Title
Influences in Voting and Growing Networks

Author
Racz, Miklos Zoltan

Publication Date
2015

Peer reviewed|Thesis/dissertation

Influences in Voting and Growing Networks

by

Miklós Zoltán Rácz

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Statistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Elchanan Mossel, Chair
Professor James W. Pitman
Professor Allan M. Sly
Professor David S. Ahn

Spring 2015

Influences in Voting and Growing Networks

Abstract

Influences in Voting and Growing Networks

by

Miklós Zoltán Rácz

Doctor of Philosophy in Statistics

University of California, Berkeley

Professor Elchanan Mossel, Chair

This thesis studies problems in applied probability using combinatorial techniques. The first part of the thesis focuses on voting, and studies the average-case behavior of voting systems with respect to manipulation of their outcome by voters. Many results in the field of voting are negative; in particular, Gibbard and Satterthwaite showed that no reasonable voting system can be strategyproof (a.k.a. nonmanipulable). We prove a quantitative version of this result, showing that the probability of manipulation is nonnegligible, unless the voting system is close to being a dictatorship. We also study manipulation by a coalition of voters, and show that the transition from being powerless to having absolute power is smooth. These results suggest that manipulation is easy on average for reasonable voting systems, and thus computational complexity cannot hide manipulations completely.

The second part of the thesis focuses on statistical inference questions in growing random graph models. In particular, we study the influence of the seed in random trees grown according to preferential attachment and uniform attachment. While the seed has no effect from a weak local limit point of view in either model, different seeds lead to different distributions of limiting trees from a total variation point of view in both models. These results open up a host of new statistical inference questions regarding the temporal dynamics of growing networks.

Mamának és Papának

Contents

Acknowledgments

First and foremost, I thank Elchanan Mossel, my advisor, for being an outstanding mentor, teacher, collaborator, and friend. I feel very lucky to have been in the right place at the right time: during my first semester at Berkeley, Elchanan taught a course called "Social Choice and Networks" which served as a springboard for my PhD. This was a fascinating class with lots of exciting topics and interesting discussions among a great group of students. I am thankful to Elchanan for letting me continue in the same spirit and for accepting me as his student. I could not have wished for a better advisor. I have grown tremendously by learning from Elchanan's advice and from his general approach to life. I have also enjoyed going hiking, climbing, playing frisbee, and just hanging out and having a good meal. Thank you, Elchanan, for everything, and I look forward to more interesting problems and fun in the future.

I am grateful to Jim Pitman, Allan Sly, and David Ahn for being on my dissertation committee. I especially thank Jim for his careful reading of my work, for his numerous comments that improved my dissertation, and for his insights and ideas that are leading to further explorations.

This thesis would not have been possible without my collaborators. I thank them for all their efforts regarding our projects and for their company which made working on these problems so much more enjoyable. I am thankful to Márton Balázs and Bálint Tóth for advising me on my first research problem, a project which we finished during the first year of my PhD; I also thank them for teaching me the fundamentals of probability theory. Another project that was finished during the beginning of my PhD was started during an unforgettable summer at IPAM with Jasmine Nirody and Susana Serna—I will always remember all our laughter. Working with Elchanan is wonderful, and I am constantly amazed at the clarity with which he sees the key issues of the problem at hand. I am very happy that I got to work on a wide variety of problems during my PhD. It was a pleasure to work with Ariel Procaccia, and I also greatly enjoyed and learned from the summer school he organized at CMU. Tonći Antunović taught me many things about Pólya urns and I also thank him for his hospitality during my visits to LA. Misha Shkolnikov's drive to know all things about interacting particle systems is infectious and I am glad that our paths intersected at Berkeley. I am thankful to Eric Friedman for taking me on board his project and to Scott Shenker for his input. Thanks to Nathan Ross for all the great times, I am happy we got to work together; thanks also to Junhyong Kim for doing his best to ground our project in biological reality. I feel very fortunate to have started working with Sébastien Bubeck when he was visiting the Simons Institute and I am thankful for the opportunity to intern with him at Microsoft Research last summer. I admire his ability to constantly ask interesting questions and I look forward to continued collaboration. MSR has a wonderful atmosphere in the summer, with many visitors present, and it is a pleasure to thank Jian Ding, Ronen Eldan, and Tselil Schramm for our collaboration which continued throughout the academic year. In addition, I thank Yuval Peres for his interest in my work, for his encouragement, and for his frequent invitations to Redmond.

While my focus has been on research, I am also very thankful to have had the chance to improve as a teacher, and this is mainly due to the excellent examples I saw around me. Zvezda Stankova does a phenomenal job of leading the Berkeley Math Circle. I am thankful to her for giving me the opportunity to lead sessions and for giving me free rein to experiment with my ideas. Philip Stark is a pioneer in online education and I feel privileged to have been a GSI for his online class based on SticiGui before MOOCs became widespread. Teaching game theory with Elchanan was a highlight of my PhD—it was so much fun. I learned to always rethink the material and to not be afraid of completely mixing things up if it is for the better. I also learned the importance of having a constant dialogue with the students, a lesson that was reaffirmed when I taught stochastic processes with Ani Adhikari. I greatly appreciate all her wisdom and encouragement concerning all things teaching. Finally, I am extremely thankful to Lajos Pósa, whose brilliance and dedication to teaching I admire, and whose example I always try to follow.

The Berkeley Statistics Department is an amazing environment to be in, providing both a vibrant academic community and a supporting group of peers, faculty, and staff. I will always cherish the five years spent here. I had the good fortune to be officemates with Miles Lopes, Christine Ho, and Linda Tran; thanks for all the great conversations and fun in the office. Thanks to Po-Ling Loh for being an amazing co-president of the SGSA and for being a good sport about all our jokes and pranks, in the good times and the bad. Many friends made the past five years memorable; in addition to those mentioned above, I thank in particular Josh Abramson, Riddhi Basu, Yuval Benjamini, Tamara Broderick, Luis Campos, Mayssa Dabaghi, Geno Guerra, Wayne Tai Lee, Joe Neeman, Karl Rohe, Ngoc Tran, Gido van de Ven, Rachel Wang, and Andre Wibisono for the great times and for the support. The staff in the department are excellent and I thank them for making life smooth, especially La Shana Porlaris, who is the best grad student advisor one could wish for.

A huge thanks to everyone I played football (soccer) with over these past five years. I always looked forward to playing in the evenings and on weekends; our matches were incredibly refreshing for me, both physically and mentally. I won't start naming names as a full list would fill a few pages. You know who you are. Thank you, and I hope to see you at Underhill or Maxwell soon.

I am thankful for the friendship and support of many people back home in Hungary, in particular Ákos, Bao, and Lili.

My family is the reason I am here today, for they have raised me and supported me throughout my life. My parents were my first teachers, my sisters my first companions. Mindent köszönök!

Jasmine's energy, laughter, smile, jokes, and unconditional love and support are what keep me going. I could not have done this without her and I look forward to the journey ahead of us.

Chapter 1

Introduction

Recent advances in technology have given us an unprecedented amount of data in a variety of disciplines throughout the social and natural sciences. This boom has greatly stimulated the fields of probability and statistics, bringing both new perspectives to classical problems and raising entirely new ones. For example, new types of data, such as time series of evolving networks, are becoming increasingly easier to collect. The availability of such data highlights the importance of considering the dynamic nature of natural and social processes.

This thesis considers if and how certain properties of a dynamic process influence its outcome. In particular, it contains contributions to two topics which have seen exciting new challenges and developments in recent years: voting and networks. The first part focuses on how easily individual voters, or a coalition of voters, can manipulate the outcome of an election by changing their votes. The second part studies statistical inference questions in growing random graphs, particularly how the initial "seed" influences the graph's structure when it is very large.

1.1 Quantitative social choice

Collective decision making has been a cornerstone of civilized society since the first democracy was established in ancient Athens. The foundations of the study of voting systems, termed social choice theory, were not laid out until 18^{th} century France by Borda and Condorcet. In particular, Condorcet noticed the following paradox [18]: when ranking three candidates, a, b, and c, it may happen that a majority of voters prefer a over b, a majority prefers b over c, and a majority prefers c over a. This produces an "irrational" circular ranking of the candidates. In the mid-20^{th} century, Arrow's impossibility theorem [4, 5] showed that this paradox holds under very natural assumptions, thus marking the basis of modern social choice theory.

Perhaps frighteningly, many results in the field of voting are negative: if there are three or more candidates, then it is impossible to design a voting system that satisfies a few desired properties all at once. For instance, a naturally desirable property of a voting system is *strategyproofness* (a.k.a. nonmanipulability): no voter should benefit from voting strategically, i.e., voting not according to her true preferences. However, Gibbard [32] and Satterthwaite [67] showed that no reasonable voting system can be strategyproof.

This has contributed to the realization that it is unlikely to expect truthfulness in voting, opening up a host of questions with regard to how voters will behave in a variety of situations. Strategizing behavior can be very complex and can lead to undesired outcomes; in fact, even the unanimously least-preferred candidate might end up winning. This problem is increasingly relevant in the area of artificial intelligence and computer science as well, because virtual elections are now an established tool in preference aggregation [27], often with a very large number of candidates, unlike for classical elections. Consequently, there have been many branches of research devoted to understanding the extent of the manipulability of voting systems, and to finding ways of circumventing the negative results.

One approach, introduced by Bartholdi, Tovey, and Trick [7], suggests computational complexity as a barrier against manipulation: if it is computationally hard for a voter to manipulate, then she would just tell the truth. This approach has been fruitful in the past 25 years, and computer scientists have shown that several voting systems are computationally resistant to manipulation. However, this is a worst-case approach, and while worst-case hardness of manipulation is a desirable property for a voting system to have, this does not tell us anything about *typical* instances of the problem—is it easy or hard to manipulate *on average*?

Understanding the average-case behavior of voting systems is the main goal of quantitative social choice, which is the topic of the first part of this thesis. We first prove a quantitative version of the classical Gibbard-Satterthwaite theorem, showing that the probability of manipulation is nonnegligible, unless the voting system is close to being a dictatorship. The main message of this result is that manipulation is easy on average for reasonable voting systems, and thus computational complexity cannot hide manipulations completely. We then study the ability of a coalition of voters to manipulate, and show that there is a smooth transition from being powerless to having absolute power. Since smooth phase transitions are often found in connection with computationally easy (polynomial) problems, this result suggests that deciding the coalitional manipulation problem may not be computationally hard in practice.

1.2 Statistical inference in networks

Data arising from real-world networks open up numerous challenging statistical inference problems. As complex systems often exhibit strong structural properties, it is natural that community detection has received a lot of attention. However, stepping away from static networks and considering *out-of-equilibrium systems*, such as growing networks, it is apparent that there are a host of challenging statistical inference problems that have not yet been considered previously. Indeed, perhaps the most important questions concern the *temporal dynamics* of these systems. For instance, given the current state of the network, what can we say about a previous state, perhaps far back in time?

A useful theoretical approach is to consider such questions on natural generative models of randomly growing graphs which capture salient features of real-world networks. Given a growing sequence of random graphs, perhaps the most basic question of this kind is whether the initial “seed” graph influences the structure of the graph at large times. If so, then in what sense? This is the problem that we tackle in the second part of the thesis, where we study the influence of the seed in two natural models of randomly growing trees: preferential attachment and uniform attachment.

We first study preferential attachment trees, where at every time step a new node is added to the tree, together with an edge connecting it to an existing node, chosen randomly with probability proportional to its degree. We show that the seed tree indeed influences the growing tree at large times: if S and T are two trees with different degree profiles, then preferential attachment trees started from these seeds and viewed at any large time n, $\mathrm{PA}(n,S)$ and $\mathrm{PA}(n,T)$, are statistically distinguishable. Subsequently, our work was extended by Curien et al. [21], who showed that the above holds also when S and T have the same degree profile but are nonisomorphic.

We then consider uniform attachment trees, where the incoming node connects to an existing one picked uniformly at random, and show that the same conclusion holds: each seed leads to a unique limiting distribution of the uniform attachment tree. We also prove, in both models, that no information can be gained about the seed by using only local statistics when the system size is large enough, and so global statistics of the tree are needed for inferring the seed.

These works are just the beginning of a larger investigation into statistical inference questions in growing networks. An immediate challenge is to understand the influence of the seed in other models of growing graphs. Another important, but more difficult, question is that of estimation: can one *find* the seed in a large random graph? Considering the effect of noise, extra information, and other variables that are present in real-world networks will make the picture clearer. Our hope is that the work presented in this thesis serves as a stepping stone for future investigations into these challenges.

1.3 Overview

Part I of the thesis studies influences in voting. In Chapter 2 we state and prove our quantitative Gibbard-Satterthwaite theorem; this chapter is joint work with Elchanan Mossel [52, 53]. Next, we study coalitional manipulation in Chapter 3, which is joint work with Elchanan Mossel and Ariel Procaccia [51]. Part II of the thesis studies influences in growing networks. After formalizing the main questions and results in Chapter 4, we prove our results for preferential attachment and uniform attachment in Chapters 5 and 6, respectively. These chapters are joint works with Sébastien Bubeck and Elchanan Mossel [14], and Sébastien Bubeck, Ronen Eldan, and Elchanan Mossel [15], respectively.

Part I

Influences in Voting

Chapter 2

A quantitative Gibbard-Satterthwaite theorem

2.1 Introduction

As discussed in Chapter 1, Part I of this thesis focuses on problems in quantitative social choice. In particular, the goal of this chapter is to prove a quantitative version of the classical Gibbard-Satterthwaite theorem. Let us begin by specifying the problem more formally.

We consider n voters electing a winner among k alternatives. We denote by $[k] := \{1, 2, \ldots, k\}$ the set of alternatives, and the voters specify their opinion by giving a strict ranking of the alternatives, also known as a total order of them. Equivalently, one can think of a vote as a permutation of the alternatives, from most-preferred to least-preferred. Accordingly, we denote the set of all possible total orderings of the k alternatives by S_k, and note that this set has cardinality $k!$. A collection of rankings by the voters is called a *ranking profile*, and is an element of S_k^n, which is the n^{th} Cartesian power of S_k and which has cardinality $(k!)^n$. The winner is determined according to some predefined *social choice function* (SCF) $f : S_k^n \to [k]$ of all the voters' rankings. We say that a SCF is *manipulable* if there exists a ranking profile where a voter can achieve a more desirable outcome of the election according to her true preferences by voting in a way that does not reflect her true preferences (see Definition 2.1 for a more detailed definition).

The Gibbard-Satterthwaite theorem states that any SCF which is not a dictatorship (i.e., not a function of a single voter), and which allows at least three alternatives to be elected, is manipulable. As discussed in Section 1.1, this has led to a lot of effort in trying to understand how voters will act. We are particularly interested in *typical* instances of the problem—is it easy or hard to manipulate *on average*?

A recent line of research with an average-case algorithmic approach has suggested that manipulation is indeed easy on average; see, e.g., Kelly [44], Conitzer and Sandholm [19], Procaccia and Rosenschein [65], and Zuckerman et al. [80] for results on certain restricted classes of SCFs (see also the survey [27]).

A different approach was taken by Friedgut, Kalai, Keller and Nisan [29, 30], who looked at the fraction of ranking profiles that are manipulable. To put it differently: assuming each voter votes independently and uniformly at random (known as the *impartial culture assumption* in the social choice literature), what is the probability that a ranking profile is manipulable? Is it perhaps exponentially small (in the parameters n and k), or is it non-negligible? Of course, if the SCF is nonmanipulable then this probability is zero. Similarly, if the SCF is "close" to being nonmanipulable in some sense, then this probability can be small. We say that a SCF f is ε-far from the family of nonmanipulable functions, if one must change the outcome of f on at least an ε-fraction of the ranking profiles in order to transform f into a nonmanipulable function. Friedgut et al. conjectured that if $k \geq 3$ and the SCF f is ε-far from the family of nonmanipulable functions, then the probability of a ranking profile being manipulable is bounded from below by a polynomial in $1/n$, $1/k$, and ε. Moreover, they conjectured that a random manipulation will succeed with nonnegligible probability, suggesting that manipulation by computational agents in this setting is easy.

Friedgut et al. proved their conjecture in the case of $k = 3$ alternatives, showing a lower bound of $C\varepsilon^6/n$ in the general setting, and $C'\varepsilon^2/n$ in the case when the SCF is *neutral* (commutes with changes made to the names of the alternatives), where C, C' are constants. Note that this result does not have any computational consequences, since when there are only $k = 3$ alternatives, a computational agent may easily try all possible permutations of the alternatives to find a manipulation (if one exists). Several follow-up works have since extended this result. First, Xia and Conitzer [75] used the proof technique of Friedgut et al. to extend their result to a constant number of alternatives, assuming several additional technical assumptions. However, this still does not have any computational consequences, since the result holds only for a constant number of alternatives. Dobzinski and Procaccia [23] proved the conjecture in the case of two voters under the assumption that the SCF is Pareto optimal. Finally, the latest work is due to Isaksson, Kindler and Mossel [40, 41], who proved the conjecture in the case of $k \geq 4$ alternatives with only the added assumption of neutrality. Moreover, they showed that a random manipulation which replaces four adjacent alternatives in the preference order of the manipulating voter by a random permutation of them succeeds with nonnegligible probability. Since this result is valid for any number of ($k \geq 4$) alternatives, it does have computational consequences, implying that for neutral SCFs, manipulation by computational agents is easy on average.

In this chapter we remove the neutrality condition of Isaksson et al. and resolve the conjecture of Friedgut et al.: if $k \geq 3$ and the SCF f is ε-far from the family of nonmanipulable functions, then the probability of a ranking profile being manipulable is bounded from below by a polynomial in $1/n$, $1/k$, and ε. We continue by first presenting our results, then discussing their implications, and finally we conclude this section by commenting on the techniques used in the proof.

2.1.1 Basic setup

Recall that our basic setup consists of n voters electing a winner among k alternatives via a SCF $f : S_k^n \to [k]$. We now define manipulability in more detail:

Definition 2.1 (Manipulation points). *Let $\sigma \in S_k^n$ be a ranking profile. Write $a \overset{\sigma_i}{>} b$ to denote that alternative a is preferred over b by voter i. A SCF $f\colon S_k^n \to [k]$ is* manipulable *at the ranking profile $\sigma \in S_k^n$ if there exists a $\sigma' \in S_k^n$ and an $i \in [n]$ such that σ and σ' only differ in the i^{th} coordinate and*

$$f(\sigma') \overset{\sigma_i}{>} f(\sigma).$$

In this case we also say that σ is a manipulation point *of f, and that (σ, σ') is* a manipulation pair *for f. We say that f is* manipulable *if it is manipulable at some point σ. We also say that σ is an r-*manipulation point *of f if f has a manipulation pair (σ, σ') such that σ' is obtained from σ by permuting (at most) r adjacent alternatives in one of the coordinates of σ. (We allow $r > k$—any manipulation point is an r-manipulation point for $r > k$.)*

Let $M(f)$ denote the set of manipulation points of the SCF f, and for a given r, let $M_r(f)$ denote the set of r-manipulation points of f. When the SCF is obvious from the context, we write simply M and M_r.

Gibbard and Satterthwaite proved the following theorem.

Theorem 2.1 (Gibbard-Satterthwaite [32, 67]). *Any SCF $f\colon S_k^n \to [k]$ which takes at least three values and is not a dictator (i.e., not a function of only one voter) is manipulable.*

This theorem is tight in the sense that *monotone* SCFs which are dictators or only have two possible outcomes are indeed nonmanipulable. A function is nonmonotone, and clearly manipulable, if for some ranking profile a voter can change the outcome from, say, a to b by moving a ahead of b in her preference. It is useful to introduce a refined notion of a dictator before defining the set of nonmanipulable SCFs.

Definition 2.2 (Dictator on a subset). *For a subset of alternatives $H \subseteq [k]$, let top_H be the SCF on one voter whose output is always the top ranked alternative among those in H.*

Definition 2.3 (Nonmanipulable SCFs). *We denote by* $\mathrm{NONMANIP} \equiv \mathrm{NONMANIP}(n,k)$ *the set of nonmanipulable SCFs, which is the following:*

$$\begin{aligned}&\mathrm{NONMANIP}(n,k)\\ &\qquad\qquad = \{f : S_k^n \to [k] \mid f(\sigma) = \mathrm{top}_H(\sigma_i) \ \textit{for some } i \in [n], H \subseteq [k], H \neq \emptyset\}\\ &\qquad\qquad \bigcup \{f : S_k^n \to [k] \mid f \textit{ is a monotone function taking on exactly two values}\}.\end{aligned}$$

When the parameters n and k are obvious from the context, we omit them.

Another important class of functions, which is larger than NONMANIP, but which has a simpler description, is the following.

Definition 2.4. *Define, for parameters n and k that remain implicit (when used the parameters will be obvious from the context):*

$$\overline{\text{NONMANIP}} = \left\{ f \colon S_k^n \to [k] \mid f \text{ only depends on one coordinate or takes at most two values} \right\}.$$

The notation should be thought of as "closure" rather than "complement". We remark that in [40, 41] the set $\overline{\text{NONMANIP}}$ is denoted by NONMANIP—but these two sets of functions should not be confused.

As discussed previously, our goal is to study manipulability from a quantitative viewpoint, and in order to do so we need to define the distance between SCFs.

Definition 2.5 (Distance between SCFs)**.** *Define the distance $\mathbf{D}(f,g)$ between two SCFs $f, g \colon S_k^n \to [k]$ as the fraction of inputs on which they differ. In other words, $\mathbf{D}(f,g) = \mathbb{P}(f(\sigma) \neq g(\sigma))$, where $\sigma \in S_k^n$ is uniformly selected. For a class G of SCFs, we write $\mathbf{D}(f,G) = \min_{g \in G} \mathbf{D}(f,g)$.*

The concepts of anonymity and neutrality of SCFs will be important to us, so we define them here.

Definition 2.6 (Anonymity)**.** *A SCF is* anonymous *if it is invariant under changes made to the names of the voters. More precisely, a SCF $f : S_k^n \to [k]$ is anonymous if for every $\sigma = (\sigma_1, \ldots, \sigma_n) \in S_k^n$ and every $\pi \in S_n$,*

$$f(\sigma_1, \ldots, \sigma_n) = f\left(\sigma_{\pi(1)}, \ldots, \sigma_{\pi(n)}\right).$$

Definition 2.7 (Neutrality)**.** *A SCF is* neutral *if it commutes with changes made to the names of the alternatives. More precisely, a SCF $f : S_k^n \to [k]$ is neutral if for every $\sigma = (\sigma_1, \ldots, \sigma_n) \in S_k^n$ and every $\pi \in S_k$,*

$$f(\pi \circ \sigma_1, \ldots, \pi \circ \sigma_n) = \pi(f(\sigma)).$$

2.1.2 Our main result

Our main result, which resolves the conjecture of Friedgut et al. [29, 30], is the following.

Theorem 2.2. *Suppose that we have $n \geq 1$ voters, $k \geq 3$ alternatives, and a SCF $f : S_k^n \to [k]$ satisfying $\mathbf{D}(f, \overline{\text{NONMANIP}}) \geq \varepsilon$. Then*

$$\mathbb{P}(\sigma \in M(f)) \geq \mathbb{P}(\sigma \in M_4(f)) \geq p\left(\varepsilon, \frac{1}{n}, \frac{1}{k}\right) \tag{2.1}$$

for some polynomial p, where $\sigma \in S_k^n$ is selected uniformly. In particular, we show a lower bound of $\frac{\varepsilon^{15}}{10^{39} n^{67} k^{166}}$.

An immediate consequence is that

$$\mathbb{P}\left((\sigma, \sigma') \text{ is a manipulation pair for } f\right) \geq q\left(\varepsilon, \frac{1}{n}, \frac{1}{k}\right)$$

for some polynomial q, where $\sigma \in S_k^n$ is uniformly selected, and σ' is obtained from σ by uniformly selecting a coordinate $i \in \{1, \ldots, n\}$, uniformly selecting $j \in \{1, \ldots, k-3\}$, and then uniformly randomly permuting the following four adjacent alternatives in σ_i: $\sigma_i(j)$, $\sigma_i(j+1)$, $\sigma_i(j+2)$, and $\sigma_i(j+3)$. In particular, the specific lower bound for $\mathbb{P}(\sigma \in M_4(f))$ implies that we can take $q\left(\varepsilon, \frac{1}{n}, \frac{1}{k}\right) = \frac{\varepsilon^{15}}{10^{41} n^{68} k^{167}}$.

2.1.3 Discussion

Our results cover all previous cases for which a quantitative Gibbard-Satterthwaite theorem has been established. In particular, the main novelty is that neutrality of the SCF is not assumed, and therefore our results hold for nonneutral SCFs as well. This solves the main open problem of Friedgut, Kalai, Keller and Nisan [29, 30], and Isaksson, Kindler and Mossel [40, 41]. The main message of our results is that the approach of masking manipulation behind computational hardness cannot hide manipulations completely even in the nonneutral setting.

Importance of nonneutrality. While neutrality seems like a very natural assumption, there are multiple reasons why removing this assumption is important:

- **Anonymity vs. neutrality.** It is known that there is a conflict between anonymity and neutrality (recall Definitions 2.6 and 2.7). In particular, there are some combinations of n and k when there exists no SCF which is both anonymous and neutral.

  Theorem 2.3. *[55, Chapter 2.4.] There exists a SCF on n voters and k alternatives which is both anonymous and neutral if and only if k cannot be written as the sum of (non-trivial) divisors of n.*

  The difficulty comes from rules governing tie-breaking. Consider the following example: suppose that $n = k = 2$, i.e., we have two voters, voter 1 and voter 2, and two alternatives, a and b. Suppose further (w.l.o.g.) that when voter 1 prefers a over b and voter 2 prefers b over a then the outcome is a. What should the outcome be when voter 1 prefers b over a and voter 2 prefers a over b? By anonymity the outcome should be a for this configuration as well, but by neutrality the outcome should be b.

  Most common voting rules (plurality, Borda count, etc.) break ties in an anonymous way, and therefore they cannot be neutral as well (or can only be neutral for special values of n and k). See Moulin [55, Chapter 2.4.] for more on anonymity and neutrality.

- **Nonneutrality in virtual elections.** As mentioned before, voting manipulation is a serious issue in artificial intelligence and computer science as well, where virtual elections are becoming more and more popular as a tool in preference aggregation (see the survey [27]). For example, consider web (meta-)search engines (see, e.g., Dwork et al. [25]), where one inputs a query and the possible outcomes ("alternatives") are the web pages (with the various search engines acting as "voters"). Here, due to various restrictions, neutrality is not a natural assumption. For example, there can be language-related restrictions: if one searches in English then the top-ranked webpage will also be in English; or safety-related restrictions: if one searches in child-safe mode, then the top-ranked webpage cannot have adult content. These restrictions imply that the appropriate aggregating function cannot be neutral.

- **Nonneutrality in real-life elections.** Although not a common occurrence, there have been cases in real-life elections when a candidate is on the ballot, but is actually ineligible—she cannot win the election no matter what. In such a case the SCF is necessarily nonneutral.

  In a recent set of local elections in Philadelphia there were actually three such occurences [70]: one of the candidates for the one open Municipal Court slot was not a lawyer, which is a prerequisite for someone elected to this position; another judicial candidate received a court order to leave the race; finally, in the race for a district seat in Philadelphia, one of the candidates had announced that he is abandoning his candidacy; yet all three of them remained on the respective ballots.

  A more curious story is that of the New York State Senate elections in 2010, where the name of a dead man appeared on the ballot (he received 828 votes) [38].

A quantitative Gibbard-Satterthwaite theorem for one voter. A major part of the work in proving Theorem 2.2 is devoted to understanding functions of a single voter, essentially proving a quantitative Gibbard-Satterthwaite theorem for one voter. This can be formulated as follows.

Theorem 2.4. *Suppose that $f : S_k \to [k]$ is a SCF on $n = 1$ voter and $k \geq 3$ alternatives which satisfies $\mathbf{D}(f, \text{NONMANIP}) \geq \varepsilon$. Then*

$$\mathbb{P}(\sigma \in M(f)) \geq \mathbb{P}(\sigma \in M_3(f)) \geq p\left(\varepsilon, \frac{1}{k}\right), \tag{2.2}$$

for some polynomial p, where $\sigma \in S_k$ is selected uniformly. In particular, we show a lower bound of $\frac{\varepsilon^3}{10^5 k^{16}}$.

We note that this is a new result, which has not been studied in the literature before.

Dobzinski and Procaccia [23] proved a quantitative Gibbard-Satterthwaite theorem for two voters, assuming that the SCF is *Pareto optimal*, i.e., if all voters rank alternative a above b, then b is not elected. The assumption of Pareto optimality is natural in the context of classical social choice, but it is a very strong assumption in the context of quantitative social choice. For one, it implies that every alternative is elected with probability at least

$1/k^2$. Second, for one voter, there exists a unique Pareto optimal SCF, while the number of nonmanipulable SCFs is exponential in k. The assumption also prevents applying the result of Dobzinski and Procaccia to SCFs obtained from a SCF on many voters when the votes of all voters but two are fixed (since even if the original SCF is Pareto optimal, the restricted function may not be so). In our proof we often deal with such restricted SCFs (where the votes of all but one or two voters are fixed), and this is also what led us to our quantitative Gibbard-Satterthwaite theorem for one voter.

On NONMANIP **versus** $\overline{\text{NONMANIP}}$**.** The quantitative Gibbard-Satterthwaite theorems of Friedgut, Kalai, Keller and Nisan [29, 30], and Isaksson, Kindler and Mossel [40, 41] involve the distance of a SCF from $\overline{\text{NONMANIP}}$. Any SCF that is not in $\overline{\text{NONMANIP}}$ is manipulable (by the Gibbard-Satterthwaite theorem), but as some SCFs in $\overline{\text{NONMANIP}}$ are manipulable as well, ideally a quantitative Gibbard-Satterthwaite theorem would involve the distance of a SCF from the set of (truly) nonmanipulable SCFs, NONMANIP. Theorem 2.2 addresses this concern, as it involves the distance of a SCF from NONMANIP. This is done via the following reduction theorem that implies that whenever one has a quantitative Gibbard-Satterthwaite theorem involving $\mathbf{D}\left(f, \overline{\text{NONMANIP}}\right)$, it can be turned into a quantitative Gibbard-Satterthwaite theorem involving $\mathbf{D}\left(f, \text{NONMANIP}\right)$.

Theorem 2.5. *Suppose that f is a SCF on n voters and $k \geq 3$ alternatives for which $\mathbf{D}\left(f, \overline{\text{NONMANIP}}\right) \leq \alpha$. Then either*

$$\mathbf{D}\left(f, \text{NONMANIP}\right) < 100 n^4 k^8 \alpha^{1/3} \tag{2.3}$$

or

$$\mathbb{P}\left(\sigma \in M\left(f\right)\right) \geq \mathbb{P}\left(\sigma \in M_3\left(f\right)\right) \geq \alpha. \tag{2.4}$$

The proof of this result also uses Theorem 2.4, our quantitative Gibbard-Satterthwaite theorem for one voter.

A note on our quantitative bounds. The lower bounds on the probability of manipulation derived in Theorems 2.2, 2.4, and various results along the way, are not tight. Moreover, we do not believe that our techniques allow us to obtain tight bounds. Consequently, we did not try to optimize these bounds, but rather focused on the qualitative result: obtaining polynomial bounds.

2.1.4 Proof techniques and ideas

In our proof we combine ideas from both Friedgut, Kalai, Keller and Nisan [29, 30] and Isaksson, Kindler and Mossel [40, 41]. In addition, we use a reverse hypercontractivity lemma that was applied in the proof of a quantitative version of Arrow's theorem by Mossel [50]. (Reverse hypercontractivity was originally proved and discussed by Borell [12], and was first applied by Mossel, O'Donnell, Regev, Steif, and Sudakov [54].) Our techniques most closely resemble those of Isaksson et al. [40, 41]; here the authors used a variant of the canonical path method to show the existence of a large interface where three bodies touch. Our goal is also to come to this conclusion, but we do so via different methods.

We first present our techniques that achieve a lower bound for the probability of manipulation that involves factors of $\frac{1}{k!}$ (see Theorem 2.9 in Section 2.3), and then describe how a refined approach leads to a lower bound which has inverse polynomial dependence on k (see Theorem 2.35 in Section 2.7).

