemma.

Lemma 3.5.4 (Ranking-matching lemma, Gkatzelis-Halpern-Shah)**.** *Let $\mathcal{E} = (V, C, \sigma)$ be an election, and let $p \in \Delta(V)$, $q \in \Delta(C)$. Then, there exists a candidate $a \in C$ whose (p, q)-domination graph $G^{\mathcal{E}}_{p,q}(a)$ admits a fractional perfect matching. Moreover, this candidate can be computed in strongly polynomial time.*

Proof. We argue by contradiction. Assume that $\mathcal{E} = (V, C, \sigma)$ is an election with the smallest possible number of voters, such that there exist $p \in \Delta(V)$ and $q \in \Delta(C)$ for which the assertion of the lemma is not true. Fix such p and q. Then, Lemma 3.5.3 shows that for any $a \in C$ we can find $S \subseteq V$ such that $q(D^{\mathcal{E}}_a(S)) < p(S)$. We shall say that S is a counterexample for a, and we shall call this counterexample minimal if there is no strict subset of S which is also a counterexample for a. For the rest of the proof, for every $a \in C$ we fix such a minimal counterexample $X_a \subseteq V$.

We consider the set of candidates C^* whose minimal counterexamples have the largest possible weight under p, i.e. C^* contains those c for which $p(X_c)$ is maximal. Next, we define a partial order R on C^* as follows: if $b, c \in C^*$ then $b\,R\,c$ if and only if $X_c \subseteq X_b$ and $b \succ_v c$ for every $v \in X_c$. If $a \in C^*$ is a maximal element for this partial order then the counterexample X_a for a has the highest weight and no other candidate b with this property is better than a with respect to R.

We fix a candidate a who is maximal for R. We set $X := X_a$ and $D := D^{\mathcal{E}}_a(X_a)$. We also set $\overline{X} = V \setminus X$ and $\overline{D} = C \setminus D$. We shall show that if a does not weakly defeat a candidate b in X_a, then the minimal counterexamples of a and b are incomparable.

Lemma 3.5.5. *For every $b \in \overline{D}$ we have that $X \setminus X_b \neq \varnothing$ and $X_b \setminus X \neq \varnothing$.*

Proof. Let $b \in \overline{D}$, Then, by the definition of D, we have that $b \succ_v a$ for all $v \in X$. Suppose that $X \setminus X_b = \varnothing$, i.e. $X \subseteq X_b$. Then, $p(X_b) \geq p(X)$. Since $p(X)$ is maximal, we must have $b \in C^*$. Moreover, since $X \subseteq X_b$ and $b \succ_v a$ for all $v \in X$, we see that $b\,R\,a$, which is a contradiction because a is a maximal element of C^* with respect to R. Next, suppose that $X_b \setminus X = \varnothing$, i.e. $X_b \subseteq X$. Since X is a minimal counterexample for a, X_b is not a counterexample for a, and hence $q(D^{\mathcal{E}}_a(X_b)) \geq p(X_b)$. On the other hand, since $b \succ_v a$ for all $v \in X$, we get $q(D^{\mathcal{E}}_b(X_b)) \geq q(D^{\mathcal{E}}_a(X_b)) \geq p(X_b)$, which is a contradiction because X_b is a counterexample for b. □

Using Lemma 3.5.5 we can prove Lemma 3.5.4. We shall use the following:

(1) $D \neq \varnothing$.

(2) $q(\overline{D}) > 0$, and hence $\overline{D} \neq \varnothing$.

(3) $p(X) > 0$, and hence $X \neq \varnothing$.

(4) $\overline{X} \neq \varnothing$.

Since $a \in D$ we clearly have (1). Next, since X is a counterexample for a, we have that $0 \leq q(D) < p(X) \leq 1$, which shows that $q(\overline{D}) > 0$ and $p(X) > 0$. For claim (4), assume that $\overline{X} = \varnothing$. This means that $X = V$, therefore $X_b \setminus X = \varnothing$ for all b, and then Lemma 3.5.4 gives $\overline{D} = \varnothing$, a contradiction by claim (2).

We show now that every $c \in \overline{D}$ the set $D_c^{\mathcal{E}}(X_c)$ of candidates that are weakly defeated by c in X_c has sufficiently large weight:

$$\begin{aligned} q(D_c^{\mathcal{E}}(X_c)) &= q(D_c^{\mathcal{E}}(X_c) \cap D) + q(D_c^{\mathcal{E}}(X_c) \cap \overline{D}) \\ &\geq q(D_c^{\mathcal{E}}(X_c \cap X) \cap D) + q(D_c^{\mathcal{E}}(X_c \cap \overline{X}) \cap \overline{D}) \\ &\geq q(D_a^{\mathcal{E}}(X_c \cap X)) + q(D_c^{\mathcal{E}}(X_c \cap \overline{X}) \cap \overline{D}) \\ &\geq p(X_c \cap X) + q(D_c^{\mathcal{E}}(X_c \cap \overline{X}) \cap \overline{D}). \end{aligned} \tag{3.14}$$

The second inequality above holds because $D_a^{\mathcal{E}}(X_c \cap X) \subseteq D_c^{\mathcal{E}}(X_c \cap X) \cap D$: every candidate who is weakly defeated by a in $X_c \cap X$ is weakly defeated by a in X and also weakly defeated by c in $X_c \cap X$ because $c \succ_v a$ for all $v \in X$. The last inequality holds because $X_c \cap X \subsetneq X$ by Lemma 3.5.5. Since X is a minimal counterexample for a, we have that $X_c \cap X$ is not a counterexample for a, and hence $q(D_a^{\mathcal{E}}(X_c \cap X)) \geq p(X_c \cap X)$.

We want to find a suitable candidate $b \in \overline{D}$ for whom we will use (3.14) to show that $q(D_b^{\mathcal{E}}(X_b)) \geq p(X_b)$, a contradiction because X_b is a counterexample for b.

We distinguish two cases. First, assume that $p(\overline{X}) = 0$. We choose an arbitrary $b \in \overline{D}$. Since $p(X) = 1$, we get $p(X_b \cap X) = p(X_b)$, and then (3.14) gives $q(D_b^{\mathcal{E}}(X_b)) \geq p(X_b)$.

Now, assume that $p(\overline{X}) > 0$, and in particular $\overline{X} \neq \varnothing$. We consider the restricted election $\overline{\mathcal{E}} = (\overline{X}, \overline{D}, \sigma|_{\overline{X},\overline{D}})$, where $\sigma|_{\overline{X},\overline{D}}$ is the preference profile σ restricted to the preferences of the voters in $\overline{X}$ over the candidates in $\overline{D}$. Note that $\overline{\mathcal{E}}$ is a valid election because $\overline{X}$ and $\overline{D}$ are non-empty. Moreover, the number of voters in $\overline{\mathcal{E}}$ is smaller, because $X \neq \varnothing$. Since $\mathcal{E}$ is assumed to be an election with minimal number of voters for which Lemma 3.5.4 fails, we have that Lemma 3.5.4 applies to $\overline{\mathcal{E}}$ for any choice of weight vectors. We choose the re-normalized weight vectors $p' \in \Delta(\overline{X})$ and $q' \in \Delta(\overline{D})$ defined by

$$p'_v = \frac{p_v}{p(\overline{X})} \quad \text{and} \quad q'_c = \frac{q_c}{q(\overline{D})}$$

where $v \in \overline{X}$ and $c \in \overline{D}$. Since Lemma 3.5.4 holds for $\overline{\mathcal{E}}$, we can find $b \in \overline{D}$ such that $p'(S) \leq q'(D_b^{\mathcal{E}}(S))$ for all $S \subseteq \overline{X}$. Choosing $S = X_b \cap \overline{X}$, we get

$$p'(X_b \cap \overline{X}) \leq q'(D_b^{\mathcal{E}}(X_b \cap \overline{X})) = q'(D_b^{\mathcal{E}}(X_b \cap \overline{X}) \cap \overline{D}),$$

because $D_b^{\mathcal{E}}(X_b \cap \overline{X}) = D_b^{\mathcal{E}}(X_b \cap \overline{X}) \cap \overline{D}$. By the definition of p' and q', this inequality gives

$$q(D_b^{\mathcal{E}}(X_b \cap \overline{X}) \cap \overline{D}) \geq p(X_b \cap \overline{X}) \cdot \frac{q(\overline{D})}{p(\overline{X})} \geq p(X_b \cap \overline{X}),$$

where the last inequality holds because X is a counterexample for a and this implies that $q(D) < p(X)$, therefore $q(\overline{D}) > p(\overline{X})$. Going back to (3.14) we get

$$q(D_b^{\mathcal{E}}(X_b)) \geq p(X_b \cap X) + p(X_b \cap \overline{X}) = p(X_b)$$

and this contradicts the fact that X_b is a counterexample for b. This contradiction proves the existence of the desired candidate a.

Finally, in order to compute the candidate a, we can iterate over all the m candidates in C, and check if the (p, q)-domination graph of each one of them admits a fractional perfect matching. By Lemma 3.5.3, this can be done in strongly polynomial time. □

Lemma 3.5.4 allows us to choose the weight vectors p and q as we wish. In this way, we obtain a family of deterministic voting rules. Given p and q, for any election $\mathcal{E} = (V, C, \sigma)$ the rule that corresponds to p and q returns an arbitrary candidate a whose (p, q)-domination graph $G_{p,q}^{\mathcal{E}}(a)$ admits a fractional perfect matching. We shall study a specific choice of p and q.

Definition 3.5.6 (PluralityMatching Rule)**.** Let $\mathcal{E} = (V, C, \sigma)$ be an election. The PluralityMatching returns a candidate a (ties broken arbitrarily) whose $(p^{\text{uni}}, q^{\text{plu}})$-domination graph admits a fractional perfect matching, where $p_v^{\text{uni}} = \frac{1}{n}$ for all $v \in V$ and $q_c^{\text{plu}} = \frac{\text{plu}(c)}{n}$ for all $c \in C$.

Lemma 3.5.3 shows that if this rule returns a candidate a then $\text{plu}(D_a^{\mathcal{E}}(S)) \geq |S|$ for every $S \subseteq V$. In other words, in any subset S of voters, a weakly defeats a set of candidates with total plurality score at least $|S|$. We can also view the domination graph in a different way. Instead of having weights $1/n$ for each voter v and $\text{plu}(c)/n$ for each candidate c, we can let the voters and candidates have integral weights 1 and $\text{plu}(c)$ respectively. Then, we can replace each node c with weight $\text{plu}(c)$ by $\text{plu}(c)$ many nodes, each representing a unique voter whose top choice is c, that have weight 1 each and are connected to the same nodes of voters as c was. In this way, we obtain a bipartite graph whose vertices on both sides correspond to the voters.

Definition 3.5.7. Let $\mathcal{E} = (V, C, \sigma)$ be an election. For each candidate $a \in C$ we define the *integral domination graph* of a to be the bipartite graph $G^{\mathcal{E}}(a) = (V, V, E_a)$, where $(v, v') \in E_a$ if and only if $a \succcurlyeq_v \text{top}(v')$.

One can check that Hall's condition for the existence of an integral perfect matching in the integral domination graph is equivalent to the condition that we described above for the candidate returned by PluralityMatching. Then, Lemma 3.5.4 shows that there always exists a candidate a such that $G^{\mathcal{E}}(a)$ admits a perfect matching, and PluralityMatching returns one such candidate.

Corollary 3.5.8 (Gkatzelis-Halpern-Shah)**.** *Let $\mathcal{E} = (V, C, \sigma)$ be an election. For every candidate $a \in C$, $G^{\mathcal{E}}(a)$ admits a perfect matching if and only if $G_{p^{\text{uni}}, q^{\text{plu}}}^{\mathcal{E}}(a)$ admits a fractional perfect matching. So, there exists a candidate $a \in C$ whose integral domination graph $G^{\mathcal{E}}(a)$ admits a perfect matching: there exists a bijection $M : V \to V$ such that $a \succcurlyeq_v \text{top}(M(v))$ for all $v \in V$.*

Now we can prove that PluralityMatching has distortion at most $2+\alpha$ for α-decisive metric spaces. Since all metric spaces are 1-decisive, this implies that the optimal deterministic rule has distortion 3, which is optimal by the lower bound of 3 established in Theorem 3.1.4.

Theorem 3.5.9 (Gkatzelis-Halpern-Shah)**.** *For every $m \geq 3$ and $0 \leq \alpha \leq 1$,* PluralityMatching *has distortion $2 + \alpha$ for α-decisive metric spaces.*

Proof. First we prove the upper bound. Let (X, d) be an α-decisive metric space and let σ be the induced preference profile. Let a be the candidate returned by PluralityMatching and let b be any other candidate. We shall show that $\text{SC}(a) \leq (2+\alpha)\,\text{SC}(b)$.

Corollary 3.5.8 shows that the integral domination graph $G(a) := G^{\mathcal{E}}(a)$ admits a perfect matching $M : V \to V$ such that $a \succcurlyeq_v \text{top}(M(v))$ for every $v \in V$. Then,

$$\begin{aligned}
\text{SC}(a) &= \sum_{v\in V} d(v,a) \\
&\leq \sum_{v\in V} d(v,\text{top}(M(v))) \\
&\leq \sum_{v\in V} \big(d(v,b) + d(b,\text{top}(M(v)))\big) \\
&= \text{SC}(b) + \sum_{v\in V} d(b,\text{top}(M(v))) \\
&= \text{SC}(b) + \sum_{v\in V} d(b,\text{top}(v)) \\
&= \text{SC}(b) + \sum_{v\in V:\text{top}(v)\neq b} d(b,\text{top}(v)) \\
&\leq \text{SC}(b) + \sum_{v\in V:\text{top}(v)\neq b} \big(d(b,v) + d(v,\text{top}(v))\big) \\
&\leq \text{SC}(b) + \sum_{v\in V:\text{top}(v)\neq b} \big(d(b,v) + \alpha\, d(v,b)\big) \\
&\leq \text{SC}(b) + \sum_{v\in V} \big(d(b,v) + \alpha\, d(v,b)\big) \\
&= (2+\alpha)\,\text{SC}(b).
\end{aligned}$$

Since the α-decisive metric space, the election and the choice of b were arbitrary, this shows that PluralityMatching has distortion at most $2+\alpha$.

Next we show that this bound cannot be improved. Consider an election with $V = \{u, v\}$ (two voters) and $C = \{x, y, w\}$ (three candidates). Let the preference profile be the following: $\sigma_u = x \succ y \succ w$ and $\sigma_v = w \succ y \succ x$. Then, the integral domination graph of b has two edges, (u, v) and (v, u), and both edges together form a perfect matching. Therefore, PluralityMatching may return b. Now, consider an α-decisive metric, aligned to the preference profile given by the following undirected graph, where the distance between any two points is the shortest distance in the graph.

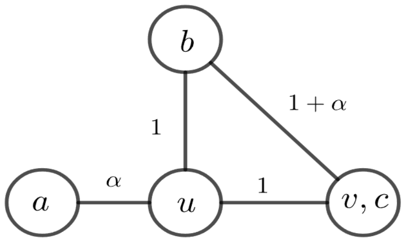

Figure 3.8: Example showing that Theorem 3.5.9 is optimal.

Then, $\text{SC}(b) = 2+\alpha$ and $\text{SC}(c) = 1$. It follows that the distortion of PluralityMatching is at least $2+\alpha$. This establishes the lower bound when $m = 3$. If $m > 3$ then we may obtain the same lower bound if we place all the other candidates sufficiently far away. ☐

In the case $m = 2$, PluralityMatching coincides with the majority rule, which chooses the top choice of a majority of voters. It is known that this rule has distortion $2 + \alpha$, which is optimal among all deterministic rules for every $\alpha \in [0, 1]$. Therefore, we get:

Corollary 3.5.10 (Gkatzelis-Halpern-Shah)**.** *For every m, the distortion of* PluralityMatching *is 3, which is the optimal distortion over all deterministic voting rules.*

This settles in the affirmative the conjecture that the optimal distortion of deterministic rules is 3 and can be achieved by a polynomial-time computable combinatorial rule.

CHAPTER 4

Learning-augmented algorithms

In this chapter we discuss an analysis framework, termed *algorithms with predictions* or *learning-augmented algorithms*, that has been defined for achieving refined bounds for the metric distortion problem using the guidance of predictions. A main question in the metric distortion problem is the information gap that the designer faces. The rankings of the voters are available, but the designer may also have historical data about the voters' choices in other matters that correlate with their preferences in the present matter, which may help in identifying their preferred outcome in the metric space. The idea is to enhance the algorithm with a prediction that it can use in order to improve its performance. The algorithm is then evaluated by its performance when the prediction is accurate (one measures the *consistency* of the algorithm) as well as when the prediction can be arbitrarily inaccurate (one measures the *robustness* of the algorithm).

Berger et al. introduced a family of algorithms, which is parameterized by $0 \leq \delta < 1$, known as BoostedSV$_\delta$, and obtained consistency and robustness bounds in terms of δ. The family BoostedSV$_\delta$ represents a learning-augmented adaptation of the SimultaneousVeto algorithm proposed in [59]. The SimultaneousVeto algorithm starts by assigning each candidate $c \in C$ a score equal to the number of voters ranking c first (its plurality score). Voters then continuously and simultaneously reduce the score of their least preferred candidate among those with remaining positive scores. The candidate whose score reaches zero last is selected as the winner. The BoostedSV$_\delta$ algorithm enhances SimultaneousVeto by boosting the initial score of the candidate $p \in C$ predicted to be optimal. The size of this boost is a carefully calibrated increasing function of δ, which also adjusts the rate at which voters reduce scores. As δ increases—indicating greater confidence in the prediction—the size of the boost grows. It is proved that the algorithm BoostedSV$_\delta$ achieves $\frac{3-\delta}{1+\delta}$-consistency and $\frac{3+\delta+13\delta^2-\delta^3}{(1+\delta)(1-\delta)^2}$-robustness. More recently, the authors presented in [21] a second family of algorithms, called LA$_\delta$, and proved that for any $\delta \in [0,1)$, the algorithm LA$_\delta$ achieves $\frac{3-\delta}{1+\delta}$-consistency and $\frac{3+\delta}{1-\delta}$-robustness. Moreover, they showed that this is the optimal trade-off. Namely, no deterministic algorithm that is $\frac{3-\delta}{1+\delta}$-consistent can be strictly better than $\frac{3+\delta}{1-\delta}$-robust, even for the line metric and just two candidates.

4.1 Metric distortion with predictions

We shall consider algorithms that receive as input a pair (σ, p), where σ is a preference profile and $p \in C$ is a *prediction* about the optimal candidate $c^*(d)$. We want to evaluate the performance of an algorithm through its *consistency* and its *robustness*. The consistency of ALG is defined as the distortion that ALG guarantees when the provided prediction is accurate, i.e. $p = c^*(d)$. More formally,

$$\text{consistency}(\text{ALG}) = \sup_{\sigma} \sup_{d: d \rhd \sigma} \frac{\text{SC}(\text{ALG}(\sigma, c^*(d)), d)}{\text{SC}(c^*(d), d)}.$$

The robustness of ALG is defined as the distortion that ALG guarantees with an arbitrary prediction, independently of how accurate this prediction may be. More formally,

$$\text{robustness}(\text{ALG}) = \sup_{\sigma} \sup_{p \in C} \sup_{d: d \rhd \sigma} \frac{\text{SC}(\text{ALG}(\sigma, p), d)}{\text{SC}(c^*(d), d)}.$$

In what follows, given a preference profile σ, a voter v, a candidate c and a subset of candidates $S \subseteq C$, we write $\text{top}(v)$ for the candidate ranked highest by v, $\text{plu}(c)$ for the number of voters who rank c as their top choice, $\text{bot}(v)$ for the candidate ranked lowest by v, and $\text{bot}_S(v)$ for the candidate in S ranked lowest by v. Note that $\text{bot}(v) = \text{bot}_C(v)$.

4.2 Tradeoff between robustness and consistency

The next theorem shows that even for instances with just two candidates, for any parameter $0 \le \delta < 1$, no deterministic algorithm can simultaneously achieve $\frac{3-\delta}{1+\delta}$-consistency and robustness strictly better than $\frac{3+\delta}{1-\delta}$.

Theorem 4.2.1 (Berger-Feldman-Gkatzelis-Tan)**.** *Let* $0 \le \delta < 1$ *and let* ALG *be a deterministic algorithm that is* $\frac{3-\delta}{1+\delta}$*-consistent. If* $\beta < \frac{3+\delta}{1-\delta}$ *then* ALG *is not* β*-robust.*

Proof. Let ALG be a $\frac{3-\delta}{1+\delta}$-consistent algorithm. For any $\epsilon > 0$ we define

$$r(n, \epsilon) := \frac{3 + \delta - \frac{4}{n} - 2\epsilon}{1 - \delta + \frac{4}{n}}.$$

Note that $r(n, \epsilon) < \frac{3+\delta}{1-\delta}$ for all $n, \epsilon > 0$. Since $\lim_{n\to\infty, \epsilon\to 0^+} r(n, \epsilon) = \frac{3+\delta}{1-\delta}$ and $\beta < \frac{3+\delta}{1-\delta}$, we may choose n, ϵ so that $\beta < r(n, \epsilon) < \frac{3+\delta}{1-\delta}$. We may also assume that $\epsilon < \frac{1}{n}$. To make this precise, if $\epsilon = \frac{1}{2n}$ then

$$r\big(n, \tfrac{1}{2n}\big) = \frac{3 + \delta - \frac{3}{n}}{1 - \delta + \frac{4}{n}} \to \frac{3 + \delta}{1 - \delta}.$$

and since $\beta < \frac{3+\delta}{1-\delta}$ we may then choose n large enough so that $\beta < r\big(n, \frac{1}{2n}\big) < \frac{3+\delta}{1-\delta}$.

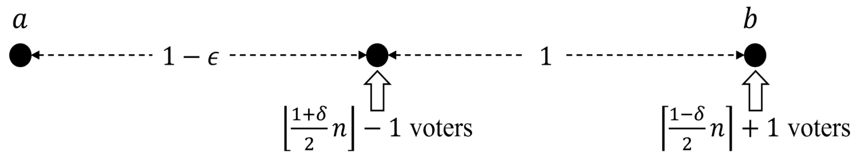

We fix n and ϵ as above. Now, consider the instance (V, C, d) where $C = \{a, b\}$ and the distance from a to b is equal to $2 - \epsilon$. We consider a set C of n voters that are located in the metric space so that $\lceil \frac{1-\delta}{2} n + 1 \rceil$ of them are placed on b and the remaining $\lfloor \frac{1+\delta}{2} n - 1 \rfloor$ of them are placed at distance 1 from b and at distance $1 - \epsilon$ from a, i,e, they are slightly closer to a than to b.

We easily check that b is the optimal candidate and the distortion of a is equal to

$$\begin{aligned}
\frac{\mathrm{SC}(a)}{\mathrm{SC}(b)} &= \frac{\lfloor \frac{1+\delta}{2}n - 1 \rfloor (1-\epsilon) + \lceil \frac{1-\delta}{2}n + 1 \rceil (2-\epsilon)}{\lfloor \frac{1+\delta}{2}n - 1 \rfloor \cdot 1} \\
&= \frac{\left(\lfloor \frac{1+\delta}{2}n - 1 \rfloor + \lceil \frac{1-\delta}{2}n + 1 \rceil\right)(1-\epsilon) + \lceil \frac{1-\delta}{2}n + 1 \rceil \cdot 1}{\lfloor \frac{1+\delta}{2}n - 1 \rfloor \cdot 1} \\
&= \frac{n(1-\epsilon) + \lceil \frac{1-\delta}{2}n + 1 \rceil}{\lfloor \frac{1+\delta}{2}n - 1 \rfloor} \geq \frac{n(1-\epsilon) + \frac{1-\delta}{2}n + 1}{\frac{1+\delta}{2}n} \\
&= \frac{(3-\delta)\frac{n}{2} + 1 - \epsilon n}{(1+\delta)\frac{n}{2}} = \frac{3 - \delta + \frac{2}{n} - 2\epsilon}{1+\delta} \\
&> \frac{3-\delta}{1+\delta} > 1.
\end{aligned}$$

This computation shows that if σ is the preference profile induced by (V, C, d) and the prediction is $p = b$, then ALG outputs b for the input (σ, b). Indeed, if it outputs a then the assumption that ALG is $\frac{3-\delta}{1+\delta}$-consistent leads to a contradiction.

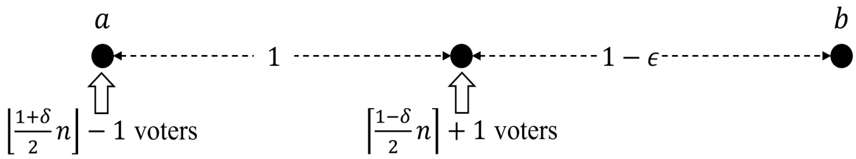

Now, we consider the following variant (V', C, d) of the previous instance. The $\lceil \frac{1-\delta}{2}n + 1 \rceil$ voters that were previously placed on b, are now placed at distance $1-\epsilon$ from b and at distance 1 from a, i,e, they are slightly closer to b than to a. The remaining $\lfloor \frac{1+\delta}{2}n - 1 \rfloor$ voters are located on a.

Note that this instance induces the same profile σ as before, therefore ALG has to output b for the input (σ, b). On the other hand, we can easily check that now a is the optimal candidate and the distortion of b is equal to

$$\begin{aligned}
\frac{\mathrm{SC}(b)}{\mathrm{SC}(a)} &= \frac{\lceil \frac{1-\delta}{2}n + 1 \rceil (1-\epsilon) + \lfloor \frac{1+\delta}{2}n - 1 \rfloor (2-\epsilon)}{\lceil \frac{1-\delta}{2}n + 1 \rceil \cdot 1} \\
&= \frac{n(1-\epsilon) + \lfloor \frac{1+\delta}{2}n - 1 \rfloor}{\lceil \frac{1-\delta}{2}n + 1 \rceil} \geq \frac{n(1-\epsilon) + \frac{1+\delta}{2}n - 2}{\frac{1-\delta}{2}n + 2} \\
&= \frac{(3+\delta)\frac{n}{2} - 2 - \epsilon n}{(1-\delta)\frac{n}{2} + 2} = \frac{3 + \delta - \frac{4}{n} - 2\epsilon}{1 - \delta + \frac{4}{n}} \\
&= r(n, \epsilon) > \beta.
\end{aligned}$$

This means that the robustness of ALG is greater than β. In other words, for any $\beta < \frac{3+\delta}{1-\delta}$ we have that ALG is not β-robust. □

4.3 Optimal tradeoff between consistency and robustness

Berger, Feldman, Gkatzelis and Tan introduced in [21] a family of algorithms, parametrized by $0 \leq \delta < 1$, which, in view of Theorem 4.2.1, achieve an optimal tradeoff between consistency and robustness. In a first version of their article the presented a family of algorithms, called BoostedSV$_\delta$, which has consistency

$$\text{consistency}(\text{BoostedSV}_\delta) \leq \frac{3-\delta}{1+\delta}$$

and robustness

$$\text{robustness}(\text{BoostedSV}_\delta) \leq T(\delta) := \frac{3+\delta+13\delta^2-\delta^3}{(1+\delta)(1-\delta)^2}$$

which is slightly weaker than $\frac{3+\delta}{1-\delta}$ for "small" values of δ.

Subsequently, the authors presented a second family of algorithms, called LA_δ, and proved that it achieves the optimal bounds.

Theorem 4.3.1 (Berger, Feldman, Gkatzelis and Tan)**.** *For every* $0 \leq \delta < 1$*,* LA_δ *achieves* $\frac{3-\delta}{1+\delta}$*-consistency and* $\frac{3+\delta}{1-\delta}$*-robustness.*

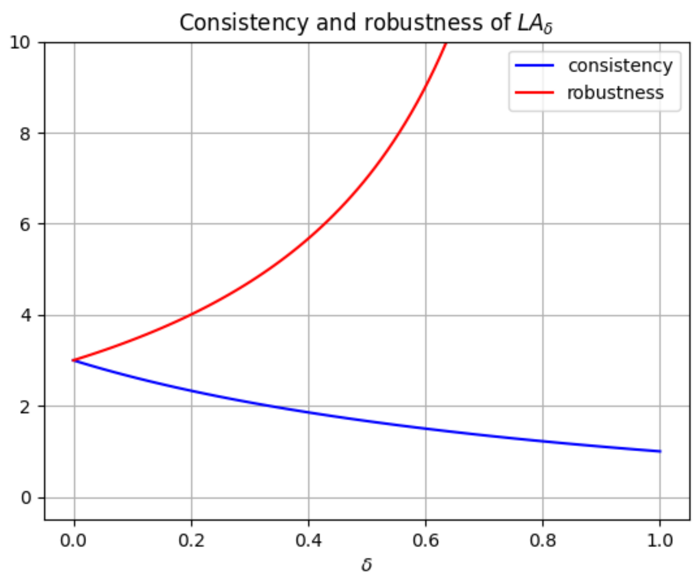

Figure 4.1: Consistency and robustness of LA_δ.

The key notion for the definition of the class LA_δ is the (p,q)-*veto core*. Recall from Chapter 3 that given an election $\mathcal{E} = (V, C, \sigma)$ and two normalized weight vectors $p \in \Delta(V)$ and $q \in \Delta(C)$, the (p,q)-domination graph of a candidate $a \in C$ is the vertex-weighted bipartite graph $G^{\mathcal{E}}_{p,q}(a) = (V, C, E_a, p, q)$, where $(v,c) \in E_a$ if and only if $a \succcurlyeq_v c$. Then, the (p,q)-veto core is the set of candidates $a \in C$ for which $G^{\mathcal{E}}_{p,q}(a)$ admits a fractional perfect matching.

Given $p \in \Delta(V)$ and $q \in \Delta(C)$, a (p,q)-algorithm is an algorithm with the property that if it receives a preference profile σ as an input then it returns a non-empty subset of candidates which is contained in the (p,q)-veto core of σ.

Now, given $0 \leq \delta < 1$, we consider a class of LA_δ-algorithms which satisfies the following: For any predicted optimal candidate $\widehat{c} \in C$ and any preference profile σ, the class LA_δ contains the $(p^{\text{uni}}, \widehat{q})$-algorithm where $p^{\text{uni}} = \left(\frac{1}{n}, \ldots, \frac{1}{n}\right)$ and $\widehat{q} \in \Delta(C)$ is defined by

$$\widehat{q}(\widehat{c}) = \frac{1-\delta}{1+\delta} \cdot \frac{\text{plu}(\widehat{c}) + \frac{2\delta n}{1-\delta}}{n}$$

and

$$\widehat{q}(c) = \frac{1-\delta}{1+\delta} \cdot \frac{\text{plu}(\widehat{c})}{n}$$

for all $c \in C \setminus \{\widehat{c}\}$.

Boosted Simultaneous Veto

In this section we study in detail the original algorithm BoostedSV$_\delta$ from [21]. The algorithm provides the candidates with some initial scores: Each candidate other than the prediction starts with his plurality score, while the initial score of the predicted candidate is boosted and equals his plurality score plus the boost $\mathsf{b} = \frac{2\delta n}{1-\delta}$. The sum of the initial scores is equal to

$$n + \mathsf{b} = \frac{1+\delta}{1-\delta} n.$$

After this initialization step, starting from time $t = 0$ and until time $t = 1$, the voters decrement continuously and simultaneously the score of their least-favorite candidate that still has positive score, at a rate of $\frac{n+\mathsf{b}}{n} = \frac{1+\delta}{1-\delta}$. We say that a voter v down-votes some candidate c if v decrements the score of c. We also say that a candidate is active if his score at that time point is positive. When the score of some candidate becomes 0, all voters that down-voted this candidate up to that point switch to down-voting their next least-favorite active candidate.

The first time point at which all candidates have reached score 0 is the end of the algorithm ($t = 1$). If the predicted candidate is among the candidates that have reached score 0 only at the end, then it is the output. Otherwise, the algorithm chooses arbitrarily one of the candidates whose score reached 0 only at the end. Therefore, we have the following.

Claim 4.3.2. For any $0 \leq \delta < 1$, the candidate output by a given execution of BoostedSV$_\delta$ is active at any time point $t < 1$.

We may also consider the following discrete implementation of BoostedSV$_\delta$. There are up to m rounds, divided by time points at which the score of some candidate reaches 0: at least one candidate is eliminated at the end of each round. Round i starts at time t_{i-1} and ends at time t_i, where $t_0 = 0$ and $t_k = 1$ if the total number of rounds is equal to k.

Given t_{i-1}, the time point t_i is determined as follows: We write score(c) for the current score of the candidate c at time t_{i-1} and denote by n_c the number of voters whose least favorite active candidate at time t_{i-1} is c. These are the voters that currently down-vote c. Let $\overline{c}$ be a candidate for whom the ratio

$$\Delta = \frac{\text{score}(c)}{n_c \cdot \frac{1+\delta}{1-\delta}}$$

is minimized. Then, $t_i = t_{i-1} + \Delta$ is the new time point, and a new round begins after appropriately updating the scores of all candidates. In the new round, all voters who were down-voting a candidate whose score reached 0 switch to down-voting their next least favorite active candidate. Note that more than one candidates may become inactive at the end of the same round. If this happens at least once, then the number k of rounds will be strictly less than m.

Algorithm 1: BoostedSV$_\delta$

Input: Preference profile σ, predicted optimal candidate $p \in C$

Output: a candidate $c \in C$

`/* All candidates start with score equal to their plurality, excepted for` p `who gets boosted. */`

$\forall c \in C \setminus \{p\},\ \text{score}(c) \leftarrow \text{plu}(c),$

$\text{score}(p) \leftarrow \text{plu}(p) + \text{b},\ \text{where b} = \frac{2\delta n}{1-\delta}.$

$i \leftarrow 1$

$t_0 \leftarrow 0$

`/* Voters decrement the score of their least-favorite candidate that still has positive score, continuously and simultaneously, at rate` $\frac{1+\delta}{1-\delta}$`. Following is a discrete implementation: */`

while $\exists c \in C$ s.t. $\text{score}(c) > 0$ **do**

- $A \leftarrow \{c \in C \mid \text{score}(c) > 0\}$ `// Set of current active candidates`
- $\forall c \in C : n_c \leftarrow |\{v \in V \mid \text{bot}_A(v) = c\}|$ `// Number of current "down-voters" for each candidate` c
- $\Delta \leftarrow \min_{c \in A} \frac{\text{score}(c)}{n \cdot \left(\frac{1+\delta}{1-\delta}\right)}$ `// Time interval until some candidate's score reaches` 0
- $t_i \leftarrow t_{i-1} + \Delta$ `// Time point when this happened`
- $\forall c \in C : \text{score}(c) \leftarrow \text{score}(c) - \Delta \cdot n_c \cdot \left(\frac{1+\delta}{1-\delta}\right)$ `// Appropriate score update`
- $i \leftarrow i + 1$

`/* All candidates now have score` 0`.` A `contains only candidates whose score reached` 0 `at` $t = 1$ `*/`

if $p \in A$ **then**

- **return** p

else

- **return** an arbitrary $c \in A$

Definition 4.3.3. Consider an execution of BoostedSV$_\delta$. We denote by $f(v, c) \in [0, 1]$ the fraction of time that the voter v spends down-voting the candidate c throughout the execution. Every voter spends the time interval $[0, 1]$ down-voting at a rate of $\frac{1+\delta}{1-\delta}$ and every candidate's score eventually reaches 0, therefore we have the following properties of f:

(i) For every $v \in V$,

$$\sum_{c \in C} f(v, c) = 1. \tag{4.1}$$

(ii) For every candidate $c \neq p$,

$$\sum_{v \in V} f(v, c) = \frac{1-\delta}{1+\delta} \cdot \text{plu}(c). \tag{4.2}$$

(iii) For the predicted candidate $c = p$,

$$\sum_{v \in V} f(v, p) = \frac{1-\delta}{1+\delta} \cdot \left(\text{plu}(p) + \frac{2\delta n}{1-\delta}\right). \tag{4.3}$$

Lemma 4.3.4. *Let $0 \leq \delta < 1$ and let a be the candidate output by a given execution of* BoostedSV$_\delta$. *Then, for every $v \in V$ and any $c \in C$ that satisfies $f(v, c) > 0$, we have $d(v, c) \geq d(v, a)$.*

Proof. From Claim 4.3.2 we know that a is active throughout the algorithm, which means that a belongs to the set A of *active candidates* in each step of the algorithm. Since $f(v,c) > 0$, we have $c = \text{bot}_A(v)$ in at least one step. Then, by the definition of bot_A we must have that $d(v,c) \geq d(v,a)$. □

We can now establish the consistency bound.

Theorem 4.3.5 (Berger-Feldman-Gkatzelis-Tan). *Let $0 \leq \delta < 1$ and let a be the candidate output by a given execution of* BoostedSV_δ *that is provided with the correct prediction $p = c^*$. Then, we have that*

$$\text{SC}(a) \leq \frac{3-\delta}{1+\delta} \cdot \text{SC}(c^*).$$

Proof. We give an upper bound for the social cost of the returned candidate a using Lemma 4.3.4 and the properties of $f(v,c)$ from Definition 4.3.3, as follows:

$$\begin{aligned}
\sum_{v\in V} d(v,a) &\leq \sum_{v\in V}\sum_{c\in C} f(v,c)\cdot d(v,a) \\
&\leq \sum_{v\in V}\sum_{c\in C} f(v,c)\cdot d(v,c) \\
&\leq \sum_{v\in V}\sum_{c\in C} f(v,c)\cdot \big(d(v,c^*) + d(c^*,c)\big) \\
&= \text{SC}(c^*) + \sum_{c\in C}\sum_{v\in V} f(v,c)\cdot d(c^*,c) \\
&= \text{SC}(c^*) + \sum_{c\in C\setminus\{c^*\}} \left(\sum_{v} f(v,c)\right)\cdot d(c^*,c) \\
&= \text{SC}(c^*) + \sum_{c\in C\setminus\{c^*\}} \frac{1-\delta}{1+\delta}\cdot \text{plu}(c)\cdot d(c^*,c) \\
&= \text{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c\in C}\sum_{\{v\in V:\, \text{top}(v)=c\}} d(c^*,c) \\
&\leq \text{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c\in C}\sum_{\{v\in V:\, \text{top}(v)=c\}} \big(d(c^*,v) + d(v,c)\big) \\
&\leq \text{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c\in C}\sum_{\{v\in V:\, \text{top}(v)=c\}} 2d(c^*,v) \\
&\leq \text{SC}(c^*) + \frac{2(1-\delta)}{1+\delta}\text{SC}(c^*) = \frac{3-\delta}{1+\delta}\text{SC}(c^*),
\end{aligned}$$

where we have also used the fact that if $c = \text{top}(v)$ then $c \succcurlyeq_v c^*$. □

Next, we prove the robustness bound in the case where the returned candidate is the prediction.

Theorem 4.3.6 (Berger-Feldman-Gkatzelis-Tan). *Let $0 \leq \delta < 1$ and assume that p is the candidate output by a given execution of* BoostedSV_δ (*the algorithm returns the prediction*). *Then, we have that*

$$\text{SC}(p) \leq \frac{3+\delta+13\delta^2-\delta^3}{(1+\delta)(1-\delta)^2}\cdot \text{SC}(c^*).$$

The main ingredient in the proof of Theorem 4.3.6 is Lemma 4.3.9 which gives an upper bound for the distance between the optimal candidate c^* and the prediction p in terms of the optimal social cost. We shall introduce a graph (the veto map) that takes into account the way in which the down-voting of each voter is distributed across the up-votes of other voters and in Lemma 4.3.8 we shall use the structure of this graph in order to obtain lower bounds for the sum of the distances of pairs of voters that form an edge from c^*.

Definition 4.3.7 (the veto map)**.** Let $0 \leq \delta < 1$ and consider an execution of BoostedSV$_\delta$. The *veto map* associated with this execution is an edge-weighted directed graph $G = (V, E)$ where each vertex corresponds to a voter in V and an edge $(v, v') \in E$ if and only if v has down-voted the top candidate of v' at some step of the execution, i.e. if $f(v, \text{top}(v')) > 0$. The weight of an edge $(v, v') \in E$ is defined by

$$w(v, v') := \frac{f(v, \text{top}(v')) \cdot \frac{1+\delta}{1-\delta}}{\text{plu}(\text{top}(v'))}. \tag{4.4}$$

For each edge $(v, v') \in E$, the numerator of the right-hand side of (4.4) is equal to the amount by which v decreased the score of top(v') throughout the execution of BoostedSV$_\delta$, i.e. the amount of time $f(v, \text{top}(v'))$ that v spent down-voting top(v') multiplied by the rate $\frac{1+\delta}{1-\delta}$. The denominator of the right-hand side (4.4) is the plurality of this candidate, i.e. the number of voters that up-voted him. So, for each v and c, the veto map "distributes" the amount by which v reduced the score of c equally among the weights of the edges from v to each voter v' such that $c = \text{top}(v')$.

In the next lemma we observe that for every edge $(v, v') \in E$ we can give a lower bound for the combined contribution of v and v' to the social cost of the optimal candidate c^*. The idea is that if v down-votes the candidate closest to v' then v must be far from that candidate and hence from v' too, and if one of v or v' is close to c^* then the other voter must be far from c^*.

Lemma 4.3.8. *Let* $0 \leq \delta < 1$ *and let* $G = (V, E)$ *be the veto map associated with a given execution of* BoostedSV$_\delta$. *If* $(v, v') \in E$ *then*

$$d(v, c^*) + d(v', c^*) \geq \frac{d(c^*, p)}{2}.$$

Proof. For any $v \in V$, by the triangle inequality we get

$$d(p, v) \geq d(c^*, p) - d(v, c^*). \tag{4.5}$$

By the definition of the veto map, if $(v, v') \in E$ then v has down-voted the top candidate $c = \text{top}(v')$ of v'. We have

$$d(c^*, c) \geq d(c, v) - d(v, c^*) \geq d(p, v) - d(v, c^*) \geq d(c^*, p) - 2d(v, c^*), \tag{4.6}$$

where the first inequality follows from the triangle inequality, the second holds because $d(c, v) \geq d(p, v)$ by Lemma 4.3.4, and for the last inequality we use (4.5). If $(v, v') \in E$, then

$$d(c^*, v') + d(v', c) \geq d(c^*, c)$$

by the triangle inequality, and since $c = \text{top}(v') \succcurlyeq_{v'} c^*$ we get $d(c^*, v') \geq d(c, v')$ we get $2d(c^*, v') \geq d(c^*, c)$. Then, using (4.6) we obtain

$$2d(c^*, v') \geq d(c^*, p) - 2d(v, c^*),$$

which is equivalent to the assertion of the lemma. □

Now, we can state and prove the following main lemma.