Rankings graph and applying the original Gibbard-Satterthwaite theorem. As in Isaksson et al. [40, 41], think of the graph $G = (V, E)$ having vertex set $V = S_k^n$, the set of all ranking profiles, and let $(\sigma, \sigma') \in E$ if and only if σ and σ' differ in exactly one coordinate. The SCF $f : S_k^n \to [k]$ naturally partitions V into k subsets. Since every manipulation point must be on the boundary between two such subsets, we are interested in the size of such boundaries.

For two alternatives a and b, and voter i, denote by $B_i^{a,b}$ the boundary between $f^{-1}(a)$ and $f^{-1}(b)$ in voter i. A lemma from Isaksson et al. [40, 41] tells us that at least two of the boundaries are large; in the following assume that these are $B_1^{a,b}$ and $B_2^{a,c}$. Now if a ranking profile σ lies on *both* of these boundaries, then applying the original Gibbard-Satterthwaite theorem to the restricted SCF on two voters where we fix all coordinates of σ except the first two, we get that there must exist a manipulation point which agrees with σ in all but the first two coordinates. Consequently, if we can show that the *intersection* of the boundaries $B_1^{a,b}$ and $B_2^{a,c}$ is large, then we have many manipulation points.

Fibers and reverse hypercontractivity. In order to have more "control" over what is happening at the boundaries, we partition the graph further—this idea is due to Friedgut et al. [29, 30]. Given a ranking profile σ and two alternatives a and b, σ induces a *vector of preferences* $x^{a,b}(\sigma) \in \{-1, 1\}^n$ between a and b. For a vector $z^{a,b} \in \{-1, 1\}^n$ we define the *fiber with respect to preferences between a and b*, denoted by $F\left(z^{a,b}\right)$, to be the set of ranking profiles for which the vector of preferences between a and b is $z^{a,b}$. We can then partition the vertex set V into such fibers, and work inside each fiber separately. Working inside a specific fiber is advantageous, because it gives us the extra knowledge of the vector of preferences between a and b.

We distinguish two types of fibers: large and small. We say that a fiber w.r.t. preferences between a and b is *large* if almost all of the ranking profiles in this fiber lie on the boundary $B_1^{a,b}$, and *small* otherwise. Now since the boundary $B_1^{a,b}$ is large, either there is big mass on the large fibers w.r.t. preferences between a and b or there is big mass on the small fibers. This holds analogously for the boundary $B_2^{a,c}$ and fibers w.r.t. preferences between a and c.

Consider the case when there is big mass on the large fibers of both $B_1^{a,b}$ and $B_2^{a,c}$. Notice that for a ranking profile σ, being in a fiber w.r.t. preferences between a and b only depends on the vector of preferences between a and b, $x^{a,b}(\sigma)$, which is a uniform bit vector. Similarly, being in a fiber w.r.t. preferences between a and c only depends on $x^{a,c}(\sigma)$. Moreover, we know the exact correlation between the coordinates of $x^{a,b}(\sigma)$ and $x^{a,c}(\sigma)$, and it is in exactly this setting where *reverse hypercontractivity* applies (see Lemma 2.8 for a precise statement), and shows that the *intersection* of the large fibers of $B_1^{a,b}$ and $B_2^{a,c}$ is also large. Finally, by the definition of a large fiber it follows that the intersection of the *boundaries* $B_1^{a,b}$ and $B_2^{a,c}$ is large as well, and we can finish the argument using the Gibbard-Satterthwaite theorem as above.

To deal with the case when there is big mass on the *small* fibers of $B_1^{a,b}$ we use various isoperimetric techniques, including the canonical path method developed for this problem by Isaksson et al. [40, 41]. In particular, we use the fact that for a small fiber for $B_1^{a,b}$, the size of the boundary of $B_1^{a,b}$ in the small fiber is comparable to the size of $B_1^{a,b}$ in the small fiber itself, up to polynomial factors.

A refined geometry. Using this approach with the rankings graph above, our bound includes $\frac{1}{k!}$ factors (see Theorem 2.9 in Section 2.3). In order to obtain inverse polynomial dependence on k (as in Theorem 2.35 in Section 2.4), we use a refined approach, similar to that in Isaksson et al. [40, 41]. Instead of the rankings graph outlined above, we use an underlying graph with a different edge structure: $(\sigma, \sigma') \in E$ if and only if σ and σ' differ in exactly one coordinate, and in this coordinate they differ by a single adjacent transposition. In order to prove the refined result, we need to show that the geometric and combinatorial quantities such as boundaries and manipulation points are roughly the same in the refined graph as in the original rankings graph. In particular, this is where we need to analyze carefully functions of one voter, and ultimately prove a quantitative Gibbard-Satterthwaite theorem for one voter.

2.1.5 Organization of the chapter

The rest of the chapter is outlined as follows. We introduce necessary preliminaries (definitions and previous technical results) in Section 2.2. We then proceed by proving Theorem 2.9 in Section 2.3, which is weaker than Theorem 2.2 in two aspects: first, the condition $\mathbf{D}(f, \mathrm{NONMANIP}) \geq \varepsilon$ is replaced with the stronger condition $\mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \geq \varepsilon$, and second, we allow factors of $\frac{1}{k!}$ in our lower bounds for $\mathbb{P}(\sigma \in M(f))$. We continue by explaining the necessary modifications we have to make in the refined setting to get inverse polynomial dependence on k in Section 2.4. Additional preliminaries necessary for the proofs of Theorems 2.35, 2.4 and 2.5 are in Section 2.5, while the remaining sections contain the proofs of these theorems. We prove Theorem 2.4 in Section 2.6, Theorem 2.35 in Section 2.7, and Theorem 2.5 and Theorem 2.2 in Section 2.8. Finally, we conclude with some open problems in Section 2.9.

2.2 Preliminaries: definitions and previous technical results

2.2.1 Boundaries and influences

For a general graph $G = (V, E)$, and a subset of the vertices $A \subseteq V$, we define the *edge boundary* of A as

$$\partial_e (A) = \{(u, v) \in E : u \in A, v \notin A\}.$$

We also define the *boundary* (or vertex boundary) of a subset of the vertices $A \subseteq V$ to be the set of vertices in A which have a neighbor that is not in A:

$$\partial (A) = \{u \in A : \text{ there exists } v \notin A \text{ such that } (u, v) \in E\}.$$

If $u \in \partial (A)$, we also say that u is *on* the edge boundary of A.

As discussed in Section 2.1.4, we can view the ranking profiles (which are elements of S_k^n) as vertices of a graph—the rankings graph—where two vertices are connected by an edge if they differ in exactly one coordinate. The SCF f naturally partitions the vertices of this graph into k subsets, depending on the value of f at a given vertex. Clearly, a manipulation point can only be on the edge boundary of such a subset, and so it is important to study these boundaries. In this spirit, we introduce the following definitions.

Definition 2.8 (Boundaries). *For a given SCF f and a given alternative $a \in [k]$, we define*

$$W^a (f) = \{\sigma \in S_k^n : f (\sigma) = a\},$$

the set of ranking profiles where the outcome of the vote is a. The edge boundary of this set is denoted by $B^a (f)$: $B^a (f) = \partial_e (W^a (f))$. This boundary can be partitioned: we say that the edge boundary of $W^a (f)$ in the direction of the i^{th} coordinate is

$$B_i^a (f) = \{(\sigma, \sigma') \in B^a (f) : \sigma_i \neq \sigma_i'\}.$$

The boundary $B^a (f)$ can be therefore written as $B^a (f) = \cup_{i=1}^n B_i^a (f)$. We can also define the boundary between two alternatives a and b in the direction of the i^{th} coordinate:

$$B_i^{a,b} (f) = \{(\sigma, \sigma') \in B_i^a (f) : f (\sigma') = b\}.$$

We also say that $\sigma \in B_i^a (f)$ is on *the boundary $B_i^{a,b} (f)$ if there exists σ' such that $(\sigma, \sigma') \in B_i^{a,b} (f)$.*

Definition 2.9 (Influences). *We define the* influence *of the i^{th} coordinate on f as*

$$\text{Inf}_i (f) = \mathbb{P} \left(f (\sigma) \neq f \left(\sigma^{(i)} \right) \right) = \mathbb{P} \left(\left(\sigma, \sigma^{(i)} \right) \in \cup_{a=1}^k B_i^a (f) \right),$$

where σ is uniform on S_k^n and $\sigma^{(i)}$ is obtained from σ by rerandomizing the i^{th} coordinate. Similarly, we define the influence of the i^{th} coordinate with respect to a single alternative $a \in [k]$ or a pair of alternatives $a, b \in [k]$ as

$$\mathrm{Inf}_i^a(f) = \mathbb{P}\left(f(\sigma) = a, f\left(\sigma^{(i)}\right) \neq a\right) = \mathbb{P}\left(\left(\sigma, \sigma^{(i)}\right) \in B_i^a(f)\right),$$

and

$$\mathrm{Inf}_i^{a,b}(f) = \mathbb{P}\left(f(\sigma) = a, f\left(\sigma^{(i)}\right) = b\right) = \mathbb{P}\left(\left(\sigma, \sigma^{(i)}\right) \in B_i^{a,b}(f)\right),$$

respectively.

Clearly

$$\mathrm{Inf}_i(f) = \sum_{a=1}^{k} \mathrm{Inf}_i^a(f) = \sum_{a,b \in [k]: a \neq b} \mathrm{Inf}_i^{a,b}(f).$$

Most of the time the specific SCF f will be clear from the context, in which case we omit the dependence on f, and write simply $B^a \equiv B^a(f)$, $B_i^a \equiv B_i^a(f)$, etc.

2.2.2 Large boundaries

The following lemma from Isaksson, Kindler and Mossel [40, 41, Lemma 3.1.] shows that there are some boundaries which are large (in the sense that they are only inverse polynomially small in n, k and ε^{-1})—our task is then to find many manipulation points on these boundaries.

Lemma 2.6. *Fix $k \geq 3$ and $f \colon S_k^n \to [k]$ satisfying $\mathbf{D}(f, \overline{\mathrm{NONMANIP}}) \geq \varepsilon$. Then there exist distinct $i, j \in [n]$ and $\{a, b\}, \{c, d\} \subseteq [k]$ such that $c \notin \{a, b\}$ and*

$$\mathrm{Inf}_i^{a,b}(f) \geq \frac{2\varepsilon}{nk^2(k-1)} \quad \text{and} \quad \mathrm{Inf}_j^{c,d}(f) \geq \frac{2\varepsilon}{nk^2(k-1)}. \tag{2.5}$$

2.2.3 General isoperimetric results

Our rankings graph is the Cartesian product of n complete graphs on $k!$ vertices. We therefore use isoperimetric results on products of graphs—see [37] for an overview. In particular, the edge-isoperimetric problem on the product of complete graphs was originally solved by Lindsey [47], implying the following result.

Corollary 2.7. *If $A \subseteq K_\ell \times \cdots \times K_\ell$ (n copies of the complete graph K_ℓ) and $|A| \leq \left(1 - \frac{1}{\ell}\right)\ell^n$, then $|\partial_e(A)| \geq |A|$.*

2.2.4 Fibers

In our proof we need to partition the graph even further—this idea is due to Friedgut, Kalai, Keller, and Nisan [29, 30].

Definition 2.10. *For a ranking profile $\sigma \in S_k^n$ define the vector*

$$x^{a,b} \equiv x^{a,b}(\sigma) = \left(x_1^{a,b}(\sigma), \dots, x_n^{a,b}(\sigma)\right)$$

of preferences between a and b, where $x_i^{a,b}(\sigma) = 1$ if $a \overset{\sigma_i}{>} b$ and $x_i^{a,b}(\sigma) = -1$ otherwise.

Definition 2.11 (Fibers)**.** *For a pair of alternatives $a, b \in [k]$ and a vector $z^{a,b} \in \{-1,1\}^n$, write*

$$F\left(z^{a,b}\right) := \left\{\sigma : x^{a,b}(\sigma) = z^{a,b}\right\}.$$

We call the $F\left(z^{a,b}\right)$ fibers *with respect to preferences between a and b.*

So for any pair of alternatives a, b, we can partition the ranking profiles according to its fibers:

$$S_k^n = \bigcup_{z^{a,b} \in \{-1,1\}^n} F\left(z^{a,b}\right).$$

Given a SCF f, for any pair of alternatives $a, b \in [k]$ and $i \in [n]$, we can also partition the boundary $B_i^{a,b}(f)$ according to its fibers. There are multiple, slightly different ways of doing this, but for our purposes the following definition is most useful. Define

$$B_i\left(z^{a,b}\right) := \left\{\sigma \in F\left(z^{a,b}\right) : f(\sigma) = a, \text{ and there exists } \sigma' \text{ s.t. } (\sigma, \sigma') \in B_i^{a,b}\right\},$$

where we omit the dependence of $B_i\left(z^{a,b}\right)$ on f. So $B_i\left(z^{a,b}\right) \subseteq F\left(z^{a,b}\right)$ is the set of vertices on the given fiber for which the outcome is a and which lies on the boundary between a and b in direction i. We call the sets of the form $B_i\left(z^{a,b}\right)$ *fibers for the boundary* $B_i^{a,b}$ (again omitting the dependence on f of both sets).

We now distinguish between small and large fibers for the boundary $B_i^{a,b}$.

Definition 2.12 (Small and large fibers)**.** *We say that the fiber $B_i\left(z^{a,b}\right)$ is* large *if*

$$\mathbb{P}\left(\sigma \in B_i\left(z^{a,b}\right) \,\middle|\, \sigma \in F\left(z^{a,b}\right)\right) \geq 1 - \frac{\varepsilon^3}{4n^3k^9}, \tag{2.6}$$

and small *otherwise.*

We denote by $\mathrm{Lg}\left(B_i^{a,b}\right)$ *the union of large fibers for the boundary $B_i^{a,b}$, i.e.,*

$$\mathrm{Lg}\left(B_i^{a,b}\right) := \left\{\sigma : B_i\left(x^{a,b}(\sigma)\right) \textit{ is a large fiber, and } \sigma \in B_i\left(x^{a,b}(\sigma)\right)\right\}$$

and similarly, we denote by $\mathrm{Sm}\left(B_i^{a,b}\right)$ *the union of small fibers.*

We remark that what is important is that the fraction appearing on the right hand side of (2.6) is a polynomial of $\frac{1}{n}$, $\frac{1}{k}$ and ε—the specific polynomial in this definition is the end result of the computation in the proof.

Finally, for a voter i and a pair of alternatives $a, b \in [k]$, we define

$$F_i^{a,b} := \left\{\sigma : B_i\left(x^{a,b}\left(\sigma\right)\right) \text{ is a large fiber}\right\}.$$

So this means that

$$\mathbb{P}\left(\sigma \in \cup_{z^{a,b}} B_i\left(z^{a,b}\right) \middle| \sigma \in F_i^{a,b}\right) \geq 1 - \frac{\varepsilon^3}{4n^3k^9}. \tag{2.7}$$

2.2.5 Boundaries of boundaries

Finally, we also look at boundaries of boundaries. In particular, for a given vector $z^{a,b}$ of preferences between a and b, we can think of the fiber $F\left(z^{a,b}\right)$ as a subgraph of the original rankings graph. When we write $\partial\left(B_i\left(z^{a,b}\right)\right)$, we mean the boundary of $B_i\left(z^{a,b}\right)$ in the subgraph of the rankings graph induced by the fiber $F\left(z^{a,b}\right)$. That is,

$$\begin{aligned}&\partial\left(B_i\left(z^{a,b}\right)\right)\\ &\quad = \{\sigma \in B_i\left(z^{a,b}\right) : \exists\, \pi \in F\left(z^{a,b}\right) \setminus B_i\left(z^{a,b}\right) \text{ s.t. } \sigma \text{ and } \pi \text{ differ in exactly one coord.}\}.\end{aligned}$$

2.2.6 Reverse hypercontractivity

We use the following lemma about reverse hypercontractivity from Mossel [50].

Lemma 2.8. *Suppose that the vectors x and y are distributed uniformly in $\{-1,1\}^n$ and that $\{(x_i, y_i)\}_{i=1}^n$ are independent. Assume further that $\left|\mathbb{E}\left(x_i y_i\right)\right| \leq \rho$. Let $B_1, B_2 \subset \{-1,1\}^n$ be two sets and assume that*

$$\mathbb{P}\left(B_1\right) \geq e^{-\alpha^2}, \qquad \mathbb{P}\left(B_2\right) \geq e^{-\beta^2}.$$

Then

$$\mathbb{P}\left(x \in B_1, y \in B_2\right) \geq \exp\left(-\frac{\alpha^2 + \beta^2 + 2\rho\alpha\beta}{1-\rho^2}\right).$$

In particular, if $\mathbb{P}\left(B_1\right) \geq \varepsilon$ and $\mathbb{P}\left(B_2\right) \geq \varepsilon$, then

$$\mathbb{P}\left(x \in B_1, y \in B_2\right) \geq \varepsilon^{\frac{2}{1-\rho}}.$$

2.2.7 Dictators and miscellaneous definitions

For a ranking profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ we sometimes write σ_{-i} for the collection of all coordinates except the i^{th} coordinate, i.e., $\sigma = (\sigma_i, \sigma_{-i})$. Furthermore, we sometimes distinguish two coordinates, e.g., we write $\sigma = \left(\sigma_1, \sigma_i, \sigma_{-\{1,i\}}\right)$.

Definition 2.13 (Induced SCF on one coordinate)**.** *Let $f_{\sigma_{-i}}$ denote the SCF on one voter induced by f by fixing all voter preferences except the i^{th} one according to σ_{-i}. That is,*

$$f_{\sigma_{-i}}(\cdot) := f(\cdot, \sigma_{-i}).$$

Recall Definition 2.2 of a dictator on a subset.

Definition 2.14 (Ranking profiles giving dictators on a subset)**.** *For a coordinate i and a subset of alternatives $H \subseteq [k]$, define*

$$D_i^H := \left\{\sigma_{-i} : f_{\sigma_{-i}}(\cdot) \equiv \mathrm{top}_H(\cdot)\right\}.$$

Also, for a pair of alternatives a and b, define

$$D_i(a,b) := \bigcup_{H:\{a,b\}\subseteq H, |H|\geq 3} D_i^H.$$

2.3 Inverse polynomial manipulability for a fixed number of alternatives

Our goal in this section is to demonstrate the proof techniques described in Section 2.1.4. We prove here the following theorem (Theorem 2.9 below), which is weaker than our main theorem, Theorem 2.2, in two aspects: first, the condition $\mathbf{D}(f, \mathrm{NONMANIP}) \geq \varepsilon$ is replaced with the stronger condition $\mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \geq \varepsilon$, and second, we allow factors of $\frac{1}{k!}$ in our lower bounds for $\mathbb{P}(\sigma \in M(f))$. The advantage is that the proof of this statement is relatively simpler. We move on to getting a lower bound with polynomial dependence on k in the following sections, and finally we replace the condition $\mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \geq \varepsilon$ with $\mathbf{D}(f, \mathrm{NONMANIP}) \geq \varepsilon$ in Section 2.8.

Theorem 2.9. *Suppose that we have $n \geq 2$ voters, $k \geq 3$ alternatives, and a SCF $f : S_k^n \to [k]$ satisfying $\mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \geq \varepsilon$. Then*

$$\mathbb{P}(\sigma \in M(f)) \geq p\left(\varepsilon, \frac{1}{n}, \frac{1}{k!}\right), \tag{2.8}$$

for some polynomial p, where $\sigma \in S_k^n$ is selected uniformly. In particular, we show a lower bound of $\frac{\varepsilon^5}{4n^7k^{12}(k!)^4}$.

An immediate consequence is that

$$\mathbb{P}((\sigma,\sigma') \text{ is a manipulation pair for } f) \geq q\left(\varepsilon, \frac{1}{n}, \frac{1}{k!}\right),$$

for some polynomial q, where $\sigma \in S_k^n$ is selected uniformly, and σ' is obtained from σ by uniformly selecting a coordinate $i \in \{1,\ldots,n\}$ and resetting the i^{th} coordinate to a random preference. In particular, the specific lower bound for $\mathbb{P}(\sigma \in M(f))$ implies that we can take $q\left(\varepsilon, \frac{1}{n}, \frac{1}{k!}\right) = \frac{\varepsilon^5}{4n^8k^{12}(k!)^5}$.

First we provide an overview of the proof of Theorem 2.9 in Section 2.3.1. In this overview we use adjectives such as "big", and "not too small" to describe probabilities—here these are all synonymous with "has probability at least an inverse polynomial of n, $k!$, and ε^{-1}".

2.3.1 Overview of proof

The tactic in proving Theorem 2.9 is roughly the following:

- By Lemma 2.6, we know that there are at least two boundaries which are big. W.l.o.g. we can assume that these are either $B_1^{a,b}$ and $B_2^{a,c}$, or $B_1^{a,b}$ and $B_2^{c,d}$ with $\{a, b\} \cap \{c, d\} = \emptyset$. Our proof works in both cases, but we continue the outline of the proof assuming the former case—this is the more interesting case, since the latter case has been solved already by Isaksson et al. [40, 41].

- We partition $B_1^{a,b}$ according to its fibers based on the preferences between a and b of the n voters, just like as described in Section 2.2. Similarly for $B_2^{a,c}$ and preferences between a and c.

- As in Section 2.2, we can distinguish small and large fibers for these two boundaries. Now since $B_1^{a,b}$ is big, either the mass of small fibers, or the mass of large fibers is big. Similarly for $B_2^{a,c}$.

- Suppose first that there is big mass on large fibers in both $B_1^{a,b}$ and $B_2^{a,c}$. In this case the probability of our random ranking σ being in $F_1^{a,b}$ is big, and similarly for $F_2^{a,c}$. Being in $F_1^{a,b}$ only depends on the vector $x^{a,b}(\sigma)$ of preferences between a and b, and similarly being in $F_2^{a,c}$ only depends on the vector $x^{a,c}(\sigma)$ of preferences between a and c. We know the correlation between $x^{a,b}(\sigma)$ and $x^{a,c}(\sigma)$ and hence we can apply reverse hypercontractivity (Lemma 2.8), which tells us that the probability that σ lies in both $F_1^{a,b}$ and $F_2^{a,c}$ is big as well. If $\sigma \in F_1^{a,b}$, then voter 1 is pivotal between alternatives a and b with big probability, and similarly if $\sigma \in F_2^{a,c}$, then voter 2 is pivotal between alternatives a and c with big probability. So now we have that the probability that both voter 1 is pivotal between a and b and voter 2 is pivotal between a and c is big, and in this case the Gibbard-Satterthwaite theorem tells us that there is a manipulation point which agrees with this ranking profile in all except for perhaps the first two coordinates. So there are many manipulation points.

- Now suppose that the mass of small fibers in $B_1^{a,b}$ is big. By isoperimetric theory, the size of the boundary of every small fiber is comparable (same order up to $\text{poly}^{-1}(\varepsilon^{-1}, n, k!)$ factors) to the size of the small fiber. Consequently, the total size of the boundaries of small fibers is comparable to the total size of small fibers, which in this case has to be big.

  We then distinguish two cases: either we are on the boundary of a small fiber in the first coordinate, or some other coordinate. If σ is on the boundary of a small fiber in some coordinate $j \neq 1$, then the Gibbard-Satterthwaite theorem tells us that there is a manipulation point which agrees with σ in all coordinates except perhaps in coordinates 1 and j. If our ranking profile σ is on the boundary of a small fiber in the first coordinate, then either there exists a manipulation point which agrees with σ in all coordinates except perhaps the first, or the SCF on one voter that we obtain from f by fixing the votes of voters 2 through n to be σ_{-1} must be a dictator on some subset of the alternatives. So either we get sufficiently many manipulation points this way, or for many votes of voters 2 through n, the restricted SCF obtained from f by fixing these votes is a dictator on coordinate 1 on some subset of the alternatives.

Finally, to deal with dictators on the first coordinate, we look at the boundary of the dictators. Since $\mathbf{D}\left(f, \overline{\text{NONMANIP}}\right) \geq \varepsilon$, the boundary is big, and we can also show that there is a manipulation point near every boundary point.

- If the mass of small fibers in $B_2^{a,c}$ is big, then we can do the same thing for this boundary.

2.3.2 Division into cases

For the remainder of Section 2.3, let us fix the number of voters $n \geq 2$, the number of alternatives $k \geq 3$, and the SCF f, which satisfies $\mathbf{D}\left(f, \overline{\text{NONMANIP}}\right) \geq \varepsilon$. Accordingly, we typically omit the dependence of various sets (e.g., boundaries between two alternatives) on f.

Our starting point is Lemma 2.6. W.l.o.g. we may assume that the two boundaries that the lemma gives us have $i = 1$ and $j = 2$, so the lemma tells us that

$$\mathbb{P}\left(\left(\sigma, \sigma^{(1)}\right) \in B_1^{a,b}\right) \geq \frac{2\varepsilon}{nk^3},$$

where σ is uniform on the ranking profiles, and $\sigma^{(1)}$ is obtained by rerandomizing the first coordinate. This also means that

$$\mathbb{P}\left(\sigma \in \cup_{z^{a,b}} B_1\left(z^{a,b}\right)\right) \geq \frac{2\varepsilon}{nk^3},$$

and similar inequalities hold for the boundary $B_2^{c,d}$. The following lemma is an immediate corollary.

Lemma 2.10. *Either*

$$\mathbb{P}\left(\sigma \in \text{Sm}\left(B_1^{a,b}\right)\right) \geq \frac{\varepsilon}{nk^3} \tag{2.9}$$

or

$$\mathbb{P}\left(\sigma \in \text{Lg}\left(B_1^{a,b}\right)\right) \geq \frac{\varepsilon}{nk^3}, \tag{2.10}$$

and the same can be said for the boundary $B_2^{c,d}$.

We distinguish cases based upon this: either (2.9) holds, or (2.9) holds for the boundary $B_2^{c,d}$, or (2.10) holds for both boundaries. We only need one boundary for the small fiber case, and we need both boundaries only in the large fiber case. So in the large fiber case we must differentiate between two cases: whether $d \in \{a, b\}$ or $d \notin \{a, b\}$. First of all, in the $d \notin \{a, b\}$ case the problem of finding a manipulation point with not too small (i.e., inverse polynomial in n, $k!$ and ε^{-1}) probability has already been solved in [40, 41]. But moreover, we will see that if $d \notin \{a, b\}$ then the large fiber case cannot occur—so this method of proof works as well.

In the rest of the section we first deal with the large fiber case, and then with the small fiber case.

2.3.3 Big mass on large fibers

We now deal with the case when

$$\mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_1^{a,b}\right)\right) \geq \frac{\varepsilon}{nk^3} \tag{2.11}$$

and also

$$\mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_2^{c,d}\right)\right) \geq \frac{\varepsilon}{nk^3}. \tag{2.12}$$

As mentioned before, we must differentiate between two cases: whether $d \in \{a, b\}$ or $d \notin \{a, b\}$.

2.3.3.1 Case 1

Suppose that $d \in \{a, b\}$, in which case we may assume w.l.o.g. that $d = a$.

Lemma 2.11. *If*

$$\mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_1^{a,b}\right)\right) \geq \frac{\varepsilon}{nk^3} \qquad \textit{and} \qquad \mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_2^{a,c}\right)\right) \geq \frac{\varepsilon}{nk^3}, \tag{2.13}$$

then

$$\mathbb{P}\left(\sigma \in M\right) \geq \frac{\varepsilon^3}{2n^3k^9\left(k!\right)^2}. \tag{2.14}$$

Proof. By (2.13) we have that

$$\mathbb{P}\left(\sigma \in F_1^{a,b}\right) \geq \frac{\varepsilon}{nk^3} \qquad \text{and} \qquad \mathbb{P}\left(\sigma \in F_2^{a,c}\right) \geq \frac{\varepsilon}{nk^3}.$$

We know that $\left|\mathbb{E}\left(x_i^{a,b}\left(\sigma\right)x_i^{a,c}\left(\sigma\right)\right)\right| = 1/3$, and so by reverse hypercontractivity (Lemma 2.8) we have that

$$\mathbb{P}\left(\sigma \in F_1^{a,b} \cap F_2^{a,c}\right) \geq \frac{\varepsilon^3}{n^3k^9}. \tag{2.15}$$

Recall that we say that σ is *on* the boundary $B_1^{a,b}$ if there exists σ' such that $(\sigma, \sigma') \in B_1^{a,b}$. If $\sigma \in F_1^{a,b}$, then with big probability σ is on the boundary $B_1^{a,b}$, and if $\sigma \in F_2^{a,c}$, then with big probability σ is on the boundary $B_2^{a,c}$. Using this and (2.15) we can show that the probability of σ lying on both the boundary $B_1^{a,b}$ and the boundary $B_2^{a,c}$ is big. Then we are done, because if σ lies on both $B_1^{a,b}$ and $B_2^{a,c}$, then by the Gibbard-Satterthwaite theorem there is a $\hat{\sigma}$ which agrees with σ on the last $n-2$ coordinates, and which is a manipulation point. Furthermore, there can be at most $(k!)^2$ ranking profiles that give the same manipulation point. Let us do the computation:

$$\begin{aligned}\mathbb{P}\left(\sigma \text{ on } B_1^{a,b}, \sigma \text{ on } B_2^{a,c}\right) &\geq \mathbb{P}\left(\sigma \text{ on } B_1^{a,b}, \sigma \text{ on } B_2^{a,c}, \sigma \in F_1^{a,b} \cap F_2^{a,c}\right)\\ &\geq \mathbb{P}\left(\sigma \in F_1^{a,b} \cap F_2^{a,c}\right) - \mathbb{P}\left(\sigma \in F_1^{a,b} \cap F_2^{a,c}, \sigma \text{ not on } B_1^{a,b}\right)\\ &\quad - \mathbb{P}\left(\sigma \in F_1^{a,b} \cap F_2^{a,c}, \sigma \text{ not on } B_2^{a,c}\right).\end{aligned}$$

The first term is bounded below via (2.15), while the other two terms can be bounded using (2.7):

$$\begin{aligned}\mathbb{P}\left(\sigma \in F_1^{a,b} \cap F_2^{a,c}, \sigma \text{ not on } B_1^{a,b}\right) &\leq \mathbb{P}\left(\sigma \in F_1^{a,b}, \sigma \text{ not on } B_1^{a,b}\right)\\ &\leq \mathbb{P}\left(\sigma \text{ not on } B_1^{a,b} \,\middle|\, \sigma \in F_1^{a,b}\right) \leq \frac{\varepsilon^3}{4n^3k^9},\end{aligned}$$

and similarly for the other term. Putting everything together gives us

$$\mathbb{P}\left(\sigma \text{ on } B_1^{a,b}, \sigma \text{ on } B_2^{a,c}\right) \geq \frac{\varepsilon^3}{2n^3k^9},$$

which, by the discussion above, implies (2.14). □

2.3.3.2 Case 2

Lemma 2.12. *If $d \notin \{a,b\}$, then* (2.11) *and* (2.12) *cannot hold simultaneously.*

Proof. Suppose on the contrary that (2.11) and (2.12) do both hold. Then

$$\mathbb{P}\left(\sigma \in F_1^{a,b}\right) \geq \frac{\varepsilon}{nk^3} \qquad \text{and} \qquad \mathbb{P}\left(\sigma \in F_2^{c,d}\right) \geq \frac{\varepsilon}{nk^3}$$

as before. Since $\{a,b\} \cap \{c,d\} = \emptyset$, $\left\{\sigma \in F_1^{a,b}\right\}$ and $\left\{\sigma \in F_2^{c,d}\right\}$ are independent events, and so

$$\mathbb{P}\left(\sigma \in F_1^{a,b} \cap F_2^{c,d}\right) = \mathbb{P}\left(\sigma \in F_1^{a,b}\right) \mathbb{P}\left(\sigma \in F_2^{c,d}\right) \geq \frac{\varepsilon^2}{n^2k^6}.$$

In the same way as before, by the definition of large fibers this implies that

$$\mathbb{P}\left(\sigma \text{ on } B_1^{a,b}, \sigma \text{ on } B_2^{c,d}\right) \geq \frac{\varepsilon^2}{2n^2k^6} > 0,$$

but it is clear that

$$\mathbb{P}\left(\sigma \text{ on } B_1^{a,b}, \sigma \text{ on } B_2^{c,d}\right) = 0,$$

since σ on $B_1^{a,b}$ and on $B_2^{c,d}$ requires $f(\sigma) \in \{a,b\} \cap \{c,d\} = \emptyset$. So we have reached a contradiction. $\square$

2.3.4 Big mass on small fibers

We now deal with the case when (2.9) holds, i.e., when we have a big mass on the small fibers for the boundary $B_1^{a,b}$. We formalize the ideas of the outline described in Section 2.3.1 in a series of statements.

First, we want to formalize that the boundaries of the boundaries are big when we are on a small fiber.

Lemma 2.13. *Fix coordinate 1 and the pair of alternatives a and b. Let $z^{a,b}$ be such that $B_1\left(z^{a,b}\right)$ is a small fiber for $B_1^{a,b}$. Then, writing $B \equiv B_1\left(z^{a,b}\right)$, we have*

$$|\partial_e(B)| \geq \frac{\varepsilon^3}{4n^3k^9}|B|$$

and

$$\mathbb{P}(\sigma \in \partial(B)) \geq \frac{\varepsilon^3}{2n^4k^9k!}\mathbb{P}(\sigma \in B), \tag{2.16}$$

where both the edge boundary $\partial_e(B)$ and the boundary $\partial(B)$ are with respect to the induced subgraph $F\left(z^{a,b}\right)$, which is isomorphic to $K_{k!/2}^n$, the Cartesian product of n complete graphs of size $k!/2$.

Proof. We use Corollary 2.7 with $\ell = k!/2$ and the set A being either B or $B^c := F\left(z^{a,b}\right) \setminus B$. Suppose first that $|B| \leq \left(1 - \frac{2}{k!}\right)(k!/2)^n$. Then $|\partial_e(B)| \geq |B|$. Suppose now that $|B| > \left(1 - \frac{2}{k!}\right)(k!/2)^n$. Since we are in the case of a small fiber, we also know that $|B| \leq \left(1 - \frac{\varepsilon^3}{4n^3k^9}\right)(k!/2)^n$. Consequently, we get

$$|\partial_e(B)| = |\partial_e(B^c)| \geq |B^c| \geq \frac{\varepsilon^3}{4n^3k^9}|B|,$$

which proves the first claim.

A ranking profile in $F\left(z^{a,b}\right)$ has $(k!/2 - 1)n \leq nk!/2$ neighbors in $F\left(z^{a,b}\right)$, which then implies (2.16). $\square$

Corollary 2.14. *If (2.9) holds, then*

$$\mathbb{P}\left(\sigma \in \bigcup_{z^{a,b}} \partial\left(B_1\left(z^{a,b}\right)\right)\right) \geq \frac{\varepsilon^4}{2n^5k^{12}k!}.$$

Proof. Using the previous lemma and (2.9) we have

$$\begin{aligned}
\mathbb{P}\left(\sigma \in \bigcup_{z^{a,b}} \partial\left(B_1\left(z^{a,b}\right)\right)\right) &= \sum_{z^{a,b}} \mathbb{P}\left(\sigma \in \partial\left(B_1\left(z^{a,b}\right)\right)\right) \\
&\geq \sum_{z^{a,b}: B_1\left(z^{a,b}\right) \subseteq \mathrm{Sm}\left(B_1^{a,b}\right)} \mathbb{P}\left(\sigma \in \partial\left(B_1\left(z^{a,b}\right)\right)\right) \\
&\geq \sum_{z^{a,b}: B_1\left(z^{a,b}\right) \subseteq \mathrm{Sm}\left(B_1^{a,b}\right)} \frac{\varepsilon^3}{2n^4k^9k!} \mathbb{P}\left(\sigma \in B_1\left(z^{a,b}\right)\right) \\
&= \frac{\varepsilon^3}{2n^4k^9k!} \mathbb{P}\left(\sigma \in \mathrm{Sm}\left(B_1^{a,b}\right)\right) \geq \frac{\varepsilon^4}{2n^5k^{12}k!}.
\end{aligned}$$

□

Next, we want to find manipulation points on the boundaries of boundaries.

Lemma 2.15. *Suppose that the ranking profile σ is on the boundary of a fiber for $B_1^{a,b}$, i.e.,*

$$\sigma \in \bigcup_{z^{a,b}} \partial\left(B_1\left(z^{a,b}\right)\right).$$

Then either $\sigma_{-1} \in D_1(a,b)$, or there exists a manipulation point $\hat{\sigma}$ which differs from σ in at most two coordinates, one of them being the first coordinate.

Proof. First of all, by our assumption that σ is on the boundary of a fiber for $B_1^{a,b}$, we know that $\sigma \in B_1\left(z^{a,b}\right)$ for some $z^{a,b}$, which means that there exists a ranking profile $\sigma' = (\sigma_1', \sigma_{-1})$ such that $(\sigma, \sigma') \in B_1^{a,b}$. We may assume that $a \overset{\sigma_1}{>} b$ and $b \overset{\sigma_1'}{>} a$, or else either σ or σ' is a manipulation point.

Now since $\sigma \in \partial\left(B_1\left(z^{a,b}\right)\right)$, we also know that there exists a ranking profile $\pi = (\pi_j, \sigma_{-j}) \in F\left(z^{a,b}\right) \setminus B_1\left(z^{a,b}\right)$ for some $j \in [k]$. We distinguish two cases: $j \neq 1$ and $j = 1$.

Case 1: $\mathbf{j \neq 1}$. What does it mean for $\pi = (\pi_j, \sigma_{-j})$ to be on the same fiber as σ, but for π to not be in $B_1\left(z^{a,b}\right)$? First of all, being on the same fiber means that σ_j and π_j both rank a and b in the same order. Now $\pi \notin B_1\left(z^{a,b}\right)$ means that

- either $f(\pi) \neq a$;
- or $f(\pi) = a$ and $f(\pi_1', \pi_{-1}) \neq b$ for every $\pi_1' \in S_k$.

If $f(\pi) = b$, then either σ or π is a manipulation point, since the order of a and b is the same in both σ_j and π_j (since σ and π are on the same fiber).

Suppose that $f(\pi) = c \notin \{a, b\}$. Then we can define a SCF function on two coordinates by fixing all coordinates except coordinates 1 and j to agree with the respective coordinates of σ—letting coordinates 1 and j vary we get a SCF function on two coordinates which takes on at least three values (a, b, and c), and does not only depend on one coordinate. Now applying the Gibbard-Satterthwaite theorem we get that this SCF on two coordinates has a manipulation point, which means that our original SCF f has a manipulation point which agrees with σ in all coordinates except perhaps in coordinates 1 and j.

So the final case is that $f(\pi) = a$ and $f(\pi'_1, \pi_{-1}) \neq b$ for every $\pi'_1 \in S_k$. In particular for $\tilde{\pi} := (\sigma'_1, \pi_{-1}) = \left(\pi_j, \sigma'_{-j}\right)$ we have $f(\tilde{\pi}) \neq b$. Now if $f(\tilde{\pi}) = a$ then either σ' or $\tilde{\pi}$ is a manipulation point, since the order of a and b is the same in both $\sigma'_j = \sigma_j$ and π_j. Finally, if $f(\tilde{\pi}) = c \notin \{a, b\}$, then we can apply the Gibbard-Satterthwaite theorem just like in the previous paragraph.

Case 2: $\mathbf{j = 1}$. We can again ask: what does it mean for $\pi = (\pi_1, \sigma_{-1})$ to be on the same fiber as σ, but for π to not be in $B_1\left(z^{a,b}\right)$? First of all, being on the same fiber means that σ_1 and π_1 both rank a and b in the same order (namely, as discussed at the beginning, ranking a above b, or else we have a manipulation point). Now $\pi \notin B_1\left(z^{a,b}\right)$ means that

- either $f(\pi) \neq a$;
- or $f(\pi) = a$ and $f(\pi'_1, \pi_{-1}) \neq b$ for every $\pi'_1 \in S_k$.

However, we know that $f(\sigma') = b$ and that σ' is of the form $\sigma' = (\sigma'_1, \sigma_{-1}) = (\sigma'_1, \pi_{-1})$, and so the only way we can have $\pi \notin B_1\left(z^{a,b}\right)$ is if $f(\pi) \neq a$.

If $f(\pi) = b$, then π is a manipulation point, since $a \overset{\pi_1}{>} b$ and $f(\sigma) = a$.

So the remaining case is if $f(\pi) = c \notin \{a, b\}$. This means that $f_{\sigma_{-1}}$ (see Definition 2.13) takes on at least three values. Denote by $H \subseteq [k]$ the range of $f_{\sigma_{-1}}$. Now either $\sigma_{-1} \in D_1^H \subseteq D_1(a, b)$, or there exists a manipulation point $\hat{\sigma}$ which agrees with σ in every coordinate except perhaps the first. $\square$

Finally, we need to deal with dictators on the first coordinate.

Lemma 2.16. *Assume that* $\mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \geq \varepsilon$. *We have that either*

$$\mathbb{P}\left(\sigma_{-1} \in D_1(a, b)\right) \leq \frac{\varepsilon^4}{4n^5 k^{12} k!},$$

or

$$\mathbb{P}(\sigma \in M) \geq \frac{\varepsilon^5}{4n^7 k^{12} (k!)^4}. \tag{2.17}$$

Proof. Suppose that $\mathbb{P}\left(\sigma_{-1} \in D_1(a, b)\right) \geq \frac{\varepsilon^4}{4n^5 k^{12} k!}$, which is the same as

$$\sum_{H : \{a,b\} \subseteq H, |H| \geq 3} \mathbb{P}\left(\sigma_{-1} \in D_1^H\right) \geq \frac{\varepsilon^4}{4n^5 k^{12} k!}. \tag{2.18}$$

Note that for every $H \subseteq [k]$ we have

$$\varepsilon \leq \mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \leq \mathbb{P}\left(f(\sigma) \neq \mathrm{top}_H(\sigma_1)\right) \leq 1 - \mathbb{P}\left(D_1^H\right),$$

and so

$$\mathbb{P}\left(D_1^H\right) \leq 1 - \varepsilon. \tag{2.19}$$

The main idea is that (2.19) implies that the size of the boundary of D_1^H is comparable to the size of D_1^H, and if we are on the boundary of D_1^H, then there is a manipulation point nearby.

So first let us establish that the size of the boundary of D_1^H is comparable to the size of D_1^H. This is done along the same lines as the proof of Lemma 2.13.

Notice that $D_1^H \subseteq S_k^{n-1}$, where S_k^{n-1} should be thought of as the Cartesian product of $n-1$ copies of the complete graph on S_k. We apply Corollary 2.7 with $\ell = k!$ and with $n-1$ copies, and we see that if $\varepsilon \geq \frac{1}{k!}$, then $\left|\partial_e\left(D_1^H\right)\right| \geq \left|D_1^H\right|$. If $\varepsilon < \frac{1}{k!}$ and $1 - \frac{1}{k!} \leq \mathbb{P}\left(D_1^H\right) \leq 1 - \varepsilon$ then

$$\left|\partial_e\left(D_1^H\right)\right| = \left|\partial_e\left(\left(D_1^H\right)^c\right)\right| \geq \left|\left(D_1^H\right)^c\right| \geq \varepsilon \left|D_1^H\right|.$$

So in any case we have $\left|\partial_e\left(D_1^H\right)\right| \geq \varepsilon\left|D_1^H\right|$. Since σ_{-1} has $(n-1)(k!-1) \leq nk!$ neighbors in S_k^{n-1}, we have that

$$\mathbb{P}\left(\sigma_{-1} \in \partial\left(D_1^H\right)\right) \geq \frac{\varepsilon}{nk!}\mathbb{P}\left(\sigma_{-1} \in D_1^H\right).$$

Consequently, by (2.18), we have

$$\begin{aligned}\mathbb{P}\left(\sigma_{-1} \in \bigcup_{H:\{a,b\}\subseteq H, |H|\geq 3} \partial\left(D_1^H\right)\right) &= \sum_{H:\{a,b\}\subseteq H,|H|\geq 3} \mathbb{P}\left(\sigma_{-1} \in \partial\left(D_1^H\right)\right) \\ &\geq \sum_{H:\{a,b\}\subseteq H,|H|\geq 3} \frac{\varepsilon}{nk!}\mathbb{P}\left(\sigma_{-1} \in D_1^H\right) \geq \frac{\varepsilon^5}{4n^6k^{12}(k!)^2}.\end{aligned}$$

Next, suppose that $\sigma_{-1} \in \partial\left(D_1^H\right)$ for some H such that $\{a, b\} \subseteq H$, and $|H| \geq 3$. We want to show that then there is a manipulation point "close" to σ_{-1} in some sense. To be more precise: for the manipulation point $\hat{\sigma}$, $\hat{\sigma}_{-1}$ will agree with σ_{-1} in all except maybe one coordinate.

If $\sigma_{-1} \in \partial\left(D_1^H\right)$, then there exist $j \in \{2, \dots, n\}$ and σ_j' such that $\sigma_{-1}' := \left(\sigma_j', \sigma_{-\{1,j\}}\right) \notin D_1^H$. That is, $f_{\sigma_{-1}'}(\cdot) \not\equiv \mathrm{top}_H(\cdot)$. There can be two ways that this can happen—the two cases are outlined below. Denote by $H' \subseteq [k]$ the range of $f_{\sigma_{-1}'}$.

Case 1: $\mathbf{H' = H}$. In this case we automatically know that there exists a manipulation point $\hat{\sigma}$ such that $\hat{\sigma}_{-1} = \sigma_{-1}'$, and so $\hat{\sigma}_{-1}$ agrees with σ_{-1} in all coordinates except coordinate j.

Case 2: $\mathbf{H' \neq H}$. W.l.o.g. suppose $H' \setminus H \neq \emptyset$, and let $c \in H' \setminus H$. (The other case when $H \setminus H' \neq \emptyset$ works in exactly the same way.) First of all, we may assume that $f_{\sigma_{-1}'}(\cdot) \equiv \mathrm{top}_{H'}(\cdot)$, because otherwise we have a manipulation point just like in Case 1.

We can define a SCF on two coordinates by fixing all coordinates except coordinate 1 and j to agree with σ_{-1}, and varying coordinates 1 and j. We know that the outcome takes on at least three different values, since $\sigma_{-1} \in D_1^H$, and $|H| \geq 3$.

Now let us show that this SCF is not a function of the first coordinate. Let σ_1 be a ranking which puts c first, and then a. Then $f(\sigma_1, \sigma_{-1}) = a$, but $f\left(\sigma_1, \sigma'_{-1}\right) = c$, which shows that this SCF is not a function of the first coordinate (since a change in coordinate j can change the outcome).

Consequently, the Gibbard-Satterthwaite theorem tells us that this SCF on two coordinates has a manipulation point, and therefore there exists a manipulation point $\hat{\sigma}$ for f such that $\hat{\sigma}_{-1}$ agrees with σ_{-1} in all coordinates except coordinate j.

Putting everything together yields (2.17). □

2.3.5 Proof of Theorem 2.9 concluded

Proof of Theorem 2.9. If (2.11) and (2.12) hold, then we are done by Lemmas 2.11 and 2.12.

If not, then either (2.9) holds, or (2.9) holds for the boundary $B_2^{c,d}$; w.l.o.g. assume that (2.9) holds.

By Corollary 2.14, we have

$$\mathbb{P}\left(\sigma \in \bigcup_{z^{a,b}} \partial\left(B_1\left(z^{a,b}\right)\right)\right) \geq \frac{\varepsilon^4}{2n^5k^{12}k!}.$$

We may assume that $\mathbb{P}(\sigma_{-1} \in D_1(a,b)) \leq \frac{\varepsilon^4}{4n^5k^{12}k!}$, since otherwise we are done by Lemma 2.16. Consequently, we then have

$$\mathbb{P}\left(\sigma \in \bigcup_{z^{a,b}} \partial\left(B_1\left(z^{a,b}\right)\right), \sigma_{-1} \notin D_1(a,b)\right) \geq \frac{\varepsilon^4}{4n^5k^{12}k!}.$$

We can then finish our argument using Lemma 2.15:

$$\mathbb{P}(\sigma \in M) \geq \frac{1}{n\,(k!)^2}\mathbb{P}\left(\sigma \in \bigcup_{z^{a,b}} \partial\left(B_1\left(z^{a,b}\right)\right), \sigma_{-1} \notin D_1(a,b)\right) \geq \frac{\varepsilon^4}{4n^6k^{12}\,(k!)^3}.$$ □

2.4 An overview of the refined proof

In order to improve on the result of Theorem 2.9—in particular to get rid of the factor of $\frac{1}{(k!)^4}$—we need to refine the methods used in the previous section. We continue the approach of Isaksson, Kindler and Mossel [40, 41], where the authors first proved a quantitative Gibbard-Satterthwaite theorem for neutral SCFs with a bound involving factors of $\frac{1}{k!}$, and then with a refined method were able to remove these factors.

The key to the refined method is to consider the so-called *refined rankings graph* instead of the general rankings graph studied in Section 2.3. The vertices of this graph are again ranking profiles (elements of S_k^n), and two vertices are connected by an edge if they differ in exactly one coordinate, and by an adjacent transposition in that coordinate. Again, the SCF

f naturally partitions the vertices of this graph into k subsets, depending on the value of f at a given vertex. Clearly a 2-manipulation point can only be on the edge boundary of such a subset in the refined rankings graph, and so it is important to study these boundaries.

One of the important steps of the proof in Section 2.3 is creating a configuration where we fix all but two coordinates, and the SCF f takes on at least three values when we vary these two coordinates—then we can define another SCF on two voters and k alternatives which must have a manipulation point by the Gibbard-Satterthwaite theorem. The advantage of the refined rankings graph is that we can create a configuration where we fix all but two coordinates, and in these two coordinates we also fix all but constantly many adjacent alternatives, and the SCF takes on at least three values when we vary these constantly many adjacent alternatives in the two coordinates. Then we can define another SCF on two voters and r alternatives, where r is a small constant, which must have a manipulation point by the Gibbard-Satterthwaite theorem. Since r is a constant, we only lose a constant factor in our estimates, not factors of $\frac{1}{k!}$.

We state the refined result in Theorem 2.35, which we also prove in Section 2.7. The proof of Theorem 2.35 follows the outline of the proof of Theorem 2.9: we know that there are at least two refined boundaries which are big (by Isaksson et al. [40, 41]); we partition them according to their fibers; we distinguish small and large fibers; and we consider two cases: the small fiber case and the large fiber case. The ideas in both cases are roughly the same as in Section 2.3, except the proofs are more involved. There is, however, one major difference in the small fiber case, which is the following.

The difficulty is dealing with the case when we are on the boundary of a small fiber in the first coordinate. Suppose $\sigma = (\sigma_1, \sigma_{-1})$ is on such a boundary. We know that there are $k!$ ranking profiles which agree with σ in coordinates 2 through n. The difficulty comes from the fact that—in order to obtain a polynomial bound in k—we are only allowed to look at a polynomial number (in k) of these ranking profiles when searching for a manipulation point. If there is an r-manipulation point among them for some small constant r, then we are done. If this is not the case then σ is what we call a *local dictator* on some subset of the alternatives in coordinate 1. We say that σ is a local dictator on some subset $H \subseteq [k]$ of the alternatives in coordinate 1 if the alternatives in H are adjacent in σ_1, and permuting the alternatives in H in every possible way in the first coordinate, the outcome of the SCF f is always the top-ranked alternative in H.

So instead of dealing with dictators on some subset in coordinate 1, as in Section 2.3, we have to deal with *local dictators* on some subset in coordinate 1. This analysis involves essentially only the first coordinate, in essence proving a quantitative Gibbard-Satterthwaite theorem for one voter. As discussed in Section 2.1.3, this has not been studied in the literature before, and, moreover, we were not able to utilize previous quantitative Gibbard-Satterthwaite theorems to solve this problem easily. Hence we separate this argument from the rest of the proof of Theorem 2.35 and formulate a quantitative Gibbard-Satterthwaite theorem for one voter, Theorem 2.4, which is proven in Section 2.6. This proof forms the backbone for the proof of Theorem 2.35, which is then proven in Section 2.7.

2.5 Refined rankings graph—introduction and preliminaries

2.5.1 Transpositions, boundaries, and influences

Definition 2.15 (Adjacent transpositions)**.** *Given two elements $a, b \in [k]$, the adjacent transposition $[a : b]$ between them is defined as follows. If $\sigma \in S_k$ has a and b adjacent, then $[a : b]\,\sigma$ is obtained from σ by exchanging a and b. Otherwise $[a : b]\,\sigma = \sigma$.*

We let T denote the set of all $k\,(k-1)\,/2$ adjacent transpositions.

For $\sigma \in S_k^n$, we let $[a : b]_i\,\sigma$ denote the ranking profile obtained by applying $[a : b]$ on the i^{th} coordinate of σ while leaving all other coordinates unchanged.

Definition 2.16 (Boundaries)**.** *For a given SCF f and a given alternative $a \in [k]$, we define*

$$W^a\,(f) = \{\sigma \in S_k^n : f\,(\sigma) = a\}\,,$$

the set of ranking profiles where the outcome of the vote is a. The edge boundary of this set (with respect to the underlying refined rankings graph) is denoted by $B^{a;T}\,(f)$: $B^{a;T}\,(f) = \partial_e\,(W^a\,(f))$. This boundary can be partitioned: we say that the edge boundary of $W^a\,(f)$ in the direction of the i^{th} coordinate is

$$B_i^{a;T}\,(f) = \left\{(\sigma, \sigma') \in B^{a;T}\,(f) : \sigma_i \neq \sigma_i'\right\}.$$

The boundary $B^a\,(f)$ can be therefore written as $B^{a;T}\,(f) = \cup_{i=1}^n B_i^{a;T}\,(f)$. We can also define the boundary between two alternatives a and b in the direction of the i^{th} coordinate:

$$B_i^{a,b;T}\,(f) = \Big\{(\sigma, \sigma') \in B_i^{a;T}\,(f) : f\,(\sigma') = b\Big\}.$$

Moreover, we can define the boundary between two alternatives a and b in the direction of the i^{th} coordinate with respect to the adjacent transposition $z \in T$:

$$B_i^{a,b;z}\,(f) = \Big\{(\sigma, \sigma') \in B_i^{a;T}\,(f) : \sigma' = z_i\sigma, f\,(\sigma') = b\Big\}.$$

We also say that σ is on *the boundary $B_i^{a,b;z}\,(f)$ if $(\sigma, z_i\sigma) \in B_i^{a,b;z}\,(f)$. Clearly we have*

$$B_i^{a,b;T}\,(f) = \bigcup_{z\in T} B_i^{a,b;z}\,(f)\,.$$

Definition 2.17 (Influences)**.** *Given $z \in T$, we define*

$$\begin{aligned}
\mathrm{Inf}_i^{a,b;z}\,(f) &= \mathbb{P}\left(f\,(\sigma) = a, f\left(\sigma^{(i)}\right) = b\right)\\
\mathrm{Inf}_i^{a;z}\,(f) &= \mathbb{P}\left(f\,(\sigma) = a, f\left(\sigma^{(i)}\right) \neq a\right)\\
\mathrm{Inf}_i^{a,b;T}\,(f) &= \sum_{z\in T} \mathrm{Inf}_i^{a,b;z}\,(f)\,,
\end{aligned}$$

where σ is uniformly distributed in S_k^n and $\sigma^{(i)}$ is obtained from σ by rerandomizing the i^{th} coordinate σ_i in the following way: with probability 1/2 we keep it as σ_i, and otherwise we replace it by $z\sigma_i$.

Note that for $a \neq b$,

$$\mathrm{Inf}_i^{a,b;z}(f) = \frac{1}{2}\mathbb{P}(f(\sigma) = a, f(z_i\sigma) = b) = \frac{1}{2}\frac{\left|B_i^{a,b;z}(f)\right|}{(k!)^n}.$$

Again, most of the time the specific SCF f will be clear from the context, in which case we omit the dependence on f.

2.5.2 Manipulation points on refined boundaries

The following two lemmas from Isaksson et al. [40, 41] identify manipulation points on (or close to) these refined boundaries.

Lemma 2.17. *[40, 41, Lemma 7.1.] Fix $f : S_k^n \to [k]$, distinct $a, b \in [k]$, and $(\sigma, \pi) \in B_i^{a,b;T}$. Then either $\sigma_i = [a : b]\pi_i$, or one of σ and π is a 2-manipulation point for f.*

Lemma 2.18. *[40, 41, Lemma 7.2.] Fix $f : S_k^n \to [k]$ and points $\sigma, \pi, \mu \in S_k^n$ such that $(\sigma, \pi) \in B_i^{a,b;T}$, and $(\mu, \pi) \in B_j^{c,b;T}$, where a, b, and c are distinct and $i \neq j$. Then there exists a 3-manipulation point $\nu \in S_k^n$ for f such that $\nu_\ell = \pi_\ell$ for $\ell \notin \{i, j\}$ and ν_i is equal to σ_i or π_i except that the position of c may be shifted arbitrarily and ν_j is equal to μ_j or π_j except that the position of a may be shifted arbitrarily.*

2.5.3 Large refined boundaries

An essential result that will be our starting point in Section 2.7 is the following lemma, again from Isaksson et al. [40, 41], which shows that there are large refined boundaries (or else we have a lot of 2-manipulation points automatically).

Lemma 2.19. *[40, 41, Lemma 7.3.] Suppose that $k \geq 3$ and that the SCF $f : S_k^n \to [k]$ satisfies $\mathbf{D}(f, \overline{\mathrm{NONMANIP}}) \geq \varepsilon$. Let σ be uniformly selected from S_k^n. Then either*

$$\mathbb{P}(\sigma \in M_2(f)) \geq \frac{4\varepsilon}{nk^7}, \tag{2.20}$$

or there exist distinct $i, j \in [n]$ and $\{a, b\}, \{c, d\} \subseteq [k]$ such that $c \notin \{a, b\}$ and

$$\mathrm{Inf}_i^{a,b;[a:b]}(f) \geq \frac{2\varepsilon}{nk^7} \quad \text{and} \quad \mathrm{Inf}_j^{c,d;[c:d]}(f) \geq \frac{2\varepsilon}{nk^7}. \tag{2.21}$$

2.5.4 Fibers

We again use fibers $F\left(z^{a,b}\right)$ as defined in Definition 2.11. However, we need more than this. We note that the following definitions only apply in Section 2.7, i.e., when we have at least two voters; in Section 2.6, when we have only one voter, things are simpler.

Given the result of Lemma 2.19, our primary interest is in the boundary $B_i^{a,b;[a:b]}$. For ranking profiles on this boundary, we know that the alternatives a and b are adjacent in coordinate i—so we know more than just the preference between a and b in coordinate i. Consequently we would like to divide the set of ranking profiles with a and b adjacent in coordinate i according to the preferences between a and b in all coordinates except coordinate i. The following definitions make this precise.

As done in Section 2.2.7 for ranking profiles, we can write $x_{-i}^{a,b} \equiv x_{-i}^{a,b}(\sigma)$ for the vector of preferences between a and b for all coordinates except coordinate i, i.e., the whole vector of preferences between a and b is $x^{a,b}(\sigma) = \left(x_i^{a,b}(\sigma), x_{-i}^{a,b}(\sigma)\right)$.

We can define $F\left(z_{-i}^{a,b}\right)$ analogously to $F\left(z^{a,b}\right)$:

$$F\left(z_{-i}^{a,b}\right) := \left\{\sigma : x_{-i}^{a,b}(\sigma) = z_{-i}^{a,b}\right\}.$$

We also define the subset of $F\left(z_{-i}^{a,b}\right)$ where a and b are adjacent in coordinate i, with a above b:

$$\bar{F}\left(z_{-i}^{a,b}\right) := \left\{\sigma \in F\left(z_{-i}^{a,b}\right) : a \text{ and } b \text{ are adjacent in coordinate } i, \text{ with } a \text{ above } b\right\}.$$

Given a SCF f, for any pair of alternatives $a, b \in [k]$ and coordinate $i \in [n]$, we can also partition the boundary $B_i^{a,b}(f)$ according to its fibers. There are multiple, slightly different ways of doing this, but for our purposes the following definition is most useful.

Define

$$B_i\left(z_{-i}^{a,b}\right) := \left\{\sigma \in \bar{F}\left(z_{-i}^{a,b}\right) : f(\sigma) = a, f\left([a:b]_i \sigma\right) = b\right\},$$

where we omit the dependence of $B_i\left(z_{-i}^{a,b}\right)$ on f. We call sets of the form $B_i\left(z_{-i}^{a,b}\right) \subseteq \bar{F}\left(z_{-i}^{a,b}\right)$ *fibers for the boundary* $B_i^{a,b;[a:b]}$.

We now distinguish between small and large fibers for the boundary $B_i^{a,b;[a:b]}$.

Definition 2.18 (Small and large fibers). *We say that the fiber* $B_i\left(z_{-i}^{a,b}\right) \subseteq \bar{F}\left(z_{-i}^{a,b}\right)$ *is* large *if*

$$\mathbb{P}\left(\sigma \in B_i\left(z_{-i}^{a,b}\right) \,\middle|\, \sigma \in \bar{F}\left(z_{-i}^{a,b}\right)\right) \geq 1 - \gamma,$$

where $\gamma = \frac{\varepsilon^3}{10^3 n^3 k^{24}}$, *and* small *otherwise.*

As before, we denote by $\mathrm{Lg}\left(B_i^{a,b;[a:b]}\right)$ *the union of large fibers for the boundary* $B_i^{a,b;[a:b]}$, *i.e.,*

$$\mathrm{Lg}\left(B_i^{a,b;[a:b]}\right) := \bigcup_{B_i\left(z_{-i}^{a,b}\right) \text{ is a large fiber}} B_i\left(z_{-i}^{a,b}\right),$$

and similarly, we denote by $\mathrm{Sm}\left(B_i^{a,b;[a:b]}\right)$ *the union of small fibers.*

As in Definition 2.12, we remark that what is important is that γ is a polynomial of $\frac{1}{n}$, $\frac{1}{k}$ and ε—the specific polynomial in this definition is the end result of the computation in the proof.

The following definition is used in Section 2.7.3 in dealing with the large fiber case in the refined setting.

Definition 2.19. *For a coordinate* i *and a pair of alternatives* a *and* b*, define* $F_i^{a,b}$ *to be the set of ranking profiles* σ *such that* $x^{a,b}(\sigma)$ *satisfies*

$$\mathbb{P}\left(f(\tilde{\sigma}) = \mathrm{top}_{\{a,b\}}(\tilde{\sigma}_i) \,\middle|\, \tilde{\sigma} \in F\left(x_{-i}^{a,b}(\sigma)\right)\right) \geq 1 - 2k\gamma.$$

Clearly $F_i^{a,b}$ is the union of fibers of the form $F\left(z^{a,b}\right)$, and also $F\left(\left(1, x_{-i}^{a,b}\right)\right) \subseteq F_i^{a,b}$ if and only if $F\left(\left(-1, x_{-i}^{a,b}\right)\right) \subseteq F_i^{a,b}$.

2.5.5 Boundaries of boundaries

In the refined graph setting, just like in the general rankings graph setting, we also look at boundaries of boundaries.

For a given vector $z_{-i}^{a,b}$ of preferences between a and b, we can think of $\bar{F}\left(z_{-i}^{a,b}\right)$ as a subgraph of the original refined rankings graph S_k^n, i.e., two ranking profiles in $\bar{F}\left(z_{-i}^{a,b}\right)$ are adjacent if they differ by one adjacent transposition in exactly one coordinate. Since both of the ranking profiles are in $\bar{F}\left(z_{-i}^{a,b}\right)$, this adjacent transposition keeps the order of a and b in all coordinates, and moreover it keeps a and b adjacent in coordinate i.

We choose to slightly modify this graph: the vertex set is still $\bar{F}\left(z_{-i}^{a,b}\right)$, but we modify the edge set by adding new edges. Suppose that $\sigma \in \bar{F}\left(z_{-i}^{a,b}\right)$ and

$$\sigma_i = \begin{pmatrix} \vdots \\ c \\ a \\ b \\ d \\ \vdots \end{pmatrix}; \qquad \sigma_i' = \begin{pmatrix} \vdots \\ a \\ b \\ c \\ d \\ \vdots \end{pmatrix}; \qquad \sigma_i'' = \begin{pmatrix} \vdots \\ c \\ d \\ a \\ b \\ \vdots \end{pmatrix}.$$

Define in this way $\sigma' = (\sigma_i', \sigma_{-i})$ and $\sigma'' = (\sigma_i'', \sigma_{-i})$, and add (σ, σ') and (σ, σ'') to the edge set. So basically, we consider the block of a and b in coordinate i as a single element, and connect two ranking profiles in $\bar{F}\left(z_{-i}^{a,b}\right)$ if they differ in an adjacent transposition in a single coordinate, allowing this transposition to move the block of a and b in coordinate i. We call this graph $G\left(z_{-i}^{a,b}\right) = \left(\bar{F}\left(z_{-i}^{a,b}\right), E\left(z_{-i}^{a,b}\right)\right)$, where $E\left(z_{-i}^{a,b}\right)$ is the edge set.