Lemma 4.3.9. *Let* $0 \le \delta < 1$ *and assume that* p *is the candidate output of a given execution of* BoostedSV_δ*, i.e. the algorithm returns the prediction. Then,*

$$\text{SC}(c^*) \ge \frac{1-\delta}{2} n \cdot \frac{d(c^*,p)}{2}.$$

Note that we would have been able to obtain Lemma 4.3.9 directly from Lemma 4.3.8 if the set E of veto map edges contained a matching of size at least $\frac{1-\delta}{2}n$. Instead of finding such a matching, we partition the set V of voters into the set V_{in} that contains the voters whose distance from c^* is less than $\frac{d(c^*,p)}{2}$ and the set V_{out} that contains the voters whose distance from c^* is at least $\frac{d(c^*,p)}{2}$. Then, we show that the subgraph of the veto map induced by the vertices in V_{in} contains in its edges a fractional matching of total size $\frac{1-\delta}{2}n - |V_{\text{out}}|$. Since each voter from V_{out} contributes at least $\frac{d(c^*,p)}{2}$ to the optimal social cost and, by Lemma 4.3.8, each edge of the fractional matching (with vertices in V_{in}) also contributes at least $\frac{d(c^*,p)}{2}$ to the optimal social cost, we will show that the combined contribution from V_{in} and V_{out} is at least $\frac{1-\delta}{2}n \cdot \frac{d(c^*,p)}{2}$, which leads to the lower bound for $\text{SC}(c^*)$ in Theorem 4.3.6.

Proof of Lemma 4.3.9. Let V_{in} denote the set of $v \in V$ who satisfy $d(v,c^*) < \frac{d(c^*,p)}{2}$ and V_{out} denote the set of voters who satisfy $d(v,c^*) \ge \frac{d(c^*,p)}{2}$. Note that if $|V_{\text{out}}| \ge \frac{1-\delta}{2}n$ then the lemma is true because

$$\text{SC}(c^*) \ge \sum_{v \in V_{\text{out}}} d(v,c^*) \ge \frac{1-\delta}{2} n \cdot \frac{d(c^*,p)}{2}.$$

So, for the rest of the proof we assume that $|V_{\text{out}}| < \frac{1-\delta}{2}n$ and consider $x > 0$ so that $|V_{\text{out}}| = \frac{1-\delta}{2}n - xn$. Then, the contribution of the voters in V_{out} is

$$\sum_{v \in V_{\text{out}}} d(v,c^*) \ge \left(\frac{1-\delta}{2}n - xn\right) \frac{d(c^*,p)}{2}, \tag{4.7}$$

and $|V_{\text{in}}| = n - |V_{\text{out}}| = \frac{1+\delta}{2}n + xn$. Our aim is to show that

$$\sum_{v \in V_{\text{in}}} d(v,c^*) \ge xn \frac{d(c^*,p)}{2}, \tag{4.8}$$

which gives the lemma.

Let $G = (V,E)$ be the veto map associated with the given execution of BoosetdSV_δ. In order to complete the proof of (4.8), and hence the proof of Lemma 4.3.9, we need a number of inequalities about the veto map edge weights.

Claim 4.3.10. *For all* $v' \in V_{\text{in}}$ *we have that*

$$\sum_{v \in V_{\text{in}}} w(v,v') \le 1.$$

Proof. Let $v' \in V_{\text{in}}$. Note that

$$\sum_{v \in V_{\text{in}}} w(v,v') = \sum_{v \in V_{\text{in}}} \frac{f(v,\text{top}(v')) \cdot \frac{1+\delta}{1-\delta}}{\text{plu}(\text{top}(v'))} \le \sum_{v \in V} \frac{f(v,\text{top}(v'))}{\text{plu}(\text{top}(v'))} \cdot \frac{1+\delta}{1-\delta} = 1.$$

Here we use the fact that $\sum_{v \in V} f(v,c) = \frac{1-\delta}{1+\delta} \cdot \text{plu}(c)$ for every candidate $c \neq p$. Note that $v' \in V_{\text{in}}$ and hence $\text{top}(v') \neq p$: the voter v' strictly prefers c^* over p, because $d(c^*,v') < d(c^*,p)/2 \le$

$\big(d(c^*, v') + d(v', p)\big)/2$ by the definition of V_{in} and the triangle inequality. ☐

Claim 4.3.11. *For all $v \in V$ we have that*

$$\sum_{v' \in V_{\text{in}}} w(v, v') \leq \frac{1+\delta}{1-\delta}.$$

Proof. For every $v \in V$ we have

$$\begin{aligned}\sum_{v' \in V_{\text{in}}} w(v, v') &= \sum_{v' \in V_{\text{in}}} \frac{f(v, \text{top}(v')) \cdot \frac{1+\delta}{1-\delta}}{\text{plu}(\text{top}(v'))} = \frac{1+\delta}{1-\delta} \sum_{c \in C} \sum_{\{v' \in V_{\text{in}} :\, \text{top}(v') = c\}} \frac{f(v, c)}{\text{plu}(\text{top}(v'))} \\ &\leq \frac{1+\delta}{1-\delta} \sum_{c \in C} \sum_{\{v' \in V :\, \text{top}(v') = c\}} \frac{f(v, c)}{\text{plu}(\text{top}(v'))} = \frac{1+\delta}{1-\delta} \sum_{c \in C} \text{plu}(c) \cdot \frac{f(v, c)}{\text{plu}(\text{top}(c))} \\ &= \frac{1+\delta}{1-\delta} \sum_{c \in C} f(v, c) = \frac{1+\delta}{1-\delta},\end{aligned}$$

where in the last step we used the fact that $\sum_{c \in C} f(v, c) = 1$ for every $v \in V$. ☐

The final claim gives a lower bound for the "amount" of up-votes from voters $v' \in V_{\text{in}}$ to candidates who are down-voted by some other voters $v \in V_{\text{in}}$. We will show that even if all the voters in V_{out} keep down-voting the up-voted candidates of all the voters in V_{in}, there are still some up-voted candidates left, and these must have been down-voted by voters in V_{in}.

Claim 4.3.12. *Let $E_{\text{in}} := \{(v, v') \in E : v, v' \in V_{\text{in}}\}$ be the set of edges in the veto map for which both vertices belong to V_{in}. Then,*

$$\sum_{(v, v') \in E_{\text{in}}} w(v, v') \geq \frac{2xn}{1-\delta}.$$

Proof. First observe that for every $v' \in V_{\text{in}}$ we have

$$\sum_{v \in V} f(v, \text{top}(v')) = \frac{1-\delta}{1+\delta} \text{plu}(\text{top}(v')).$$

This is because, by (4.2), for every $c \neq p$ we have

$$\sum_{v \in V} f(v, c) = \frac{1-\delta}{1+\delta} \text{plu}(c)$$

and $\text{top}(v') \neq p$ for all $v' \in V_{\text{in}}$ because for every $v' \in V_{\text{in}}$ we have that

$$2d(v', c^*) \leq d(x^*, p) \leq d(v', p) + d(v', c^*)$$

which implies $d(v', p) > d(v', c^*)$, and this shows that p cannot be the top candidate of v'. It follows that

$$\sum_{v' \in V_{\text{in}}} \sum_{v \in V} w(v, v') = \sum_{v' \in V_{\text{in}}} \sum_{v \in V} \frac{f(v, \text{top}(v')) \cdot \frac{1+\delta}{1-\delta}}{\text{plu}(\text{top}(v'))} = \sum_{v' \in V_{\text{in}}} \frac{\text{plu}(\text{top}(v'))}{\text{plu}(\text{top}(v'))} = |V_{\text{in}}|. \tag{4.9}$$

From Claim 4.3.11 we also have

$$\sum_{v' \in V_{\text{in}}} \sum_{v \in V_{\text{out}}} w(v,v') = \sum_{v \in V_{\text{out}}} \sum_{v' \in V_{\text{in}}} w(v,v') \leq \sum_{v \in V_{\text{out}}} \frac{1+\delta}{1-\delta} = \frac{1+\delta}{1-\delta} |V_{\text{out}}|. \tag{4.10}$$

We write

$$\sum_{(v,v') \in E_{\text{in}}} w(v,v') = \sum_{v' \in V_{\text{in}}} \sum_{v \in V_{\text{in}}} w(v,v') = \sum_{v' \in V_{\text{in}}} \sum_{v \in V} w(v,v') - \sum_{v' \in V_{\text{in}}} \sum_{v \in V_{\text{out}}} w(v,v')$$

and using (4.9) and (4.10) we obtain

$$\begin{aligned}\sum_{(v,v') \in E_{\text{in}}} w(v,v') \geq |V_{\text{in}}| - \frac{1+\delta}{1-\delta} |V_{\text{out}}| &= \left(\frac{1+\delta}{2} n + xn\right) - \frac{1+\delta}{1-\delta} \left(\frac{1-\delta}{2} n - xn\right) \\ &= \frac{2xn}{1-\delta},\end{aligned}$$

as claimed. □

Using Lemma 4.3.8 and Claims 4.3.10–4.3.12 we give a lower bound for the contribution of voters in V_{in} to the optimal score. If we could find a matching $M \subseteq E_{\text{in}}$ of size $|M| \geq xn$, then by Lemma 4.3.8 we would immediately get (4.8) and hence the lemma. We can define a fractional matching with the property that for every edge $(v,v') \in E_{\text{in}}$ it includes a fraction of this edge equal to

$$m(v,v') := w(v,v') \cdot \frac{1-\delta}{2}.$$

In order to show that this is a valid fractional matching, we first show that for any vertex $v \in V_{\text{in}}$ the total fraction of the (incoming and outgoing) edges adjacent to it that are included in the matching is at most 1. From Claim 4.3.10 and Claim 4.3.11 we see that

$$\begin{aligned}\sum_{v' \in V_{\text{in}}} m(v',v) + \sum_{v' \in V_{\text{in}}} m(v,v') &= \frac{1-\delta}{2} \left(\sum_{v' \in V_{\text{in}}} w(v',v) + \sum_{v' \in V_{\text{in}}} w(v,v') \right) \\ &\leq \frac{1-\delta}{2} \cdot \left(1 + \frac{1+\delta}{1-\delta}\right) = 1.\end{aligned} \tag{4.11}$$

Also, from Claim 4.3.12 we get

$$\sum_{(v,v') \in E_{\text{in}}} m(v,v') = \sum_{(v,v') \in E_{\text{in}}} w(v,v') \cdot \frac{1-\delta}{2} \geq \frac{1-\delta}{2} \cdot \frac{2xn}{1-\delta} = xn. \tag{4.12}$$

Using (4.11) and rearranging the sum, we write

$$\begin{aligned}\sum_{v \in V_{\text{in}}} d(v,c^*) &\geq \sum_{v \in V_{\text{in}}} \left(\sum_{v' \in V_{\text{in}}} m(v',v) + \sum_{v' \in V_{\text{in}}} m(v,v') \right) \cdot d(v,c^*) \\ &= \sum_{(v,v') \in E} m(v,v') \cdot \big(d(v,c^*) + d(v',c^*)\big).\end{aligned}$$

Then, Lemma 4.3.8 and (4.12) give us that

$$\sum_{v\in V_{\text{in}}} d(v,c^*) \geq \sum_{(v,v')\in E} m(v,v')\cdot \frac{d(c^*,p)}{2} \geq xn\cdot\frac{d(c^*,p)}{2}.$$

This estimate, together with the lower bound (4.7) for the contribution of voters from V_{out} to the optimal social cost, allow us to prove the lower bound for the optimal cost. We have that

$$\begin{aligned}\text{SC}(c^*) &= \sum_{v\in V_{\text{out}}} d(v,c^*) + \sum_{v\in V_{\text{in}}} d(v,c^*)\\ &\geq \left(\frac{1-\delta}{2}n - xn\right)\frac{d(c^*,p)}{2} + xn\,\frac{d(c^*,p)}{2} \geq \frac{1-\delta}{2}n\cdot\frac{d(c^*,p)}{2},\end{aligned}$$

and the proof of Lemma 4.3.9 is complete. □

Now, we are able to prove the main theorem.

Proof of Theorem 4.3.6. We give an upper bound for the social cost of the returned candidate p using Lemma 4.3.4, the properties of $f(v,c)$ from Definition 4.3.3 and finally Lemma 4.3.8, as follows:

$$\begin{aligned}\text{SC}(p) = \sum_{v\in V} d(v,p) &\leq \sum_{v\in V}\sum_{c\in C} f(v,c)\cdot d(v,p) \leq \sum_{v\in V}\sum_{c\in C} f(v,c)\cdot d(v,c)\\ &\leq \sum_{v\in V}\sum_{c\in C} f(v,c)\cdot \big(d(v,c^*)+d(c^*,c)\big)\\ &= \text{SC}(c^*) + \sum_{c\in C}\sum_{v\in V} f(v,c)\cdot d(c^*,c)\\ &= \text{SC}(c^*) + \sum_{c\in C\setminus\{p\}}\sum_{v\in V} f(v,c)\cdot d(c^*,c) + \sum_{v\in V} f(v,p)\cdot d(c^*,p)\\ &= \text{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c\in C}\text{plu}(c)\cdot d(c^*,c) + \frac{2\delta n}{1-\delta}\,d(c^*,p)\\ &= \text{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c\in C}\sum_{\{v\in V:\,\text{top}(v)=c\}} d(c^*,c) + \frac{2\delta n}{1-\delta}\,d(c^*,p).\end{aligned}$$

Recalling (4.2) and (4.3) we get

$$\begin{aligned}\text{SC}(p) &\leq \text{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c\in C}\sum_{\{v\in V:\,\text{top}(v)=c\}} \big(d(c^*,v)+d(v,c)\big) + \frac{2\delta n}{1-\delta}\,d(c^*,p)\\ &\leq \text{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c\in C}\sum_{\{v\in V:\,\text{top}(v)=c\}} 2d(c^*,v) + \frac{2\delta n}{1-\delta}d(c^*,p)\end{aligned}$$

where we have also used the fact that if $c = \text{top}(v)$ then $d(v,c)\leq d(v,c^*)$. Therefore,

$$\begin{aligned}\text{SC}(p) &\leq \text{SC}(c^*) + \frac{2(1-\delta)}{1+\delta}\,\text{SC}(c^*) + \frac{2\delta n}{1-\delta}d(c^*,p)\\ &\leq \frac{3-\delta}{1+\delta}\,\text{SC}(c^*) + \frac{2\delta n}{1-\delta}\,d(c^*,p)\\ &\leq \frac{3-\delta}{1+\delta}\,\text{SC}(c^*) + \frac{2\delta n}{1-\delta}\,\frac{2\,\text{SC}(c^*)}{\frac{1-\delta}{2}n}\end{aligned}$$

$$= \left(\frac{3-\delta}{1+\delta} + \frac{8\delta}{(1-\delta)^2} \right) \mathrm{SC}(c^*)$$
$$= \frac{3+\delta+13\delta^2-s^3}{(1+\delta)(1-\delta)} \mathrm{SC}(c^*) = T(\delta)\,\mathrm{SC}(c^*),$$

□

Next, we prove the same robustness bound even when BoostedSV$_\delta$ does not output the prediction.

Theorem 4.3.13 (Berger-Feldman-Gkatzelis-Tan). *Let* $0 \le \delta < 1$ *and assume that the candidate output by a given execution of* BoostedSV$_\delta$ *is a candidate* $a \in C \setminus \{p\}$*, i.e. the algorithm does not return the prediction. Then,*

$$\mathrm{SC}(a) \le T(\delta) \cdot \mathrm{SC}(c^*).$$

We note that the approach that we used for the proof of Theorem 4.3.6 cannot be useful because the prediction can be arbitrarily far away from the optimal candidate. What can be shown is that if $a \neq p$ is returned, then a would have been again returned in the instance which is identical with the given one except that a is the prediction. In other words, one can reduce instances in which the prediction is not the output to instances in which it is, while maintaining the distortion.

Proposition 4.3.14 (Berger-Feldman-Gkatzelis-Tan). *Let* $0 \le \delta < 1$ *and assume that the candidate output by a given execution of* BoostedSV$_\delta$ *is a candidate* $a \in C\setminus\{p\}$*. Then, if we execute* BoostedSV$_\delta$ *on the same preference profile but having* a *as the prediction, we get* a *as an output.*

Before presenting the proof of Proposition 4.3.14, we first show how it implies Theorem 4.3.13.

Proof of Theorem 4.3.13. By Proposition 4.3.14, a is returned by BoostedSV$_\delta$ when executed on the same preference profile as in the given execution, with candidate a also being the prediction. That is, the new execution returns the new prediction (candidate a). The claim follows by Theorem 4.3.6. □

It remains to prove Proposition 4.3.14. Consider two executions of BoostedSV$_\delta$. The first one (called "p execution") has p as the prediction and a as the output, while the second one (called "a execution") has a as the prediction. We shall show that a is also the output in the a execution.

For any candidate $c \in C$, any $t \in [0, 1]$ and any $b \in \{a, p\}$ we define the following:

- $\mathrm{dec}_b(c, t)$ is the amount of score decremented from c up to time t, in the continuous interpretation of the b execution.
- $\mathrm{score}_b(c, t)$ is the score of c at time t in the continuous interpretation of the b execution. Therefore, $\mathrm{score}_b(c, t) = s_c^b - \mathrm{dec}_b(c, t)$, where s_c^b is the initial score of c in the b execution.

Claim 4.3.15. For any $c \in C$ and $b \in \{a, p\}$ the functions $\mathrm{dec}_b(c, \cdot)$ and $\mathrm{score}_b(c, \cdot)$ are piece-wise linear and in particular continuous in t.

Proof. For any $b \in \{a, p\}$ and any fixed $c \in C$, the function $\mathrm{dec}_b(c, \cdot)$ is piece-wise linear (in t) and can be computed using the round boundary time points t_i from the discrete implementation of BoostedSV$_\delta$ (see Algorithm 1). More precisely, $\mathrm{dec}_b(c, 0) = 0$, and if there are k-rounds in the While loop then for each $i = 1, \dots, k$ and for any $t \in [t_{i-1}, t_i]$ (recall that $t_0 = 0$ and $t_k = 1$) we have that

$$\mathrm{dec}_b(c, t) = \mathrm{dec}_b(c, t_{i-1}) + n_c^i \cdot \frac{1+\delta}{1-\delta} \cdot (t - t_{i-1}),$$

where n_c^i is the number of voters who down-voted c in round i (recall also that $\frac{1+\delta}{1-\delta}$ is the down-voting rate per voter). □

Our next claim is that both executions are equivalent up to the time point where the score of p reaches 0 in the a execution.

Claim 4.3.16. *For any* $c \in C$ *and for any* t *such that* $\text{score}_a(p,t) > 0$ *we have* $\text{dec}_p(c,t) = \text{dec}_a(c,t)$.

Proof. Consider the implementation of BoostedSV$_\delta$ as given by Algorithm 1. In the initialization phase of the two executions, the only difference is in the initial scores of p and a, where in the p execution p gets the boost and in the a execution a gets the boost. Since a is the output in the p execution, it has positive score throughout the p execution, and hence it always belongs to the set of active candidates A. Therefore, before the score of p reaches 0 in the a execution, every step of the While loop in Algorithm 1 is computed in the same way in both executions. □

The next lemma shows that when p does not get the boost, the rest of the candidates can only be down-voted faster.

Lemma 4.3.17. *For all* $c \in C \setminus \{p\}$ *and* $t \in [0,1]$ *we have that* $\text{dec}_p(c,t) \leq \text{dec}_a(c,t)$.

Proof. Assume that there is a candidate $c \neq p$ such that for some $t \in [0,1]$ we have $\text{dec}_a(c,t) < \text{dec}_p(c,t)$. We define

$$T := \{t \in [0,1] : \text{ there exists } c \in C \setminus \{p\} \text{ such that } \text{dec}_a(c,t) < \text{dec}_p(c,t)\}.$$

Note that if $t \in T$ then $t > 0$, because the amount of score decremented from any candidate at time $t = 0$ equals 0 at both executions of the algorithm.

Let $t_{\text{inf}} = \inf(T)$. We can check that $t_{\text{inf}} \notin T$, because for any $t \in T$ and any candidate c for which $\text{dec}_a(c,t) < \text{dec}_p(c,t)$, by the continuity of the functions $\text{dec}_a(c,\cdot)$ and $\text{dec}_p(c,\cdot)$ we can find $\epsilon > 0$ small enough so that $\text{dec}_a(c,t-\epsilon) < \text{dec}_p(c,t-\epsilon)$, and hence $t - \epsilon \in S$. Furthermore, we can check that $t_{\text{inf}} < 1$ because otherwise T would be empty, a contradiction. Finally, the continuity of the functions $\text{dec}_a(c,\cdot)$ and $\text{dec}_p(c,\cdot)$ also implies that at time $t = t_{\text{inf}}$ there must exist a candidate $c \neq p$ for which the following hold true:

(i) $\text{dec}_a(c, t_{\text{inf}}) = \text{dec}_p(c, t_{\text{inf}})$.