When we write $\partial_e\left(B_i\left(z_{-i}^{a,b}\right)\right)$, we mean the edge boundary of $B_i\left(z_{-i}^{a,b}\right)$ in the graph $G\left(z_{-i}^{a,b}\right)$, and similarly when we write $\partial\left(B_i\left(z_{-i}^{a,b}\right)\right)$, we mean the vertex boundary of $B_i\left(z_{-i}^{a,b}\right)$ in the graph $G\left(z_{-i}^{a,b}\right)$.

2.5.6 Local dictators, conditioning and miscellaneous definitions

In the general rankings graph setting we defined a dictator on a subset of the alternatives, but in the refined rankings graph setting we need to define so-called *local dictators*.

Definition 2.20 (Local dictators). *For a coordinate i and a subset of alternatives $H \subseteq [k]$, define LD_i^H to be the set of ranking profiles σ such that the alternatives in H form an adjacent block in σ_i, and permuting them among themselves in any order, the outcome of the SCF f is always the top ranked alternative among those in H. If $\sigma \in \mathrm{LD}_i^H$, then we call σ a local dictator on H in coordinate i.*

Also, for a pair of alternatives a and b, define

$$\mathrm{LD}_i\left(a,b\right) := \bigcup_{c \notin \{a,b\}} \mathrm{LD}_i^{\{a,b,c\}},$$

the set of local dictators on three alternatives, two of which are a and b, in coordinate i.

In dealing with local dictators, we will condition on the top of a particular coordinate being fixed. We therefore introduce the following notation.

Definition 2.21 (Conditioning). *For any coordinate $i \in [n]$ and any vector $\mathbf{v}$ of alternatives, we define*

$$\mathbb{P}_i^{\mathbf{v}}\left(\cdot\right) := \mathbb{P}\left(\cdot \,|\, \left(\sigma_i\left(1\right), \ldots, \sigma_i\left(|\mathbf{v}|\right)\right) = \mathbf{v}\right),$$

where $|\mathbf{v}|$ denotes the length of the vector $\mathbf{v}$. For example, $\mathbb{P}_1^{(a)}\left(\cdot\right) = \mathbb{P}\left(\cdot \,|\, \sigma_1\left(1\right) = a\right)$ and

$$\mathbb{P}_1^{(a,b,c)} = \mathbb{P}\left(\cdot \,|\, \left(\sigma_1\left(1\right), \sigma_1\left(2\right), \sigma_1\left(3\right)\right) = \left(a,b,c\right)\right).$$

We use the following notation in the proof of Theorem 2.5.

Definition 2.22 (Majority function). *For a function f whose domain X is finite and whose range is the set $\{a,b\}$, define $\mathrm{Maj}\left(f\right)$ by*

$$\mathrm{Maj}\left(f\right) = \begin{cases} a & \text{if} \quad \#\left\{x \in X : f\left(x\right) = a\right\} \geq \#\left\{x \in X : f\left(x\right) = b\right\}, \\ b & \text{if} \quad \#\left\{x \in X : f\left(x\right) = a\right\} < \#\left\{x \in X : f\left(x\right) = b\right\}. \end{cases}$$

2.6 Quantitative Gibbard-Satterthwaite theorem for one voter

In this section we prove our quantitative Gibbard-Satterthwaite theorem for one voter, Theorem 2.4. As mentioned before, we present this proof before proving Theorem 2.35, because the proof of Theorem 2.35 follows the lines of this proof, with slight modifications needed to deal with having $n > 1$ coordinates.

For the remainder of this section, let us fix the number of voters to be 1, the number of alternatives $k \geq 3$, and the SCF f, which satisfies $\mathbf{D}(f, \text{NONMANIP}) \geq \varepsilon$. Accordingly, we typically omit the dependence of various sets (e.g., boundaries between two alternatives) on f.

An additional notational remark: since our SCF is on one voter only, we omit the subscripts that denote the coordinate we are on. For example, we write simply $\text{Inf}^{a,b}$ instead of $\text{Inf}_1^{a,b}$, etc.

We present the proof in several steps.

2.6.1 Large boundary between two alternatives

The first thing we have to establish is a large boundary between two alternatives. This can be done just like in Lemma 2.19, except there are two small differences. On the one hand, the assumption of the lemma, namely that $\mathbf{D}(f, \text{NONMANIP}) \geq \varepsilon$, is weaker than that of the original lemma. On the other hand, here we only need one big boundary, unlike in Lemma 2.19, where Isaksson et al. [40, 41] showed that there are two big boundaries in two different coordinates. The following lemma formulates what we need.

Lemma 2.20. *Recall that f is a SCF on 1 voter and $k \geq 3$ alternatives which satisfies $\mathbf{D}(f, \text{NONMANIP}) \geq \varepsilon$. Let $\sigma \in S_k$ be selected uniformly. Then either*

$$\mathbb{P}(\sigma \in M_2) \geq \frac{4\varepsilon}{k^6} \tag{2.22}$$

or there exist alternatives $a, b \in [k]$ such that $a \neq b$ and

$$\text{Inf}^{a,b;[a:b]} \geq \frac{2\varepsilon}{k^6}. \tag{2.23}$$

Proof. The proof is just like the proof of Lemma 2.19. First, suppose that $\text{Inf}^{a,b;z} \geq \frac{2\varepsilon}{k^6}$ for some pair of alternatives $a \neq b$, and some transposition $z \neq [a:b]$. Then by Lemma 2.17, for any point $(\sigma, \sigma') \in B^{a,b;z}$, at least one of σ or $\sigma' = z\sigma$ is a 2-manipulation point. Then

$$|M_2| \geq \left|B^{a,b;z}\right| = 2 \cdot k! \cdot \text{Inf}^{a,b;z} \geq \frac{4\varepsilon}{k^6} k!,$$

and dividing with $k!$ gives (2.22). So for the remainder of the proof we may assume that $\text{Inf}^{a,b;z} < \frac{2\varepsilon}{k^6}$ for every $a \neq b$ and $z \neq [a:b]$.

For every $a \in [k]$, we have $\mathbf{D}\left(f, \mathrm{top}_{\{a\}}\right) \geq \varepsilon$, so $\mathbb{P}\left(f\left(\sigma\right) = a\right) \leq 1 - \varepsilon$. On the other hand, there exists an alternative, say $a \in [k]$, such that $\mathbb{P}\left(f\left(\sigma\right) = a\right) \geq \frac{1}{k}$. So for this alternative we have

$$\mathrm{Var}\left(\mathbf{1}\left[f\left(\sigma\right) = a\right]\right) \geq \frac{\varepsilon}{k},$$

and consequently using [40, 41, Corollary 6.5] and [40, 41, Proposition 2.3] we have

$$\sum_{w \in T} \sum_{b \neq a} \mathrm{Inf}^{a,b;w} = \sum_{w \in T} \mathrm{Inf}^{a;w} \geq \frac{1}{k^2} \mathrm{Var}\left(\mathbf{1}\left[f\left(\sigma\right) = a\right]\right) \geq \frac{\varepsilon}{k^3}.$$

Hence there must exist some $w \in T$ and $b \neq a$ such that $\mathrm{Inf}^{a,b;w} \geq \frac{2\varepsilon}{k^6}$, but by our assumption we must have $w = [a:b]$. $\square$

If (2.22) holds, then we are done, so in the following we assume that (2.23) holds.

We know that σ is on $B^{a,b;[a:b]}$ if $f\left(\sigma\right) = a$ and $f\left(\left[a:b\right]\sigma\right) = b$. We know that if $b \overset{\sigma}{>} a$, then σ is a 2-manipulation point, so if this happens in more than half of the cases when σ is on $B^{a,b;[a:b]}$, then we have

$$\mathbb{P}\left(\sigma \in M_2\right) \geq \frac{2\varepsilon}{k^6},$$

in which case we are again done. So we may assume in the following that

$$\mathbb{P}\left(\sigma \in B\right) \geq \frac{2\varepsilon}{k^6}, \tag{2.24}$$

where

$$B := \left\{\sigma : f\left(\sigma\right) = a, f\left(\left[a:b\right]\sigma\right) = b, a \overset{\sigma}{>} b\right\}.$$

2.6.2 Division into cases

We again divide into two cases.

We introduce the set $\bar{F}$ of permutations where a is directly above b:

$$\bar{F} := \left\{\sigma \in S_k : a \overset{\sigma}{>} b \text{ and } b \overset{\sigma'}{>} a, \text{ where } \sigma' = \left[a:b\right]\sigma\right\}.$$

One of the following two cases must hold.

Case 1: Small fiber case. We have

$$\mathbb{P}\left(\sigma \in B \,\middle|\, \sigma \in \bar{F}\right) \leq 1 - \frac{\varepsilon}{4k}. \tag{2.25}$$

Case 2: Large fiber case. We have

$$\mathbb{P}\left(\sigma \in B \,\middle|\, \sigma \in \bar{F}\right) > 1 - \frac{\varepsilon}{4k}. \tag{2.26}$$

2.6.3 Small fiber case

In this subsection we assume that (2.25) holds.

We first formalize that the boundary $\partial(B)$ of B is big (recall the definition of $\partial(B)$ from Section 2.5.5). The proof uses the canonical path method, as successfully adapted to this setting by Isaksson et al. [40, 41].

Lemma 2.21. *If* (2.25) *holds, then*

$$\mathbb{P}(\sigma \in \partial(B)) \geq \frac{\varepsilon}{2k^4}\mathbb{P}(\sigma \in B). \tag{2.27}$$

Proof. Let $B^c = \bar{F} \setminus B$. For every $(\sigma, \sigma') \in B \times B^c$, we define a canonical path from σ to σ', which has to pass through at least one edge in $\partial_e(B)$. Then if we show that every edge in $\partial_e(B)$ lies on at most r canonical paths, then it follows that $|\partial_e(B)| \geq |B| |B^c| / r$.

So let $(\sigma, \sigma') \in B \times B^c$. We apply the path construction of [40, 41, Proposition 6.4.], but considering the block formed by a and b as a single element. Since this path goes from σ (which is in B) to σ' (which is in B^c), it must pass through at least one edge in $\partial_e(B)$.

For a given edge $(\pi, \pi') \in \partial_e(B)$, at most how many possible $(\sigma, \sigma') \in B \times B^c$ pairs are there such that the canonical path between σ and σ' defined above passes through (π, π')? We learn from [40, 41, Proposition 6.4.] that there are at most $(k-1)^2 (k-1)!/2 < k^2 (k-1)!/2$ possibilities for the pair (σ, σ').

Recall that $|\bar{F}| = (k-1)!$. By our assumption we have $|B| \leq \left(1 - \frac{\varepsilon}{4k}\right)(k-1)!$, and so $|B^c| \geq \frac{\varepsilon}{4k}(k-1)!$. Therefore

$$|\partial_e(B)| \geq \frac{|B| |B^c|}{\frac{k^2}{2}(k-1)!} \geq \frac{\varepsilon}{2k^3} |B|.$$

Now in G every ranking profile has $k-2 < k$ neighbors, which implies (2.27). $\square$

Corollary 2.22. *If* (2.25) *holds, then*

$$\mathbb{P}(\sigma \in \partial(B)) \geq \frac{\varepsilon^2}{k^{10}}. \tag{2.28}$$

Proof. Combine Lemma 2.21 and (2.24). $\square$

Next we want to find manipulation points on the boundary $\partial(B)$. The next lemma tells us that if we are on the boundary $\partial(B)$, then either we can find manipulation points easily, or we are at a local dictator on three alternatives.

Lemma 2.23. *Suppose that $\sigma \in \partial(B)$. Then*

- *either $\sigma \in \mathrm{LD}(a, b)$,*
- *or there exists $\hat{\sigma} \in M_3$ such that $\hat{\sigma}$ is equal to σ or $[a : b]\sigma$ except that the position of a third alternative c might be shifted arbitrarily.*

Proof. Since $\sigma \in \partial(B) \subseteq B$, we know that $f(\sigma) = a$, and if $\sigma' = [a : b]\,\sigma$, then $f(\sigma') = b$. Let $\pi \in B^c$ denote the ranking profile such that $(\sigma, \pi) \in \partial_e(B)$, and let $\pi' = [a : b]\,\pi$. Since $\pi \notin B$, $(f(\pi), f(\pi')) \neq (a, b)$. Then, by Lemma 2.17, if $f(\pi) \neq f(\pi')$, then one of π and π' is a 2-manipulation point. So assume $f(\pi) = f(\pi')$.

There are two cases to consider: either σ and π differ by an adjacent transposition not involving the block of a and b, or they differ by an adjacent transposition that moves the block of a and b.

In the former case, it is not hard to see that one of σ, σ', π, or π' is a 2-manipulation point, by Lemma 2.17.

If σ and π differ by an adjacent transposition that involves the block of a and b, then there are again two cases to consider: either this transposition moves the block of a and b up in the ranking, or it moves it down.

If the block of a and b is moved up to get from σ to π, then we must have $f(\pi) = a$, or else σ or π is a 3-manipulation point. Then we must have $f(\pi') = f(\pi) = a$, in which case π' is a 3-manipulation point, since $f(\sigma') = b$.

The final case is when the block of a and b is moved down to get from σ to π, and a third alternative, call it c, is moved up, directly above the block of a and b. Now if $f(\pi) = d \notin \{a, b, c\}$, then σ or π is a 3-manipulation point. If $f(\pi) = f(\pi') = a$, then π' is a 3-manipulation point, whereas if $f(\pi) = f(\pi') = b$, then π is a 3-manipulation point. The remaining case is when $f(\pi) = f(\pi') = c$. Now if $f([b : c]\,\sigma) \neq a$ or $f([a : c]\,\sigma') \neq b$, then we again have a 3-manipulation point close to σ. Otherwise $\sigma \in \mathrm{LD}(a, b)$. ☐

The following corollary then tells us that either we have found many 3-manipulation points, or we have many local dictators on three alternatives.

Corollary 2.24. *If* (2.25) *holds, then either*

$$\sum_{c \notin \{a,b\}} \mathbb{P}\left(\sigma \in \mathrm{LD}^{\{a,b,c\}}\right) = \mathbb{P}(\sigma \in \mathrm{LD}(a, b)) \geq \frac{\varepsilon^2}{2k^{10}} \tag{2.29}$$

or

$$\mathbb{P}(\sigma \in M_3) \geq \frac{\varepsilon^2}{4k^{12}}.$$

2.6.3.1 Dealing with local dictators

So the remaining case we have to deal with in this small fiber case is when (2.29) holds, i.e., we have many local dictators on three alternatives.

Lemma 2.25. *Suppose that $\sigma \in \mathrm{LD}^{\{a,b,c\}}$ for some alternative $c \notin \{a, b\}$. Let σ' be equal to σ except that the block of a, b, and c is moved to the top of the coordinate. Then*

- *either $\sigma' \in \mathrm{LD}^{\{a,b,c\}}$,*

- *or there exists a 3-manipulation point $\hat{\sigma}$ which agrees with σ except that the positions of a, b, and c might be shifted arbitrarily.*

Proof. W.l.o.g. we may assume that in σ alternative a is ranked above b, which is ranked above c. Now move a to the top using a sequence of adjacent transpositions, all involving a; we call this procedure "bubbling" a to the top. If at any point during this the outcome of f is not a, then we have found a 2-manipulation point. Now bubble up b to right below a, and then bubble up c to be right below b. Again, if at any point during this the outcome of f is not a, then there is a 2-manipulation point. Otherwise we now have a, b, and c at the top (in this order), with the outcome of f being a. Now permuting alternatives a, b, and c at the top, we either have a 3-manipulation point, or $\sigma' \in \mathrm{LD}^{\{a,b,c\}}$. □

Corollary 2.26. *If* (2.29) *holds, then either*

$$\sum_{c \notin \{a,b\}} \mathbb{P}\left(\sigma \in \mathrm{LD}^{\{a,b,c\}}, \{\sigma(1), \sigma(2), \sigma(3)\} = \{a,b,c\}\right) \geq \frac{\varepsilon^2}{4k^{11}} \tag{2.30}$$

or

$$\mathbb{P}(\sigma \in M_3) \geq \frac{\varepsilon^2}{4k^{13}}.$$

Proof. Lemma 2.25 tells us that when we move the block of a, b, and c up to the top, we either encounter a 3-manipulation point, or we get a local dictator on $\{a,b,c\}$ at the top.

If we get a 3-manipulation point, by the describtion of this manipulation point in the lemma, there can be at most k^3 ranking profiles that give the same manipulation point.

If we arrive at a local dictator at the top, then there could have been at most k different places where the block of a, b, and c could have come from. □

Now (2.30) is equivalent to

$$\sum_{c \notin \{a,b\}} \mathbb{P}\left(\sigma \in \mathrm{LD}^{\{a,b,c\}}, (\sigma(1), \sigma(2), \sigma(3)) = (a,b,c)\right) \geq \frac{\varepsilon^2}{24k^{11}}. \tag{2.31}$$

We know that

$$\mathbb{P}((\sigma(1), \sigma(2), \sigma(3)) = (a,b,c)) = \frac{1}{k(k-1)(k-2)} \leq \frac{6}{k^3},$$

and so (2.31) implies (recall Definition 2.21)

$$\sum_{c \notin \{a,b\}} \mathbb{P}^{(a,b,c)}\left(\sigma \in \mathrm{LD}^{\{a,b,c\}}\right) \geq \frac{\varepsilon^2}{144k^8}. \tag{2.32}$$

Now fix an alternative $c \notin \{a,b\}$ and define the graph $G_{(a,b,c)} = \left(V_{(a,b,c)}, E_{(a,b,c)}\right)$ to have vertex set

$$V_{(a,b,c)} := \{\sigma \in S_k : (\sigma(1), \sigma(2), \sigma(3)) = (a,b,c)\}$$

and for $\sigma, \pi \in V_{(a,b,c)}$ let $(\sigma, \pi) \in E_{(a,b,c)}$ if and only if σ and π differ by an adjacent transposition. So $G_{(a,b,c)}$ is the subgraph of the refined rankings graph induced by the vertex set $V_{(a,b,c)}$. (If $k = 3$ or $k = 4$, then this graph consists of only one vertex, and no edges.)

Let

$$T\left(a, b, c\right) := V_{(a,b,c)} \cap \mathrm{LD}^{\{a,b,c\}},$$

and let $\partial_e \left(T\left(a, b, c\right)\right)$ and $\partial \left(T\left(a, b, c\right)\right)$ denote the edge- and vertex-boundary of $T\left(a, b, c\right)$ in $G_{(a,b,c)}$, respectively.

The next lemma shows that unless $T\left(a, b, c\right)$ is almost all of $V_{(a,b,c)}$, the size of the boundary $\partial \left(T\left(a, b, c\right)\right)$ is comparable to the size of $T\left(a, b, c\right)$. The proof uses a canonical path argument, just like in Lemma 2.21.

Lemma 2.27. *Let $c \notin \{a, b\}$ be arbitrary. Write $T \equiv T\left(a, b, c\right)$ for simplicity. If*

$$\mathbb{P}^{(a,b,c)}\left(\sigma \in T\right) \leq 1 - \delta,$$

then

$$\mathbb{P}^{(a,b,c)}\left(\sigma \in \partial\left(T\right)\right) \geq \frac{\delta}{k^3} \mathbb{P}^{(a,b,c)}\left(\sigma \in T\right). \tag{2.33}$$

Proof. Let $T^c = V_{(a,b,c)} \backslash T\left(a, b, c\right)$. For every $(\sigma, \sigma') \in T \times T^c$, we define a canonical path from σ to σ', which has to pass through at least one edge in $\partial_e\left(T\right)$. Then if we show that every edge in $\partial_e\left(T\right)$ lies on at most r canonical paths, then it follows that $\left|\partial_e\left(T\right)\right| \geq |T| \left|T^c\right| / r$.

So let $(\sigma, \sigma') \in T \times T^c$. We apply the path construction of [40, 41, Proposition 6.4.], but only to alternatives in $[k] \setminus \{a, b, c\}$.

The analysis of this construction is done in exactly the same way as in Lemma 2.21; in the end we get that there are at most $k^2 \left(k - 3\right)!$ paths that pass through a given edge in $\partial_e\left(T\right)$.

Recall that $\left|V_{(a,b,c)}\right| = \left(k - 3\right)!$ and that by our assumption $|T| \leq \left(1 - \delta\right)\left(k - 3\right)!$, so $\left|T^c\right| \geq \delta \left(k - 3\right)!$. Therefore

$$\left|\partial_e\left(T\right)\right| \geq \frac{|T| \left|T^c\right|}{k^2 \left(k - 3\right)!} \geq \frac{\delta}{k^2} |T|.$$

Now every vertex in $V_{(a,b,c)}$ has $k - 4 < k$ neighbors, which implies (2.33). $\square$

The next lemma tells us that if σ is on the boundary of a set of local dictators on $\{a, b, c\}$ for some alternative $c \notin \{a, b\}$, then there is a 2-manipulation point $\hat{\sigma}$ which is close to σ.

Lemma 2.28. *Suppose that $\sigma \in \partial\left(T\left(a, b, c\right)\right)$ for some $c \notin \{a, b\}$. Then there exists $\hat{\sigma} \in M_2$ which equals $z\sigma$ for some adjacent transposition z that does not involve a, b, or c, except that the order of the block of a, b, and c might be rearranged.*

Proof. Let π be the ranking profile such that $(\sigma, \pi) \in \partial_e\left(T\left(a, b, c\right)\right)$, and let z be the adjacent transposition in which they differ, i.e., $\pi = z\sigma$. Since $\pi \notin T\left(a, b, c\right)$, there exists a reordering of the block of a, b, and c at the top of π such that the outcome of f is not the top ranked alternative. Call the resulting vector π'. W.l.o.g. let us assume that $\pi'\left(1\right) = a$. Let us also define $\sigma' := z\pi'$. Now π' is a 2-manipulation point, since $f\left(\sigma'\right) = a$. $\square$

The next corollary puts together Corollary 2.26 and Lemmas 2.27 and 2.28.

Corollary 2.29. *Suppose that* (2.30) *holds. Then if for every $c \notin \{a,b\}$ we have*

$$\mathbb{P}^{(a,b,c)}\left(\sigma \in T\left(a,b,c\right)\right) \le 1 - \frac{\varepsilon}{100k},$$

then

$$\mathbb{P}\left(\sigma \in M_2\right) \ge \frac{\varepsilon^3}{10^5 k^{16}}.$$

Proof. We know that (2.30) implies

$$\sum_{c \notin \{a,b\}} \mathbb{P}^{a,b,c}\left(\sigma \in T\left(a,b,c\right)\right) \ge \frac{\varepsilon^2}{144k^8}.$$

Now using the assumptions, Lemma 2.27 with $\delta = \frac{\varepsilon}{100k}$, and Lemma 2.28, we have

$$\begin{aligned}\mathbb{P}\left(\sigma \in M_2\right) &\ge \sum_{c \neq \{a,b\}} \frac{1}{k^3}\mathbb{P}^{(a,b,c)}\left(\sigma \in M_2\right) \ge \sum_{c \notin \{a,b\}} \frac{1}{6k^4}\mathbb{P}^{(a,b,c)}\left(\sigma \in \partial\left(T\left(a,b,c\right)\right)\right)\\ &\ge \sum_{c \notin \{a,b\}} \frac{\varepsilon}{600k^8}\mathbb{P}^{(a,b,c)}\left(\sigma \in T\left(a,b,c\right)\right) \ge \frac{\varepsilon^3}{86400k^{16}} \ge \frac{\varepsilon^3}{10^5 k^{16}}.\end{aligned}$$

□

So again we are left with one case to deal with: if there exists an alternative $c \notin \{a,b\}$ such that $\mathbb{P}^{(a,b,c)}\left(\sigma \in T\left(a,b,c\right)\right) > 1 - \frac{\varepsilon}{100k}$. Define a subset of alternatives $K \subseteq [k]$ in the following way:

$$K := \{a,b\} \cup \left\{c \in [k] \setminus \{a,b\} : \mathbb{P}^{(a,b,c)}\left(\sigma \in T\left(a,b,c\right)\right) > 1 - \frac{\varepsilon}{100k}\right\}.$$

In addition to a and b, K contains those alternatives that whenever they are at the top with a and b, they form a local dictator with high probability.

So our assumption now is that $|K| \ge 3$.

Our next step is to show that unless we have many manipulation points, for any alternative $c \in K$, conditioned on c being at the top, the outcome of f is c with probability close to 1.

Lemma 2.30. *Let $c \in K$. Then either*

$$\mathbb{P}^{(c)}\left(f\left(\sigma\right) = c\right) \ge 1 - \frac{\varepsilon}{50k}, \tag{2.34}$$

or

$$\mathbb{P}\left(\sigma \in M_2\right) \ge \frac{\varepsilon}{100k^4}. \tag{2.35}$$

Proof. First assume that $c \notin \{a, b\}$.

Let σ be uniform according to $\mathbb{P}^{(c)}$, i.e., uniform on S_k conditioned on $\sigma(1) = c$. Define σ', where σ' is constructed from σ by first bubbling up alternative a to just below c, using adjacent transpositions, and then bubbling up b to just below a. Clearly σ' is distributed according to $\mathbb{P}^{(c,a,b)}$, i.e., it is uniform on S_k conditioned on $(\sigma(1), \sigma(2), \sigma(3)) = (c, a, b)$.

Since $c \in K$, we know that $\mathbb{P}^{(c,a,b)}\left(\sigma \in \mathrm{LD}^{\{a,b,c\}}\right) > 1 - \frac{\varepsilon}{100k}$. This also means that

$$\mathbb{P}^{(c)}\left(\sigma' \in \mathrm{LD}^{\{a,b,c\}}\right) > 1 - \frac{\varepsilon}{100k}.$$

Now we can partition the ranking profiles into three parts, based on the outcome of the SCF f at σ and σ':

$$\begin{aligned}
I_1 &= \{\sigma : f(\sigma) = c, f(\sigma') = c\}, \\
I_2 &= \{\sigma : f(\sigma) \neq c, f(\sigma') = c\}, \\
I_3 &= \{\sigma : f(\sigma') \neq c\}.
\end{aligned}$$

If $\mathbb{P}^{(c)}(I_1) \geq 1 - \frac{\varepsilon}{50k}$, then (2.34) holds. Otherwise we have $\mathbb{P}^{(c)}(I_2 \cup I_3) \geq \frac{\varepsilon}{50k}$, and since $\mathbb{P}^{(c)}(I_3) \leq \frac{\varepsilon}{100k}$, we have $\mathbb{P}^{(c)}(I_2) \geq \frac{\varepsilon}{100k}$.

Now if $\sigma \in I_2$, then we know that there is a 2-manipulation point along the way as we go from σ to σ'. That is, to every $\sigma \in I_2$ there exists $\hat{\sigma} \in M_2$ such that $\hat{\sigma}$ is equal to σ except perhaps a and b are shifted arbitrarily. So there can be at most k^2 ranking profiles σ giving the same 2-manipulation point $\hat{\sigma}$, and so we have

$$\mathbb{P}(\sigma \in M_2) \geq \frac{1}{k}\mathbb{P}^{(c)}(\sigma \in M_2) \geq \frac{1}{k^3}\mathbb{P}^{(c)}(I_2) \geq \frac{\varepsilon}{100k^4},$$

showing (2.35).

Now suppose that $c \in \{a, b\}$, w.l.o.g. assume $c = a$. We know that $|K| \geq 3$ and so there exists an alternative $d \in K \setminus \{a, b\}$. We can then do the same thing as above, but we now bubble up b and d. □

We now deal with alternatives that are not in K: either we have many manipulation points, or for any alternative $d \notin K$, the outcome of f is *not* d with probability close to 1.

Lemma 2.31. *Let $d \notin K$. If $\mathbb{P}(f(\sigma) = d) \geq \frac{\varepsilon}{4k}$, then*

$$\mathbb{P}(\sigma \in M_2) \geq \frac{\varepsilon^2}{10^6 k^9}.$$

Proof. Let σ be such that $f(\sigma) = d$. Bubble up d to the top, and call this ranking profile σ'. Now if $f(\sigma') \neq d$, then we know that there exists a 2-manipulation point $\hat{\sigma}$ along the way, i.e., a 2-manipulation $\hat{\sigma}$ which agrees with σ except perhaps d is shifted arbitrarily. Consequently, either

$$\mathbb{P}(\sigma \in M_2) \geq \frac{\varepsilon}{8k^2},$$

in which case we are done, or

$$\mathbb{P}\left(\sigma : f\left(\sigma\right) = f\left(\sigma'\right) = d\right) \geq \frac{\varepsilon}{8k}.$$

Next, let us bubble up a to just below d, and then bubble up b to just below d. Denote this ranking profile by $\sigma^{(d,b,a)}$, and analogously define $\sigma^{(d,a,b)}$, $\sigma^{(a,b,d)}$, $\sigma^{(a,d,b)}$, $\sigma^{(b,a,d)}$, and $\sigma^{(b,d,a)}$. Either we encounter a 2-manipulation point $\hat{\sigma}$ along the way of bubbling up to $\sigma^{(d,b,a)}$ ($\hat{\sigma}$ agrees with σ except d is at the top, and a and b might be arbitrarily shifted), or the outcome of the SCF f is d all along. So we have that either

$$\mathbb{P}\left(\sigma \in M_2\right) \geq \frac{\varepsilon}{16k^3},$$

in which case we are done, or

$$\mathbb{P}\left(\sigma : f\left(\sigma\right) = f\left(\sigma'\right) = f\left(\sigma^{(d,b,a)}\right) = f\left(\sigma^{(d,a,b)}\right) = d\right) \geq \frac{\varepsilon}{16k}.$$

Now start from $\sigma^{(d,a,b)}$. First swap a and d to get $\sigma^{(a,d,b)}$, then swap d and b to get $\sigma^{(a,b,d)}$, and finally bubble d and b down to their original positions in σ, except for the fact that a is now at the top of the coordinate. Call this profile $\bar{\sigma}$. Since σ is uniformly distributed, $\bar{\sigma}$ is distributed according to $\mathbb{P}_1^{(a)}$, i.e., uniformly conditional on $\bar{\sigma}\left(1\right) = a$. Now note that one of the following three events has to happen. (These events are not mutually exclusive.)

$$\begin{aligned}
I_1 &= \left\{f\left(\sigma^{(a,d,b)}\right) = f\left(\sigma^{(a,b,d)}\right) = a\right\},\\
I_2 &= \left\{f\left(\bar{\sigma}\right) \neq a\right\},\\
I_3 &= \{\sigma : \exists\ \hat{\sigma} \in M_2 \text{ which is equal to } \sigma \text{ except } a \text{ is shifted}\\
&\qquad\qquad \text{to the top, and } b \text{ and } d \text{ may be shifted arbitrarily}\}.
\end{aligned}$$

Since $a \in K$, we know by Lemma 2.30 that (unless we already have enough manipulation points by the lemma) we must have

$$\mathbb{P}\left(f\left(\bar{\sigma}\right) \neq a\right) = \mathbb{P}^{(a)}\left(f\left(\bar{\sigma}\right) \neq a\right) \leq \frac{\varepsilon}{50k}.$$

Consequently

$$\mathbb{P}\left(I_1 \cup I_3, f\left(\sigma\right) = f\left(\sigma'\right) = f\left(\sigma^{(d,b,a)}\right) = f\left(\sigma^{(d,a,b)}\right) = d\right) \geq \frac{\varepsilon}{16k} - \frac{\varepsilon}{50k} = \frac{17\varepsilon}{400k},$$

and so either

$$\mathbb{P}\left(\sigma \in M_2\right) \geq \frac{17\varepsilon}{800k^3},$$

in which case we are done, or

$$\mathbb{P}\left(\sigma : f\left(\sigma^{(d,b,a)}\right) = f\left(\sigma^{(d,a,b)}\right) = d, f\left(\sigma^{(a,b,d)}\right) = f\left(\sigma^{(a,d,b)}\right) = a\right) \geq \frac{17\varepsilon}{800k}.$$

Next, we can do the same thing with b on top, and we ultimately get that either

$$\mathbb{P}\left(\sigma \in M_2\right) \geq \frac{\varepsilon}{1600k^3},$$

in which case we are done, or

$$\mathbb{P}^{(a,b,d)}\left(\sigma^{(a,b,d)} \in \mathrm{LD}^{\{a,b,d\}}\right) = \mathbb{P}\left(\sigma : \sigma^{(a,b,d)} \in \mathrm{LD}^{\{a,b,d\}}\right) \geq \frac{\varepsilon}{1600k}. \tag{2.36}$$

Define $G_{(a,b,d)}$ and $T_{(a,b,d)}$ analogously to $G_{(a,b,c)}$ and $T_{(a,b,c)}$, respectively.