(ii) There exists $\epsilon_0 > 0$ small enough so that for any $0 < \epsilon < \epsilon_0$ we have $\text{dec}_a(c, t_{\text{inf}} + \epsilon) < \text{dec}_p(c, t_{\text{inf}} + \epsilon)$.

In other words, the amount of score decremented from c until time t_{inf} is the same in both executions, but also slightly beyond time t_{inf}, score(c) is decremented more in the p execution. It follows that at time t_{inf} we have $\text{score}_p(c, t_{\text{inf}}) > 0$. This also implies that

$$\text{score}_a(c, t_{\text{inf}}) \geq \text{plu}(c) - \text{dec}_a(c, t_{\text{inf}}) = \text{plu}(c) - \text{dec}_p(c, t_{\text{inf}}) = \text{score}_p(c, t_{\text{inf}}) > 0,$$

where the first inequality holds by Claim 4.3.15 since the initial score of each candidate is at least its plurality score, the first equality follows from (i), and the second equality holds because $c \neq p$ and the initial score of each candidate other than p equals its plurality score in the p execution.

We have explained that the score of c at t_{inf} is positive in both executions, but the decrement rate is higher in the p execution. Then, we must have that at time t_{inf} the number of voters who are down-voting c is larger in the p execution than in the a execution. In particular, there exists a voter v who at time t_{inf} down-votes c in the p execution, but in the a execution she down-votes some other candidate $c' \neq c$, although c is also active in the a execution at that time. So, v ranks c' lower than c and we also have the following:

(iii) $\text{score}_a(c', t_{\text{inf}}) > 0$, because v down-votes c' in the a execution at time t_{inf} (a candidate must have positive score if he is being down-voted at any given time).

(iv) $\text{score}_p(c', t_{\text{inf}}) = 0$, because otherwise v would have down-voted c' (or some other active candidate that v ranks lower than c) in the p execution at time t_{inf}.

Note that $c' \neq a$ since a is the output in the p execution (and in particular $\text{score}_p(a, t_{\text{inf}}) > 0$). We also note that $c' \neq p$, because for any t (and in particular for $t = t_{\text{inf}}$) we have

$$\text{score}_p(p, t) \geq \text{score}_a(p, t). \tag{4.13}$$

To see this, note first that it is trivial for any t for which $\text{score}_a(p, t) = 0$. For times t in which $\text{score}_a(p, t) > 0$ we have

$$\text{score}_p(p, t) = \text{plu}(p) + \text{b} - \text{dec}_p(p, t) = \text{plu}(p) + \text{b} - \text{dec}_a(p, t) = \text{score}_a(p, t) + \text{b} \geq \text{score}_a(p, t),$$

where the first equality holds by Claim 4.3.15, the second one holds by Claim 4.3.16 and the third holds again by Claim 4.3.15.

We have checked that $c' \notin \{a, p\}$, and this means that the initial score of c' in both executions is $\text{plu}(c')$. By Claim 4.3.15 we get

$$\text{dec}_a(c', t_{\text{inf}}) = \text{plu}(c') - \text{score}_a(c', t_{\text{inf}}) < \text{plu}(c') - \text{score}_p(c', t_{\text{inf}}) = \text{dec}_p(c', t_{\text{inf}}),$$

using (iii) and (iv). Then, by definition, we get that $t_{\text{inf}} \in T$, which is a contradiction. □

Lemma 4.3.17 shows that if p does not get the boost then all other candidates can only be down-voted faster. The next lemma provides an upper bound for the extent to which this can happen.

Lemma 4.3.18. *For any $c \in C \setminus \{p\}$, and for any $t \in [0, 1]$, we have* $\text{dec}_a(c, t) \leq \text{dec}_p(c, t) + \text{b}$.

Proof. Let $t \in [0, 1]$. Our first observation is that

$$\text{dec}_p(p, t) \leq \text{dec}_a(p, t) + \text{b}. \tag{4.14}$$

This holds because

$$\text{dec}_p(p, t) = \text{plu}(p) + \text{b} - \text{score}_p(p, t) \leq \text{plu}(p) + \text{b} - \text{score}_a(p, t) = \text{dec}_a(p, t) + \text{b},$$

where the two equalities follow from Claim 4.3.15 and the inequality holds by (4.13). It follows that

$$\begin{aligned}\sum_{c \in C} \text{dec}_a(c, t) = \sum_{c \in C} \text{dec}_p(c, t) &= \left(\sum_{c \in C \setminus \{p\}} \text{dec}_p(c, t) \right) + \text{dec}_p(p, t) \\ &\leq \left(\sum_{c \in C \setminus \{p\}} \text{dec}_p(c, t) \right) + \text{dec}_a(p, t) + \text{b},\end{aligned}$$

where the first equality holds since the total amount of score decremented from all candidates up to time t equals $t \cdot n \cdot \frac{1+\delta}{1-\delta}$ for any prediction, and the inequality follows from (4.14). We rewrite the above inequality in the form

$$\sum_{c \in C \setminus \{p\}} \text{dec}_a(c, t) \leq \left(\sum_{c \in C \setminus \{p\}} \text{dec}_p(c, t) \right) + \text{b},$$

or equivalently,

$$\sum_{c\in C\setminus\{p\}} (\text{dec}_a(c,t) - \text{dec}_p(c,t)) \leq \text{b}.$$

Lemma 4.3.17 shows that all the terms in the sum are non-negative, and hence each of them is also bounded by b. This proves the lemma. □

Proof of Proposition 4.3.14. It suffices to show that for any $0 \leq t < 1$ we have $\text{score}_a(a,t) > 0$. This implies that a belongs to the set A at the end of Algorithm 1, and therefore a is the output in the a execution. We check that

$$\text{score}_a(a,t) = \text{plu}(a) + \text{b} - \text{dec}_a(a,t) \geq \text{plu}(a) - \text{dec}_p(a,t) = \text{score}_p(a,t) > 0,$$

where the first equality holds by Claim 4.3.15, the first inequality holds by Lemma 4.3.18, the second equality holds again by Claim 4.3.15, and the second inequality holds by our assumption that a is the output in the p execution. □

4.4 Prediction error

Consistency and robustness bounds focus on the case where the prediction is perfect (consistency) and on the case where it can be arbitrarily bad (robustness). A way to understand better the distortion that BoostedSV_δ achieves as a function of the prediction quality, is to define a natural measure of error, which is proportional to the distance $d(p,c^*)$ between the prediction p and the optimal candidate c^* normalized by the average optimal distance $\text{SC}(c^*)/n$. The *prediction error* is the parameter

$$\eta := \frac{nd(p,c^*)}{\text{SC}(c^*)}.$$

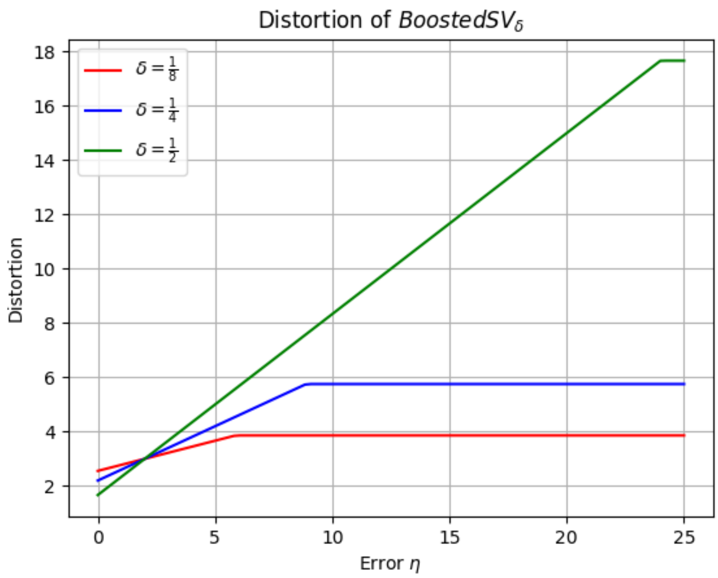

Figure 4.2: Distortion of BoostedSV_δ.

The next proposition shows that for each choice of $\delta \in [0,1)$, the BoostedSV_δ algorithm guarantees distortion at most $\frac{3-\delta+2\delta\eta}{1+\delta}$ when the error of the prediction is η. This bound is equal to the consistency bound when the prediction is correct (which corresponds to $\eta = 0$) and then grows linearly as a function of η, with a slope that is increasing in δ.

This is a reasonable behavior, consistent with our intuition that trusting the prediction more results into a worse dependence on the error. It is useful to keep in mind that the distortion never exceeds the robustness bound, so it transitions from the consistency bound to the robustness bound as a function of the error.

Proposition 4.4.1 (Berger-Feldman-Gkatzelis-Tan)**.** *For any* $\delta \in [0,1)$*, the distortion achieved by* BoostedSV_δ *on instances where the prediction* p *has error* $\eta = n \cdot d(p, c^*)/\text{SC}(c^*)$ *is at most*

$$\min\left\{\frac{3-\delta+2\delta\eta}{1+\delta}, \frac{3+\delta+13\delta^2-\delta^3}{(1+\delta)(1-\delta)^2}\right\}.$$

Proof. Let $\delta \in [0,1)$ and denote by a the candidate output of a given execution of BoostedSV_δ, when provided with a prediction p with error $\eta = nd(p,c^*)/\text{SC}(c^*)$. From Theorem 4.3.5 we know that the worst-case distortion of BoostedSV_δ is at most $\frac{3+\delta+13\delta^2-\delta^3}{(1+\delta)(1-\delta)^2}$ for any δ, so we just need to prove that it is also at most $\frac{3-\delta+2\delta\eta}{1+\delta}$.

Using Lemma 4.3.4 and the properties of $f(v,c)$ from Definition 4.3.3 we can give an upper bound for the social cost of a as follows:

$$\begin{aligned}
\sum_v d(v,a) &\le \sum_v \sum_c f(v,c) \cdot d(v,a) \\
&\le \sum_v \sum_c f(v,c) \cdot d(v,c) \\
&\le \sum_v \sum_c f(v,c) \cdot (d(v,c^*) + d(c^*,c)) \\
&= \text{SC}(c^*) + \sum_c \sum_v f(v,c) \cdot d(c^*,c).
\end{aligned}$$

Isolating $c = p$ and using again the properties of $f(v,c)$ from Definition 4.3.3 we get:

$$\begin{aligned}
\sum_v d(v,a) &= \text{SC}(c^*) + \sum_v f(v,p) \cdot d(c^*,p) + \sum_{c \neq p} \sum_v f(v,c) \cdot d(c^*,c) \\
&= \text{SC}(c^*) + \frac{1-\delta}{1+\delta} \cdot \left(\text{plu}(p) + \frac{2\delta n}{1-\delta}\right) d(c^*,p) + \frac{1-\delta}{1+\delta} \sum_{c \neq p} \text{plu}(c) \cdot d(c^*,c) \\
&= \text{SC}(c^*) + \frac{1-\delta}{1+\delta} \cdot \frac{2\delta n}{1-\delta} \cdot d(c^*,p) + \frac{1-\delta}{1+\delta} \sum_c \text{plu}(c) \cdot d(c^*,c).
\end{aligned}$$

Finally, since $d(c^*,p) = \eta \cdot \text{SC}(c^*)/n$, rearranging the sum and taking into account the fact that if $c = \text{top}(v)$ then $d(v,c) \le d(v,c^*)$, we see that

$$\begin{aligned}
\sum_v d(v,a) &= \text{SC}(c^*) + \frac{2\delta\eta}{1+\delta}\text{SC}(c^*) + \frac{1-\delta}{1+\delta} \sum_c \text{plu}(c) \cdot d(c^*,c) \\
&= \frac{1+\delta+2\delta\eta}{1+\delta}\text{SC}(c^*) + \frac{1-\delta}{1+\delta} \sum_c \sum_{v:\text{top}(v)=c} d(c^*,c) \\
&\le \frac{1+\delta+2\delta\eta}{1+\delta}\text{SC}(c^*) + \frac{1-\delta}{1+\delta} \sum_c \sum_{v:\text{top}(v)=c} d(c^*,v) + d(v,c)
\end{aligned}$$

$$\leq \frac{1+\delta+2\delta\eta}{1+\delta}\mathrm{SC}(c^*) + \frac{1-\delta}{1+\delta}\sum_{c}\sum_{v:\mathrm{top}(v)=c} 2d(c^*, v)$$

$$\leq \frac{1+\delta+2\delta\eta}{1+\delta}\mathrm{SC}(c^*) + \frac{2(1-\delta)}{1+\delta}\mathrm{SC}(c^*) \leq \left(\frac{3-\delta+2\delta\eta}{1+\delta}\right)\mathrm{SC}(c^*).$$

This concludes the proof. □

CHAPTER 5

Distortion of committee election on the real line

In the previous chapters we considered the problem of single-winner election and its metric distortion. In this chapter we study the metric distortion of multi-winner election, where we aim to elect a committee of $k \geq 2$ (out of $m \geq k+1$) candidates based on ordinal preferences provided by n voters. Similarly to the single-winner case, the voters and candidates are associated with locations in a metric space and the voters' cardinal preferences correspond to their distance from the candidate locations. Here we consider the setting where the cost of each voter for a committee is defined as her distance to the nearest member. We focus on the simplest case of *linear preferences*, where the voters and candidates are embedded on the real line. The results of [30] imply that in this setting, the distortion is unbounded for all $k \geq 3$ even in the linear case.

Despite this result, we present the work of Fotakis, Gourvès and Patsilinakos from [45], where it is shown that one can use a restricted amount of cardinal distance queries to achieve bounded or even constant distortion in k-committee election with linear preferences, for $k \geq 3$. Specifically, we will see that at least $\Omega(k)$ distance queries are needed in order to bound the distortion and we shall present a greedy algorithm which achieves $O(n)$ distortion with $O(k)$ queries.

5.1 The model

We consider a set $C = \{c_1, \ldots, c_m\}$ of m candidates and a set $V = \{v_1, \ldots, v_n\}$ of n voters. All candidates and voters are assumed to be located on the real line: each candidate c_i is associated with a location $x(c_i) \in \mathbb{R}$ and each voter v_j is associated with a location $x(v_j) \in \mathbb{R}$. For simplicity we denote by c_i both the candidate c_i and its location $x(c_i)$, and similarly we denote by v_j both the voter v_j and its location $x(v_j)$. We index the candidates according to the order of their real coordinates, i.e. we assume that $x(c_1) < \cdots < x(c_m)$ and write $c_1 < \cdots < c_m$. For any $c < c'$ in $\mathbb{R}$ we define $C[c, c'] = C \cap [c, c']$; this is the set of all candidates in C that lie between c and c'.

For each voter v we define the cost of v for being represented by a candidate c as follows:

$$\text{cost}_v(c) = d(v, c) = |v - c| = |x(v) - x(c)|.$$

Then, for any set $S \subseteq C$ of candidates, we set

$$d(v, S) = \min_{c \in S}\{d(v, c)\} = \min_{c \in S}\{|v - c|\}.$$

We assume that each voter is represented by the nearest candidate in any given set S of elected candidates. In other words, the cost of v being represented by a set S of elected candidates is the quantity

$$\text{cost}_v(S) = d(v, S) = \min_{c \in S}\{d(v, c)\}.$$

The **problem of the k-committee election** is to choose a committee (that is, a set of candidates) $S \subseteq C$ of cardinality $|S| = k \leq m - 1$, that minimizes the (utilitarian) *social cost*

$$\text{SC}(S) = \sum_{v \in V} \text{cost}_v(S)$$

of the voters. One can also consider the *egalitarian cost*

$$\text{EC}(S) = \max_{v \in V}\{\text{cost}_v(S)\}$$

for a k-committee S of elected candidates. Whenever we use the egalitarian cost, this will be explicitly mentioned. An *instance* of a k-committee election is a pair (C, V), where C is the set of candidates and V is the set of voters, equipped with their locations on the real line (these are assumed fixed, but unknown to the voting rule).

In our model, each voter v provides only a ranking $\succ_v$ on the set C of candidates, which is aligned with the function $\text{cost}_v : C \to \mathbb{R}^+$. This means that for any pair of candidates c and c' we have that $c \succ_v c'$ (then we say that v prefers c to c') if and only if $d(v, c) < d(v, c')$. We assume that for every voter v and any pair of candidates $c \neq c'$ we have that $d(v, c) \neq d(v, c')$, and hence $\succ_v$ is a strict total order.

A *deterministic rule* R for k-committee election receives as an input a linear ranking profile $\vec{\succ} = (\succ_1, \ldots, \succ_n)$ over a set C of m candidates, the committee size k and a non-negative integer q. Then, using $\vec{\succ}$ and information about the distance of at most q candidate pairs on the real line, the rule R computes a committee $R(\vec{\succ}, k, q) = S \subseteq C$ consisting of k candidates. We consider committee election rules that assume availability of a strict ordering of the candidates on the real line, from left to right, which we call the *candidate axis*. We may obtain such an ordering from a given linear profile $\vec{\succ}$ using the algorithm of Elkind and Faliszewski [38]. We evaluate committee election rules with their distortion, whose definition in Chapter 3 extends naturally for this setting.

Below we use the following notation. We write top(v) for the top candidate of voter v with respect to $\succ_v$. The cluster Cluster(c) of a candidate c is the set of all voters v in V that have c as their top candidate. A candidate c is called *active* if Cluster$(c) \neq \varnothing$, equivalently if there exists a voter v with c as her top choice. We assume that all inactive candidates in C are removed before given as an input to the algorithms. For convenience we sometimes consider *candidate-restricted* instances in our analysis, where all candidates are active and all voters are assumed to be at the same location with their top candidate.

We shall discuss three different types of distance queries:

- **Regular queries.** Given a voter $v \in V$ and a candidate $c \in C$, we ask for the distance $d(v, c) = |v - c|$.
- **Candidate queries.** Given two candidates $c, c' \in C$, we ask for the distance $d(c, c') = |c - c'|$.
- **Voter queries.** Given two voters $v, v' \in V$, we ask for the distance $d(v, v') = |v - v'|$.

It is shown in [45] that in the linear case we can simulate candidate queries and voter queries with six and two regular queries, respectively. In the following rules we assume access to candidate queries but as long as we care about the asymptotics of the number of queries used, these three query types can be used interchangeably.

5.2 Lower bound for the number of queries required for bounded distortion

A lower bound on the number of queries that are required to ensure bounded distortion is given by the next theorem.

Theorem 5.2.1 (Fotakis-Gourvés-Patsilinakos). *For any $k \geq 3$, the distortion of any deterministic k-committee election rule that uses at most $k-3$ distance queries and selects k out of at least $2(k-1)$ candidates on the real line cannot be bounded by any function of n, m and k. This applies to both the social and the egalitarian cost.*

Proof. Let $k \geq 3$ and consider $m = 2(k-1)$ candidates $c_1 < c_2 < \cdots < c_{2k-3} < c_{2k-2}$. We fix D large enough so that $D^2 \gg \max\{2D+1, k\}$ and $\epsilon \in \left(0, \frac{1}{k}\right)$ small enough. First, we define a basic instance, where $d(c_{2i-1}, c_{2i}) = 1$ for all $i \in [k-1]$ and $d(c_{2i}, c_{2i+1}) = D^2 + (i-1)\epsilon$ for all $i \in [k-2]$ (see Fig.5.1). There are $n = m$ voters and each one of them has a different top candidate. Both in the basic instance and in the variants that we shall define, each voter has the same location with its top candidate.

Now, we construct $2(k-1)$ variants of the basic instance as follows. For the j-th instance we move c_j by D and keep all the other candidates at their original locations: If j is odd, we increase the distance $d(c_j, c_{j+1})$ from 1 to $D+1$ (moving c_j by D on the left), while if j is even, we increase the distance $d(c_{j-1}, c_j)$ from 1 to $D+1$ (moving c_j by D on the right). All the other candidates stay at the locations they had in the basic instance (see Fig.5.2). This means that, in the j-th variant, if j is odd then all distances $d(c_i, c_j)$, $1 \leq i \leq j-1$, decrease by D and all distances $d(c_j, c_i)$, $j+1 \leq i \leq 2(k-1)$, increase by D. On the other hand, if j is even then all distances $d(c_i, c_j)$, $1 \leq i \leq j-1$, increase by D and all distances $d(c_j, c_i)$, $j+1 \leq i \leq 2(k-1)$, decrease by D. Note that the distance of c_j to all other candidates is affected, however the distances between any other pair of candidates remain the same.

With this construction, there are $k-1$ "isolated" pairs of candidates, and the same is true for all the variants. In each variant, there exists exactly one pair of candidates who are far from each other, and this means that if a voting rule has a bounded distortion then it must identify this pair and elect both candidates of this pair. On the other hand, the candidates of the remaining pairs are very close to each other, and this means that a voting rule needs only to elect one candidate from each such pair in order to achieve bounded distortion. Apart from this, the $2(k-1)$ variants are completely symmetric and hence any distance query that can discover that the distance between a pair of candidates is the same as in the basic instance can exclude at most two variants of the basic instance. This implies that any deterministic rule requires at least $k-2$ distance queries, in the worst case, in order to identify the pair of candidates that have distance D.