Suppose that (2.36) holds. We also know that $d \notin K$, so Lemma 2.27 applies, and then Lemma 2.28 shows us how to find manipulation points. We can put these arguments together, just like in the proof of Corollary 2.29, to show what we need:

$$\begin{aligned}\mathbb{P}\left(\sigma \in M_2\right) &\geq \frac{1}{k^3}\mathbb{P}^{(a,b,d)}\left(\sigma \in M_2\right) \geq \frac{1}{6k^4}\mathbb{P}^{(a,b,d)}\left(\sigma \in \partial\left(T\left(a,b,d\right)\right)\right) \\ &\geq \frac{\varepsilon}{600k^8}\mathbb{P}^{(a,b,d)}\left(\sigma \in T\left(a,b,d\right)\right) \geq \frac{\varepsilon^2}{10^6k^9}.\end{aligned}$$

□

Putting together the results of the previous lemmas, there is only one case to be covered, which is covered by the following final lemma. Basically, this lemma says that unless there are enough manipulation points, our function is close to a dictator on the subset of alternatives K.

Lemma 2.32. *Recall that we assume that* $\mathbf{D}\left(f, \mathrm{NONMANIP}\right) \geq \varepsilon$. *Furthermore assume that* $|K| \geq 3$, *that for every* $c \in K$ *we have*

$$\mathbb{P}^{(c)}\left(f\left(\sigma\right) = c\right) \geq 1 - \frac{\varepsilon}{50k}, \tag{2.37}$$

and that for every $d \notin K$ *we have*

$$\mathbb{P}\left(f\left(\sigma\right) = d\right) \leq \frac{\varepsilon}{4k}.$$

Then

$$\mathbb{P}\left(\sigma \in M_2\right) \geq \frac{\varepsilon}{4k^2}. \tag{2.38}$$

Proof. First note that

$$\mathbb{P}\left(f\left(\sigma\right) \neq \mathrm{top}_K\left(\sigma\right)\right) = \mathbb{P}\left(f\left(\sigma\right) \notin K\right) + \mathbb{P}\left(f\left(\sigma\right) \neq \mathrm{top}_K\left(\sigma\right), f\left(\sigma\right) \in K\right).$$

We know that

$$\varepsilon \leq \mathbf{D}\left(f, \mathrm{NONMANIP}\right) \leq \mathbb{P}\left(f\left(\sigma\right) \neq \mathrm{top}_K\left(\sigma\right)\right)$$

and also that

$$\mathbb{P}\left(f\left(\sigma\right) \notin K\right) \leq \left(k - |K|\right)\frac{\varepsilon}{4k} \leq \frac{\varepsilon}{2},$$

which together imply that

$$\mathbb{P}\left(f\left(\sigma\right) \neq \operatorname{top}_K\left(\sigma\right), f\left(\sigma\right) \in K\right) \geq \frac{\varepsilon}{2}.$$

Let σ be such that $f\left(\sigma\right) \neq \operatorname{top}_K\left(\sigma\right)$ and $f\left(\sigma\right) \in K$. Now bubble $\operatorname{top}_K\left(\sigma\right)$ up to the top in σ, call this ranking profile $\bar{\sigma}$. Clearly then $\operatorname{top}_K\left(\bar{\sigma}\right) = \operatorname{top}_K\left(\sigma\right)$.

There are two cases to consider. If $f\left(\sigma\right) \neq f\left(\bar{\sigma}\right)$, then there is a 2-manipulation point along the way from σ to $\bar{\sigma}$, i.e., a 2-manipulation point $\hat{\sigma}$ such that $\hat{\sigma}$ agrees with σ except perhaps some alternative c is arbitrarily shifted. Otherwise $f\left(\sigma\right) = f\left(\bar{\sigma}\right)$, and so $f\left(\bar{\sigma}\right) \neq \operatorname{top}_K\left(\bar{\sigma}\right)$.

Consequently we have that either (2.38) holds, or that

$$\mathbb{P}\left(\sigma : f\left(\bar{\sigma}\right) \neq \operatorname{top}_K\left(\bar{\sigma}\right)\right) \geq \frac{\varepsilon}{4}. \tag{2.39}$$

By the construction of $\bar{\sigma}$, we know that $\bar{\sigma}$ is uniformly distributed conditional on $\bar{\sigma}\left(1\right) \in K$. Consequently, by (2.37), we have that

$$\mathbb{P}\left(\sigma : f\left(\bar{\sigma}\right) \neq \operatorname{top}_K\left(\bar{\sigma}\right)\right) \leq \frac{\varepsilon}{50k},$$

which contradicts with (2.39) since $\frac{\varepsilon}{50k} < \frac{\varepsilon}{4}$. □

This concludes the proof of the small fiber case.

2.6.4 Large fiber case

In this subsection we assume that (2.26) holds. We show that we either have a lot 2-manipulation points or we have a lot of local dictators on three alternatives.

Our first step towards this is the following lemma.

Lemma 2.33. *Suppose that* (2.26) *holds. Then*

$$\mathbb{P}^{(a,b)}\left(\sigma \in B\right) \geq 1 - \frac{\varepsilon}{4}. \tag{2.40}$$

Proof. Let $B^c = \bar{F} \setminus B$. Our assumption (2.26) implies that $\mathbb{P}\left(\sigma \in B^c \,\middle|\, \sigma \in \bar{F}\right) \leq \frac{\varepsilon}{4k}$, which means that $\left|B^c\right| \leq \frac{\varepsilon(k-1)!}{4k}$, and so

$$\mathbb{P}^{(a,b)}\left(\sigma \notin B\right) \leq \frac{\varepsilon\left(k-1\right)!}{4k\left(k-2\right)!} < \frac{\varepsilon}{4},$$

which is equivalent to (2.40). □

The next lemma (together with Section 2.6.3.1) concludes the proof in the large fiber case.

Lemma 2.34. *Suppose that* (2.26) *holds and recall that our SCF f satisfies the condition* $\mathbf{D}(f, \mathrm{NONMANIP}) \geq \varepsilon$. *Then either*

$$\mathbb{P}(\sigma \in M_2) \geq \frac{\varepsilon}{4k^2} \tag{2.41}$$

or

$$\mathbb{P}(\sigma \in \mathrm{LD}(a,b)) \geq \frac{\varepsilon}{4k^2}. \tag{2.42}$$

Proof. By Lemma 2.33 we know that (2.40) holds.

Let $\sigma \in S_k$ be uniform. Define σ' by being the same as σ except alternatives a and b are moved to the top of the coordinate: $\sigma'(1) = a$ and $\sigma'(2) = b$. Clearly σ' is distributed according to $\mathbb{P}^{(a,b)}(\cdot)$. Also define $\sigma'' = [a : b]\, \sigma'$.

We partition the set of ranking profiles S_k into three parts:

$$\begin{aligned}
I_1 &:= \left\{\sigma \in S_k : f(\sigma) = \mathrm{top}_{\{a,b\}}(\sigma), (f(\sigma'), f(\sigma'')) = (a,b)\right\},\\
I_2 &:= \left\{\sigma \in S_k : f(\sigma) \neq \mathrm{top}_{\{a,b\}}(\sigma), (f(\sigma'), f(\sigma'')) = (a,b)\right\},\\
I_3 &:= \left\{\sigma \in S_k : (f(\sigma'), f(\sigma'')) \neq (a,b)\right\}.
\end{aligned}$$

By (2.40) we know that $\mathbb{P}(\sigma \in I_3) \leq \frac{\varepsilon}{4}$. We also know that $\mathbb{P}(\sigma \in I_1) \leq 1 - \varepsilon$, since $\mathbf{D}(f, \mathrm{NONMANIP}) \geq \varepsilon$. Therefore we must have

$$\mathbb{P}(\sigma \in I_2) \geq \frac{3\varepsilon}{4} > \frac{\varepsilon}{2}.$$

Let us partition I_2 further, and write it as $I_2 = I_2' \cup \left(\cup_{c \notin \{a,b\}} I_{2,c}\right)$, where

$$I_2' := \left\{\sigma \in I_2 : f(\sigma) \neq \mathrm{top}_{\{a,b\}}(\sigma), f(\sigma) \in \{a,b\}\right\}$$

and for any $c \notin \{a,b\}$,

$$I_{2,c} := \{\sigma \in I_2 : f(\sigma) = c\}.$$

Suppose that $\sigma \in I_2'$. W.l.o.g. let us assume that a is ranked higher than b by σ, and therefore $f(\sigma) = b$, since $\sigma \in I_2'$. Then we can get from σ to σ' by first bubbling up a to the top, and then bubbling up b to just below a. Since $f(\sigma) = b$ and $f(\sigma') = a$, there must be a 2-manipulation point $\hat{\sigma}$ along the way, which is equal to σ except perhaps the positions of a and b are arbitrarily shifted.

Now suppose that $\sigma \in I_{2,c}$ for some $c \notin \{a,b\}$. We distinguish two cases: either c is ranked above both a and b in σ, or it is not.

If not, then say a is ranked above c in σ. Bubble a all the way to the top, and then bubble b as well, all the way to the top, just below a. Since $f(\sigma) = c$ and $f(\sigma') = a$, there must be a 2-manipulation point $\hat{\sigma}$ along the way, which is equal to σ except perhaps the positions of a and b are arbitrarily shifted.

If c is ranked above both a and b in σ, then the argument is similar. First bubble up a and b to just below c, and denote this ranking profile by $\tilde{\sigma}$, then permute these three alternatives arbitrarily, and then bubble a and b to the top. It is not hard to think through that either there is a 2-manipulation $\hat{\sigma}$ along the way, which is then equal to σ except perhaps the positions of a and b are arbitrarily shifted, or else $\tilde{\sigma} \in \mathrm{LD}^{\{a,b,c\}}$.

Combining these cases we see that either (2.41) or (2.42) must hold. □

So if (2.41) holds then we are done, and if (2.42) holds, then we refer back to Section 2.6.3.1, where we deal with the case of local dictators on three alternatives.

2.6.5 Proof of Theorem 2.4 concluded

Proof of Theorem 2.4. Our starting point is Lemma 2.20, which implies that (2.24) holds (unless we already have many 2-manipulation points). We then consider two cases, as indicated in Section 2.6.2.

We deal with the small fiber case, when (2.25) holds, in Section 2.6.3. First, we have that Lemma 2.21, Corollary 2.22, Lemma 2.23 and Corollary 2.24 show that either there are many 3-manipulation points, or there are many local dictators on three alternatives. We then deal with the case of many local dictators in Section 2.6.3.1. Lemma 2.25, Corollary 2.26, Lemmas 2.27 and 2.28, Corollary 2.29, and Lemmas 2.30, 2.31 and 2.32 together show that there are many 3-manipulation points if there are many local dictators on three alternatives, and the SCF is ε-far from the family of nonmanipulable functions.

We deal with the large fiber case—when (2.26) holds—in Section 2.6.4. Here Lemma 2.34 shows that either we have many 2-manipulation points, or we have many local dictators on three alternatives. In this latter case we refer back to Section 2.6.3.1 to conclude the proof. □

2.7 Inverse polynomial manipulability for any number of alternatives

In this section we prove the theorem below, which is the same as our main theorem, Theorem 2.2, except that the condition of $\mathbf{D}(f, \mathrm{NONMANIP}) \geq \varepsilon$ from Theorem 2.2 is replaced with the stronger condition $\mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \geq \varepsilon$.

Theorem 2.35. *Suppose that we have $n \geq 2$ voters, $k \geq 3$ alternatives, and a SCF $f : S_k^n \to [k]$ satisfying $\mathbf{D}\left(f, \overline{\mathrm{NONMANIP}}\right) \geq \varepsilon$. Then*

$$\mathbb{P}(\sigma \in M(f)) \geq \mathbb{P}(\sigma \in M_4(f)) \geq p\left(\varepsilon, \frac{1}{n}, \frac{1}{k}\right), \tag{2.43}$$

for some polynomial p, where $\sigma \in S_k^n$ is selected uniformly. In particular, we show a lower bound of $\frac{\varepsilon^5}{10^9 n^7 k^{46}}$.

An immediate consequence is that

$$\mathbb{P}((\sigma, \sigma') \text{ is a manipulation pair for } f) \geq q\left(\varepsilon, \frac{1}{n}, \frac{1}{k}\right),$$

for some polynomial q, where $\sigma \in S_k^n$ is uniformly selected, and σ' is obtained from σ by uniformly selecting a coordinate $i \in \{1, \ldots, n\}$, uniformly selecting $j \in \{1, \ldots, k-3\}$, and then uniformly randomly permuting the following four adjacent alternatives in σ_i: $\sigma_i(j)$, $\sigma_i(j+1)$, $\sigma_i(j+2)$, and $\sigma_i(j+3)$. In particular, the specific lower bound for $\mathbb{P}(\sigma \in M_4(f))$ implies that we can take $q\left(\varepsilon, \frac{1}{n}, \frac{1}{k}\right) = \frac{\varepsilon^5}{10^{11} n^8 k^{47}}$.

For the remainder of the section, let us fix the number of voters $n \geq 2$, the number of alternatives $k \geq 3$, and the SCF f, which satisfies $\mathbf{D}\left(f, \overline{\text{NONMANIP}}\right) \geq \varepsilon$. Accordingly, we typically omit the dependence of various sets (e.g., boundaries between two alternatives) on f.

2.7.1 Division into cases

Our starting point in proving Theorem 2.35 is Lemma 2.19. Clearly if (2.20) holds then we are done, so in the rest of Section 2.7 we assume that this is not the case. Then Lemma 2.19 tells us that (2.21) holds, and w.l.o.g. we may assume that the two boundaries that the lemma gives us have $i = 1$ and $j = 2$. That is, we have

$$\mathbb{P}\left(\sigma \text{ on } B_1^{a,b;[a:b]}\right) \geq \frac{4\varepsilon}{nk^7} \quad \text{and} \quad \mathbb{P}\left(\sigma \text{ on } B_2^{c,d;[c:d]}\right) \geq \frac{4\varepsilon}{nk^7},$$

where recall that σ is on $B_1^{a,b;[a:b]}$ if $f(\sigma) = a$ and $f([a:b]_1 \sigma) = b$. If σ is on $B_1^{a,b;[a:b]}$ and $b \overset{\sigma_1}{>} a$, then σ is a 2-manipulation point, so if this happens in more than half of the cases when σ is on $B_1^{a,b;[a:b]}$, then we have

$$\mathbb{P}(\sigma \in M_2) \geq \frac{2\varepsilon}{nk^7},$$

and we are done. Similarly in the case of the boundary between c and d in coordinate 2. So we may assume from now on that

$$\mathbb{P}\left(\sigma \in \cup_{z_{-1}^{a,b}} B_1\left(z_{-1}^{a,b}\right)\right) \geq \frac{2\varepsilon}{nk^7} \quad \text{and} \quad \mathbb{P}\left(\sigma \in \cup_{z_{-2}^{c,d}} B_2\left(z_{-2}^{c,d}\right)\right) \geq \frac{2\varepsilon}{nk^7}.$$

The following lemma is an immediate corollary.

Lemma 2.36. *Either*

$$\mathbb{P}\left(\sigma \in \text{Sm}\left(B_1^{a,b;[a:b]}\right)\right) \geq \frac{\varepsilon}{nk^7} \tag{2.44}$$

or

$$\mathbb{P}\left(\sigma \in \text{Lg}\left(B_1^{a,b;[a:b]}\right)\right) \geq \frac{\varepsilon}{nk^7}, \tag{2.45}$$

and the same can be said for the boundary $B_2^{c,d;[c:d]}$.

We distinguish cases based upon this: either (2.44) holds, or (2.44) holds for the boundary $B_2^{c,d;[c:d]}$, or (2.45) holds for both boundaries. We only need one boundary for the small fiber case, and we need both boundaries only in the large fiber case. So in the large fiber case we must differentiate between two cases: whether $d \in \{a, b\}$ or $d \notin \{a, b\}$. First of all, in the $d \notin \{a, b\}$ case the problem of finding a manipulation point with not too small (i.e., inverse polynomial in n, k and ε^{-1}) probability has already been solved by Isaksson et al. [40, 41], so we are primarily interested in the $d \in \{a, b\}$ case. But moreover, we will see that our method of proof works in both cases.

In the rest of the section we first deal with the small fiber case, and then with the large fiber case.

2.7.2 Small fiber case

We now deal with the case when (2.44) holds. We formalize the ideas of the outline in a series of statements.

First, we want to formalize that the boundaries of the boundaries are big in this refined graph setting as well, when we are on a small fiber. The proof uses the canonical path method, as successfully adapted to this setting by Isaksson, Kindler and Mossel [40, 41], and is very similar to the proof of Lemma 2.21, with some necessary modifications due to the fact that we now have n coordinates.

Lemma 2.37. *Fix a coordinate and a pair of alternatives—for simplicity we choose coordinate 1 and alternatives a and b, but we note that this lemma holds in general, we do not assume anything special about these choices. Let $z_{-1}^{a,b}$ be such that $B_1\left(z_{-1}^{a,b}\right)$ is a small fiber for $B_1^{a,b;[a:b]}$. Then, writing $B \equiv B_1\left(z_{-1}^{a,b}\right)$ for simplicity, we have*

$$\mathbb{P}\left(\sigma \in \partial\left(B\right)\right) \geq \frac{\gamma}{2nk^5}\mathbb{P}\left(\sigma \in B\right). \tag{2.46}$$

Proof. Let $B^c = \bar{F}\left(z_{-1}^{a,b}\right) \setminus B$. For every $(\sigma, \sigma') \in B \times B^c$, we define a canonical path from σ to σ', which has to pass through at least one edge in $\partial_e\left(B\right)$. Then if we show that every edge in $\partial_e\left(B\right)$ lies on at most r canonical paths, then it follows that $\left|\partial_e\left(B\right)\right| \geq \left|B\right| \left|B^c\right| / r$.

So let $(\sigma, \sigma') \in B \times B^c$. We define a path from σ to σ' by applying a path construction in each coordinate one by one, and then concatenating these paths: first in the first coordinate we get from σ_1 to σ_1', while leaving all other coordinates unchanged, then in the second coordinate we get from σ_2 to σ_2', while leaving all other coordinates unchanged, and so on, finally in the last coordinate we get from σ_n to σ_n'. In the first coordinate we apply the path construction of [40, 41, Proposition 6.4.], but considering the block formed by a and b as a single element; in all other coordinates we apply the path construction of [40, 41, Proposition 6.6.]. Since this path goes from σ (which is in B) to σ' (which is in B^c), it must pass through at least one edge in $\partial_e\left(B\right)$.

For a given edge $(\pi, \pi') \in \partial_e(B)$, at most how many possible $(\sigma, \sigma') \in B \times B^c$ pairs are there such that the canonical path between σ and σ' defined above passes through (π, π')? Let us differentiate two cases.

Suppose that π and π' differ in the first coordinate. Then coordinates 2 through n of σ must agree with the respective coordinates of π, while coordinates 2 through n of σ' can be anything (up to the restriction given by $\sigma' \in B^c \subseteq \bar{F}\left(z_{-1}^{a,b}\right)$), giving $\left(\frac{k!}{2}\right)^{n-1}$ possibilities. Now fixing all coordinates except the first, [40, 41, Proposition 6.4.] tells us that there are at most $(k-1)^2 (k-1)!/2 < k^2 (k-1)!$ possibilities for the pair (σ_1, σ_1'). So altogether there are at most $k^2 (k-1)! \left(\frac{k!}{2}\right)^{n-1}$ paths that pass through a given edge in $\partial_e(B)$ in this case.

Suppose now that π and π' differ in the i^{th} coordinate, where $i \neq 1$. Then the first $i-1$ coordinates of σ' must agree with the first $i-1$ coordinates of π, while coordinates $i+1, \dots, n$ of σ must agree with the respective coordinates of π. The first $i-1$ coordinates of σ, and coordinates $i+1, \dots, n$ of σ' can be anything (up to the restriction given by $\sigma, \sigma' \in \bar{F}\left(z_{-1}^{a,b}\right)$), giving $(k-1)! \left(\frac{k!}{2}\right)^{n-2}$ possibilities. Now fixing all coordinates except the i^{th} coordinate, [40, 41, Proposition 6.6.] tells us that there are at most $k^4 k!$ possibilities for the pair (σ_i, σ_i'). So altogether there are at most $2k^4 (k-1)! \left(\frac{k!}{2}\right)^{n-1}$ paths that pass through a given edge in $\partial_e(B)$ in this case.

So in any case, there are at most $2k^4 (k-1)! \left(\frac{k!}{2}\right)^{n-1}$ paths that pass through a given edge in $\partial_e(B)$.

Recall that $\left|\bar{F}\left(z_{-1}^{a,b}\right)\right| = (k-1)! \left(\frac{k!}{2}\right)^{n-1}$, and also $|B^c| \geq \gamma (k-1)! \left(\frac{k!}{2}\right)^{n-1}$ since B is a small fiber. Therefore

$$|\partial_e(B)| \geq \frac{|B| |B^c|}{2k^4 (k-1)! \left(\frac{k!}{2}\right)^{n-1}} \geq \frac{\gamma}{2k^4} |B|.$$

Now in $G\left(z_{-1}^{a,b}\right)$ every ranking profile has no more than nk neighbors, which implies (2.46). $\square$

Corollary 2.38. *If* (2.44) *holds, then*

$$\mathbb{P}\left(\sigma \in \bigcup_{z_{-1}^{a,b}} \partial\left(B_1\left(z_{-1}^{a,b}\right)\right)\right) \geq \frac{\gamma\varepsilon}{2n^2k^{12}}.$$

Proof. Using the previous lemma and (2.44) we have

$$\begin{aligned}\mathbb{P}\left(\sigma \in \bigcup_{z_{-1}^{a,b}} \partial\left(B_1\left(z_{-1}^{a,b}\right)\right)\right) &= \sum_{z_{-1}^{a,b}} \mathbb{P}\left(\sigma \in \partial\left(B_1\left(z_{-1}^{a,b}\right)\right)\right) \\ &\geq \sum_{z_{-1}^{a,b}: B_1\left(z_{-1}^{a,b}\right) \subseteq \mathrm{Sm}\left(B_1^{a,b;[a:b]}\right)} \mathbb{P}\left(\sigma \in \partial\left(B_1\left(z_{-1}^{a,b}\right)\right)\right)\end{aligned}$$

$$\geq \sum_{z_{-1}^{a,b}: B_1\left(z_{-1}^{a,b}\right) \subseteq \mathrm{Sm}\left(B_1^{a,b;[a:b]}\right)} \frac{\gamma}{2nk^5} \mathbb{P}\left(\sigma \in B_1\left(z^{a,b}\right)\right)$$
$$= \frac{\gamma}{2nk^5} \mathbb{P}\left(\sigma \in \mathrm{Sm}\left(B_1^{a,b}\right)\right) \geq \frac{\gamma\varepsilon}{2n^2k^{12}}. \qquad \square$$

Next, we want to find manipulation points on the boundaries of boundaries.

Before we do this, let us divide the boundaries of the boundaries according to which direction they are in. If $\sigma \in \partial\left(B_1\left(z_{-1}^{a,b}\right)\right)$ for some $z_{-1}^{a,b}$, then we know that there exists a ranking profile π such that $(\sigma,\pi) \in \partial_e\left(B_1\left(z_{-1}^{a,b}\right)\right)$. We know that σ and π differ in exactly one coordinate, say coordinate j; in this case we say that σ is on the boundary of $B_1\left(z_{-1}^{a,b}\right)$ in direction j, and we write $\sigma \in \partial_j\left(B_1\left(z_{-1}^{a,b}\right)\right)$. (This notation should not be confused with that of the edge boundary.)

We can write the boundary of $B_1\left(z_{-1}^{a,b}\right)$ as a union of boundaries in the different directions:

$$\partial\left(B_1\left(z_{-1}^{a,b}\right)\right) = \cup_{j=1}^{n} \partial_j\left(B_1\left(z_{-1}^{a,b}\right)\right),$$

but note that this is not (necessarily) a disjoint union, as a ranking profile σ for which $\sigma \in \partial\left(B_1\left(z_{-1}^{a,b}\right)\right)$ might lie on the boundary in multiple directions.

In particular, we differentiate between the boundary in direction 1 and the boundary in all other directions. To this end we introduce the notation

$$\partial_{-1}\left(B_1\left(x_{-1}^{a,b}\right)\right) := \cup_{j=2}^{n} \partial_j\left(B_1\left(x_{-1}^{a,b}\right)\right).$$

With this notation we have the following corollary of Corollary 2.38.

Corollary 2.39. *If* (2.44) *holds, then either*

$$\mathbb{P}\left(\sigma \in \cup_{z_{-1}^{a,b}} \partial_{-1}\left(B_1\left(z_{-1}^{a,b}\right)\right)\right) \geq \frac{\gamma\varepsilon}{4n^2k^{12}} \tag{2.47}$$

or

$$\mathbb{P}\left(\sigma \in \cup_{z_{-1}^{a,b}} \partial_{1}\left(B_1\left(z_{-1}^{a,b}\right)\right)\right) \geq \frac{\gamma\varepsilon}{4n^2k^{12}}. \tag{2.48}$$

Lemma 2.40. *Suppose that the ranking profile σ is on the boundary of a fiber for $B_1^{a,b;[a:b]}$ in direction $j \neq 1$, i.e.,*

$$\sigma \in \cup_{z_{-1}^{a,b}} \partial_{-1}\left(B_1\left(z_{-1}^{a,b}\right)\right).$$

Then there exists a 3-manipulation point $\hat{\sigma}$ which agrees with σ in all coordinates except perhaps coordinate 1 and some coordinate $j \neq 1$; furthermore $\hat{\sigma}_1$ is equal to σ_1 or $[a:b]\,\sigma_1$, except that the position of a third alternative c might be shifted arbitrarily, and $\hat{\sigma}_j$ is equal to σ_j or $z\sigma_j$ for some adjacent transposition $z \in T$, except the position of b might be shifted arbitrarily.

Proof. Suppose that $x_{-1}^{a,b}(\sigma) = z_{-1}^{a,b}$. Since $\sigma \in \partial\left(B_1\left(z_{-1}^{a,b}\right)\right) \subseteq B_1\left(z_{-1}^{a,b}\right)$, we know that $f(\sigma) = a$, and if $\sigma' = [a:b]_1 \sigma$, then $f(\sigma') = b$.

Let $\pi = (\pi_j, \sigma_{-j})$ denote the ranking profile such that $(\sigma, \pi) \in \partial_e\left(B_1\left(z_{-1}^{a,b}\right)\right)$. Let $\pi' := [a:b]_1 \pi$. Since $\pi \notin B_1\left(z_{-1}^{a,b}\right)$, $(f(\pi), f(\pi')) \neq (a,b)$. Then, by Lemma 2.17, if $f(\pi) \neq f(\pi')$, then one of π and π' is a 2-manipulation point.

So let us suppose that $f(\pi) = f(\pi')$. If $f(\pi') = a$, then one of σ' and π' is a 2-manipulation point by Lemma 2.17, since $\pi' = z_j\sigma'$ for some adjacent transposition $z \neq [a:b]$. If $f(\pi) = b$, then similarly one of σ and π is a 2-manipulation point.

Finally, let us suppose that $f(\pi) = c$ for some $c \notin \{a,b\}$. In this case Lemma 2.18 tells us that there exists an appropriate 3-manipulation point $\hat{\sigma}$. □

Corollary 2.41. *If* (2.47) *holds, then*

$$\mathbb{P}(\sigma \in M_3) \geq \frac{\gamma\varepsilon}{8n^3k^{16}}. \tag{2.49}$$

Proof. Lemma 2.40 tells us that for every ranking profile σ which is on the boundary of a fiber for $B_1^{a,b;[a:b]}$ in some direction $j \neq 1$, there is a 3-manipulation point $\hat{\sigma}$ "nearby"; the lemma specifies what "nearby" means.

How many ranking profiles σ may give the same $\hat{\sigma}$? At most $2nk^4$, which comes from the following: σ and $\hat{\sigma}$ agree in all coordinates except maybe two, one of which is the first coordinate; there are $n-1 < n$ possibilities for the other coordinate; in the first coordinate, $\hat{\sigma}_1$ is either σ_1 or $[a:b]\,\sigma_1$ (giving 2 possibilities), while some alternative c ($k-2 < k$ possibilities) might be shifted arbitrarily (at most k possibilities); in the other coordinate $j \neq 1$, $\hat{\sigma}_j$ is equal to σ_j or $z\sigma_j$ for some adjacent transposition $z \in T$ (at most k possibilities), except b might be shifted arbitrarily (k possibilities).

So putting this result from Lemma 2.40 together with (2.47) yields (2.49). □

The remaining case we have to deal with is when (2.48) holds.

Lemma 2.42. *Suppose that the ranking profile σ is on the boundary of a fiber for $B_1^{a,b;[a:b]}$ in direction* 1*, i.e.,*

$$\sigma \in \cup_{z_{-1}^{a,b}} \partial_1\left(B_1\left(z_{-1}^{a,b}\right)\right).$$

Then either $\sigma \in \mathrm{LD}_1(a,b)$, or there exists a 3*-manipulation point $\hat{\sigma}$ which agrees with σ in all coordinates except perhaps in coordinate* 1*; furthermore $\hat{\sigma}_1$ is equal to σ_1, or $[a:b]\,\sigma_1$ except that the position of a third alternative c might be shifted arbitrarily.*

Proof. Just like the proof of Lemma 2.23. □

The following corollary then tells us that either we have found many 3-manipulation points, or we have many local dictators on three alternatives in coordinate 1.

Corollary 2.43. *Suppose that* (2.48) *holds. Then either*

$$\sum_{c\notin\{a,b\}} \mathbb{P}\left(\sigma \in \mathrm{LD}_1^{\{a,b,c\}}\right) = \mathbb{P}\left(\sigma \in \mathrm{LD}_1\left(a,b\right)\right) \geq \frac{\gamma\varepsilon}{8n^2k^{12}} \tag{2.50}$$

or

$$\mathbb{P}\left(\sigma \in M_3\right) \geq \frac{\gamma\varepsilon}{16n^2k^{14}}.$$

2.7.2.1 Dealing with local dictators

So the remaining case we have to deal with in this small fiber case is when (2.50) holds, i.e., we have many local dictators in coordinate 1.

Lemma 2.44. *Suppose that $\sigma \in \mathrm{LD}_1^{\{a,b,c\}}$ for some alternative $c \notin \{a,b\}$. Define $\sigma' := (\sigma_1', \sigma_{-1})$ by letting σ_1' be equal to σ_1 except that the block of a, b, and c is moved to the top of the coordinate. Then*

- *either $\sigma' \in \mathrm{LD}_1^{\{a,b,c\}}$,*
- *or there exists a 3-manipulation point $\hat{\sigma}$ which agrees with σ in all coordinates except perhaps in coordinate 1; furthermore $\hat{\sigma}_1$ is equal to σ_1 except that the position of a, b, and c might be shifted arbitrarily.*

Proof. Just like the proof of Lemma 2.25. □

Corollary 2.45. *If* (2.50) *holds, then either*

$$\sum_{c\notin\{a,b\}} \mathbb{P}\left(\sigma \in LD_1^{\{a,b,c\}}, \left\{\sigma_1\left(1\right), \sigma_1\left(2\right), \sigma_1\left(3\right)\right\} = \left\{a,b,c\right\}\right) \geq \frac{\gamma\varepsilon}{16n^2k^{13}} \tag{2.51}$$

or

$$\mathbb{P}\left(\sigma \in M_3\right) \geq \frac{\gamma\varepsilon}{16n^2k^{15}}.$$

Proof. Just like the proof of Corollary 2.26. □

Now (2.51) is equivalent to

$$\sum_{c\notin\{a,b\}} \mathbb{P}\left(\sigma \in LD_1^{\{a,b,c\}}, \left(\sigma_1\left(1\right), \sigma_1\left(2\right), \sigma_1\left(3\right)\right) = \left(a,b,c\right)\right) \geq \frac{\gamma\varepsilon}{96n^2k^{13}}. \tag{2.52}$$

We know that

$$\mathbb{P}\left(\left(\sigma_1\left(1\right), \sigma_1\left(2\right), \sigma_1\left(3\right)\right) = \left(a,b,c\right)\right) = \frac{1}{k\left(k-1\right)\left(k-2\right)} \leq \frac{6}{k^3},$$

and so (2.52) implies (recall Definition 2.21)

$$\sum_{c \notin \{a,b\}} \mathbb{P}_1^{(a,b,c)} \left(\sigma \in LD_1^{\{a,b,c\}} \right) \geq \frac{\gamma \varepsilon}{576 n^2 k^{10}}. \tag{2.53}$$

Now fix an alternative $c \notin \{a, b\}$ and define the graph $G_{(a,b,c)} = \left(V_{(a,b,c)}, E_{(a,b,c)}\right)$ to have vertex set

$$V_{(a,b,c)} := \{\sigma \in S_k^n : (\sigma_1(1), \sigma_1(2), \sigma_1(3)) = (a, b, c)\}$$

and for $\sigma, \pi \in V_{(a,b,c)}$ let $(\sigma, \pi) \in E_{(a,b,c)}$ if and only if σ and π differ in exactly one coordinate, and by an adjacent transposition in this coordinate. So $G_{(a,b,c)}$ is the subgraph of the refined rankings graph induced by the vertex set $V_{(a,b,c)}$.

Let

$$T_1(a, b, c) := V_{(a,b,c)} \cap LD_1^{\{a,b,c\}},$$

and let $\partial_e (T_1(a, b, c))$ and $\partial (T_1(a, b, c))$ denote the edge and vertex boundary of $T_1(a, b, c)$ in $G_{(a,b,c)}$, respectively.

The next lemma shows that unless $T_1(a, b, c)$ is almost all of $V_{(a,b,c)}$, the size of the boundary $\partial (T_1(a, b, c))$ is comparable to the size of $T_1(a, b, c)$.

Lemma 2.46. *Let $c \notin \{a, b\}$ be arbitrary. Write $T \equiv T_1(a, b, c)$ for simplicity. If*

$$\mathbb{P}_1^{(a,b,c)} (\sigma \in T) \leq 1 - \delta,$$

then

$$\mathbb{P}_1^{(a,b,c)} (\sigma \in \partial (T)) \geq \frac{\delta}{nk^3} \mathbb{P}_1^{(a,b,c)} (\sigma \in T). \tag{2.54}$$

Proof. The proof is essentially the same as the proof of Lemma 2.27, with a slight modification to deal with n coordinates. Let $T^c = V_{(a,b,c)} \setminus T(a, b, c)$. For every $(\sigma, \sigma') \in T \times T^c$ we define a canonical path from σ to σ' by applying a path construction in each coordinate one by one, and then concatenating these paths. In all coordinates we apply the path construction of [40, 41, Proposition 6.4.], but in the first coordinate we only apply it to alternatives in $[k] \setminus \{a, b, c\}$.

The analysis of this construction is done in exactly the same way as in Lemma 2.37; in the end we get that $|\partial_e (T)| \geq \frac{\delta}{k^2} |T|$. Now every vertex in $V_{(a,b,c)}$ has no more than nk neighbors, which implies (2.54). □

The next lemma tells us that if σ is on the boundary of a set of local dictators on $\{a, b, c\}$ for some alternative $c \notin \{a, b\}$ in coordinate 1, then there is a 4-manipulation point $\hat{\sigma}$ which is close to σ. The proof is similar to that of Lemma 2.28, but we have to take care of all n coordinates.

Lemma 2.47. *Suppose that $\sigma \in \partial (T_1(a, b, c))$ for some $c \notin \{a, b\}$. We distinguish two cases, based on the number of alternatives.*

If $k = 3$, then there exists a (3-)manipulation point $\hat{\sigma}$ which differs from σ in at most two coordinates, one of them being the first coordinate.

If $k \geq 4$, then there exists a 4-manipulation point $\hat{\sigma}$ which differs from σ in at most two coordinates, one of them being the first coordinate; furthermore, $\hat{\sigma}_1$ is equal to σ_1 except that the order of the block of a, b, and c might be rearranged and an additional alternative d might be shifted arbitrarily; and in the other coordinate, call it j, $\hat{\sigma}_j$ is equal to σ_j except perhaps a, b, and c are shifted arbitrarily.

Proof. Let π be the ranking profile such that $(\sigma, \pi) \in \partial_e (T_1 (a, b, c))$, let j be the coordinate in which they differ, and let z be the adjacent transposition in which they differ, i.e., $\pi = z_j \sigma$. Since $\pi \notin T_1 (a, b, c)$, there exists a reordering of the block of a, b, and c at the top of π_1 such that the outcome of f is not the top ranked alternative in coordinate 1. Call the resulting vector π_1', and let $\pi' := (\pi_1', \pi_{-1})$. W.l.o.g. let us assume that $\pi_1' (1) = a$. Let us also define $\sigma' := z_j \pi'$. We distinguish two cases: $j = 1$ or $j \neq 1$.

If $j = 1$ (in which case we must have $k \geq 5$), π' is a 2-manipulation point, since $f (\sigma') = a$.

If $j \neq 1$, then there are various cases to consider. If the adjacent transposition z does not move a, then either π' or σ' is a 2-manipulation point. So let us suppose that $z = [a : d]$ for some $d \neq a$.

Clearly we must have $f (\pi') = d$, or else π' or σ' is a 2-manipulation point. Suppose first that $d \in \{b, c\}$. W.l.o.g. suppose that $d = b$.

Then take alternative c in coordinate j of both σ' and π', and bubble it to the block of a and b simultaneously in the two ranking profiles. If along the way the value of the outcome of the SCF f changes from a or b, respectively, then we have a 2-manipulation point by Lemma 2.17. Otherwise, we now have a, b, and c adjacent in both coordinates 1 and j. Now rearranging the order of the blocks of a, b, and c in these two coordinates (which can be done using adjacent transpositions), we either get a 2-manipulation point by Lemma 2.17, or we can define a new SCF on two voters and three alternatives, a, b, and c. This SCF takes on three values and it is also not hard to see that the outcome is not only a function of the first coordinate, so by the Gibbard-Satterthwaite theorem we know that this SCF has a manipulation point, which is a 3-manipulation point of the original SCF f.

Now let us look at the case when $d \notin \{b, c\}$. In this case we do something similar to what we just did in the previous paragraph. In both σ' and π', first bubble up alternative d in coordinate 1 up to the block of a, b, and c, and then bubble b and c in coordinate j to the block of a and d. All of this using adjacent transpositions. If the value of the outcome of the SCF f changes from a or d, respectively, at any time along the way, then we have a 2-manipulation point by Lemma 2.17. Otherwise, we now have a, b, c, and d adjacent in both coordinates 1 and j, and we can apply the same trick to find a 4-manipulation point, using the Gibbard-Satterthwaite theorem. $\square$

The next corollary puts together Corollary 2.45 and Lemmas 2.46 and 2.47.

Corollary 2.48. *Suppose that* (2.51) *holds. Then if for every $c \notin \{a,b\}$ we have*

$$\mathbb{P}_1^{(a,b,c)}\left(\sigma \in T_1\left(a,b,c\right)\right) \leq 1 - \frac{\varepsilon}{100k},$$

then

$$\mathbb{P}\left(\sigma \in M_4\right) \geq \frac{\gamma\varepsilon^2}{345600 n^4 k^{22}}.$$

Proof. We know that (2.51) implies

$$\sum_{c \notin \{a,b\}} \mathbb{P}_1^{a,b,c}\left(\sigma \in T_1\left(a,b,c\right)\right) \geq \frac{\gamma\varepsilon}{576 n^2 k^{10}}.$$

Now then using the assumptions, Lemma 2.46 with $\delta = \frac{\varepsilon}{100k}$ and Lemma 2.47, we have

$$\begin{aligned}\mathbb{P}\left(\sigma \in M_4\right) &\geq \sum_{c \neq \{a,b\}} \frac{1}{k^3}\mathbb{P}_1^{(a,b,c)}\left(\sigma \in M_4\right) \geq \sum_{c \notin \{a,b\}} \frac{1}{6nk^8}\mathbb{P}_1^{(a,b,c)}\left(\sigma \in \partial\left(T_1\left(a,b,c\right)\right)\right) \\ &\geq \sum_{c \notin \{a,b\}} \frac{\varepsilon}{600 n^2 k^{12}}\mathbb{P}_1^{(a,b,c)}\left(\sigma \in T_1\left(a,b,c\right)\right) \geq \frac{\gamma\varepsilon^2}{345600 n^4 k^{22}}.\end{aligned}$$

□

So again we are left with one case to deal with: if there exists an alternative $c \notin \{a,b\}$ such that $\mathbb{P}_1^{(a,b,c)}\left(\sigma \in T_1\left(a,b,c\right)\right) > 1 - \frac{\varepsilon}{100k}$. Define a subset of alternatives $K \subseteq [k]$ in the following way:

$$K := \{a,b\} \cup \left\{c \in [k] \setminus \{a,b\} : \mathbb{P}_1^{(a,b,c)}\left(\sigma \in T_1\left(a,b,c\right)\right) > 1 - \frac{\varepsilon}{100k}\right\}.$$

In addition to a and b, K contains those alternatives that whenever they are at the top of coordinate 1 with a and b, they form a local dictator with high probability.

So our assumption now is that $|K| \geq 3$.

Our next step is to show that unless we have many manipulation points, for any alternative $c \in K$, conditioned on c being at the top of the first coordinate, the outcome of f is c with probability close to 1.

Lemma 2.49. *Let $c \in K$. Then either*

$$\mathbb{P}_1^{(c)}\left(f\left(\sigma\right) = c\right) \geq 1 - \frac{\varepsilon}{50k}, \tag{2.55}$$

or

$$\mathbb{P}\left(\sigma \in M_2\right) \geq \frac{\varepsilon}{100k^4}. \tag{2.56}$$

Proof. Just like the proof of Lemma 2.30. □

We now deal with alternatives that are not in K: either we have many manipulation points, or for any alternative $d \notin K$, the outcome of f is *not* d with probability close to 1.

Lemma 2.50. *Let $d \notin K$. If $\mathbb{P}(f(\sigma) = d) \geq \frac{\varepsilon}{4k}$, then*

$$\mathbb{P}(\sigma \in M_4) \geq \frac{\varepsilon^2}{10^6 n^2 k^{13}}.$$

Proof. The proof is very similar to that of Lemma 2.31: we do the same steps in the first coordinate as done in the proof of Lemma 2.31, and the fact that we have n coordinates only matters at the very end.

Let σ be such that $f(\sigma) = d$. We will keep coordinates 2 through n to be fixed as σ_{-1} throughout the proof. By bubbling alternatives d, a, and b in the first coordinate, we can define σ', $\sigma^{(d,b,a)}$, $\sigma^{(d,a,b)}$, $\sigma^{(a,b,d)}$, $\sigma^{(a,d,b)}$, $\sigma^{(b,a,d)}$, and $\sigma^{(b,d,a)}$ just as in the proof of Lemma 2.31. Again, we can show that either

$$\mathbb{P}(\sigma \in M_2) \geq \frac{\varepsilon}{1600k^3},$$

in which case we are done, or

$$\mathbb{P}_1^{(a,b,d)}\left(\sigma^{(a,b,d)} \in LD_1^{\{a,b,d\}}\right) = \mathbb{P}\left(\sigma : \sigma^{(a,b,d)} \in LD_1^{\{a,b,d\}}\right) \geq \frac{\varepsilon}{1600k}. \tag{2.57}$$

Define $G_{(a,b,d)}$ and $T_{(a,b,d)}$ analogously to $G_{(a,b,c)}$ and $T_{(a,b,c)}$, respectively.

Suppose that (2.57) holds. We also know that $d \notin K$, so Lemma 2.46 applies, and then Lemma 2.47 shows us how to find manipulation points. We can put these arguments together, just like in the proof of Corollary 2.48, to show what we need:

$$\begin{aligned}\mathbb{P}(\sigma \in M_4) &\geq \frac{1}{k^3}\mathbb{P}_1^{(a,b,d)}(\sigma \in M_4) \geq \frac{1}{6nk^8}\mathbb{P}_1^{(a,b,d)}(\sigma \in \partial(T_1(a,b,d))) \\ &\geq \frac{\varepsilon}{600n^2k^{12}}\mathbb{P}_1^{(a,b,d)}(\sigma \in T_1(a,b,d)) \geq \frac{\varepsilon^2}{10^6 n^2 k^{13}}.\end{aligned}$$

□

Putting together the results of the previous lemmas, there is only one case to be covered, which is covered by the following final lemma. Basically, this lemma says that unless there are enough manipulation points, our function is close to a dictator in the first coordinate, on the subset of alternatives K.

Lemma 2.51. *Recall that we assume that $\mathbf{D}\left(f, \overline{\text{NONMANIP}}\right) \geq \varepsilon$. Furthermore assume that $|K| \geq 3$, that for every $c \in K$ we have*

$$\mathbb{P}_1^{(c)}(f(\sigma) = c) \geq 1 - \frac{\varepsilon}{50k}, \tag{2.58}$$

and that for every $d \notin K$ we have

$$\mathbb{P}(f(\sigma) = d) \leq \frac{\varepsilon}{4k}.$$

Then

$$\mathbb{P}(\sigma \in M_2) \geq \frac{\varepsilon}{4k^2}. \tag{2.59}$$

Proof. Just like the proof of Lemma 2.32. □

To conclude the proof in the small fiber case, inspect all the lower bounds for $\mathbb{P}(\sigma \in M_4)$ obtained in Section 2.7.2, and recall that $\gamma = \frac{\varepsilon^3}{10^3 n^3 k^{24}}$.

2.7.3 Large fiber case

We now deal with the large fiber case, when (2.45) holds for both boundaries, i.e., when

$$\mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_1^{a,b;[a:b]}\right)\right) \geq \frac{\varepsilon}{nk^7}$$

and

$$\mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_2^{c,d;[c:d]}\right)\right) \geq \frac{\varepsilon}{nk^7}.$$

We differentiate between two cases: whether $d \in \{a, b\}$ or $d \notin \{a, b\}$.

2.7.3.1 Case 1

Suppose that $d \in \{a, b\}$, in which case w.l.o.g. we may assume that $d = a$. That is, in the rest of this case we may assume that

$$\mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_1^{a,b;[a:b]}\right)\right) \geq \frac{\varepsilon}{nk^7} \tag{2.60}$$

and

$$\mathbb{P}\left(\sigma \in \mathrm{Lg}\left(B_2^{a,c;[a:c]}\right)\right) \geq \frac{\varepsilon}{nk^7}. \tag{2.61}$$

First, let us look at only the boundary between a and b in direction 1. Let us fix a vector $z_{-1}^{a,b}$ which gives a large fiber $B_1\left(z_{-1}^{a,b}\right)$ for the boundary $B_1^{a,b;[a:b]}$, i.e., we know that

$$\mathbb{P}\left(\sigma \in B_1\left(z_{-1}^{a,b}\right) \,\middle|\, \sigma \in \bar{F}\left(z_{-1}^{a,b}\right)\right) \geq 1 - \gamma. \tag{2.62}$$

Our basic goal in the following will be to show that conditional on the ranking profile σ being in the fiber $F\left(z_{-1}^{a,b}\right)$ (but not necessarily in $\bar{F}\left(z_{-1}^{a,b}\right)$), with high probability the outcome of the vote is $\mathrm{top}_{\{a,b\}}\left(\sigma_1\right)$, or else we have a lot of 2-manipulation points or local dictators on three alternatives in coordinate 1.

Our first step towards this is the following.

Lemma 2.52. *Suppose that $z_{-1}^{a,b}$ gives a large fiber $B_1\left(z_{-1}^{a,b}\right)$ for the boundary $B_1^{a,b;[a:b]}$. Then*

$$\mathbb{P}_1^{(a,b)}\left(\sigma \in B_1\left(z_{-1}^{a,b}\right) \,\middle|\, \sigma \in F\left(z_{-1}^{a,b}\right)\right) \geq 1 - k\gamma. \tag{2.63}$$

Proof. We know that

$$\mathbb{P}\left(\left(\sigma_1\left(1\right), \sigma_1\left(2\right)\right) = \left(a, b\right) \,\middle|\, \sigma \in \bar{F}\left(z_{-1}^{a,b}\right)\right) = \frac{1}{k-1},$$

and so

$$\begin{aligned}\mathbb{P}_1^{(a,b)}\left(\sigma \notin B_1\left(z_{-1}^{a,b}\right) \,\middle|\, \sigma \in F\left(z_{-1}^{a,b}\right)\right) &= \mathbb{P}_1^{(a,b)}\left(\sigma \notin B_1\left(z_{-1}^{a,b}\right) \,\middle|\, \sigma \in \bar{F}\left(z_{-1}^{a,b}\right)\right) \\ &= \left(k-1\right)\mathbb{P}\left(\sigma \notin B_1\left(z_{-1}^{a,b}\right), \left(\sigma_1\left(1\right), \sigma_1\left(2\right)\right) = \left(a, b\right) \,\middle|\, \sigma \in \bar{F}\left(z_{-1}^{a,b}\right)\right) \leq \left(k-1\right)\gamma < k\gamma. \square\end{aligned}$$

The next lemma formalizes our goal mentioned above.

Lemma 2.53. *Suppose that $z_{-1}^{a,b}$ gives a large fiber $B_1\left(z_{-1}^{a,b}\right)$ for the boundary $B_1^{a,b;[a:b]}$. Then either*

$$\mathbb{P}\left(f(\sigma)=\mathrm{top}_{\{a,b\}}(\sigma_1)\,\middle|\,\sigma\in F\left(z_{-1}^{a,b}\right)\right)\geq 1-2k\gamma \tag{2.64}$$

or

$$\mathbb{P}\left(\sigma\in M_2\,\middle|\,\sigma\in F\left(z_{-1}^{a,b}\right)\right)\geq\frac{\gamma}{2k} \tag{2.65}$$

or

$$\mathbb{P}\left(\sigma\in LD_1(a,b)\,\middle|\,\sigma\in F\left(z_{-1}^{a,b}\right)\right)\geq\frac{\gamma}{2k}. \tag{2.66}$$

Proof. The proof of this lemma is essentially the same as that of Lemma 2.34, there are only two slight differences. First, we use Lemma 2.52 to know that (2.63) holds. Second, we take $\sigma\in F\left(z_{-1}^{a,b}\right)$ to be uniform, and we stay on the fiber $F\left(z_{-1}^{a,b}\right)$ throughout the proof: we modify only the first coordinate throughout the proof, in the same way as we did for Lemma 2.34. We omit the details. □

Now this lemma holds for all vectors $z_{-1}^{a,b}$ which give a large fiber $B_1\left(z_{-1}^{a,b}\right)$ for the boundary $B_1^{a,b;[a:b]}$. By (2.60) we know that

$$\mathbb{P}\left(\sigma : B_1\left(x_{-1}^{a,b}(\sigma)\right)\text{ is a large fiber}\right)\geq\frac{\varepsilon}{nk^7}.$$

Now if (2.65) holds for at least a third of the vectors $z_{-1}^{a,b}$ that give a large fiber $B_1\left(z_{-1}^{a,b}\right)$, then it follows that

$$\mathbb{P}(\sigma\in M_2)\geq\frac{\gamma\varepsilon}{6nk^8}$$

and we are done. If (2.66) holds for at least a third of the vectors $z_{-1}^{a,b}$ that give a large fiber $B_1\left(z_{-1}^{a,b}\right)$, then similarly we have

$$\mathbb{P}(\sigma\in LD_1(a,b))\geq\frac{\gamma\varepsilon}{6nk^8},$$

which means that (2.50) also holds, and so we are done by the argument in Section 2.7.2.1.

So the remaining case to consider is when (2.64) holds for at least a third of the vectors $z_{-1}^{a,b}$ that give a large fiber $B_1\left(z_{-1}^{a,b}\right)$.

We can go through this same argument for the boundary between a and c in direction 2 as well, and either we are done because

$$\mathbb{P}(\sigma\in M_2)\geq\frac{\gamma\varepsilon}{6nk^8}$$

or

$$\mathbb{P}(\sigma\in LD_2(a,c))\geq\frac{\gamma\varepsilon}{6nk^8},$$

or for at least a third of the vectors $z_{-2}^{a,c}$ that give a large fiber $B_2\left(z_{-2}^{a,c}\right)$ we have

$$\mathbb{P}\left(f(\sigma)=\operatorname{top}_{\{a,c\}}(\sigma_2)\,\middle|\,\sigma\in F\left(z_{-2}^{a,c}\right)\right)\geq 1-2k\gamma.$$

So basically our final case is if

$$\mathbb{P}\left(\sigma\in F_1^{a,b}\right)\geq\frac{\varepsilon}{3nk^7} \tag{2.67}$$

and also

$$\mathbb{P}\left(\sigma\in F_2^{a,c}\right)\geq\frac{\varepsilon}{3nk^7}. \tag{2.68}$$

Notice that being in the set $F_1^{a,b}$ only depends on the vector $x^{a,b}(\sigma)$ of preferences between a and b, and similarly being in the set $F_2^{a,c}$ only depends on the vector $x^{a,c}(\sigma)$ of preferences between a and c. We know that $\left\{\left(x_i^{a,b}(\sigma),x_i^{a,c}(\sigma)\right)\right\}_{i=1}^n$ are independent, and for any given i we know that $\left|\mathbb{E}\left(x_i^{a,b}(\sigma)\,x_i^{a,c}(\sigma)\right)\right|=\frac{1}{3}$. Hence we can apply reverse hypercontractivity (Lemma 2.8) to get the following result.

Lemma 2.54. *If* (2.67) *and* (2.68) *hold, then also*

$$\mathbb{P}\left(\sigma\in F_1^{a,b}\cap F_2^{a,c}\right)\geq\frac{\varepsilon^3}{27n^3k^{21}}. \tag{2.69}$$

Proof. See above. ∎

The next and final lemma then concludes that we have lots of manipulation points.

Lemma 2.55. *Suppose that* (2.69) *holds. Then*

$$\mathbb{P}\left(\sigma\in M_3\right)\geq\frac{\varepsilon^3}{54n^3k^{27}}-\frac{9\gamma}{k^3}. \tag{2.70}$$

Proof. First let us define two events:

$$\begin{aligned}I_1&:=\left\{\sigma:f(\sigma)=\operatorname{top}_{\{a,b\}}(\sigma_1)\right\},\\ I_2&:=\left\{\sigma:f(\sigma)=\operatorname{top}_{\{a,c\}}(\sigma_2)\right\}.\end{aligned}$$

Using similar estimates as previously in Lemma 2.11, we have

$$\begin{aligned}\mathbb{P}\left(\sigma\in I_1\cap I_2\cap F_1^{a,b}\cap F_2^{a,c}\right)&\geq\mathbb{P}\left(\sigma\in F_1^{a,b}\cap F_2^{a,c}\right)\\&\quad-\mathbb{P}\left(\sigma\notin I_1,\sigma\in F_1^{a,b}\cap F_2^{a,c}\right)-\mathbb{P}\left(\sigma\notin I_2,\sigma\in F_1^{a,b}\cap F_2^{a,c}\right).\end{aligned}$$

The first term is bounded below via (2.69), while the other two terms can be bounded using the definition of $F_1^{a,b}$ and $F_2^{a,c}$, respectively:

$$\mathbb{P}\left(\sigma\notin I_1,\sigma\in F_1^{a,b}\cap F_2^{a,c}\right)\leq\mathbb{P}\left(\sigma\notin I_1,\sigma\in F_1^{a,b}\right)\leq\mathbb{P}\left(\sigma\notin I_1\,\middle|\,\sigma\in F_1^{a,b}\right)\leq 2k\gamma,$$

and similarly for the other term. Putting everything together gives us

$$\mathbb{P}\left(\sigma \in I_1 \cap I_2 \cap F_1^{a,b} \cap F_2^{a,c}\right) \geq \frac{\varepsilon^3}{27n^3k^{21}} - 4k\gamma.$$

If $\sigma \in I_1 \cap I_2 \cap F_1^{a,b} \cap F_2^{a,c}$, then clearly we must have $f(\sigma) = a$, and therefore $x_1^{a,b}(\sigma) = 1$ and $x_2^{a,c}(\sigma) = 1$. Now define σ' from σ by bubbling up b in coordinate 1 to just below a, and bubbling up c in coordinate 2 to just below a. Either we encounter a 2-manipulation point along the way, or the outcome is still a: $f(\sigma') = a$. If we encounter a 2-manipulation point along the way for at least half of such ranking profiles, then we are done:

$$\mathbb{P}(\sigma \in M_2) \geq \frac{1}{k^2}\left(\frac{\varepsilon^3}{54n^3k^{21}} - 2k\gamma\right) = \frac{\varepsilon^3}{54n^3k^{23}} - \frac{2\gamma}{k}.$$

Otherwise, we may assume that

$$\mathbb{P}\left(\sigma \in I_1 \cap I_2 \cap F_1^{a,b} \cap F_2^{a,c}, f(\sigma') = a\right) \geq \frac{\varepsilon^3}{54n^3k^{21}} - 2k\gamma.$$

In this case define $\tilde{\sigma}' := [a:b]_1 \sigma'$ and $\tilde{\sigma}'' := [a:c]_2 \sigma'$. If $f(\tilde{\sigma}') \notin \{a,b\}$ or $f(\tilde{\sigma}'') \notin \{a,c\}$, then we automatically have that one of σ', $\tilde{\sigma}'$, or $\tilde{\sigma}''$ is a 2-manipulation point. If $f(\tilde{\sigma}') = b$ and $f(\tilde{\sigma}'') = c$, then by Lemma 2.18 we know that there exists a 3-manipulation point $\hat{\sigma}$ which agrees with σ except perhaps a, b, and c could be arbitrarily shifted in the first two coordinates. The final case is when $a \in \{f(\tilde{\sigma}'), f(\tilde{\sigma}'')\}$. But we now show that this has small probability, and therefore (2.70) follows.

First let us look at the case of $f(\tilde{\sigma}') = a$. We have

$$\begin{aligned}
&\mathbb{P}\left(\sigma \in I_1 \cap I_2 \cap F_1^{a,b} \cap F_2^{a,c}, f(\sigma') = a, f(\tilde{\sigma}') = a\right) \\
&\quad= \sum_{z_{-1}^{a,b} : F\left(z_{-1}^{a,b}\right) \subseteq F_1^{a,b}} \mathbb{P}\left(\sigma \in I_1 \cap I_2 \cap F\left(\left(1, z_{-1}^{a,b}\right)\right) \cap F_2^{a,c}, f(\sigma') = a, f(\tilde{\sigma}') = a\right) \\
&\quad= \sum_{z_{-1}^{a,b} : F\left(z_{-1}^{a,b}\right) \subseteq F_1^{a,b}} \mathbb{P}\left(\sigma \in I_1 \cap I_2 \cap F_2^{a,c}, f(\sigma') = a, f(\tilde{\sigma}') = a \,\middle|\, \sigma \in F\left(\left(1, z_{-1}^{a,b}\right)\right)\right) \times \\
&\qquad\qquad \times \mathbb{P}\left(\sigma \in F\left(\left(1, z_{-1}^{a,b}\right)\right)\right) \\
&\quad\leq \sum_{z_{-1}^{a,b} : F\left(z_{-1}^{a,b}\right) \subseteq F_1^{a,b}} \mathbb{P}\left(\sigma : f(\tilde{\sigma}') = a \,\middle|\, \sigma \in F\left(\left(1, z_{-1}^{a,b}\right)\right)\right) \mathbb{P}\left(\sigma \in F\left(\left(1, z_{-1}^{a,b}\right)\right)\right).
\end{aligned}$$

Now we know that $\tilde{\sigma}' \in F\left(\left(-1, z_{-1}^{a,b}\right)\right) \subseteq F_1^{a,b}$, and we also know that

$$\mathbb{P}\left(f(\sigma) \neq b \,\middle|\, \sigma \in F\left(\left(-1, z_{-1}^{a,b}\right)\right)\right) \leq 4k\gamma.$$

The number of σ's that give the same $\tilde{\sigma}'$ is at most k^2, and so we can conclude that

$$\mathbb{P}\left(\sigma \in I_1 \cap I_2 \cap F_1^{a,b} \cap F_2^{a,c}, f\left(\sigma'\right) = a, f\left(\tilde{\sigma}'\right) = a\right) \leq 4k^3\gamma,$$

and similarly

$$\mathbb{P}\left(\sigma \in I_1 \cap I_2 \cap F_1^{a,b} \cap F_2^{a,c}, f\left(\sigma'\right) = a, f\left(\tilde{\sigma}''\right) = a\right) \leq 4k^3\gamma,$$

which shows that

$$\mathbb{P}\left(\sigma \in M_3\right) \geq \frac{1}{k^6}\left(\frac{\varepsilon^3}{54n^3k^{21}} - 2k\gamma - 8k^3\gamma\right) \geq \frac{\varepsilon^3}{54n^3k^{27}} - \frac{9\gamma}{k^3}. \qquad \square$$

To conclude the proof in this case, recall that we have chosen $\gamma = \frac{\varepsilon^3}{10^3n^3k^{24}}$.

2.7.3.2 Case 2

First, as in the previous case, we can look at simply the boundary between a and b in direction 1, and conclude that either there are many manipulation points, or there are many local dictators, or (2.67) holds. This holds similarly for the boundary between c and d in direction 2. Finally, just as in Section 2.3.3.2, we can show that (2.67) and (2.68) cannot hold at the same time. We omit the details.

2.7.4 Proof of Theorem 2.35 concluded

Proof of Theorem 2.35. The starting point for the proof is Lemma 2.19, which directly implies Lemma 2.36 (unless there are many 2-manipulation points, in which case we are done). We then consider two cases, as indicated in Section 2.7.1.

We deal with the small fiber case in Section 2.7.2. First, Lemmas 2.37, 2.40, and 2.42, and Corollaries 2.38, 2.39, 2.41, and 2.43 imply that either there are many 3-manipulation points, or there are many local dictators on three alternatives in coordinate 1. We then deal with the case of many local dictators in Section 2.7.2.1. Lemma 2.44, Corollary 2.45, Lemmas 2.46 amd 2.47, Corollary 2.48, and Lemmas 2.49, 2.50, and 2.51 together show that there are many 4-manipulation points if there are many local dictators on three alternatives, and the SCF is ε-far from the family of nonmanipulable functions.

We deal with the large fiber case in Section 2.7.3. Here Lemmas 2.52, 2.53, 2.54, and 2.55 show that if there are not many local dictators on three alternatives, then there are many 3-manipulation points. In the case when there are many local dictators, we refer back to Section 2.7.2.1 to conclude the proof. $\square$

2.8 Reduction to distance from truly nonmanipulable SCFs

Proof of Theorem 2.5. Our assumption means that there exists a SCF $g \in \overline{\text{NONMANIP}}$ such that $\mathbf{D}(f,g) \leq \alpha$. We distinguish two cases: either g is a function of one coordinate, or g takes on at most two values.

Case 1. g **is a function of one coordinate.** In this case we can assume w.l.o.g. that g is a function of the first coordinate, i.e., there exists a SCF $h : S_k \to [k]$ on one coordinate such that for every ranking profile σ, we have $g(\sigma) = h(\sigma_1)$.

We know from the quantitative Gibbard-Satterthwaite theorem for one voter that for any β either $\mathbf{D}(h, \text{NONMANIP}(1,k)) \leq \beta$, or $\mathbb{P}(\sigma \in M_3(h)) \geq \frac{\beta^3}{10^5 k^{16}}$.

In the former case, we have that

$$\mathbf{D}(g, \text{NONMANIP}(n,k)) \leq \mathbf{D}(h, \text{NONMANIP}(1,k)) \leq \beta,$$

and so consequently

$$\mathbf{D}(f, \text{NONMANIP}(n,k)) \leq \alpha + \beta.$$

In the latter case, we have that

$$\mathbb{P}(\sigma \in M_3(g)) = \mathbb{P}(\sigma \in M_3(h)) \geq \frac{\beta^3}{10^5 k^{16}},$$

and so consequently

$$\mathbb{P}(\sigma \in M_3(f)) \geq \frac{\beta^3}{10^5 k^{16}} - 6nk\alpha,$$

since changing the outcome of a SCF at one ranking profile can change the number of 3-manipulation points by at most $6nk$. Now choosing $\beta = 100nk^6\alpha^{1/3}$ shows that either (2.3) or (2.4) holds.

Case 2. g **is a function which takes on at most two values.** W.l.o.g. we may assume that the range of g is $\{a,b\} \subset [k]$, i.e., for every ranking profile $\sigma \in S_k^n$ we have $g(\sigma) \in \{a,b\}$.

There is one thing we have to be careful about: even though g takes on at most two values, it is not necessarily a Boolean function, since the value of $g(\sigma)$ does not necessarily depend only on the Boolean vector $x^{a,b}(\sigma)$.