We pass now to the details. The optimal committee, for both the social cost and the egalitarian cost, for the j-th variant of the basic instance is to select the candidates at distance D from each other, namely c_j and c_{j+1} (if j is odd) or c_{j-1} and c_j (if j is even) and any candidate from each of the remaining pairs (c_{2i-1}, c_{2i}) of candidates. With this choice, the social cost is $k-2$ and the egalitarian cost is 1. Any other committee will have social and egalitarian cost at least D, which is arbitrarily larger than $k-2$ and 1 by our choice for D. This shows that a deterministic committee election rule can have bounded distortion only if it can identify the pair of candidates with distance D and then add both these candidates to the remaining $k-2$ candidates elected by the rule.

We shall show that this can be done only if the algorithm asks for at least $k-2$ queries. First we observe that the ranking of every voter in the basic instance is the same with the ranking of this voter in each of the $2(k-1)$ variants. Therefore, a voting rule cannot distinguish between the variants just by looking at the rankings of the voters. In the j-th variant, the distances of c_j to all other candidates differ from the corresponding distances in the basic instance, but the distances between all other pairs of candidates remain the same as in the basic instance. So, whenever we query the distance between a pair of candidates and check that it remains the same as in the basic instance, we can exclude at most two from the $2(k-1)$ variants. Therefore, any deterministic committee election rule needs at least $k-2$ distance queries, in the worst case, before it identifies the pair (c_{2i-1}, c_{2i}) of candidates that are at distance D in the input variant and must be elected both. □

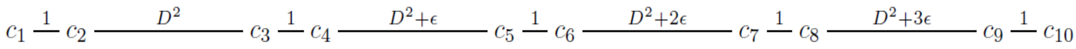

Figure 5.1: The basic instance used in the proof of Theorem 5.2.1 for $k=6$.

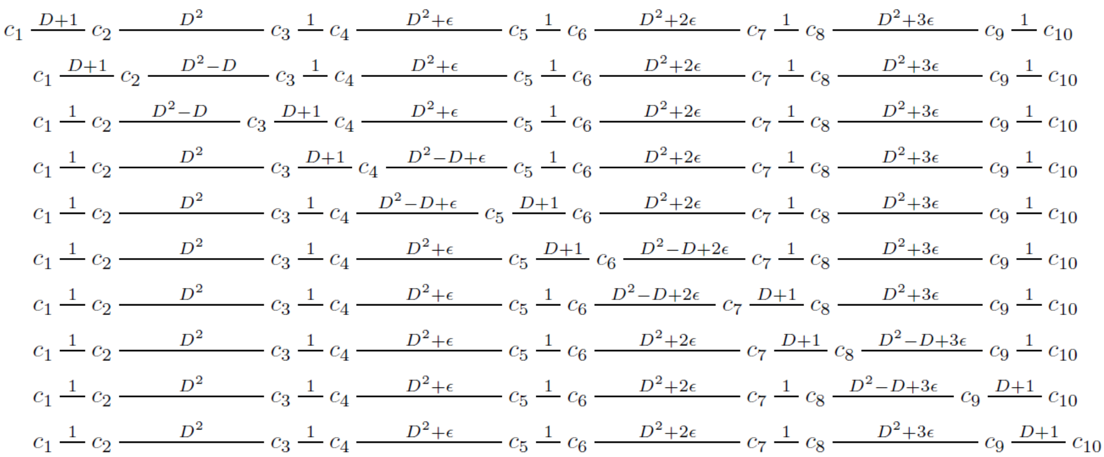

Figure 5.2: The $2(k-1)=10$ variants obtained from the basic instance used in the proof of Theorem 5.2.1 for $k=6$.

5.3 Bounded distortion with $\Theta(k)$ queries

In this section we show that we can achieve bounded distortion using $\Theta(k)$ distance queries, thus asymptotically matching the lower bound of Theorem 5.2.1. This is achieved by a query efficient implementation of the 2-approximate greedy algorithm for k-center from Williamson and Shmoys [81].

We show that the classical 2-approximate greedy algorithm for k-center can be implemented with few distance queries. The greedy algorithm iteratively maintains a set S of candidates, starting with any candidate, and adding the candidate c with maximum distance $d(c,S)$ to the current set S in each iteration. When applied to linear instances, the greedy algorithm starts with the leftmost candidate c_1 and the rightmost candidate c_m. Then, for the next $k-2$ iterations, it adds to S the candidate $c \in C$ with maximum $d(c,S)$.

In order to implement the greedy algorithm in our setting with distance queries (Algorithm 2), we need to compute the most distant candidate to a set $S = \{c_1, \ldots, c_\ell\} \subseteq C$, where the candidates are indexed as they appear on the candidate axis, from left to right, and c_1, c_ℓ are the leftmost and rightmost candidate in S respectively. For every $1 \leq i \leq \ell - 1$ we define $\hat{c}_i$ as the most distant candidate in the interval $C[c_i, c_{i+1}]$ to its endpoints $c_i, c_{i+1} \in S$, i.e.

$$d(\hat{c}_i, \{c_i, c_{i+1}\}) = \max_{c \in C[c_i, c_{i+1}]} d(c, \{c_i, c_{i+1}\}), \tag{5.1}$$

and let

$$\delta_i := d(\hat{c}_i, \{c_i, c_{i+1}\})$$

the distance of $\hat{c}_i$ to the endpoints of $C[c_i, c_{i+1}]$. This information is provided by the Distant-Candidate algorithm (Algorithm 3). We can show that the most distant candidate to S is the candidate $\hat{c}_i$ with the maximum distance δ_i.

Proposition 5.3.1. *Let S be the set of currently elected candidates in Algorithm* 2*, and let $\hat{c}_1, \dots, \hat{c}_{\ell-1}$ be defined as in* (5.1)*. Then,*

$$\max_{c \in C} d(c, S) = \max_{1 \leq i \leq \ell - 1} d(\hat{c}_i, \{c_i, c_{i+1}\}).$$

Proof. Note that, since c_1 is the leftmost candidate and c_ℓ is the rightmost candidate in C, every candidate who belongs to $C \setminus S$ lies in one of the intervals $C[c_1, c_2], \dots, C[c_{\ell-1}, c_\ell]$. Suppose that the farthest candidate c to S is in the interval $C[c_i, c_{i+1}]$. Then, c has to be the candidate $\hat{c}_i \in C[c_i, c_{i+1}]$ with maximum distance to the endpoints c_i, c_{i+1} and $d(\hat{c}_i, S) = d(\hat{c}_i, \{c_i, c_{i+1}\})$. □

For any interval $C[c, c']$ with $|C[c, c']| \geq 3$, which is defined by two consecutive candidates in S, Algorithm 3 computes the candidate $\hat{c} \in C[c, c']$ with the maximum distance to the endpoints c, c' as well as the distance $d(\hat{c}, \{c, c'\}) = d(\hat{c}, S)$. Moreover, each application of Algorithm 3 uses at most 3 distance queries.

Lemma 5.3.2. *Let $c, c' \in C$ with $c < c'$. Then, Algorithm* 3 *correctly returns the candidate $\hat{c} \in C[c, c']$ with maximum distance to the interval's endpoints c, c', i.e.*

$$d(\hat{c}, \{c, c'\}) = \max_{c'' \in C[c, c']} d(c'', \{c, c'\})$$

as well as the distance $d(\hat{c}, \{c, c'\}) = d(\hat{c}, S)$ to S.

Proof. The case $|C[c, c']| = 3$ is clear because there is only one $c'' \in C[c, c'] \setminus \{c, c'\}$ and this is necessarily the most distant candidate. Then, Algorithm 3 returns c'' and computes the distance $d(c'', \{c, c'\})$ using 2 distance queries.

Assume that $|C[c, c']| \geq 4$. We first consider candidate-restricted instances, where for any candidate $c'' \in C[c, c']$ all voters $v \in \text{Cluster}(c'')$ are collocated with c''. Then, the ranking $\succ_v$ submitted by a voter $v \in \text{Cluster}(c'')$ in Step 6 of Algorithm 3 is the same with the ranking $\succ_{c''}$ where all candidates $c \in C$ appear in increasing order of their distance to c''.

Let $m = (c + c')/2$ be the midpoint of the interval $[c, c']$, which is the point in this interval with maximum distance $d(m, \{c, c'\})$ to the endpoints c, c'. In Algorithm 3, c_r is the leftmost candidate in $C[c, c']$ that is closer to the right endpoint c' than to the left endpoint c. By the definition of c_r, we have that c_ℓ is the rightmost candidate in $C[c, c']$ that is closer to the left endpoint c than to the right endpoint c'. Note that $d(c_r, \{c, c'\}) = d(c_r, c')$ and $d(c_\ell, \{c, c'\}) = d(c_\ell, c)$. It follows that $c_\ell \leq m \leq c_r$, with at least one of the inequalities strict, and there are no other candidates between c_ℓ and c_r. Therefore, c_ℓ and c_r are the candidates in $C[c, c']$ that are closest to m. Then, comparing

$d(c_\ell, c) = d(c_\ell, S)$ and $d(c_r, c') = d(c_r, S)$ we determine the candidate in $C[c, c']$ with maximum distance to its endpoints.

Next, we remove the assumption that Algorithm 3 has access to $\succ_{c''}$. Let $\hat{c}$ be the candidate in $C[c, c']$ with maximum distance to $\{c, c'\}$. Then, $\hat{c}$ is the closest candidate to the midpoint m. Without loss of generality we assume that $\hat{c} \leq m$ and show that $\hat{c} \in \{c_\ell, c_r\}$ and that Algorithm 3 correctly returns $\hat{c}$. The case where $\hat{c} > m$ is similar by symmetry.

Let $\hat{v} \in$ Cluster($\hat{c}$) be the voter whose preference list $\succ_{\hat{v}}$ is used in place of $\succ_{\hat{c}}$ in Algorithm 3. Let c_a and c_b be the next candidate on the left and respectively on the right of $\hat{c}$ in $C[c, c']$ (note that we might have $c_a = c$ or $c_b = c'$ but not both). Let v_a and v_b be the voters in Cluster(c_a) and Cluster(c_b), respectively, whose preference lists $\succ_{v_a}$ and $\succ_{v_b}$ are used in place of $\succ_{c_a}$ and $\succ_{c_b}$.

Since $\hat{c} \leq m$, we have that $c \succ_{v_a} c'$, because $c_a < \hat{c} \leq m$ and $d(v_a, c_a) < d(v_a, \hat{c})$. It follows that $v_a < \hat{c} \leq m$. Moreover, since $\hat{c}$ is the closest candidate to m, $\hat{c} \leq m < c_b$ and $d(v_b, c_b) < d(v_b, \hat{c})$, we have that $c' \succ_{v_b} c$. Therefore, $m < v_b$, because if $v_b \leq m$ then using the fact that m is closer to $\hat{c}$ we would get $d(\hat{c}, v_b) < d(c_b, v_b)$ and v_b would be in Cluster($\hat{c}$).

Now, since $c \succ_{v_a} c'$ and $c' \succ_{v_b} c$, which imply that $c_a \neq c_r$ and c_r is either $\hat{c}$ or c_b, we get that $\hat{c}$ is either c_ℓ or c_r. We distinguish two cases depending on the placement of $\hat{v}$ with respect to m.

If $\hat{v} \leq m$ and $c \succ_{\hat{v}} c'$ then $c_r = c_b$ and $c_\ell = \hat{c}$. Moreover, $d(\hat{c}, c) \geq d(c_b, c')$ because $\hat{c}$ is the farthest candidate to $\{c, c'\}$. Therefore, Algorithm 3 returns $\hat{c}$ as the farthest candidate and the distance $d(\hat{c}, c) = d(\hat{c}, S)$.

If $m < \hat{v}$ and $c' \succ_{\hat{v}} c$ then $c_r = \hat{c}$ and $c_\ell = c_a$. Moreover, $d(\hat{c}, c') > d(c_a, c)$ because $c \leq c_a < \hat{c} \leq m < c'$. Therefore, Algorithm 3 returns $\hat{c}$ as the farthest candidate and the distance $d(\hat{c}, c) = d(\hat{c}, S)$. □

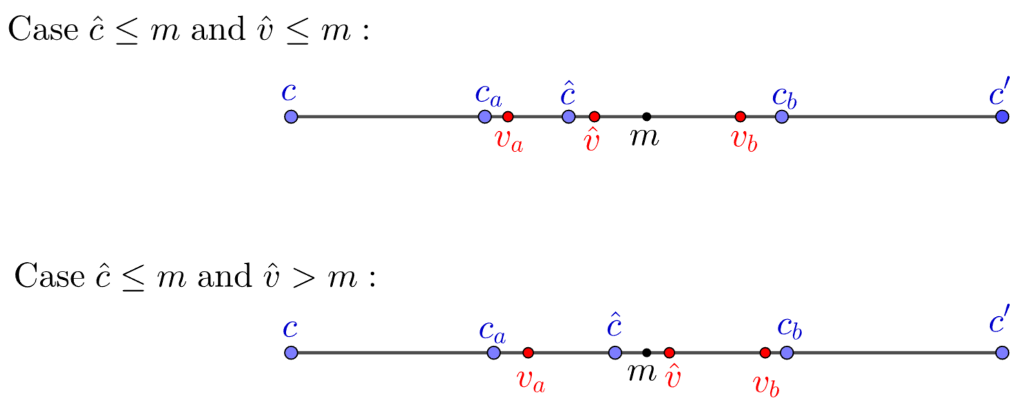

Figure 5.3: The two cases for the location of $\hat{v}$ in the proof of Lemma 5.3.2.

Algorithm 2: Query-efficient implementation of Greedy

Input: Candidates $C = \{c_1, \ldots, c_m\}$, $k \in \{2, \ldots, m-1\}$, voter ranking profile $\vec{\succ} = (\vec{\succ}_1, \ldots, \vec{\succ}_n)$
Output: Set $S \subseteq C$ of k candidates

$S \leftarrow \{c_1, c_m\}$ {pick leftmost and rightmost candidates}
$\hat{C} \leftarrow \{\text{Distant} - \text{Candidate}(C[c_1, c_m])\}$
while $|S| < k$ **do**
  Let c be s.t. $(c, \delta) \in \hat{C}$ and $\delta \geq \delta'$ for all $(c', \delta') \in \hat{C}$
  $S \leftarrow S \cup \{c\}$
  $\hat{C} \leftarrow \hat{C} \setminus \{(c, \delta)\}$
  if $|S| < k$ **then**
    Let c_i be the rightmost candidate in S on c's left
    Let c_{i+1} be the leftmost candidate in S on c's right
    $\hat{C} \leftarrow \hat{C} \cup \{\text{Distant} - \text{Candidate}(C[c_i, c])\} \cup \{\text{Distant} - \text{Candidate}(C[c, c_{i+1}])\}$
return S

Algorithm 3: The Distant-Candidate algorithm

Input: Candidate interval $C[c, c']$, a voter $v \in \text{Cluster}(c'')$ for every $c'' \in C[c, c']$
Output: Candidate $\hat{c} \in C[c, c']$ with maximum $d(\hat{c}, \{c, c'\})$

if $|C[c, c']| = 3$ **then**
  $c'' \leftarrow C[c, c'] \setminus \{c, c'\}$
  return $(c'', \min\{d(c'', c), d(c'', c')\})$
Let c'' be the leftmost candidate in $C[c, c'] \setminus \{c\}$
while $c'' \in C[c, c']$ **do**
  Let $\succ_{c''}$ be the ranking $\succ_v$ of any $v \in \text{Cluster}(c'')$
  if $c' \succ_{c''} c$ **then**
    Let c_r be c'' and c_ℓ be next candidate on c'''s left
     {c_ℓ and c_r found, while-loop terminates}
    break
  else
    $c'' \leftarrow$ the next candidate on c'''s right {proceed to the next candidate on the right}
if $d(c, c_\ell) \geq d(c_r, c')$ **then**
  return $(c_\ell, \min\{d(c, c_\ell), d(c', c_\ell)\})$
else
  return $(c_r, \min\{d(c_r, c), d(c_r, c')\})$

We shall also use the next theorem.

Theorem 5.3.3 (Fotakis-Gourvés-Patsilinakos). *Let (C, V) be an instance of the k-committee election. Let $S \subseteq C$ (respectively, $S^* \subseteq C$) be a β-approximate (respectively, an optimal) k-committee with respect to the egalitarian cost for the candidate-restricted instance (respectively, the original instance). Then,*

$$\text{EC}(S) \leq (1 + 2\beta)\, \text{EC}(S^*).$$

Proof. If $\text{top}(v) \in C$ is the top candidate of $v \in V$ then, by the triangle inequality, we have that

$d(v, S) \leq d(v, \text{top}(v)) + d(\text{top}(v), S)$ and taking maximum with respect to all $v \in V$ we see that

$$\text{EC}(S) \leq \text{EC}(C) + \text{EC}(C_{\text{cr}}, S), \tag{5.2}$$

where

$$\text{EC}(C) = \max\{d(v, \text{top}(v)) : v \in V\} = \max\{d(v, C) : v \in V\}$$

and

$$\text{EC}(C_{\text{cr}}, S) = \max\{d(\text{top}(v), S) : v \in V\}.$$

Note that $\text{EC}(C_{\text{cr}}, S)$ is the egalitarian cost of S for the candidate-restricted instance C_{cr} induced by C, and

$$\text{EC}(C_{\text{cr}}, S) \leq \beta\, \text{EC}(C_{\text{cr}}, S^{\#}) \leq \beta\, \text{EC}(C_{\text{cr}}, S^{*}). \tag{5.3}$$

The first inequality holds true because S is a β-approximate k-committee for C_{cr}. For the second inequality we use the fact that the β-approximation ratio of S is established against an optimal solution $S^{\#}$ for C_{cr} that may include inactive candidates from C, and hence S^* is also a feasible option. Since $S^{\#}$ is an optimal solution for C_{cr} with respect to the egalitarian cost, it follows that $\text{EC}(C_{\text{cr}}, S^{\#}) \leq \text{EC}(C_{\text{cr}}, S^{*})$.

Next, since $d(\text{top}(v), S^*) \leq d(\text{top}(v), v) + d(v, S^*)$ for all $v \in V$, taking the maximum over all v we see that

$$\text{EC}(C_{\text{cr}}, S^*) \leq \text{EC}(C) + \text{EC}(S^*).$$

Combining these observations with (5.2) and using (5.3) we get

$$\text{EC}(S) \leq (1 + \beta)\, \text{EC}(C) + \beta\, \text{EC}(S^*) \leq (1 + 2\beta)\, \text{EC}(S^*).$$

□

Theorem 5.3.4 (Fotakis-Gourvés-Patsilinakos)**.** *For any $k \geq 3$, Algorithm 2 achieves a distortion of at most $5n$ for the social cost, and at most 5 for the egalitarian cost, for k-Committee Election on the real line using at most $6k - 15$ candidate distance queries.*

Proof. In Algorithm 2, the distant-candidate algorithm is called once in Step 2 and $2(k-3)$ times in Step 10, twice in each iteration of the while-loop, from the iteration where $|S| = 3$ to the iteration where $|S| = k - 1$. So, the total number of distance queries is at most $6(k-3) + 3$. Lemma 5.3.2 and Proposition 5.3.1 establish the correctness of Algorithm 2, i.e. the fact that in each iteration the candidate c with maximum $d(c, S)$ is added to S. The distortion bound for the egalitarian cost follows from the fact that the algorithm of Williamson and Shmoys is 2-approximate for the egalitarian cost of candidate-restricted instances (see [81, Theorem 2.3]). Then, Theorem 5.3.3 implies an upper bound of 5 on the distortion for the egalitarian cost of the original instance. The upper bound of $5n$ on the distortion for the social cost holds because for any committee $S \subseteq C$ we have that $\text{EC}(S) \leq \text{SC}(S) \leq n\, \text{EC}(S)$. Thus, if S^* is the optimal committee with respect to the egalitarian cost and S^{**} is the optimal committee with respect to the social cost, for any committee $S \subseteq C$ we have that

$$\frac{\text{SC}(S)}{\text{SC}(S^{**})} \leq n\, \frac{\text{EC}(S)}{\text{EC}(S^{**})} = n\, \frac{\text{EC}(S)}{\text{EC}(S^{*})}\, \frac{\text{EC}(S^*)}{\text{EC}(S^{**})} \leq 5n\, \frac{\text{EC}(S^*)}{\text{EC}(S^{**})} \leq 5n,$$

because $\text{EC}(S^*) \leq \text{EC}(S^{**})$. □

CHAPTER 6

Voting rules with predictions for single-winner and committee elections

In this chapter, we modify some of the voting rules that we studied in previous chapters by augmenting them with predictions and bound their consistency and robustness.

Our first contribution is a discussion of consistency and robustness bounds for suitably defined boosted versions of the Plurality and Borda rules. In the case of the plurality rule, for any $\delta \in [0, 1)$ we define the algorithm BoostedPlurality$_\delta$ that uses a prediction p for the optimal candidate and depends on the confidence parameter δ. It elects either the candidate with the highest plurality score or p if his plurality score is high enough. The learning-augmented version of the Borda rule, BoostedBorda$_\delta$ is similar. We also provide bounds on the distortion achieved by BoostedPlurality$_\delta$ and BoostedBorda$_\delta$ on instances where the prediction p has a given error η. Our results indicate that predictions are not useful when the only information available is the plurality score or the Borda score of the candidates.