We now define a function $h : S_k^n \to \{a,b\}$ that is close in some sense to g and which can be viewed as a Boolean function $h : \{a,b\}^n \to \{a,b\}$ because $h(\sigma)$ depends on σ only through $x^{a,b}(\sigma)$. (The vector $x^{a,b}(\sigma) \in \{-1,1\}^n$ encodes which of a and b is preferred in each coordinate, and a vector in $\{a,b\}^n$ can encode the same information.) For a given ranking profile σ, let us consider the fiber on which it is on, $F\left(x^{a,b}(\sigma)\right)$, and let us define $g|_{F\left(x^{a,b}(\sigma)\right)}$ to be the restriction of g to ranking profiles in the fiber $F\left(x^{a,b}(\sigma)\right)$. Then define (see Definition 2.22)

$$h(\sigma) := \text{Maj}\left(g|_{F\left(x^{a,b}(\sigma)\right)}\right).$$

By definition, $h(\sigma)$ depends on σ only through $x^{a,b}(\sigma)$, so we may also view h as a Boolean function $h : \{a,b\}^n \to \{a,b\}$.

For any given $0 < \delta < 1$, we either have $\mathbf{D}(g,h) \leq \delta$, in which case $\mathbf{D}(f,h) \leq \alpha + \delta$, or if $\mathbf{D}(g,h) > \delta$, then we show presently that

$$\mathbb{P}(\sigma \in M_2(f)) \geq \frac{\delta}{4nk^5} - nk\alpha. \tag{2.71}$$

Choosing $\delta = 8n^2k^6\alpha$ then shows that either (2.4) holds, or $\mathbf{D}(f,h) \leq 9n^2k^6\alpha$.

Let us now show (2.71). We use a canonical path argument again, but first we divide the ranking profiles according to the fibers with respect to preference between a and b.

Let us consider an arbitrary fiber $F\left(z^{a,b}\right)$, and divide it into two disjoint sets: into those ranking profiles for which the outcome of g and h agree, and those for which these outcomes are different. That is,

$$F\left(z^{a,b}\right) = F^{\mathrm{maj}}\left(z^{a,b}\right) \cup F^{\mathrm{min}}\left(z^{a,b}\right),$$

where

$$\begin{aligned} F^{\mathrm{maj}}\left(z^{a,b}\right) &= \left\{\sigma \in F\left(z^{a,b}\right) : g(\sigma) = h(\sigma)\right\}, \\ F^{\mathrm{min}}\left(z^{a,b}\right) &= \left\{\sigma \in F\left(z^{a,b}\right) : g(\sigma) \neq h(\sigma)\right\}. \end{aligned}$$

By construction, we know that

$$\left|F^{\mathrm{min}}\left(z^{a,b}\right)\right| \leq \frac{1}{2}\left|F\left(z^{a,b}\right)\right| = \frac{1}{2}\left(\frac{k!}{2}\right)^n.$$

Now for every pair of profiles $(\sigma,\sigma') \in F^{\mathrm{min}}\left(z^{a,b}\right) \times F^{\mathrm{maj}}\left(z^{a,b}\right)$ define a canonical path from σ to σ' by applying a path construction in each coordinate one by one, and then concatenating these paths. In each coordinate we apply the path construction of [40, 41, Proposition 6.6.]: we bubble up everything except a and b, and then finally bubble up the last two alternatives as well.

For a given edge $(\pi,\pi') \in F^{\mathrm{min}}\left(z^{a,b}\right) \times F^{\mathrm{maj}}\left(z^{a,b}\right)$ there are at most $2k^4\left(\frac{k!}{2}\right)^n$ possible pairs $(\sigma,\sigma') \in F^{\mathrm{min}}\left(z^{a,b}\right) \times F^{\mathrm{maj}}\left(z^{a,b}\right)$ such that the canonical path between σ and σ' defined above passes through (π,π'). (This can be shown just like in the previous lemmas, e.g., Lemma 2.37.) Consequently we have

$$\left|\partial_e\left(F^{\mathrm{min}}\left(z^{a,b}\right)\right)\right| \geq \frac{\left|F^{\mathrm{min}}\left(z^{a,b}\right)\right|\left|F^{\mathrm{maj}}\left(z^{a,b}\right)\right|}{2k^4\left(\frac{k!}{2}\right)^n} \geq \frac{\left|F^{\mathrm{min}}\left(z^{a,b}\right)\right|}{4k^4},$$

where the edge boundary $\partial_e\left(F^{\mathrm{min}}\left(z^{a,b}\right)\right)$ is defined via the refined rankings graph restricted to the fiber $F\left(z^{a,b}\right)$. Summing this over all fibers we have that

$$\sum_{z^{a,b}}\left|\partial_e\left(F^{\mathrm{min}}\left(z^{a,b}\right)\right)\right| \geq \sum_{z^{a,b}}\frac{\left|F^{\mathrm{min}}\left(z^{a,b}\right)\right|}{4k^4} \geq \frac{\delta}{4k^4}(k!)^n, \tag{2.72}$$

using the fact that $\mathbf{D}(g,h) > \delta$.

Now it is easy to see that if $(\sigma, \sigma') \in \partial_e \left(F^{\min}\left(z^{a,b}\right)\right)$ for some $z^{a,b}$, then either σ or σ' is a 2-manipulation point for g. In the refined rankings graph every vertex (ranking profile) has $n(k-1) < nk$ neighbors, so each 2-manipulation point can be counted at most nk times in the sum on the left hand side of (2.72), showing that

$$\mathbb{P}(\sigma \in M_2(g)) \geq \frac{\delta}{4nk^5},$$

from which (2.71) follows immediately, since changing the outcome of a SCF at one ranking profile can change the number of 2-manipulation points by at most nk.

So either we are done because (2.4) holds, or $\mathbf{D}(f,h) \leq 9n^2k^6\alpha$; suppose the latter case. Our final step is to look at h as a Boolean function, and use a result on testing monotonicity [33].

Denote by $\tilde{\mathbf{D}}$ the distance of h when viewed as a Boolean function from the set of monotone Boolean functions. Let $0 < \varepsilon < 1$ be arbitrary. Then either $\tilde{\mathbf{D}} \leq \varepsilon$, in which case $\mathbf{D}(h, \text{NONMANIP}) \leq \tilde{\mathbf{D}} \leq \varepsilon$ and therefore $\mathbf{D}(f, \text{NONMANIP}) \leq 9n^2k^6\alpha + \varepsilon$, or $\tilde{\mathbf{D}} > \varepsilon$. In the latter case we show that then

$$\mathbb{P}(\sigma \in M_2(f)) \geq \frac{2\varepsilon}{nk} - 9n^3k^7\alpha. \tag{2.73}$$

Choosing $\varepsilon = 5n^4k^8\alpha$ then shows that either (2.3) or (2.4) holds.

Let us now show (2.73). Let us view h as a Boolean function, and denote by $p(h)$ the fraction of pairs of strings, differing on one coordinate, that violate the monotonicity condition. Goldreich, Goldwasser, Lehman, Ron, and Samorodnitsky showed in [33, Theorem 2] that $p(h) \geq \frac{\tilde{\mathbf{D}}}{n}$.

Now going back to viewing h as a SCF on k alternatives, this tells us that there are at least $\frac{\varepsilon}{2}2^n$ pairs of fibers, which differ on one coordinate, that violate monotonicity. For each such pair of fibers, whenever a and b are adjacent in the coordinate where the two fibers differ, we get a 2-manipulation point. Such a 2-manipulation point can be counted at most n times in this way (since there are n coordinates where a and b can be adjacent). Consequently, we have

$$|M_2(h)| \geq \frac{\varepsilon}{2} \cdot 2^n \cdot 2(k-1)! \left(\frac{k!}{2}\right)^{n-1} \cdot \frac{1}{n} = \frac{2\varepsilon}{nk}(k!)^n,$$

i.e.,

$$\mathbb{P}(\sigma \in M_2(h)) \geq \frac{2\varepsilon}{nk},$$

from which (2.73) follows immediately, since changing the outcome of a SCF at one ranking profile can change the number of 2-manipulation points by at most nk. □

Proof of Theorem 2.2. First we argue without specific bounds. Suppose on the contrary that our SCF f does not have many 4-manipulation points. Then f is close to $\overline{\text{NONMANIP}}$

by Theorem 2.35. Consequently, by Theorem 2.5, f is close to NONMANIP, which is a contradiction.

Now we argue with specific bounds. Assume on the contrary that

$$\mathbb{P}\left(\sigma \in M_4\left(f\right)\right) < \frac{\varepsilon^{15}}{10^{39} n^{67} k^{166}}.$$

Then by Theorem 2.35 we have that $\mathbf{D}\left(f, \overline{\text{NONMANIP}}\right) < \frac{\varepsilon^3}{10^6 n^{12} k^{24}}$, and consequently by Theorem 2.5 we have $\mathbf{D}\left(f, \text{NONMANIP}\right) < \varepsilon$, which is a contradiction. □

2.9 Open problems

We conclude with a few open problems that arise naturally, some of which have already been asked by Isaksson et al. [40, 41].

- In Section 2.1.3 we mentioned that our techniques do not lead to tight bounds. It would be interesting to find the correct tight bounds. When discussing tight bounds there are various different ways to measure the manipulability of a function: in terms of the probability of having manipulating voters, in terms of the expected number of manipulating voters, in terms of the number of manipulative edges (either in the refined or non-refined graph), etc.

- A related question is to find, in some natural subsets of functions, the one that minimizes manipulation. For example, among anonymous SCFs, which function minimizes the expected number of manipulating voters? For example, for plurality, the probability that a ranking profile is manipulable is $\Theta\left(1/\sqrt{n}\right)$, and if it is manipulable, then $\Theta\left(n\right)$ voters can manipulate, so consequently the expected value of the number of voters who can manipulate individually is $\Theta\left(\sqrt{n}\right)$. Is it true that for all anonymous SCFs, this expectation is $\Omega\left(\sqrt{n}\right)$?

- What if the distribution over rankings is not i.i.d. uniform? It would be interesting to consider a quantitative Gibbard-Satterthwaite theorem, and also the questions asked above, in this setting.

Chapter 3

Coalitional manipulation: a smooth transition from powerlessness to absolute power

3.1 Introduction

While in the previous chapter we focused on a single voter manipulating the outcome of the election, in this chapter we are interested in the *coalitional manipulation* problem, where a group of voters can change their votes in unison, with the goal of making a given candidate win. Various variations of this problem are known to be $\mathcal{NP}$-hard under many of the common SCFs [20, 78, 10].

Crucially, this line of work focuses on worst-case complexity, but, as mentioned before, a recent line of research on average-case manipulability has been questioning the relevance of such worst-case complexity results. The goal of this alternative line of work is to show that there are no "reasonable" voting rules that are computationally hard to manipulate *on average*. Specifically, the goal is to rule out the following informal statement: there are "good" voting rules that are hard to manipulate on average under any "sufficiently rich" distribution over votes.

Taking this point of view, showing easiness of manipulation under a restricted class of distributions—such as i.i.d. votes or even uniform votes (the impartial culture assumption)—is interesting, even if these do not necessarily capture all possible real-world elections. Specifically, if we show that manipulation is easy under such distributions, then any average-case hardness result would necessarily have to make some unnatural technical assumptions to avoid these distributions. Studying such restricted distributions over votes is indeed exactly what some recent papers have done.

For the coalitional manipulation problem, Procaccia and Rosenschein [64] first suggested that it is trivial to determine whether manipulation is possible for most coalitional manipulation instances, from a typical-case computational point of view; one can make a highly informed guess purely based on the number of manipulators. Specifically, they studied a setting where there is a distribution over votes (which satisfies some conditions), and concentrated on a family of SCFs known as positional scoring rules. They showed that if the size of the coalition is $o(\sqrt{n})$, then with probability converging to 1 as $n \to \infty$, the coalition is powerless, i.e., it cannot change the outcome of the election. In contrast, if the size of the coalition is $\omega(\sqrt{n})$ (and $o(n)$), then with probability converging to 1 as $n \to \infty$, the coalition is all-powerful, i.e., it can elect any candidate. Later Xia and Conitzer [77] proved an analogous result for so-called generalized scoring rules, a family that contains almost all common voting rules. See also related work by Peleg [61], Slinko [69], Pritchard and Slinko [62], and Pritchard and Wilson [63]. We discuss additional related work in Section 3.1.2.

Our primary interest in this chapter is to understand the critical window that these papers leave open, when the size of the coalition is $\Theta(\sqrt{n})$. Specifically, we are interested in the *phase transition* in the probability of coalitional manipulation, when the size of the coalition is $c\sqrt{n}$ and c varies from zero to infinity, i.e., the *transition from powerlessness to absolute power*.

In the past few decades there has been much research on the connection between phase transitions and computationally hard problems (see, e.g., [31, 17, 3]). In particular, it is often the case that the computationally hardest problems can be found at critical values of a sharp phase transition (see, e.g., [34] for an overview). On the other hand, smooth phase transitions are often found in connection with computationally easy (polynomial) problems, such as 2-coloring [2] and 1-in-2 SAT [71]. Thus understanding the phase transition in this critical window may shed light on where the computationally hardest problems lie.

Recently, Walsh [72] *empirically* analyzed two well-known voting rules—veto and single transferable vote (STV)—and found that there is a smooth phase transition between the two regimes. Specifically, Walsh studied coalitional manipulation with unweighted votes for STV and weighted votes for veto, and sampled from a number of distributions in his experiments, including i.i.d. distributions, correlated distributions, and votes sampled from real-world elections. Our main result complements and improves upon Walsh's analysis in two ways; while Walsh's results show how the phase transition looks like concretely for veto and STV, we *analytically* show that the phase transition is indeed smooth for *any generalized scoring rule* (including veto and STV) when the votes are i.i.d. This suggests that deciding the coalitional manipulation problem may not be computationally hard in practice.

3.1.1 Our results

Before presenting our results, we first formally specify the setup of the problem. Again we consider n voters electing a winner among k alternatives via a SCF $f : S_k^n \to [k]$. We denote a ranking profile by $\sigma = (\sigma_1, \dots, \sigma_n) \in S_k^n$, and for a candidate a, define $W_a = \{\sigma \in S_k^n \mid f(\sigma) = a\}$, the set of ranking profiles where the outcome of f is a. Our setup and assumptions are the following.

Assumption 3.1. *We assume that the number of candidates, k, is constant.*

Assumption 3.2. *We assume that the SCF f is* anonymous*, i.e., that it treats each voter equally.*

Assumption 3.3. *We assume that the votes of voters are i.i.d., according to some distribution p on S_k. Furthermore, we assume that there exists $\delta > 0$ such that for every $\pi \in S_k$, we have $p(\pi) \geq \delta$ (necessarily $\delta \leq 1/k!$).*

If we were to assume only these, then our setup would include uninteresting cases, such as when f is a constant—i.e., no matter what the votes are, a specific candidate wins. Another less interesting case is when the probability of a given candidate winning vanishes as $n \to \infty$—we can then essentially forget about this candidate for large n (in the sense that a coalition of size $\Omega(n)$ would be necessary to make this candidate win). To exclude these and focus on the interesting cases, we make an additional assumption which concerns both the SCF and the distribution of votes.

Assumption 3.4. *We assume that there exists $\varepsilon > 0$ such that for every n and for every candidate $a \in [k]$, the probability of a being elected is at least $\varepsilon > 0$, i.e., $\mathbb{P}(W_a) \geq \varepsilon$ (necessarily $\varepsilon \leq 1/k$).*

All four assumptions are satisfied when the distribution is uniform (i.e., under the impartial culture assumption) and the SCF is close to being neutral (i.e., neutral up to some tie-breaking rules); in particular, they hold for all commonly used SCFs. The assumptions are somewhat more general than this, although the i.i.d. assumption remains a restrictive one. However, as discussed earlier, even showing easiness of manipulation under such a restricted class of distributions is interesting.

As mentioned before, we are interested in the case when the coalition has size $c\sqrt{n}$ for some constant c. Define the probabilities

$$\begin{aligned}
\underline{q}_n(c) &:= \mathbb{P}\left(\text{some coalition of size } c\sqrt{n} \text{ can elect any candidate} \right), \\
\overline{q}_n(c) &:= \mathbb{P}\left(\text{some coalition of size } c\sqrt{n} \text{ can change the outcome of the election} \right), \\
\underline{r}_n(c) &:= \mathbb{P}\left(\text{a specific coalition of size } c\sqrt{n} \text{ can elect any candidate} \right), \\
\overline{r}_n(c) &:= \mathbb{P}\left(\text{a specific coalition of size } c\sqrt{n} \text{ can change the outcome of the election} \right),
\end{aligned}$$

and let

$$\underline{q}(c) := \lim_{n\to\infty} \underline{q}_n(c), \quad \overline{q}(c) := \lim_{n\to\infty} \overline{q}_n(c), \quad \underline{r}(c) := \lim_{n\to\infty} \underline{r}_n(c), \quad \overline{r}(c) := \lim_{n\to\infty} \overline{r}_n(c),$$

provided these limits exist. Clearly $\underline{q}_n(c) \leq \overline{q}_n(c)$, $\underline{r}_n(c) \leq \overline{r}_n(c)$, $\underline{r}_n(c) \leq \underline{q}_n(c)$, and $\overline{r}_n(c) \leq \overline{q}_n(c)$.

Before we describe our results, which deal with these quantities, we first explain how these relate to the various variants of the coalitional manipulation problem. In the coalitional manipulation problem the coalition is fixed, and thus the relevant quantities are $\underline{r}_n(c)$ and $\overline{r}_n(c)$. Closely related is the problem of determining the *margin of victory*, which is the minimum number of voters who need to change their votes to change the outcome of the election. Also related is the problem of *bribery*, the minimum number of voters who need to change their votes to make a given candidate win. The main difference between these problems is that in coalitional manipulation the coalition is fixed, whereas in the latter two problems the coalition is not fixed. Hence the relevant quantities for studying the latter two are $\underline{q}_n(c)$ and $\overline{q}_n(c)$. Our tools also allow us to deal with other related quantities (such as *microbribery* [28]), but we focus our attention on the four quantities described above.

Our first result analyzes the case when the size of the coalition is $c\sqrt{n}$ for large c. We show that if c is large enough, then with probability close to 1, a specific coalition of size $c\sqrt{n}$ can elect any candidate. This holds for any SCF that satisfies the above (mild) restrictions.

Theorem 3.1. *Assume that Assumptions 3.1, 3.2, 3.3, and 3.4 hold. For any $\eta > 0$ there exists a constant $c = c(\eta, \delta, \varepsilon, k)$ such that $\underline{r}_n(c) \geq 1 - \eta$ for every n. In particular, we can choose*

$$c = (4/\delta)\log(2k!/\eta)\left[\sqrt{\log(2k/\eta)} + \sqrt{\log(2/\varepsilon)}\right].$$

It follows that

$$\lim_{c\to\infty}\liminf_{n\to\infty}\underline{r}_n(c) = 1.$$

This result extends previous theorems of Procaccia and Rosenschein [64], and Xia and Conitzer [77], from scoring rules and generalized scoring rules, respectively, to anonymous SCFs.

Our second result deals with the case when the size of the coalition is $c\sqrt{n}$ for small c, and the transition as c goes from 0 to ∞. Here we assume additionally that f is a generalized scoring rule (to be defined in Section 3.3.1.1); this is needed because there exist (pathological) anonymous SCFs for which the result below does not hold (see the beginning of Section 3.3 for an example).

Theorem 3.2. *Assume that Assumptions 3.1, 3.2, 3.3, and 3.4 hold, and furthermore that f is a generalized scoring rule. Then:*

(1) The limits $\underline{q}(c)$, $\overline{q}(c)$, $\underline{r}(c)$, and $\overline{r}(c)$ exist.

(2) There exists a constant $K = K(f, \delta) < \infty$ such that $\overline{q}(c) \leq Kc$; in particular,

$$\lim_{c\to 0}\overline{q}(c) = 0.$$

(3) For all $0 < c < \infty$, we have $0 < \underline{q}(c) \leq \overline{q}(c) < 1$ and $0 < \underline{r}(c) \leq \overline{r}(c) < 1$. Furthermore, the functions $\underline{q}(\cdot)$, $\overline{q}(\cdot)$, $\underline{r}(\cdot)$, and $\overline{r}(\cdot)$ are all continuously differentiable with bounded derivative.

In words, Part 2 means that if c is small enough then with probability close to 1 no coalition of size $c\sqrt{n}$ can change the outcome of the election, and the statements about $\overline{r}$ and $\underline{r}$ in Part 3 mean that the coalitional manipulation problem has a smooth phase transition: as the number of manipulators increases, the probabilities that a coalition has some power, and that it has absolute power, increase smoothly. Parts 1 and 2 of the theorem simply make a result of Xia and Conitzer [77] more precise, by extending the analysis to the $\Theta(\sqrt{n})$ regime. More importantly, in the proofs of these statements we introduce the machinery needed to establish Part 3, which is our main result.

Since the coalitional manipulation problem does *not* have a sharp phase transition, Theorem 3.2 can be interpreted as suggesting that realistic distributions over votes are likely to yield coalitional manipulation instances that are tractable in practice, even if the size of the coalition concentrates on the previously elusive $\Theta(\sqrt{n})$ regime; this is true for any generalized scoring rule, and in particular for almost all common social choice functions (an exception is Dodgson's rule). This interpretation has a negative flavor in further strengthening the conclusion that worst-case complexity is a poor barrier to manipulation.

However, the complexity glass is in fact only half empty. The probability that the margin of victory is at most $c\sqrt{n}$ is captured by the quantity $\overline{q}_n$, hence Part 3 of Theorem 3.2 also implies that the margin of victory problem has a smooth phase transition. As recently pointed out by Xia [73], efficiently solving the margin of victory problem could help in *post-election audits*—used to determine whether electronic elections have resulted in an incorrect outcome due to software or hardware bugs—and its tractability is in fact desirable.

The methods we use are flexible, and can be extended to various setups of interest that do not directly satisfy our assumptions above, for instance single-peaked preferences. Consider a one-dimensional political spectrum represented by the interval $[0, 1]$, and fix the location of the candidates. Assume that voters are uniformly distributed on the interval, independently of each other. For technical reasons, this distribution does not satisfy our assumptions, since there will be rankings $\pi \in S_k$ such that $p(\pi) = 0$; however, our tools allow us to deal with this setting as well. For instance, if the locations of the k candidates are $\left\{\frac{1}{2k}, \frac{3}{2k}, \ldots, \frac{2k-1}{2k}\right\}$, then our results hold (with appropriate quantitative modifications). Similarly, if the locations were something else, then there would exist a subset of candidates who have an asymptotically nonvanishing probability of winning, and the same results hold restricted to this subset of candidates.

Finally, we discuss the role of tie-breaking in our setup, since this is often an important issue when studying manipulation. However, since we consider manipulation by coalitions of size $c\sqrt{n}$, ties where there exist a constant number of voters such that if their votes are changed appropriately there is no longer a tie, are not relevant. Indeed, our tools allow us to extend the results of Theorem 3.2 to a class of SCFs slightly beyond generalized scoring rules, and, in particular, these allow for arbitrary tie-breaking rules (see Section 3.3.2.1 for details).

3.1.2 Additional related work

As discussed at the beginning of Chapter 2, there is a recent line of research with an average-case algorithmic flavor that suggests that manipulation is indeed typically easy. A different approach, initiated by Friedgut, Kalai, Keller and Nisan [29, 30], and culminating in the work presented in Chapter 2, studies the fraction of ranking profiles that are manipulable, and also suggests that manipulation is easy on average. We refer to the survey by Faliszewski and Procaccia [27] for a detailed history of the surrounding literature. See also related literature in economics, e.g., the work of [35, 16, 56].

Recent work by Xia [73] is independent from, and closely related to, our work. As mentioned above, Xia's paper is concerned with computing the margin of victory in elections. He focuses on computational complexity questions and approximation algorithms, but one of his results is similar to Parts 1 and 2 of Theorem 3.2. However, our analysis is completely different; our approach facilitates the proof of Part 3 of the theorem, which is our main contribution. An even more recent (and also independent) manuscript by Xia [74] considers similar questions for generalized scoring rules and captures additional types of strategic behavior (such as control), but again, crucially, this work does not attempt to understand the phase transition (nor does it subsume our Theorem 3.1).

3.2 Large coalitions

Without further ado, we prove Theorem 3.1. The main idea is to observe that for i.i.d. distributions, the Hamming distance of a random ranking profile from a fixed subset of ranking profiles concentrates around its mean. The theorem follows from standard concentration inequalities.

Proof of Theorem 3.1. For $\sigma, \sigma' \in S_k^n$, define

$$d(\sigma, \sigma') = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}[\sigma_i \neq \sigma'_i],$$

i.e., $d(\sigma, \sigma')$ is $1/n$ times the Hamming distance of σ and σ'. If U is a subset of ranking profiles and σ is a specific ranking profile then define $d_U(\sigma) = \min_{\sigma' \in U} d(\sigma, \sigma')$. The function d_U is Lipschitz with constant $1/n$, and therefore by McDiarmid's inequality we have the following concentration inequality:

$$\mathbb{P}(|d_U(\sigma) - \mathbb{E}d_U| \geq c) \leq 2\exp\left(-2c^2 n\right) \tag{3.1}$$

for any $c > 0$ and $U \subseteq S_k^n$. Suppose $U \subseteq S_k^n$ has measure at least ε, i.e., U is such that $\mathbb{P}(\sigma \in U) \geq \varepsilon$, and take γ such that $2\exp(-2\gamma^2 n) < \varepsilon$, e.g., let $\gamma = \sqrt{\log(2/\varepsilon)}/\sqrt{n}$. Then (3.1) implies that there exists $\sigma \in U$ such that $|d_U(\sigma) - \mathbb{E}d_U| \leq \gamma$, but since $d_U(\sigma) = 0$, this means that $\mathbb{E}d_U \leq \gamma$. So for such a set U, we have

$$\mathbb{P}(d_U(\sigma) > \gamma + c) \leq \exp\left(-2c^2 n\right)$$

for any $c > 0$. Choosing $c = B/\sqrt{n}$ and defining $B' = B + \sqrt{\log(2/\varepsilon)}$ we get that

$$\mathbb{P}\left(d_U(\sigma) > B'/\sqrt{n}\right) \leq \exp\left(-2B^2\right). \tag{3.2}$$

In the language of the usual Hamming distance, this means that the probability that the ranking profile needs to be changed in at least $B'\sqrt{n}$ coordinates to be in U is at most $\exp\left(-2B^2\right)$, which can be made arbitrarily small by choosing B large enough.

By our assumption, $\mathbb{P}(\sigma \in W_a) \geq \varepsilon$ for every a, and therefore by (3.2) and a union bound we get

$$\mathbb{P}\left(\exists a : d_{W_a}(\sigma) > B'/\sqrt{n}\right) \leq k \exp\left(-2B^2\right).$$

By choosing $B = \sqrt{\log(2k/\eta)}$, this probability is at most $\eta/2$.

Consider a specific coalition of size $DB'\sqrt{n}$, where $D = D(\delta, k)$ will be chosen later. Using Chernoff's bound and a union bound, with probability close to one, for every possible ranking π the coalition has at least $B'\sqrt{n}$ voters with the ranking π:

$$\begin{aligned}\mathbb{P}&\left(\exists \pi \in S_k : \text{ coalition of size } DB'\sqrt{n} \text{ has less than } B'\sqrt{n} \text{ voters with ranking } \pi\right)\\ &\leq k!\mathbb{P}\left(\mathrm{Bin}\left(DB'\sqrt{n}, \delta\right) < B'\sqrt{n}\right) \leq k! \exp\left(-\left(1 - 1/D\delta\right)^2 DB'\sqrt{n}\delta/2\right)\\ &\leq k! \exp\left(-\left(1 - 1/D\delta\right)^2 D\delta/2\right),\end{aligned}$$

where $\mathrm{Bin}\left(DB'\sqrt{n}, \delta\right)$ denotes a binomial random variable with parameters $DB'\sqrt{n}$ and δ, and where we used our assumption that for every voter the probability for every ranking is at least $\delta > 0$. Choosing $D = (4/\delta)\log(2k!/\eta)$, this probability is at most $\eta/2$.

By the anonymity of f, the outcome only depends on the number of voters voting according to each ranking. Consequently, if σ is such that it is at a distance of at most $B'/\sqrt{n}$ away from each W_a, and where for each ranking π there are at least $B'\sqrt{n}$ voters in the coalition with ranking π, then the coalition is able to achieve any outcome. Using the above and a union bound this happens with probability at least $1 - \eta$. □

3.3 Small coalitions and the phase transition

This section is almost entirely devoted to the proof of Theorem 3.2, but it also includes some helpful definitions, examples, and intuitions.

Consider the following example of a SCF. For $a \in [k]$ let $n_a(\sigma)$ denote the number of voters who ranked candidate a on top in the ranking profile σ. Define the SCF f by $f(\sigma) = \sum_{a=1}^{k} a n_a(\sigma) \mod k$. This SCF is clearly anonymous (since it only depends on the number of voters voting according to specific rankings), and moreover it is easy to see that *any* single voter can *always* elect any candidate.

This example shows that, in general, we cannot have a matching lower bound for the size of the manipulating coalition on the order of $\sqrt{n}$. However, this is an artificial example (one would not consider such a voting system in real life), and we expect that a matching lower bound holds for most reasonable SCFs.

Xia and Conitzer [77] introduced a large class of SCFs called generalized scoring rules, which include most commonly occurring SCFs. In the following we introduce this class of SCFs, provide an alternative way of looking at them (as so-called "hyperplane rules"), and show that for this class of SCFs if the coalition has size $c\sqrt{n}$ for small enough c, then the probability of being able to change the outcome of the election can be arbitrarily close to zero. At the end of the section we then prove the smooth transition as stated in Part 3 of Theorem 3.2.

3.3.1 Generalized scoring rules, hyperplane rules, and their equivalence

In this subsection we introduce generalized scoring rules and hyperplane rules, and show their equivalence.

3.3.1.1 Generalized scoring rules

We now define generalized scoring rules.

Definition 3.1. *For any $y, z \in \mathbb{R}^m$, we say that y and z are* equivalent*, denoted by $y \sim z$, if for every $i, j \in [m]$, we have $y_i \geq y_j$ if and only if $z_i \geq z_j$.*

Definition 3.2. *A function $g : \mathbb{R}^m \to [k]$ is* compatible *if for any $y \sim z$, we have $g(y) = g(z)$.*

That is, for any function g that is compatible, $g(y)$ is completely determined by the total preorder of $\{y_1, \ldots, y_m\}$ (a total preorder is an ordering in which ties are allowed).

Definition 3.3 (Generalized scoring rules)**.** *Let $m \in \mathbb{N}$ be a natural number, let $f : S_k \to \mathbb{R}^m$ be a function mapping a ranking to an m-tuple of real numbers (called a* generalized scoring function*), and let $g : \mathbb{R}^m \to [k]$ be a function mapping an m-tuple of reals to an integer in $[k]$, where g is compatible (g is called a* decision function*). The functions f and g determine the* generalized scoring rule $\mathrm{GS}(f,g)$ *as follows: for $\sigma \in S_k^n$, let*

$$\mathrm{GS}(f,g)(\sigma) := g\left(\sum_{i=1}^{n} f(\sigma_i)\right).$$

From the definition it is clear that every generalized scoring rule (GSR) is anonymous.

3.3.1.2 Hyperplane rules

Preliminaries and notation. In the following, for a SCF let us write $f \equiv f_n$, i.e., let us explicitly note that f is a function on n voters; also let us write $\sigma \equiv \sigma^n$. Since the SCF f_n is anonymous, the outcome only depends on the numbers of voters who vote according to particular rankings. Let D_n denote the set of points in the probability simplex $\Delta^{k!}$ for which all coordinates are integer multiples of $1/n$. Let us denote a typical element of the probability simplex $\Delta^{k!}$ by $x = \{x_\pi\}_{\pi \in S_k}$. For a ranking profile σ^n, let us denote the corresponding element of the probability simplex by $x(\sigma^n)$, i.e., for all $\pi \in S_k$, we have

$$x(\sigma^n)_\pi = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}[\sigma_i = \pi].$$

By our assumptions the outcome of f_n only depends on $x(\sigma^n)$, so by abuse of notation we may write that $f_n : \Delta^{k!}|_{D_n} \to [k]$ with $f_n(\sigma^n) = f_n(x(\sigma^n))$.

We are now ready to define hyperplane rules.