Subsequently, we focus on the committee election problem. We assume that our algorithms have access to a prediction $\mathcal{P} = \{p_1, \dots, p_k\} \subseteq C$ for the optimal committee. The first algorithm that we examine, is a learning-augmented version of the Greedy algorithm from [45], parameterized by $\delta \in [0, 1)$. At each iteration it elects in the committee either the most distant candidate thus far or the most distant *predicted* candidate thus far. It uses $\Theta(k)$ distance queries, but for all $\delta \in [0, 1)$ its consistency is still $\Omega(n)$.

Afterwards, we introduce an algorithm which, using a prediction $\mathcal{P}$ and the Greedy algorithm, computes a **good** representative set of candidates and elects the optimal k-committee in the restricted instance induced by this set. Our algorithm achieves constant consistency and linear robustness with $O(k)$ distance queries.

6.1 Boosted versions of Plurality and Borda

In this section we discuss consistency and robustness of boosted versions of the Plurality and Borda rules, parameterized by $\delta \in [0, 1)$. We assume that we have access to the Plurality or Borda score of every candidate, in addition to a prediction about the optimal candidate. With this information, it is natural to elect the predicted candidate, if his score is sufficiently high when compared to the highest score. The threshold of the score, above which the prediction is chosen, depends on the confidence parameter δ.

Starting with the Plurality rule, recall that if $\sigma := (\sigma_v)_{v \in V}$ is a preference profile which is induced by an instance (V, C, d) then, for any $c \in C$, the plurality score plu(c) of c is the number of its first positions in the rankings. Formally,

$$\text{plu}(c) = |\{v \in V : \sigma_v(c) = 1\}|.$$

We assume that the candidates $c \in C$ are labeled according to their plurality score: c_i is the candidate placed i-th with respect to the plurality score. Therefore, $\text{plu}(c_1) \geq \text{plu}(c_2) \geq \cdots \geq \text{plu}(c_m)$. Let $\delta \in [0, 1)$. We define BoostedPlurality$_\delta$ as follows. Given a prediction $p \in C$ about the optimal candidate $c^* = c_1$, we set $a = p$ if $\text{plu}(p) > (1 - \delta)\text{plu}(c_1)$ and $a = c_1$ otherwise. The algorithm outputs a. Therefore, when δ is close to 1 the algorithm tends to elect the predicted candidate.

Algorithm 4: BoostedPlurality$_\delta$

Input: Preference profile σ, predicted optimal candidate $p \in C$
Output: a candidate $c \in C$

$c_1 \leftarrow \text{argmax}_{c \in C}\text{plu}(c)$
$\forall\, c \in C \setminus \{p\},\ \text{score}(c) \leftarrow \text{plu}(c),$
$\text{score}(p) \leftarrow \text{plu}(p) + \delta \cdot \text{plu}(c_1)$ `// Boost score of predicted candidate in proportion to δ`
if score(p) $>$ score(c_1) **then**
  return p
else
  return c_1

The next proposition provides an upper bound for the robustness of BoostedPlurality$_\delta$.

Proposition 6.1.1. *For every $\delta \in [0, 1)$ the algorithm* BoostedPlurality$_\delta$ *has robustness*

$$\text{robustness}(\text{BoostedPlurality}_\delta) \leq \frac{2m}{1 - \delta} - 1.$$

Proof. From Lemma 3.1.6 we know that,

$$\frac{\text{SC}(a)}{\text{SC}(c^*)} \leq \frac{2n}{|ac^*|} - 1.$$

Recall that c^* is the optimal candidate. Note that

$$|ac^*| \geq \text{plu}(a) \geq (1 - \delta)\text{plu}(c_1).$$

It follows that

$$\frac{\text{SC}(a)}{\text{SC}(c^*)} \leq \frac{2n}{(1 - \delta)\text{plu}(c_1)} - 1.$$

Since $n = \text{plu}(c_1) + \cdots + \text{plu}(c_m) \leq m \cdot \text{plu}(c_1)$, we know that $\text{plu}(c_1) \geq n/m$. It follows that

$$\frac{\text{SC}(a)}{\text{SC}(c^*)} \leq \frac{2m}{1 - \delta} - 1$$

as claimed. □

On the other hand, we show that the consistency of BoostedPlurality$_\delta$ exceeds $2m - 1 - 2\delta$.

Proposition 6.1.2. *For every $\delta \in [0, 1)$ the algorithm* BoostedPlurality$_\delta$ *has consistency*

$$\text{consistency}(\text{BoostedPlurality}_\delta) \geq 2m - 1 - 2\delta.$$

Proof. Consider the example shown in Figure 6.1.

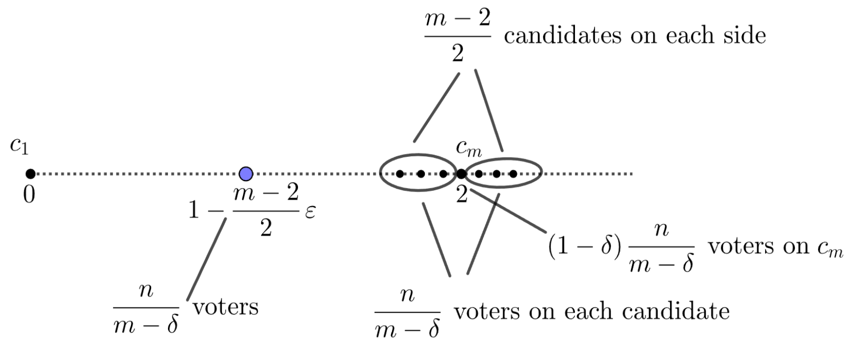

Figure 6.1: Example for Proposition 6.1.2.

All candidates and voters are located on the real axis and the distance is the Euclidean distance. Candidate c_1 is located at the point $x = 0$, candidate c_m is located at the point $x = 2$, and the remaining $m-2$ candidates are very close to c_m. Specifically, there are $\frac{m-2}{2}$ candidates to the left of c_m and $\frac{m-2}{2}$ candidates to the right of c_m, at distances $\epsilon, 2\epsilon, \ldots, \frac{m-2}{2}\epsilon$ respectively, where $\epsilon > 0$ is a small constant. There are $\frac{n}{m-\delta}$ voters at $x = 1 - \frac{m-2}{2}\epsilon$, $(1-\delta)\frac{n}{m-\delta}$ voters collocated with c_m, and $\frac{n}{m-\delta}$ voters collocated with each of the candidates that are close to c_m.

We easily check that c_m is the optimal candidate with respect to the social cost, therefore when we consider the consistency, the prediction p is c_m. However,

$$\text{plu}(p) = \text{plu}(c_m) = (1-\delta)\frac{n}{m-\delta} = (1-\delta)\text{plu}(c_1)$$

and the algorithm BoostedPlurality$_\delta$ outputs candidate c_1.

The social cost of c_1 is

$$\begin{aligned}
\text{SC}(c_1) &= \frac{n}{m-\delta} \cdot (1 - \frac{m-2}{2}\epsilon) + (1-\delta)\frac{n}{m-\delta} \cdot 2 \\
&\quad + \frac{n}{m-\delta}\left(2 - \frac{m-2}{2}\epsilon + 2 + \frac{m-2}{2}\epsilon + \cdots + 2 - \epsilon + 2 + \epsilon\right) \\
&\to \frac{n}{m-\delta}\big(1 + 2(1-\delta) + (m-2)\cdot 2\big) = \frac{n}{m-\delta}(2m - 1 - 2\delta)
\end{aligned}$$

when $\epsilon \to 9$, and the social cost of c_m is

$$\begin{aligned}
\text{SC}(c_m) &= \frac{n}{m-\delta} \cdot 1 + \frac{n}{m-\delta} \cdot 2\left(\epsilon + 2\epsilon + \cdots + \frac{m-2}{2}\epsilon\right) \\
&= \frac{n}{m-\delta} + \frac{n}{m-\delta}\epsilon\frac{m(m-2)}{4}) \\
&\to \frac{n}{m-\delta}
\end{aligned}$$

when $\epsilon \to 0$. It follows that the consistency of the algorithm is greater than $2m - 1 - 2\delta$. $\square$

Note that if $\delta = 0$ then the lower bound of Proposition 6.1.2 matches the distortion of the simple plurality algorithm. For $\delta \to 1$ the lower bound for the consistency becomes $2m - 3$, meaning that the prediction does not give nearly any improvement in the distortion of the plurality algorithm.

We pass now to the Borda rule. Recall that, in this case, the scoring vector is $\vec{s} = (m-1, m-2, \ldots, 1, 0)$. Let σ be a preference profile and $p \in C$ be a prediction about the optimal candidate c^*. The Borda score of a candidate c is

$$\text{borda}(c) = \sum_{v \in V} (m - \sigma_v(c)).$$

Let $c_1 \in C$ be a candidate with maximal score. For any $\delta \in \left[0, \frac{1}{m-1}\right)$ we set $a = p$ if $\text{borda}(p) > (1-\delta)\text{borda}(c_1)$ and $a = c_1$ otherwise. The algorithm outputs a.

Algorithm 5: BoostedBorda$_\delta$

Input: Preference profile σ, predicted optimal candidate $p \in C$
Output: a candidate $c \in C$

$c_1 \leftarrow \text{argmax}_{c \in C} \text{borda}(c)$
$\forall\, c \in C \setminus \{p\},\ \text{score}(c) \leftarrow \text{borda}(c),$
$\text{score}(p) \leftarrow \text{borda}(p) + \delta \cdot \text{borda}(c_1)$ // Boost score of predicted candidate in proportion to δ
if $\text{score(p)} > \text{score(c}_1)$ **then**
  return p
else
  return c_1

Proposition 6.1.3. *For every* $\delta \in \left[0, \frac{1}{m-1}\right)$ *the algorithm* BoostedBorda$_\delta$ *has robustness*

$$\text{robustness}(\text{BoostedBorda}_\delta) \leq \frac{2m}{1 - \delta(m-1)} - 1.$$

Proof. Consider the case $a = p$. We observe that

$$\begin{aligned}\text{borda}(p) - \text{borda}(c^*) &= \sum_{v \in V} (\sigma_v(c^*) - \sigma_v(p)) \\ &\leq (-1) \cdot |c^* p| + (m-1) \cdot |pc^*| = -n + m|pc^*|.\end{aligned}$$

Since $a = p$, we have

$$\begin{aligned}\text{borda}(p) - \text{borda}(c^*) &> (1-\delta)\text{borda}(c_1) - \text{borda}(c^*) \\ &= \text{borda}(c_1) - \text{borda}(c^*) - \delta \text{borda}(c_1) \\ &\geq -\delta \cdot \text{borda}(c_1).\end{aligned}$$

Combining the above, as well as the simple observation that $\text{borda}(c_1) \leq (m-1)n$, we see that

$$|pc^*| \geq \frac{n - \delta\, \text{borda}(c_1)}{m} \geq \frac{(1 - \delta(m-1))n}{m}.$$

From Lemma 3.1.6 we know that

$$\frac{\text{SC}(a)}{\text{SC}(c^*)} \leq \frac{2n}{|ac^*|} - 1.$$

It follows that

$$\frac{\text{SC}(a)}{\text{SC}(c^*)} \leq \frac{2m}{1 - \delta(m-1)} - 1.$$

We also know that if $a = c_1$ then

$$\frac{\mathrm{SC}(a)}{\mathrm{SC}(c^*)} \leq 2m - 1.$$

So, we get the claimed bound. □

Moreover, we can also provide a lower bound for the consistency of BoostedBorda$_\delta$.

Proposition 6.1.4. *For every $\delta \in [0, 1)$ the algorithm* BoostedBorda$_\delta$ *has consistency*

$$\text{consistency}(\text{BoostedBorda}_\delta) \geq 2(m-1)\frac{1-\delta}{1+\delta(m-2)} + 1.$$

Proof. Consider the example shown in Figure 6.2.

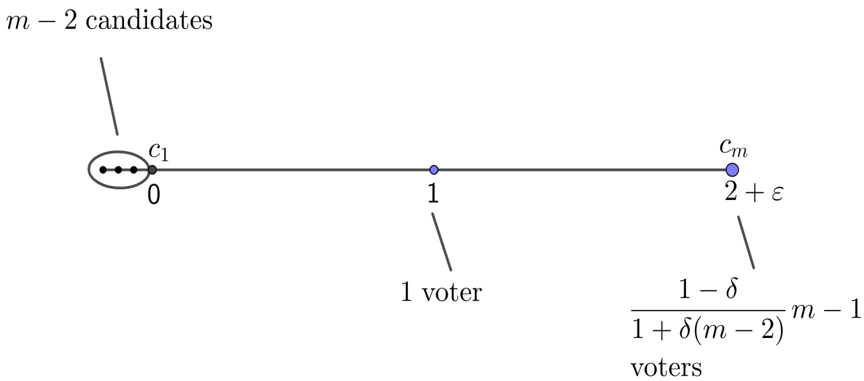

Figure 6.2: Example for Proposition 6.1.4.

All candidates and voters are located on the real axis and the distance is the Euclidean distance. Candidate c_1 is located at the point $x = 0$, candidate c_m is located at the point $x = 2+\epsilon$ (where $\epsilon > 0$ is a small constant), and the remaining $m-2$ candidates are very close to c_1 and to his left. There is one voter at $x = 1$ (she ranks c_1 first and c_m last) and $\frac{1-\delta}{1+\delta(m-2)}(m-1)$ voters collocated with c_m.

We easily check that c_m is the optimal candidate with respect to the social cost, therefore when we consider the consistency, the prediction p is c_m. However,

$$\begin{aligned}\text{borda}(c_1) &= (m-1) + \frac{1-\delta}{1+\delta(m-2)}(m-1)(m-2)\\ &= (m-1)\left[(m-2)\frac{1-\delta}{1+\delta(m-2)} + 1\right] = \frac{(m-1)^2}{1+\delta(m-2)}\end{aligned}$$

and

$$\text{borda}(p) = \text{borda}(c_m) = \frac{1-\delta}{1+\delta(m-2)}(m-1)^2 = (1-\delta)\,\text{borda}(c_1)$$

and the algorithm BoostedPlurality$_\delta$ outputs candidate c_1.

The social cost of c_1 satisfies

$$\mathrm{SC}(c_1) \to 1 + 2\cdot\frac{1-\delta}{1+\delta(m-2)}(m-1)$$

when $\epsilon \to 0$, and

$$\mathrm{SC}(c_m) \to 1$$

when $\epsilon \to 0$. It follows that the consistency of the algorithm is greater than $2 \cdot \frac{1-\delta}{1+\delta(m-2)}(m-1)+1$. □

Note that if $\delta = 0$ then the lower bound of Proposition 6.1.4 matches the distortion of the simple Borda rule. For $\delta < \frac{1}{m-1}$ the lower bound for the consistency becomes $1 + \frac{(m-1)(m-2)}{m-\frac{3}{2}}$. This is close to m for large values of m, which is better than the lower bound for the distortion of the simple Borda rule. Still, we showed that BoostedPlurality$_\delta$ cannot achieve sublinear consistency while having bounded robustness, meaning that predictions do not offer significant improvement.

Finally, we shall also provide bounds on the distortion achieved by BoostedPlurality$_\delta$ and BoostedBorda$_\delta$ on instances where the prediction p has a given error η. We consider first the boosted plurality algorithm.

Proposition 6.1.5. *For any $\delta \in [0, 1)$, the distortion achieved by* BoostedPlurality$_\delta$ *on instances where the prediction p has error $\eta = n \cdot d(p, c^*)/\mathrm{SC}(c^*)$ is at most*

$$\min\left\{\max\left\{1 + \left(1 - \frac{1-\delta}{m}\right)\eta, 2m-1\right\}, \frac{2m}{1-\delta} - 1\right\}.$$

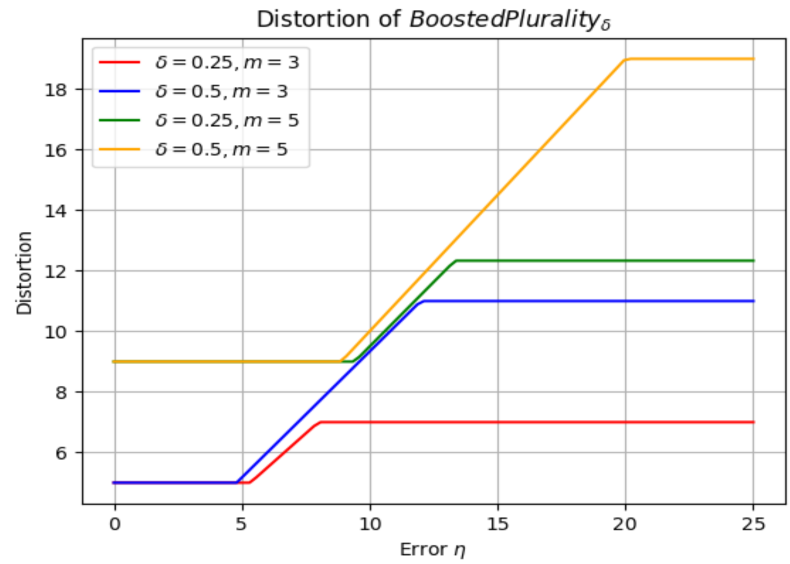

Figure 6.3: Distortion of BoostedPlurality$_\delta$ as a function of error η

.

Proof. Proposition 6.1.1 already shows that

$$\frac{\mathrm{SC}(a)}{\mathrm{SC}(c^*)} \leq \frac{2m}{1-\delta} - 1.$$

Recall from (3.1) that

$$\frac{\mathrm{SC}(c)}{\mathrm{SC}(c^*)} \leq 1 + \frac{(n - |cc^*|) \cdot d(c, c^*)}{\mathrm{SC}(c^*)}$$

for every candidate c. For any $\delta \in [0, 1)$ we have set $a = p$ if $\mathrm{plu}(p) > (1-\delta)\mathrm{plu}(c_1)$ and $a = c_1$ otherwise. In the proof of Proposition 6.1.1 we have also seen that, in the case $a = p$,

$$|pc^*| \geq \mathrm{plu}(p) > (1-\delta)n/m.$$

Combining the above we get

$$\frac{\mathrm{SC}(p)}{\mathrm{SC}(c^*)} \leq 1 + \frac{(n-(1-\delta)n/m)\cdot d(p,c^*)}{\mathrm{SC}(c^*)}$$
$$= 1 + \left(1 - \frac{1-\delta}{m}\right)\frac{n\cdot d(p,c^*)}{\mathrm{SC}(c^*)}$$
$$= 1 + \left(1 - \frac{1-\delta}{m}\right)\eta.$$

On the other hand, if $a = c_1$ we have seen that

$$\frac{\mathrm{SC}(c_1)}{\mathrm{SC}(c^*)} \leq 2m - 1.$$

So, we have proved that

$$\frac{\mathrm{SC}(a)}{\mathrm{SC}(c^*)} \leq \max\left\{1 + \left(1 - \frac{1-\delta}{m}\right)\eta, 2m-1\right\},$$

which completes the proof. □

For the boosted Borda algorithm we work in the same way.

Proposition 6.1.6. *For any* $\delta \in \left[0, \frac{1}{m-1}\right)$*, the distortion achieved by* BoostedBorda$_\delta$ *on instances where the prediction* p *has error* $\eta = n\cdot d(p,c^*)/\mathrm{SC}(c^*)$ *is at most*

$$\min\left\{\max\left\{1 + \left(1 - \frac{1-\delta(m-1)}{m}\right)\eta, 2m-1\right\}, \frac{2m}{1-\delta(m-1)} - 1\right\}.$$

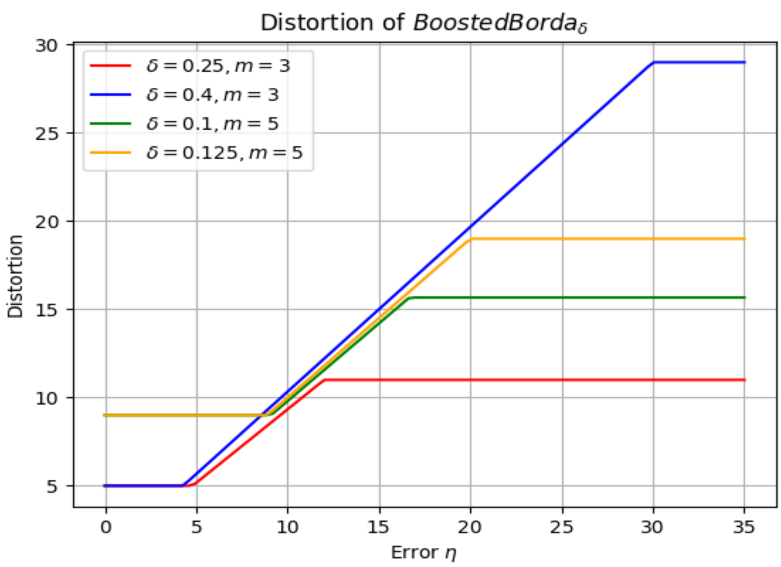

Figure 6.4: Distortion of BoostedBorda$_\delta$ as a function of error η

.

Proof. Proposition 6.1.3 already shows that

$$\frac{\mathrm{SC}(a)}{\mathrm{SC}(c^*)} \leq \frac{2m}{1-\delta(m-1)} - 1.$$

Again we have that

$$\frac{\text{SC}(c)}{\text{SC}(c^*)} \leq 1 + \frac{(n - |cc^*|) \cdot d(c, c^*)}{\text{SC}(c^*)}$$

for every candidate c. For any $\delta \in \left[0, \frac{1}{m-1}\right)$ we have set $a = p$ if $\text{borda}(p) > (1-\delta)\text{score}(c_1)$ and $a = c_1$ otherwise. In the case $a = p$ we have also seen that

$$|pc^*| \geq \frac{(1 - \delta(m-1))n}{m}.$$

Combining the above we get

$$\begin{aligned}
\frac{\text{SC}(p)}{\text{SC}(c^*)} &\leq 1 + \frac{(n - (1 - \delta(m-1))n/m) \cdot d(p, c^*)}{\text{SC}(c^*)} \\
&= 1 + \left(1 - \frac{1 - \delta(m-1)}{m}\right) \frac{n \cdot d(p, c^*)}{\text{SC}(c^*)} \\
&= 1 + \left(1 - \frac{1 - \delta(m-1)}{m}\right) \eta.
\end{aligned}$$

On the other hand, if $a = c_1$ we have seen that

$$\frac{\text{SC}(c_1)}{\text{SC}(c^*)} \leq 2m - 1.$$

So, we have proved that

$$\frac{\text{SC}(a)}{\text{SC}(c^*)} \leq \max\left\{1 + \left(1 - \frac{1 - \delta(m-1)}{m}\right)\eta, 2m - 1\right\},$$

which completes the proof. □

6.2 Greedy algorithm for k-committee election with predictions

In the k-committee election problem, we assume that predictions are given in the form of a subset $\mathcal{P} = \{p_1, \ldots, p_k\} \subseteq C$ of size k, and that $p_1, \ldots, p_k$ are ordered according to their location on the real line. We define $\mathcal{P}[c, c'] = \mathcal{P} \cap [c, c'] \cup \{c, c'\}$. This is the set of candidates in the predicted committee who are located between c and c', including the candidates at the endpoints of the interval.

We modify the greedy algorithm that we studied in Chapter 5 (Algorithm 2) enhanced by predictions about the optimal committee. The algorithm is parameterized by $\delta \in [0, 1)$, which indicates our confidence in the prediction. At each step of the algorithm, we add to the committee S either the most distant candidate thus far or the most distant candidate *in* $\mathcal{P}$ thus far. The latter is chosen if his distance from S is greater than $1 - \delta$ times the distance of the former from S. The same applies for the selection of the leftmost and rightmost candidates in the first steps. The algorithm maintains a set $\hat{\mathcal{P}}$ similar to $\hat{C}$, with the most distant candidates from S in $\mathcal{P}$ in each interval as well as their distances. See Section 5.3 for the analysis of the simple greedy algorithm, and see also the comments regarding Algorithm 6 for implementation details needed for its correctness.

Regarding the number of queries used by Algorithm 6, note that 3 queries are used at the start. Distant-candidate algorithm is called at most $4(k-3)+2$ times, and it uses at most 3 queries each time. Therefore, in total, the max number of distance queries used is at most $12k - 27$, which is $\Theta(k)$.

Algorithm 6: Greedy algorithm with predictions, $\delta \in [0, 1)$

Input: Candidates $C = \{c_1, \ldots, c_m\}$, $k \in \{2, \ldots, m-1\}$, voter ranking profile $\vec{\succ} = (\vec{\succ}_1, \ldots, \vec{\succ}_n)$, prediction $\mathcal{P} = \{p_1, \ldots, p_k\} \subseteq C$

Output: Set $S \subseteq C$ of k candidates

if $d(p_1, c_m) > (1-\delta)d(c_1, c_m)$ **then**
    $S \leftarrow \{p_1\}$
    $\hat{C} \leftarrow \{(c_1, d(c_1, p_1))\}$ // c_1 is the most distant candidate left of p_1 and can be selected later
    $c_\ell \leftarrow p_1$
else
    $S \leftarrow \{c_1\}$
    $c_\ell \leftarrow c_1$
if $d(c_1, p_k) > (1-\delta)d(c_1, c_m)$ **then**
    $S \leftarrow S \cup \{p_k\}$
    $\hat{C} \leftarrow \hat{c} \cup \{(c_m, d(p_k, c_m))\}$ // c_m is the most distant candidate right of p_k and can be selected later
    $c_r \leftarrow p_k$
else
    $S \leftarrow S \cup \{c_m\}$
    $c_r \leftarrow c_m$
$\hat{C} \leftarrow \{\text{Distant} - \text{Candidate}\,(C[c_\ell, c_r])\}$
$\hat{\mathcal{P}} \leftarrow \{\text{Distant} - \text{Candidate}\,(\mathcal{P}[c_\ell, c_r])\}$
while $|S| < k$ **do**
    Let c be s.t. $(c, t) \in \hat{C}$ and $t \geq t'$ for all $(c', t') \in \hat{C}$
    Let p be s.t. $(p, r) \in \hat{\mathcal{P}}$ and $r \geq r'$ for all $(p', r') \in \hat{\mathcal{P}}$
    if $r > (1-\delta)t$ **then**
        $S \leftarrow S \cup \{p\}$
        $c \leftarrow p$
        Let $(c'', t'') \in \hat{C}$ and $c_i < c'' < c_{i+1}$ // c_i, c_{i+1} defined later, c'' is unique
        $\hat{C} \leftarrow \hat{C} \setminus \{(c'', t'')\}$ // The interval $[c_i, c_{i+1}]$ will split, so we remove its most distant candidate from $\hat{C}$
        $\hat{\mathcal{P}} \leftarrow \hat{\mathcal{P}} \setminus \{(p, r)\}$
    else
        $S \leftarrow S \cup \{c\}$
        Let $(p'', r'') \in \hat{\mathcal{P}}$ and $c_i < p'' < c_{i+1}$ // c_i, c_{i+1} defined later, p'' is unique if it exists
        $\hat{C} \leftarrow \hat{C} \setminus \{(c, t)\}$
        $\hat{\mathcal{P}} \leftarrow \{(p'', r'')\}$ // The interval $[c_i, c_{i+1}]$ will split, so we remove its most distant candidate from $\hat{P}$
    if $|S| < k$ **then**
        Let c_i be the rightmost candidate in S on c's left // if $c = c_1$ then c_i does not exist and in the following step we ignore it
        Let c_{i+1} be the leftmost candidate in S on c's right // if $c = c_m$ then c_{i+1} does not exist and in the following step we ignore it
        $\hat{C} \leftarrow \hat{C} \cup \{\text{Distant} - \text{Candidate}(C[c_i, c])\} \cup \{\text{Distant} - \text{Candidate}(C[c, c_{i+1}])\}$
        $\hat{\mathcal{P}} \leftarrow \hat{\mathcal{P}} \cup \{\text{Distant} - \text{Candidate}(\mathcal{P}[c_i, c])\} \cup \{\text{Distant} - \text{Candidate}(C[c, c_{i+1}])\}$
return S

We shall give an upper bound for the robustness of the Greedy algorithm with respect to the egalitarian cost (EC) and to the social cost (SC).

Proposition 6.2.1. *The algorithm achieves robustness*

$$\text{robustness}(\text{Greedy})_{\text{EC}} \leq 1 + \frac{4}{1-\delta} \quad \textit{and} \quad \text{robustness}(\text{Greedy})_{\text{SC}} \leq \left(1 + \frac{4}{1-\delta}\right) n.$$

Proof. Let $S^* = \{j_1, \ldots, j_k\}$ denote the optimal k-committee with respect to the egalitarian cost for the candidate restricted instance C_{cr}, and let OPT denote the optimal egalitarian cost. This solution partitions the candidates (and their corresponding voters) into clusters $C_1, \ldots, C_k$, where each $c \in C$ is placed in C_i if his closest committee member is j_i. By the triangle inequality, two points in the same cluster are at most 2 OPT apart.

Now, consider the committee $S \subseteq C$ selected by the algorithm. If some candidate in S is selected from each cluster of the optimal solution S^*, then S is 2-approximate.

Suppose that the algorithm selects two candidates in the same cluster. In some iteration step it selects $j \in C_i$ although it had selected $j' \in C_i$ in an earlier iteration step. Then, $d(j, j') \leq 2\,\text{OPT}$.

If (j, t_ℓ) was in $\hat{C}$ then j would be the farthest candidate from S up to this point (and $t_\ell \leq d(j, j') \leq 2\,\text{OPT}$). So, all voters are within a distance of 2 OPT from S, and that remains true.

If (j, r_ℓ) was in $\hat{\mathcal{P}}$ then we should have $r_\ell > (1-\delta)t_\ell$. Also, $r_\ell \leq 2\,\text{OPT}$. The distance of the farthest candidate from S is less than

$$t_\ell < \frac{r_\ell}{1-\delta} \leq \frac{2}{1-\delta}\,\text{OPT}.$$

So, all voters are within a distance of $\dfrac{2}{1-\delta}$ OPT from S and S is $\frac{2}{1-\delta}$-approximate.

From Theorem 5.3.3 we get the following bound for the robustness of S for the original instance with respect to the egalitarian cost:

$$\text{EC(S)} \leq \left(1 + 2 \cdot \frac{2}{1-\delta}\right) \text{EC(S}^*) = 1 + \frac{4}{1-\delta}\text{EC(S}^*).$$

Also,

$$\text{SC(S)} \leq n\text{EC(S)} \leq \left(1 + \frac{4}{1-\delta}\right) n\,\text{EC(S}^*) \leq \left(1 + \frac{4}{1-\delta}\right) n\,\text{SC(S}^*).$$

This concludes the proof. □

Next, we provide a lower bound for the consistency of the Greedy algorithm.

Proposition 6.2.2. *For every $\delta \in [0, 1/3)$ the Greedy algorithm has consistency*

$$\text{consistency}(\text{Greedy}) \geq \frac{n-1}{3}$$

and for $\delta \in [1/3, 1)$

$$\text{consistency}(\text{Greedy}) \geq (n-1)\frac{1-\delta}{6\delta}$$

Proof. Consider the example shown in Figure 6.5.

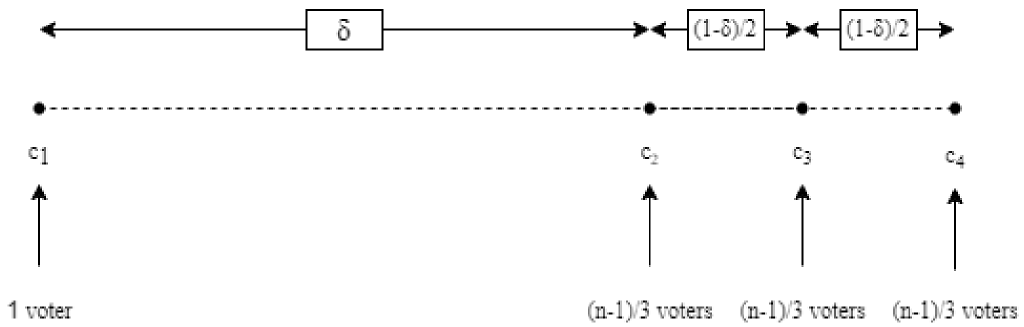

Figure 6.5: Example for Proposition 6.2.2.

For a fixed value of δ we choose $n > \frac{6\delta}{1-\delta} + 1$ and construct the instance above, where $m = 4$ and $k = 3$.

There are 4 possible committees. The following table summarizes their social cost.

Candidate omitted	Social cost
c_1	δ
c_2	$\frac{n-1}{3} \min\left\{\delta, \frac{1-\delta}{2}\right\}$
c_3	$\frac{n-1}{3}\frac{1-\delta}{2}$
c_4	$\frac{n-1}{3}\frac{1-\delta}{2}$

Since $n > \frac{6\delta}{1-\delta} + 1$ we have that $\delta < \frac{n-1}{3}\frac{1-\delta}{2}$. Then, the table shows that the optimal committee is $S^* = \{c_2, c_3, c_4\}$ and for the consistency case we assume that $P = S^* = \{c_2, c_3, c_4\}$.

In the first two iterations, the algorithm selects c_1 and c_4. So far, $S = \{c_1, c_4\}$. In the third iteration we have $d(c_2, S) = \min\{\delta, 1-\delta\}$ (which is δ if $\delta < \frac{1}{2}$ and $1-\delta$ if $\delta \geq \frac{1}{2}$) and $d(c_3, S) = \frac{1-\delta}{2}$. So, we have two cases for the farthest candidate:

- If $\delta \in \left[0, \frac{1}{3}\right)$ then $S = \{c_1, c_3, c_4\}$ and $\text{SC} = \delta\,\frac{n-1}{3}$.
- If $\delta \in \left[\frac{1}{3}, 1\right)$ then $S = \{c_1, c_2, c_4\}$ and $\text{SC} = \frac{1-\delta}{2}\frac{n-1}{3}$.

The optimal social cost is equal to δ, for $S^* = \{c_2, c_3, c_4\}$, therefore:

- In the first case the algorithm has distortion $\frac{n-1}{3}$.
- In the second case the algorithm has distortion $(n-1)\frac{1-\delta}{6\delta}$.

The result follows. □

In both cases, the consistency is $\Omega(n)$, therefore the improvement given by the predictions is not significant, especially for large values of n.

6.3 An algorithm achieving constant consistency with $\Theta(k)$ queries

Before presenting the algorithm, we define the notion of a (ℓ, β)-**good** set.

Let $\beta \geq 1$ and $\ell \geq k$ be an integer. A subset $C' \subseteq C$ of candidates is called (ℓ, β)-**good** if $|C'| = \ell$ and

$$\text{SC}(C') \leq \beta\,\text{SC}(S^*)$$

where S^* is an optimal k-committee for the original instance. We say that C' is ℓ-*sparse*, in the sense that there are $\ell \leq m$ candidates in C' (ideally, $\ell \ll m$) and that C' is β-*good*, in the sense that if we represent each voter by its top candidate in C' then we impose a cost at most β times the optimal social cost. In this terminology, the original set C of candidates is $(m, 1)$-good and any k-committee with distortion β is (k, β)-good.

Algorithm 7: Good set computation with predictions

Input: $C, k, \vec{\succ}, \mathcal{P}$

Output: Set $S \subseteq C$ of k candidates

1 $S' \leftarrow$ Greedy$\{C, k, \vec{\succ}\}$

2 $C' \leftarrow S' \cup \mathcal{P}$ // $|C'| \leq 2k$

3 Let C'_{cr} be the candidate-restricted instance induced by C'

4 Compute the distances between all active candidates in C' using distance queries

5 $S \leftarrow DP\{C'_{\text{cr}}, \text{distances}, k\}$ // $|S| = k, S$ optimal for C'_{cr}

6 **return** S

This algorithm computes a **good** set C' of at most $2k$ candidates by taking the union of the set S' given by the Greedy Algorithm 2 and the predicted set $\mathcal{P}$. The DP-algorithm in Step 5, proposed in Hassin and Tamir in [53], runs in $O(k^2)$ time and computes the optimal k-committee for the candidate restricted instance C'_{cr} when having access to all distances between candidates.

This algorithm uses at most $6k - 15$ queries in Step 1 by calling Algorithm 2 and at most $2k - 1$ queries in Step 4 by asking the distance between each pair of consecutive candidates in C'. So, in total, the algorithm uses $\Theta(k)$ candidate distance queries.

To prove consistency and robustness bounds for Algorithm 7 we need a theorem from [45]. Let C' be an (ℓ, β)-good set of candidates and let

$$C'_{\text{cr}} = \{(c_1, n_1), \ldots, (c_\ell, n_\ell)\}$$

be the candidate-restricted instance induced by C'. Then, $c_1 < \cdots < c_\ell$ are the locations of the candidates from C' on the real line, and $n_i = |\text{Cluster}(c_i)|$ is the number of voters that have c_i as their top candidate in C'. We always have $n_1 + \cdots + n_\ell = n$ and we may assume that all n_i are strictly positive, by removing the inactive candidates from C'.

The next theorem of Fotakis, Gourvés and Patsilinakos [45] shows that an optimal k-committee for the candidate-restricted instance C'_{cr} induced by an (ℓ, β)-good set C' achieves a distortion of $1 + 2\beta$ for the original instance.

Theorem 6.3.1 (Fotakis-Gourvés-Patsilinakos)**.** *Let (C, V) be an instance of the k-committee election. Let $C' \subseteq C$ be an (ℓ, β)-good set, let C'_{cr} be the candidate-restricted instance induced by C' and let S, respectively S^*, be an optimal k-committee for C'_{cr}, respectively for (C, V). Then,*

$$\text{SC}(S) \leq (1 + 2\beta)\, \text{SC}(S^*).$$

For the proof of Theorem 6.3.1 we need the next proposition.

Proposition 6.3.2 (Fotakis-Gourvés-Patsilinakos)**.** *Let C be the set of all candidates and let $\tilde{C}$ be the set of active candidates. For any instance in which every voter has the same location with its top candidate and for every committee S that includes inactive candidates, i.e. $S \setminus \tilde{C} \neq \varnothing$, we can find another committee $S' \subseteq \tilde{C}$ such that $\text{SC}(S') \leq \text{SC}(S)$.*

Proof. Let S be a committee in which there exists some candidate $c \notin \tilde{C}$. If there are no voters represented by c, then we can remove c from S to get a new committee with $k - 1$ candidates, less inactive candidates, and $\mathrm{SC}(S \setminus \{c\}) = \mathrm{SC}(S)$. Adding any candidate from $\tilde{C} \setminus S$ to $S \setminus \{c\}$, we obtain a k-committee with at most the same cost.

If c represents some voters, let V_{left} be the set of voters on the left of c that are represented by c under the k-committee S, and let V_{right} be the set of voters on the right of c that are represented by c under the k-committee S. By our hypothesis, every voter in $V_{\text{left}} \cup V_{\text{right}}$ has the same location with some candidate. If $|V_{\text{left}}| > |V_{\text{right}}|$ then we replace c in S with a candidate c' which has the same location as the rightmost voter of V_{left}. Otherwise, we replace c in S with the candidate c' that has the same location as the leftmost voter of V_{right}. In this way, we obtain a new k-committee $S' = (S \setminus \{c\}) \cup \{c'\}$ with less inactive candidates than S and $\mathrm{SC}(S') \leq \mathrm{SC}(S)$. Applying the same argument repeatedly, we obtain a k-committee $S' \subseteq \tilde{C}$ with $\mathrm{SC}(S') \leq \mathrm{SC}(S)$. □

Proof of Theorem 6.3.1. For every voter v we set $\mathrm{top}'(v) \in C'$ to be the top candidate of v in C'. By the triangle inequality we have that

$$d(v, S) \leq d(v, \mathrm{top}'(v)) + d(\mathrm{top}'(v), S).$$

Taking the sum over all voters $v \in V$ we get

$$\mathrm{SC}(S) \leq \mathrm{SC}(C') + \mathrm{SC}(C'_{\text{cr}}, S), \tag{6.1}$$

where

$$\mathrm{SC}(C') = \sum_{v \in V} d(v, \mathrm{top}'(v)) = \sum_{v \in V} d(v, C')$$

and

$$\mathrm{SC}(C'_{\text{cr}}, S) = \sum_{v \in V} d(\mathrm{top}'(v), S) = \sum_{i=1}^{\ell} n_i d(c_i, S)$$

is the social cost of S for the candidate-restricted instance C'_{cr}.

From Proposition 6.3.2 we know that we can replace candidates in $S^* \setminus C'$ with candidates in S in C'_{cr} without increasing the social cost, therefore

$$\mathrm{SC}(C'_{\text{cr}}, S) \leq \mathrm{SC}(C'_{\text{cr}}, S^*).$$

By the triangle inequality we also have

$$d(\mathrm{top}'(v), S^*) \leq d(\mathrm{top}'(v), v) + d(v, S^*),$$

and so, we finally get

$$\mathrm{SC}(C'_{\text{cr}}, S^*) \leq \mathrm{SC}(C') + \mathrm{SC}(S^*).$$

Then, from (6.1) we see that

$$\mathrm{SC}(S) \leq 2\mathrm{SC}(C') + \mathrm{SC}(S^*) \leq (1 + 2\beta)\, \mathrm{SC}(S^*),$$

where the second inequality follows from the hypothesis that C' is an (ℓ, β)-good set of candidates. □

Now, we can prove upper bounds for the consistency and the robustness of Algorithm 7.

Proposition 6.3.3. *Algorithm 7 achieves a consistency of at most* 3.

Proof. For the consistency case, we assume that $\mathcal{P}$ is optimal for (C, V) and C' is a superset of $\mathcal{P}$, so $\text{SC}(C') \leq \text{SC}(\mathcal{P}) = \text{SC}(S^*)$, where S^* is an optimal k-committee. Therefore, C' is a $(|C'|, 1)$-good set.

From Theorem 6.3.1 we have that

$$\text{SC}(S) \leq (1 + 2 \cdot 1)\text{SC}(S^*) = 3\text{SC}(S^*).$$

This shows that the consistency of Algorithm 7 is at most 3. □

Proposition 6.3.4. *Algorithm 7 achieves a robustness of at most* $10n + 1$.

Proof. From Theorem 5.3.4 we have that the Greedy Algorithm has a distortion of at most $5n$. Also, C' is a superset of S', so $\text{SC}(C') \leq \text{SC}(S') \leq 5n\,\text{SC}(S^*)$, where S^* is an optimal k-committee. Therefore, C' is a $(|C'|, 5n)$-good set.

From Theorem 6.3.1 we have that

$$\text{SC}(S) \leq (1 + 2 \cdot 5n)\text{SC}(S^*) = (10n + 1)\text{SC}(S^*).$$

This shows that the robustness of Algorithm 7 is at most $\Theta(n)$. □

Bibliography

[1] B. Abramowitz, E. Anshelevich, and W. Zhu, "Awareness of voter passion greatly improves the distortion of metric social choice," in *Web and Internet Economics - 15th International Conference, WINE 2019, New York, NY, USA, December 10-12, 2019, Proceedings*, ser. Lecture Notes in Computer Science, I. Caragiannis, V. S. Mirrokni, and E. Nikolova, Eds., vol. 11920. Springer, 2019, pp. 3–16. [Online]. Available: https://doi.org/10.1007/978-3-030-35389-6_1

[2] P. Agrawal, E. Balkanski, V. Gkatzelis, T. Ou, and X. Tan, "Learning-augmented mechanism design: Leveraging predictions for facility location," *CoRR*, vol. abs/2204.01120, 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2204.01120

[3] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, and A. A. Voudouris, "Peeking behind the ordinal curtain: Improving distortion via cardinal queries," *Artif. Intell.*, vol. 296, p. 103488, 2021. [Online]. Available: https://doi.org/10.1016/j.artint.2021.103488

[4] ——, "Don't roll the dice, ask twice: The two-query distortion of matching problems and beyond," in *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., 2022. [Online]. Available: http://papers.nips.cc/paper_files/paper/2022/hash/c5ec22711f3a4a2f4a0a8ffd92167190-Abstract-Conference.html

[5] ——, "A few queries go a long way: Information-distortion tradeoffs in matching," *J. Artif. Intell. Res.*, vol. 74, 2022. [Online]. Available: https://doi.org/10.1613/jair.1.12690

[6] I. Anagnostides, D. Fotakis, and P. Patsilinakos, "Dimensionality and coordination in voting: The distortion of STV," in *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*. AAAI Press, 2022, pp. 4776–4784. [Online]. Available: https://doi.org/10.1609/aaai.v36i5.20404

[7] ——, "Metric-distortion bounds under limited information," *J. Artif. Intell. Res.*, vol. 74, pp. 1449–1483, 2022. [Online]. Available: https://doi.org/10.1613/jair.1.13338

[8] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron, "Approximating optimal social choice under metric preferences," *Artif. Intell.*, vol. 264, pp. 27–51, 2018. [Online]. Available: https://doi.org/10.1016/j.artint.2018.07.006

[9] E. Anshelevich, A. Filos-Ratsikas, N. Shah, and A. A. Voudouris, "Distortion in social choice problems: The first 15 years and beyond," *CoRR*, vol. abs/2103.00911, 2021. [Online]. Available: https://arxiv.org/abs/2103.00911

[10] E. Anshelevich, A. Filos-Ratsikas, and A. A. Voudouris, "The distortion of distributed metric social choice," *CoRR*, vol. abs/2107.05456, 2021. [Online]. Available: https://arxiv.org/abs/2107.05456

[11] E. Anshelevich and J. Postl, "Randomized social choice functions under metric preferences," *J. Artif. Intell. Res.*, vol. 58, pp. 797–827, 2017. [Online]. Available: https://doi.org/10.1613/jair.5340

[12] A. Antoniadis, T. Gouleakis, P. Kleer, and P. Kolev, “Secretary and online matching problems with machine learned advice,” *Discret. Optim.*, vol. 48, no. Part 2, p. 100778, 2023. [Online]. Available: https://doi.org/10.1016/j.disopt.2023.100778

[13] K. Arrow, *Advances in the Spatial Theory of Voting*, J. M. Enelow and M. J. Hinich, Eds. Cambridge University Press, 1990.

[14] K. J. Arrow, *Social Choice and Individual Values*. Yale University Press, 2012. [Online]. Available: http://www.jstor.org/stable/j.ctt1nqb90

[15] Y. Azar, D. Panigrahi, and N. Touitou, “Online graph algorithms with predictions,” in *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, J. S. Naor and N. Buchbinder, Eds. SIAM, 2022, pp. 35–66. [Online]. Available: https://doi.org/10.1137/1.9781611977073.3

[16] E. Balkanski, V. Gkatzelis, and X. Tan, “Strategyproof scheduling with predictions,” in *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, ser. LIPIcs, Y. T. Kalai, Ed., vol. 251. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pp. 11:1–11:22. [Online]. Available: https://doi.org/10.4230/LIPIcs.ITCS.2023.11

[17] E. Balkanski, V. Gkatzelis, X. Tan, and C. Zhu, “Online mechanism design with predictions,” *CoRR*, vol. abs/2310.02879, 2023. [Online]. Available: https://doi.org/10.48550/arXiv.2310.02879

[18] É. Bamas, A. Maggiori, and O. Svensson, “The primal-dual method for learning augmented algorithms,” in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., 2020. [Online]. Available: https://proceedings.neurips.cc/paper/2020/hash/e834cb114d33f729dbc9c7fb0c6bb607-Abstract.html

[19] S. Banerjee, V. Gkatzelis, A. Gorokh, and B. Jin, “Online nash social welfare maximization with predictions,” in *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, J. S. Naor and N. Buchbinder, Eds. SIAM, 2022, pp. 1–19. [Online]. Available: https://doi.org/10.1137/1.9781611977073.1

[20] S. Banerjee, V. Gkatzelis, S. Hossain, B. Jin, E. Micha, and N. Shah, “Proportionally fair online allocation of public goods with predictions,” in *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*. ijcai.org, 2023, pp. 20–28. [Online]. Available: https://doi.org/10.24963/ijcai.2023/3

[21] B. Berger, M. Feldman, V. Gkatzelis, and X. Tan, “Optimal metric distortion with predictions,” *CoRR*, vol. abs/2307.07495, 2023. [Online]. Available: https://doi.org/10.48550/arXiv.2307.07495

[22] A. Borodin and R. El-Yaniv, *Online computation and competitive analysis*. Cambridge University Press, 1998.

[23] A. Borodin, D. Halpern, M. Latifian, and N. Shah, “Distortion in voting with top-t preferences,” in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, ser. Proceedings of the International Joint Conference on Artificial Intelligence, L. De Raedt, Ed. IJCAI Organization, Jul. 2022, pp. 116–122, the 31st International Joint Conference on Artificial Intelligence and the 25th European Conference on Artificial Intelligence, IJCAI-ECAI 2022 ; Conference date: 23-07-2022 Through 29-07-2022. [Online]. Available: https://ijcai-22.org/

[24] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet, “Optimal social choice functions: A utilitarian view,” *Artif. Intell.*, vol. 227, pp. 190–213, 2015. [Online]. Available: https://doi.org/10.1016/j.artint.2015.06.003

[25] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, Eds., *Handbook of Computational Social Choice*. Cambridge University Press, 2016. [Online]. Available: https://doi.org/10.1017/CBO9781107446984

[26] J. Burkhardt, I. Caragiannis, K. Fehrs, M. Russo, C. Schwiegelshohn, and S. Shyam, “Low-distortion clustering with ordinal and limited cardinal information,” *CoRR*, vol. abs/2402.04035, 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2402.04035

[27] I. Caragiannis and G. Kalantzis, “Randomized learning-augmented auctions with revenue guarantees,” *CoRR*, vol. abs/2401.13384, 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2401.13384

[28] I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah, “Subset selection via implicit utilitarian voting,” *J. Artif. Intell. Res.*, vol. 58, pp. 123–152, 2017. [Online]. Available: https://doi.org/10.1613/jair.5282

[29] I. Caragiannis and A. D. Procaccia, “Voting almost maximizes social welfare despite limited communication,” *Artif. Intell.*, vol. 175, no. 9-10, pp. 1655–1671, 2011. [Online]. Available: https://doi.org/10.1016/j.artint.2011.03.005

[30] I. Caragiannis, N. Shah, and A. A. Voudouris, “The metric distortion of multiwinner voting,” *Artif. Intell.*, vol. 313, p. 103802, 2022. [Online]. Available: https://doi.org/10.1016/j.artint.2022.103802

[31] J. R. Chamberlin and P. N. Courant, “Representative deliberations and representative decisions: Proportional representation and the borda rule,” *American Political Science Review*, vol. 77, no. 3, p. 718–733, 1983.

[32] M. Charikar and P. Ramakrishnan, “Metric distortion bounds for randomized social choice,” in *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, J. S. Naor and N. Buchbinder, Eds. SIAM, 2022, pp. 2986–3004. [Online]. Available: https://doi.org/10.1137/1.9781611977073.116

[33] M. Charikar, K. Wang, P. Ramakrishnan, and H. Wu, “Breaking the metric voting distortion barrier,” in *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, D. P. Woodruff, Ed. SIAM, 2024, pp. 1621–1640. [Online]. Available: https://doi.org/10.1137/1.9781611977912.65

[34] X. Chen, M. Li, and C. Wang, “Favorite-candidate voting for eliminating the least popular candidate in a metric space,” 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*. AAAI Press, 2020, pp. 1894–1901. [Online]. Available: https://doi.org/10.1609/aaai.v34i02.5558

[35] S. Cho and J. W. Endersby, “Issues, the spatial theory of voting, and british general elections: A comparison of proximity and directional models,” *Public Choice*, vol. 114, no. 3/4, pp. 275–293, 2003. [Online]. Available: http://www.jstor.org/stable/30025956

[36] A. Copeland, “A reasonable social welfare function,” 1951.

[37] P. Dütting, S. Lattanzi, R. P. Leme, and S. Vassilvitskii, “Secretaries with advice,” *Math. Oper. Res.*, vol. 49, no. 2, pp. 856–879, 2024. [Online]. Available: https://doi.org/10.1287/moor.2023.1384

[38] E. Elkind and P. Faliszewski, “Recognizing 1-euclidean preferences: An alternative approach,” in *Algorithmic Game Theory - 7th International Symposium, SAGT 2014, Haifa, Israel, September 30 - October 2, 2014. Proceedings*, ser. Lecture Notes in Computer Science, R. Lavi, Ed., vol. 8768. Springer, 2014, pp. 146–157. [Online]. Available: https://doi.org/10.1007/978-3-662-44803-8_13

[39] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko, “Properties of multiwinner voting rules,” *Soc. Choice Welf.*, vol. 48, no. 3, pp. 599–632, 2017. [Online]. Available: https://doi.org/10.1007/s00355-017-1026-z

[40] J. M. Enelow and M. J. Hinich, *The Spatial Theory of Voting*, ser. Cambridge Books. Cambridge University Press, 1984, no. 9780521275156. [Online]. Available: https://ideas.repec.org/b/cup/cbooks/9780521275156.html

[41] B. Fain, A. Goel, K. Munagala, and N. Prabhu, "Random dictators with a random referee: constant sample complexity mechanisms for social choice," in *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, ser. AAAI'19/IAAI'19/EAAI'19. AAAI Press, 2019. [Online]. Available: https://doi.org/10.1609/aaai.v33i01.33011893

[42] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon, *Multiwinner voting: A new challenge for social choice theory*, ser. Trends in computational social choice. Lulu Publisher, 2017, vol. 74, pp. 27–47.

[43] M. Feldman, A. Fiat, and I. Golomb, "On voting and facility location," *CoRR*, vol. abs/1512.05868, 2015. [Online]. Available: http://arxiv.org/abs/1512.05868

[44] A. Filos-Ratsikas, E. Micha, and A. A. Voudouris, "The distortion of distributed voting," *Artif. Intell.*, vol. 286, p. 103343, 2020. [Online]. Available: https://doi.org/10.1016/j.artint.2020.103343

[45] D. Fotakis, L. Gourvès, and P. Patsilinakos, "On the distortion of committee election with 1-euclidean preferences and few distance queries," *CoRR*, vol. abs/2408.11755, 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2408.11755

[46] A. Gibbard, "Manipulation of voting schemes," *Econometrica*, vol. 41, pp. 587–602, 1973.

[47] V. Gkatzelis, D. Halpern, and N. Shah, "Resolving the optimal metric distortion conjecture," in *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, S. Irani, Ed. IEEE, 2020, pp. 1427–1438. [Online]. Available: https://doi.org/10.1109/FOCS46700.2020.00134

[48] V. Gkatzelis, K. Kollias, A. Sgouritsa, and X. Tan, "Improved price of anarchy via predictions," in *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, D. M. Pennock, I. Segal, and S. Seuken, Eds. ACM, 2022, pp. 529–557. [Online]. Available: https://doi.org/10.1145/3490486.3538296

[49] V. Gkatzelis, M. Latifian, and N. Shah, "Best of both distortion worlds," in *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023, London, United Kingdom, July 9-12, 2023*, K. Leyton-Brown, J. D. Hartline, and L. Samuelson, Eds. ACM, 2023, pp. 738–758. [Online]. Available: https://doi.org/10.1145/3580507.3597739

[50] A. Goel, R. Hulett, and A. K. Krishnaswamy, "Relating metric distortion and fairness of social choice rules," in *Proceedings of the 13th Workshop on Economics of Networks, Systems and Computation, NetEcon@SIGMETRICS 2018, Irvine, CA, USA, June 18, 2018*. ACM, 2018, p. 4:1. [Online]. Available: https://doi.org/10.1145/3230654.3230658

[51] A. Goel, A. K. Krishnaswamy, and K. Munagala, "Metric distortion of social choice rules: Lower bounds and fairness properties," *CoRR*, vol. abs/1612.02912, 2016. [Online]. Available: http://arxiv.org/abs/1612.02912

[52] S. Gross, E. Anshelevich, and L. Xia, "Vote until two of you agree: Mechanisms with small distortion and sample complexity," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, Feb. 2017. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/10587

[53] R. Hassin and A. Tamir, "Improved complexity bounds for location problems on the real line," *Oper. Res. Lett.*, vol. 10, no. 7, pp. 395–402, 1991. [Online]. Available: https://doi.org/10.1016/0167-6377(91)90041-M

[54] S. Im, R. Kumar, M. M. Qaem, and M. Purohit, "Online knapsack with frequency predictions," in *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan, Eds., 2021, pp. 2733–2743. [Online]. Available: https://proceedings.neurips.cc/paper/2021/hash/161c5c5ad51fcc884157890511b3c8b0-Abstract.html

[55] G. Istrate and C. Bonchis, “Mechanism design with predictions for obnoxious facility location,” *CoRR*, vol. abs/2212.09521, 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2212.09521

[56] D. Kempe, “An analysis framework for metric voting based on ℓ_p duality,” in *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 2020, pp. 2079–2086. [Online]. Available: https://doi.org/10.1609/aaai.v34i02.5581

[57] ——, “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*. AAAI Press, 2020, pp. 2087–2094. [Online]. Available: https://doi.org/10.1609/aaai.v34i02.5582

[58] F. E. Kizilkaya and D. Kempe, “Plurality veto: A simple voting rule achieving optimal metric distortion,” in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, L. D. Raedt, Ed. ijcai.org, 2022, pp. 349–355. [Online]. Available: https://doi.org/10.24963/ijcai.2022/50

[59] ——, “Generalized veto core and a practical voting rule with optimal metric distortion,” in *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023, London, United Kingdom, July 9-12, 2023*, K. Leyton-Brown, J. D. Hartline, and L. Samuelson, Eds. ACM, 2023, pp. 913–936. [Online]. Available: https://doi.org/10.1145/3580507.3597798

[60] A. Lindermayr and N. Megow, “Alps.” [Online]. Available: https://algorithms-with-predictions.github.io/

[61] P. Lu, Z. Wan, and J. Zhang, “Competitive auctions with imperfect predictions,” *CoRR*, vol. abs/2309.15414, 2023. [Online]. Available: https://doi.org/10.48550/arXiv.2309.15414

[62] T. Lykouris and S. Vassilvitskii, “Competitive caching with machine learned advice,” *J. ACM*, vol. 68, no. 4, pp. 24:1–24:25, 2021. [Online]. Available: https://doi.org/10.1145/3447579

[63] D. Mandal, A. D. Procaccia, N. Shah, and D. P. Woodruff, “Efficient and thrifty voting by any means necessary,” in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett, Eds., 2019, pp. 7178–7189. [Online]. Available: https://proceedings.neurips.cc/paper/2019/hash/c09f9caf5e08836d4673ccdd69bb041e-Abstract.html

[64] D. Mandal, N. Shah, and D. P. Woodruff, “Optimal communication-distortion tradeoff in voting,” in *EC ’20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, P. Biró, J. D. Hartline, M. Ostrovsky, and A. D. Procaccia, Eds. ACM, 2020, pp. 795–813. [Online]. Available: https://doi.org/10.1145/3391403.3399510

[65] K. O. May, “A set of independent necessary and sufficient conditions for simple majority decision,” *Econometrica*, vol. 20, p. 680, 1952. [Online]. Available: https://api.semanticscholar.org/CorpusID:14254458

[66] S. Merrill, III and B. Grofman, *A Unified Theory of Voting: Directional and Proximity Spatial Models*. Cambridge University Press, 1999.

[67] M. Mitzenmacher and S. Vassilvitskii, “Algorithms with predictions,” *CoRR*, vol. abs/2006.09123, 2020. [Online]. Available: https://arxiv.org/abs/2006.09123

[68] B. L. Monroe, “Fully proportional representation,” *American Political Science Review*, vol. 89, no. 4, p. 925–940, 1995.

[69] K. Munagala and K. Wang, “Improved metric distortion for deterministic social choice rules,” in *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, A. R. Karlin, N. Immorlica, and R. Johari, Eds. ACM, 2019, pp. 245–262. [Online]. Available: https://doi.org/10.1145/3328526.3329550

[70] N. Nisan, *Introduction to Mechanism Design (for Computer Scientists)*. Cambridge University Press, 2007, p. 209–242.

[71] S. Prasad, M. Balcan, and T. Sandholm, "Bicriteria multidimensional mechanism design with side information," in *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., 2023. [Online]. Available: http://papers.nips.cc/paper_files/paper/2023/hash/8039ca1e9860daab3a79e45d010d5398-Abstract-Conference.html

[72] A. D. Procaccia and J. S. Rosenschein, "The distortion of cardinal preferences in voting," in *Cooperative Information Agents X, 10th International Workshop, CIA 2006, Edinburgh, UK, September 11-13, 2006, Proceedings*, ser. Lecture Notes in Computer Science, M. Klusch, M. Rovatsos, and T. R. Payne, Eds., vol. 4149. Springer, 2006, pp. 317–331. [Online]. Available: https://doi.org/10.1007/11839354_23

[73] H. Pulyassary, "Algorithm design for ordinal settings," Master's thesis, University of Waterloo, 2022.

[74] M. Purohit, Z. Svitkina, and R. Kumar, "Improving online algorithms via ML predictions," in *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., 2018, pp. 9684–9693. [Online]. Available: https://proceedings.neurips.cc/paper/2018/hash/73a427badebe0e32caa2e1fc7530b7f3-Abstract.html

[75] M. Satterthwaite, "Strategyproofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions," *Journal of Economic Theory*, vol. 10, pp. 187–217, 1975.

[76] N. Schofield, *The Spatial Model of Politics*, ser. Routledge frontiers of political economy. Routledge, 2008. [Online]. Available: https://books.google.gr/books?id=rVSamKrb_LEC

[77] P. K. Skowron and E. Elkind, "Social choice under metric preferences: Scoring rules and STV," in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, S. Singh and S. Markovitch, Eds. AAAI Press, 2017, pp. 706–712. [Online]. Available: https://doi.org/10.1609/aaai.v31i1.10591

[78] T. N. Tideman, "Independence of clones as a criterion for voting rules," *Social Choice and Welfare*, vol. 4, no. 3, pp. 185–206, 1987. [Online]. Available: http://www.jstor.org/stable/41105866

[79] V. V. Vazirani, *Approximation algorithms*. Springer, 2001. [Online]. Available: http://www.springer.com/computer/theoretical+computer+science/book/978-3-540-65367-7

[80] J. von Neumann, O. Morgenstern, and A. Rubinstein, *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press, 1944. [Online]. Available: http://www.jstor.org/stable/j.ctt1r2gkx

[81] D. P. Williamson and D. B. Shmoys, *The Design of Approximation Algorithms*. Cambridge University Press, 2011. [Online]. Available: http://www.cambridge.org/de/knowledge/isbn/item5759340/?site_locale=de_DE

[82] C. Xu and P. Lu, "Mechanism design with predictions," in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, L. D. Raedt, Ed. ijcai.org, 2022, pp. 571–577. [Online]. Available: https://doi.org/10.24963/ijcai.2022/81

[83] H. P. Young, "Condorcet's theory of voting," *American Political Science Review*, vol. 82, no. 4, p. 1231–1244, 1988.

[84] W. S. Zwicker, "Introduction to the theory of voting," in *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, Eds. Cambridge University Press, 2016, pp. 23–56. [Online]. Available: https://doi.org/10.1017/CBO9781107446984.003