Definition 3.4 (Hyperplane rules)**.** *Fix a finite set of affine hyperplanes of the simplex $\Delta^{k!}$: $H_1, \ldots, H_\ell$. Each affine hyperplane partitions the simplex into three parts: the affine hyperplane itself and two open halfspaces on either side of the affine hyperplane. Thus the affine hyperplanes $H_1, \ldots, H_\ell$ partition the simplex into finitely many (at most 3^ℓ) regions. Let $F : \Delta^{k!} \to [k]$ be a function which is constant on each such region. Then the sequence of SCFs $\{f_n\}_{n \geq 1}$, $f_n : S_k^n \to [k]$, defined by*

$$f_n(\sigma^n) = F(x(\sigma^n))$$

is called a hyperplane rule induced by the affine hyperplanes $H_1, \ldots, H_\ell$ and the function F.

A function $F : \Delta^{k!} \to [k]$ naturally partitions the simplex $\Delta^{k!}$ into k parts based on the outcome of F. (For hyperplane rules this partition is coarser than the partition of $\Delta^{k!}$ induced by the affine hyperplanes $H_1, \ldots, H_\ell$.) We abuse notation and denote these parts by $\{W_a\}_{a \in [k]}$. The following definition will be useful for us.

Definition 3.5 (Interior and boundaries of a partition induced by F)**.** *We say that $x \in \Delta^{k!}$ is an interior point of the partition $\{W_a\}_{a \in [k]}$ induced by F if there exists $\alpha > 0$ such that for all $y \in \Delta^{k!}$ for which $|x - y| \leq \alpha$, we have $F(x) = F(y)$. Otherwise, we say that $x \in \Delta^{k!}$ is on the boundary of the partition, which we denote by B.*

For a hyperplane rule the boundary B is contained in the union of the corresponding affine hyperplanes. Conversely, suppose $F : \Delta^{k!} \to [k]$ is an arbitrary function and the sequence of (anonymous) SCFs $\{f_n\}_{n \geq 1}$, $f_n : S_k^n \to [k]$ is defined by $f_n(\sigma^n) = F(x(\sigma^n))$. If the boundary B of F is contained in the union of finitely many affine hyperplanes of $\Delta^{k!}$, then F is not necessarily a hyperplane rule, but there exists a hyperplane rule $\hat{F}$ such that F and $\hat{F}$ agree everywhere except perhaps on the union of the finitely many affine hyperplanes.

3.3.1.3 Equivalence

Xia and Conitzer [76] gave a characterization of generalized scoring rules: a SCF is a generalized scoring rule if and only if it is anonymous and finitely locally consistent (see Xia and Conitzer [76, Definition 5]). This characterization is related to saying that generalized scoring rules are the same as hyperplane rules, yet we believe that spelling this out explicitly is important, because the geometric viewpoint of hyperplane rules is somewhat different, and in this probabilistic context it is also more flexible.

Lemma 3.3. *The class of generalized scoring rules coincides with the class of hyperplane rules.*

Proof. First, let us show that every hyperplane rule is a generalized scoring rule. Let us consider the hyperplane rule induced by affine hyperplanes $H_1, \ldots, H_\ell$ of the simplex $\Delta^{k!}$, and the function $F : \Delta^{k!} \to [k]$. The affine hyperplanes of $\Delta^{k!}$ can be thought of as hyperplanes of $\mathbb{R}^{k!}$ that go through the origin—abusing notation we also denote these by $H_1, \ldots, H_\ell$. Let $u_1, \ldots, u_\ell$ denote unit normal vectors of these hyperplanes.

We need to define functions f and g such that for every ranking profile $\sigma^n \in S_k^n$, $\mathrm{GS}(f,g)(\sigma^n) = F(x(\sigma^n))$. We will have $f : S_k \to \mathbb{R}^{\ell+1}$ and $g : \mathbb{R}^{\ell+1} \to [k]$. Coordinates $1, \ldots, \ell$ of f correspond to hyperplanes $H_1, \ldots, H_\ell$, while the last coordinate of f will always be 0 (this is a technical necessity to make sure that the function g is compatible). Let us look at the coordinate corresponding to hyperplane H_j with normal vector u_j. For $\pi \in S_k$ define

$$(f(\pi))_j \equiv (f(\pi))_{H_j} \equiv (f(\pi))_{u_j} := (u_j)_\pi \,,$$

where the coordinates of $\mathbb{R}^{k!}$ are indexed by elements of S_k. Then

$$(f(\sigma^n))_j := \sum_{i=1}^{n} (f(\sigma_i))_j = \sum_{i=1}^{n} (u_j)_{\sigma_i} = n\,(u_j \cdot x(\sigma^n))\,.$$

The sign of $(f(\sigma^n))_j$ thus tells us which side of the hyperplane H_j the point $x(\sigma^n)$ lies on. We define $g(y)$ for all $y \in \mathbb{R}^{\ell+1}$ such that $y_{\ell+1} = 0$; then the requirement that g be compatible defines g for all $y \in \mathbb{R}^{\ell+1}$. For $x \in \mathbb{R}$, define $\mathrm{sgn}(x)$ to be 1 if $x > 0$, -1 if $x < 0$, and 0 if $x = 0$.

To define $g(y_1, \ldots, y_\ell, 0)$, look at the vector $(\mathrm{sgn}(y_1), \ldots, \mathrm{sgn}(y_\ell))$. This vector determines a region in $\Delta^{k!}$ in the following way: if $\mathrm{sgn}(y_j) = 1$, then the region lies in the same open halfspace as u_j, if $\mathrm{sgn}(y_j) = -1$, then the region lies in the open halfspace which does not contain u_j, and finally, if $y_j = 0$, then the region lies in the hyperplane H_j. Now we define $g(y_1, \ldots, y_\ell, 0)$ to be the value of F on the region of $\Delta^{k!}$ defined by $(\mathrm{sgn}(y_1), \ldots, \mathrm{sgn}(y_\ell))$. The value of $g(y_1, \ldots, y_\ell, 0)$ is well-defined, since F is constant in each such region. Moreover, if we take $y \sim z$ with $y_{\ell+1} = z_{\ell+1} = 0$, then necessarily $(\mathrm{sgn}(y_1), \ldots, \mathrm{sgn}(y_\ell)) = (\mathrm{sgn}(z_1), \ldots, \mathrm{sgn}(z_\ell))$, and thus $g(y) = g(z)$: so g is compatible (this is where we used the extra coordinate).

Now let us show that every generalized scoring rule is a hyperplane rule. Suppose a generalized scoring rule is given by functions $f : S_k \to \mathbb{R}^m$ and $g : \mathbb{R}^m \to [k]$. For a ranking profile $\sigma^n \in S_k^n$, define $f(\sigma^n) := \sum_{i=1}^n f(\sigma_i) = n \sum_{\pi \in S_k} f(\pi)(x(\sigma^n))_\pi$; in this way we can view f as a function mapping $\mathbb{N}_{\geq 0}^{k!} \setminus \{0\}$ to $\mathbb{R}^m$ (and hence can also view $\mathrm{GS}(f,g)$ as a function mapping $\mathbb{N}_{\geq 0}^{k!} \setminus \{0\}$ to $[k]$). Since this mapping is homogeneous, we may extend the domain of f (and hence that of $\mathrm{GS}(f,g)$) to $\mathbb{Q}_{\geq 0}^{k!} \setminus \{0\}$ in the natural way.

For a total preorder $\mathcal{O}$, let $R_{\mathcal{O}} = \left\{x \in \mathbb{Q}_{\geq 0}^{k!} \setminus \{0\} : f(x) \sim \mathcal{O}\right\}$. By definition, if $x, y \in R_{\mathcal{O}}$ then $g(f(x)) = g(f(y))$, i.e., $\mathrm{GS}(f, g)$ is constant in each region $R_{\mathcal{O}}$. Each region $R_{\mathcal{O}}$ is a $\mathbb{Q}$-convex cone, i.e., if $x, y \in R_{\mathcal{O}}$ and $\lambda \in \mathbb{Q} \cap [0, 1]$, then $\lambda x + (1 - \lambda) y \in R_{\mathcal{O}}$, and furthermore if $\mu \in \mathbb{Q}_{>0}$, then $\mu x \in R_{\mathcal{O}}$ (both of these properties follow immediately from Definition 3.1). Thus we can write $\mathbb{Q}_{\geq 0}^{k!} \setminus \{0\}$ as the disjoint union of the $\mathbb{Q}$-convex cones $\{R_{\mathcal{O}}\}_{\mathcal{O}}$. The only way to do this is by taking finitely many hyperplanes of $\mathbb{R}^{k!}$ and cutting $\mathbb{Q}_{\geq 0}^{k!} \setminus \{0\}$ using these hyperplanes; a precise statement of this can be found in Section 3.4. This essentially follows from a result by Kemperman [45, Theorem 2]—to keep the chapter self-contained we reproduce in Section 3.4 his results and proof, and show how the statement above follows from his results. Since our function is homogeneous, we need only look at the values of $\mathrm{GS}(f, g)$ on the simplex $\Delta^{k!}$. By the above, the simplex is divided into regions $\left\{R_{\mathcal{O}} \cap \Delta^{k!}\right\}_{\mathcal{O}}$ via affine hyperplanes of $\Delta^{k!}$, and the function $\mathrm{GS}(f, g)$ is constant on $R_{\mathcal{O}} \cap \Delta^{k!}$ for each total preorder $\mathcal{O}$, so $\mathrm{GS}(f, g)$ is indeed a hyperplane rule. $\square$

3.3.1.4 Examples

Most commonly used SCFs are generalized scoring rules / hyperplane rules, including all positional scoring rules, instant-runoff voting, Coombs' method, contingent vote, the Kemény-Young method, Bucklin voting, Nanson's method, Baldwin's method, Copeland's method, maximin, and ranked pairs. Some of these examples were already shown by Xia and Conitzer [77, 76], but nevertheless in Section 3.5 we detail explanations of many of these examples. The main reason for this is that the perspective of a hyperplane rule arguably makes these explanations simpler and clearer. A rule that does not fit into this framework is Dodgson's rule, which is not homogeneous (see, e.g., [13]), and therefore it is not a hyperplane rule.

3.3.2 Small coalitions for generalized scoring rules

We now show that for generalized scoring rules, a coalition of size $c\sqrt{n}$ for small enough c can only change the outcome of the election with small probability. By the equivalence above, we can work in the framework of hyperplane rules.

We consider two metrics on $\Delta^{k!}$: the L^1 metric, denoted by d_1 or $\|\cdot\|_1$, and the L^2 metric, denoted by d_2 or $\|\cdot\|_2$. The L^1 metric is important in this setting, since changing the votes of voters corresponds to moving in the L^1 metric on $\Delta^{k!}$; this connection is formalized in the following lemma.

Lemma 3.4. *Let $\sigma^n, \tau^n \in S_k^n$. Then $d_1(x(\sigma^n), x(\tau^n)) \leq \frac{2}{n} d_H(\sigma^n, \tau^n)$, where d_H denotes Hamming distance, i.e., $d_H(\sigma^n, \tau^n) = \sum_{i=1}^n \mathbf{1}[\sigma_i \neq \tau_i]$. Furthermore, if $y \in D_n$, then there exists $\hat{\tau}^n \in S_k^n$ such that $x(\hat{\tau}^n) = y$ and $d_1(x(\sigma^n), y) = \frac{2}{n} d_H(\sigma^n, \hat{\tau}^n)$.*

Proof. Let $\pi^0 = \sigma^n$, and for $i = 1, \ldots, n$, define the ranking profile π^i as

$$\pi^i = (\tau_1, \ldots, \tau_i, \sigma_{i+1}, \ldots, \sigma_n).$$

By definition, $\pi^n = \tau^n$. The desired inequality then follows from the triangle inequality:

$$d_1\left(x\left(\sigma^n\right), x\left(\tau^n\right)\right) = d_1\left(x\left(\pi^0\right), x\left(\pi^n\right)\right) \leq \sum_{i=1}^{n} d_1\left(x\left(\pi^{i-1}\right), x\left(\pi^i\right)\right)$$
$$= \sum_{i=1}^{n} \frac{2}{n}\mathbf{1}\left[\sigma_i \neq \tau_i\right] = \frac{2}{n} d_H\left(\sigma^n, \tau^n\right).$$

For the second part of the lemma, construct $\hat{\tau}^n$ as follows. For each $\pi \in S_k$, let $I_\pi := \{i \in [n] : \sigma_i = \pi\}$. If $x\left(\sigma^n\right)_\pi \leq y_\pi$, then for every $i \in I_\pi$, let $\hat{\tau}_i = \pi$. If $x\left(\sigma^n\right)_\pi > y_\pi$, then choose a subset of indices $I'_\pi \subset I_\pi$ of size $|I'_\pi| = n y_\pi$, and for every $i \in I'_\pi$, let $\hat{\tau}_i = \pi$. Finally, define the rest of the coordinates of $\hat{\tau}^n$ so that $x\left(\hat{\tau}^n\right) = y$. The construction guarantees that then $d_1\left(x\left(\sigma^n\right), y\right) = \frac{2}{n} d_H\left(\sigma^n, \hat{\tau}^n\right)$. □

It is therefore natural to define distances from the boundary B using the L^1 metric:

Definition 3.6 (Blowup of boundary)**.** *For $\alpha > 0$, we define the blowup of the boundary B by α to be*

$$B^{+\alpha} = \left\{y \in \Delta^{k!} : \exists x \in B \text{ such that } \|x - y\|_1 \leq \alpha\right\}.$$

In order for some coalition to be able to change the outcome of the election at a given ranking profile, the point on the simplex corresponding to this ranking profile needs to be sufficiently close to the boundary B; this is formulated in the following lemma.

Lemma 3.5. *Suppose we have n voters, a coalition of size m, and the ranking profile is $\sigma^n \in S_k^n$, which corresponds to the point $x\left(\sigma^n\right) \in \Delta^{k!}$ on the probability simplex. A necessary condition for the coalition to be able to change the outcome of the election from this position is that $x\left(\sigma^n\right) \in B^{+2m/n}$. Conversely, if $x\left(\sigma^n\right) \in B^{+(2m-k!)/n}$, then there exists a coalition of size m that can change the outcome of the election.*

Proof. For any ranking profile τ^n that the coalition can reach, we have $d_H\left(\sigma^n, \tau^n\right) \leq m$, and so by Lemma 3.4 we have $d_1\left(x\left(\sigma^n\right), x\left(\tau^n\right)\right) \leq \frac{2m}{n}$. If $x\left(\sigma^n\right) \notin B^{+2m/n}$, then for every ranking profile τ^n which the coalition can reach, $x\left(\sigma^n\right)$ and $x\left(\tau^n\right)$ are in the same region determined by the hyperplanes, and so $F\left(x\left(\tau^n\right)\right) = F\left(x\left(\sigma^n\right)\right)$, i.e., the coalition cannot change the outcome of the election.

Now suppose that $x\left(\sigma^n\right) \in B^{+(2m-k!)/n}$. Then there exists $y \in B$ such that we have $d_1\left(x\left(\sigma^n\right), y\right) \leq \frac{2m-k!}{n}$. Since $y \in B$, there exists $\hat{y} \in D_n$ such that $d_1\left(y, \hat{y}\right) \leq \frac{k!}{n}$ and $F\left(\hat{y}\right) \neq F\left(x\left(\sigma^n\right)\right)$. By the triangle inequality, $d_1\left(x\left(\sigma^n\right), \hat{y}\right) \leq \frac{2m}{n}$, and then by the second part of Lemma 3.4 there exists $\hat{\tau}^n \in S_k^n$ such that $x\left(\hat{\tau}^n\right) = \hat{y}$ and $d_H\left(\sigma^n, \hat{\tau}^n\right) \leq m$. The coalition consisting of voters with indices in $I := \{i \in [n] : \sigma_i \neq \hat{\tau}_i\}$ can thus change the outcome of the election. □

Corollary 3.6. *If we have n voters, the probability that some coalition of size m can change the outcome of the election is bounded from below by $\mathbb{P}\left(x\left(\sigma^n\right) \in B^{+(2m-k!)/n}\right)$ and from above by $\mathbb{P}\left(x\left(\sigma^n\right) \in B^{+2m/n}\right)$, where σ^n is drawn according to the probability distribution satisfying the conditions of the setup.*

Gaussian limit. Due to the i.i.d.-ness of the votes, the multinomial random variable $x(\sigma^n)$ concentrates around its expectation, and the rescaled random variable

$$\tilde{x}(\sigma^n) := \sqrt{n}\left(x(\sigma^n) - \mathbb{E}\left(x(\sigma^n)\right)\right)$$

converges to a normal distribution, with zero mean and specific covariance structure. For our analysis it is better to use this Gaussian picture, and thus we reformulate the preliminaries above in this limiting setting. First, let us determine the limiting distribution.

Lemma 3.7. *We have $\tilde{x}(\sigma^n) \Rightarrow_n N(0,\Sigma)$, where the covariance structure is given by $\Sigma = \operatorname{diag}(p) - pp^T$, where recall that p is the distribution of a vote.*

Proof. It is clear that $\mathbb{E}(\tilde{x}(\sigma^n)) = 0$. Computing the covariance structure, we first have that

$$\mathbb{E}\left(x_\pi^2\right) = \frac{1}{n^2}\sum_{i,j=1}^{n}\mathbb{P}\left(\sigma_i = \pi, \sigma_j = \pi\right) = \left(1 - \frac{1}{n}\right)p(\pi)^2 + \frac{1}{n}p(\pi),$$

from which we have $\operatorname{Var}(x_\pi) = \frac{1}{n}\left(p(\pi) - p(\pi)^2\right)$ and thus $\operatorname{Var}(\tilde{x}_\pi) = p(\pi) - p(\pi)^2$. Then similarly for $\pi \neq \pi'$ we have

$$\mathbb{E}\left(x_\pi x_{\pi'}\right) = \frac{1}{n^2}\sum_{i,j=1}^{n}\mathbb{P}\left(\sigma_i = \pi, \sigma_j = \pi'\right) = \frac{1}{n^2}\sum_{i\neq j}p(\pi)\,p(\pi') = \left(1 - \frac{1}{n}\right)p(\pi)\,p(\pi'),$$

from which we have that $\operatorname{Cov}(x_\pi, x_{\pi'}) = -\frac{1}{n}p(\pi)\,p(\pi')$ and thus $\operatorname{Cov}(\tilde{x}_\pi, \tilde{x}_{\pi'}) = -p(\pi)\,p(\pi')$. □

Note that, because of the concentration of $x(\sigma^n)$ around its mean, and our assumption that for every n and for every candidate $a \in [k]$, we have $\mathbb{P}(f(\sigma^n) = a) \geq \varepsilon$, it is necessary that for every $\alpha > 0$ and for every candidate $a \in [k]$ there exists $y \in \Delta^{k!}$ such that $\left\|y - \mathbb{E}(x(\sigma_1))\right\|_1 \leq \alpha$ and $F(y) = a$.

Denote by μ the distribution of $N(0,\Sigma)$ and let $\tilde{X}$ denote a random variable distributed according to μ. Note that μ is a degenerate multivariate normal distribution, as the support of μ concentrates on the hyperplane H_0 where the coordinates sum to zero. (This is because $\sum_{\pi\in S_k}\tilde{x}(\sigma^n)_\pi = 0$.)

The underlying function $F : \Delta^{k!} \to [k]$ corresponds to a function $\tilde{F} : \mathbb{R}^{k!}|_{H_0} \to [k]$ in the Gaussian limit, and this function $\tilde{F}$ partitions $\mathbb{R}^{k!}|_{H_0}$ into k parts based on the outcome of $\tilde{F}$. We denote these parts by $\left\{\tilde{W}_a\right\}_{a\in[k]}$. We need the following definitions and properties of boundaries, analogous to those above.

Definition 3.7 (Interior and boundaries of a partition)**.** *We say that $\tilde{x} \in \mathbb{R}^{k!}|_{H_0}$ is an interior point of the partition $\left\{\tilde{W}_a\right\}_{a\in[k]}$ induced by $\tilde{F}$ if there exists $\alpha > 0$ such that for all $\tilde{y} \in \mathbb{R}^{k!}|_{H_0}$ for which $\|\tilde{x} - \tilde{y}\|_1 \leq \alpha$, we have $\tilde{F}(\tilde{x}) = \tilde{F}(\tilde{y})$. Otherwise, we say that $\tilde{x} \in \mathbb{R}^{k!}|_{H_0}$ is on the boundary of the partition, which we denote by $\tilde{B}$.*

Lemma 3.8. *If the boundary B comes from a hyperplane rule, i.e., B is contained in the union of ℓ affine hyperplanes in $\Delta^{k!}$, then $\tilde{B}$ is contained in the union of $\tilde{\ell}$ hyperplanes of $\mathbb{R}^{k!}|_{H_0}$, where $\tilde{\ell} \leq \ell$.*

Proof. Two things can happen to an affine hyperplane H of $\Delta^{k!}$ when we take the Gaussian limit: (1) if $\mathbb{E}(x(\pi)) \in H$, then translation by $\mathbb{E}(x(\pi))$ takes H into a hyperplane $\tilde{H}$ of $\mathbb{R}^{k!}|_{H_0}$, and since $\tilde{H}$ goes through the origin, scaling (in particular by $\sqrt{n}$) does not move this hyperplane; (2) if $\mathbb{E}(x(\pi)) \notin H$, then translation by $\mathbb{E}(x(\pi))$ takes H into an affine hyperplane $\tilde{H}$ of $\mathbb{R}^{k!}|_{H_0}$ that does not go through the origin, and then scaling by $\sqrt{n}$ moves $\tilde{H}$ to an affine hyperplane of $\mathbb{R}^{k!}|_{H_0}$ whose L^2 distance from the origin is proportional to $\sqrt{n}$, so in the $n \to \infty$ limit this affine hyperplane "vanishes". $\square$

Definition 3.8 (Blowup of boundary)**.** *For $\alpha > 0$, we define the blowup of the boundary $\tilde{B}$ by α to be*

$$\tilde{B}^{+\alpha} = \left\{ \tilde{y} \in \mathbb{R}^{k!}|_{H_0} : \exists \tilde{x} \in \tilde{B} \text{ such that } \left\| \tilde{x} - \tilde{y} \right\|_1 \leq \alpha \right\}.$$

Let us focus specifically on a coalition of size $c\sqrt{n}$ for some (small) constant c. Corollary 3.6 implies the following.

Corollary 3.9. *For hyperplane rules the limit of the probability that in an election with n voters some coalition of size $c\sqrt{n}$ can change the outcome of the election is $\mu\left(\tilde{X} \in \tilde{B}^{+2c}\right)$.*

The following claim, together with Corollary 3.9, tells us that for hyperplane rules a coalition of size $c\sqrt{n}$ can change the outcome of the election with only small probability, given that c is sufficiently small, proving Part 2 of Theorem 3.2.

Claim 3.10. *Suppose that our SCF is a hyperplane rule, and in particular let $\left\{\tilde{H}_i\right\}_{i=1}^{M}$ be a collection of hyperplanes in $\mathbb{R}^{k!}|_{H_0}$ such that $\tilde{B} \subseteq \bigcup_{i=1}^{M} \tilde{H}_i$. Then*

$$\mu\left(\tilde{X} \in \tilde{B}^{+c}\right) \leq \sqrt{\frac{2}{\pi}} \frac{Mc}{\sqrt{\delta}}.$$

Proof. By our condition and a union bound we have

$$\mu\left(\tilde{X} \in \tilde{B}^{+c}\right) \leq \sum_{i=1}^{M} \mu\left(\tilde{X} \in \tilde{H}_i^{+c}\right).$$

For a hyperplane $\tilde{H}$ in $\mathbb{R}^{k!}|_{H_0}$, denote (one of) the corresponding unit normal vector(s) (in the hyperplane H_0) by u. Then

$$\tilde{H} = \left\{ \tilde{x} \in \mathbb{R}^{k!}|_{H_0} : u \cdot \tilde{x} = 0 \right\}$$

and since L^1 distance is always greater than L^2 distance, we have

$$\tilde{H}^{+c} \subseteq \left\{ \tilde{x} \in \mathbb{R}^{k!}|_{H_0} : \exists \tilde{y} \in \tilde{H} \text{ such that } \left\| \tilde{x} - \tilde{y} \right\|_2 \leq c \right\} = \left\{ \tilde{x} \in \mathbb{R}^{k!}|_{H_0} : |u \cdot \tilde{x}| \leq c \right\}.$$

Since $\tilde{X}$ is a multidimensional Gaussian r.v., $u \cdot \tilde{X}$ is a one-dimensional Gaussian r.v. (which is centered). Therefore

$$\mu \left(\tilde{X} \in \tilde{H}^{+c} \right) \leq \mu \left(u \cdot \tilde{X} \in [-c, c] \right) \leq \frac{2c}{\sqrt{2\pi \operatorname{Var}\left(u \cdot \tilde{X} \right)}}.$$

We have that

$$\operatorname{Var}\left(u \cdot \tilde{X} \right) = \mathbb{E}\left(u \cdot \tilde{X} \right)^2 = \mathbb{E}\left(u^T \tilde{X} \tilde{X}^T u \right) = u^T \Sigma u,$$

and so all that remains to show is that

$$\min_{u : \|u\| = 1, u \perp \mathbf{1}} u^T \Sigma u \geq \delta,$$

where $\mathbf{1}$ is the $k!$-dimensional vector having 1 in every coordinate.

Let $\lambda_1(\Sigma) \geq \lambda_2(\Sigma) \geq \cdots \geq \lambda_{k!}(\Sigma)$ denote the eigenvalues of Σ. Since Σ is positive semidefinite, all eigenvalues are nonnegative. We know that 0 is an eigenvalue of Σ (the corresponding eigenvector is $\mathbf{1}$), so $\lambda_{k!}(\Sigma) = 0$. By the variational characterization of eigenvalues we have

$$\min_{u : \|u\| = 1, u \perp \mathbf{1}} u^T \Sigma u = \lambda_{k!-1}(\Sigma),$$

and so we need to show that $\lambda_{k!-1}(\Sigma) \geq \delta$. To do this we use Weyl's inequalities.

Lemma 3.11 (Weyl's inequalities). *For an $m \times m$ matrix M, let $\lambda_1(M) \geq \lambda_2(M) \geq \cdots \geq \lambda_m(M)$ denote its eigenvalues. If A and C are $m \times m$ symmetric matrices then*

$$\begin{aligned} \lambda_j(A + C) &\leq \lambda_i(A) + \lambda_{j-i+1}(C) \qquad \text{if} \quad i \leq j, \\ \lambda_j(A + C) &\geq \lambda_i(A) + \lambda_{j-i+m}(C) \qquad \text{if} \quad i \geq j. \end{aligned}$$

We use Weyl's inequality for $A = \operatorname{diag}(p)$ and $C = -pp^T$. The eigenvalues of A are $\{p(\pi)\}_{\pi \in S_k}$, all of which are no less than δ. Since C has rank 1, all its eigenvalues but one are zero, and the single nonzero eigenvalue is $\lambda_{k!}(C) = -p^T p$. Since $\Sigma = \operatorname{diag}(p) - pp^T = A + C$, Weyl's inequality tells us that

$$\lambda_{k!-1}(\Sigma) \geq \lambda_{k!}(\operatorname{diag}(p)) + \lambda_{k!-1}\left(-pp^T \right) \geq \delta + 0 = \delta. \qquad \square$$

This implies that we have a lower bound of $\Omega(\sqrt{n})$ for the size of the coalition needed in order to change the outcome of the election for hyperplane rules. As mentioned before, most commonly occurring SCFs are in this class of rules: see Section 3.5 for many examples.

3.3.2.1 "Almost" hyperplane rules

Furthermore, the Gaussian limiting setting above is not sensitive to small changes to the voting rule for finite n. Consequently, for SCFs that are "almost" hyperplane rules (in a sense we make precise below), the same conclusion holds: a coalition of size $\Omega\left(\sqrt{n}\right)$ is needed in order to be able to change the outcome of the election with non-negligible probability. In particular, the same result holds for SCFs with arbitrary tie-breaking rules for ranking profiles which lie on one of the hyperplanes (e.g., the tie-breaking rule can depend on the number of voters n).

Definition 3.9 ("Almost" hyperplane rules)**.** *Fix a finite set of affine hyperplanes of the simplex $\Delta^{k!}$: $H_1, \ldots, H_\ell$. These partition the simplex into finitely many regions. Let $F : \Delta^{k!} \to [k]$ be a function which is constant on each such region, and let B denote the induced boundary. Then the sequence of SCFs $\{f_n\}_{n\geq 1}$, $f_n : S_k^n \to [k]$, is called an "almost" hyperplane rule if for every σ^n such that $x\left(\sigma^n\right) \notin B^{+o\left(1/\sqrt{n}\right)}$, we have*

$$f_n\left(\sigma^n\right) = F\left(x\left(\sigma^n\right)\right).$$

This SCF is called an "almost" hyperplane rule induced by the affine hyperplanes $H_1, \ldots, H_\ell$ and the function F.

Lemma 3.12. *Suppose that the sequence of SCFs $\{f_n\}_{n\geq 1}$, $f_n : S_k^n \to [k]$, is an "almost" hyperplane rule defined by ℓ hyperplanes. Then in the Gaussian limiting setting the boundary $\tilde{B}$ is contained in the union of $\tilde{\ell}$ hyperplanes of $\mathbb{R}^{k!}|_{H_0}$, where $\tilde{\ell} \leq \ell$.*

Proof. For finite n, the induced boundary of f_n in the simplex $\Delta^{k!}$ is contained in $B^{+o\left(1/\sqrt{n}\right)}$, by definition. Since in the Gaussian limit we scale by $\sqrt{n}$, the blowup by $o\left(1/\sqrt{n}\right)$ of the boundary B disappears in the limit, and hence we are back to the situation of Lemma 3.8. Consequently, the affine hyperplanes corresponding to our "almost" hyperplane rule either "disappear to infinity" or become hyperplanes of $\mathbb{R}^{k!}|_{H_0}$. □

Corollary 3.13. *Corollary 3.9 and Claim 3.10 hold for "almost" hyperplane rules as well.*

3.3.3 Smoothness of the phase transition

In this final subsection our goal is to show Parts 1 and 3 of Theorem 3.2. The existence of the limits in Part 1 follows immediately from the Gaussian limit described above; we do not detail this, but rather give formulas for these limiting probabilities. These then imply the properties described in Part 3 of the theorem.

In the following let the hyperplane rule be given by affine hyperplanes $H_1, \ldots, H_\ell$ of $\Delta^{k!}$ and the function $F : \Delta^{k!} \to [k]$; in the limiting setting denote by $\tilde{H}_1, \ldots, \tilde{H}_{\tilde{\ell}}$ the corresponding hyperplanes of $\mathbb{R}^{k!}|_{H_0}$ and denote by $\tilde{F} : \mathbb{R}^{k!}|_{H_0} \to [k]$ the corresponding function.

3.3.3.1 The quantities $\overline{q}$ and $\underline{q}$

For $\tilde{x} \in \mathbb{R}^{k!}|_{H_0}$ define

$$\alpha(\tilde{x}) := \inf_{\tilde{y}:\tilde{F}(\tilde{y}) \neq \tilde{F}(\tilde{x})} d_1(\tilde{x}, \tilde{y}), \qquad \beta(\tilde{x}) := \max_{a \in [k]} \inf_{\tilde{y}:\tilde{F}(\tilde{y})=a} d_1(\tilde{x}, \tilde{y}).$$

From the previous subsection it is then immediate that we can write

$$\overline{q}(c) = \mu\left(\tilde{X} : \alpha\left(\tilde{X}\right) \leq 2c\right),$$
$$\underline{q}(c) = \mu\left(\tilde{X} : \beta\left(\tilde{X}\right) \leq 2c\right).$$

It is important to note that the boundary $\tilde{B}$ is contained in the union of finitely many hyperplanes, $\tilde{H}_1, \ldots, \tilde{H}_{\tilde{\ell}}$, and thus the regions where $\tilde{F}$ is constant are convex cones which are the intersection of finitely many halfspaces. Consequently $\alpha(\tilde{x})$ is either $d_1(\tilde{x}, 0)$, where 0 denotes the origin of $\mathbb{R}^{k!}$, or it is $d_1\left(\tilde{x}, \tilde{H}_j\right)$ for some $1 \leq j \leq \tilde{\ell}$, where $d_1\left(\tilde{x}, \tilde{H}_j\right) = \inf_{\tilde{y} \in \tilde{H}_j} d_1(\tilde{x}, \tilde{y})$. If we scale $\tilde{x}$ by some positive constant λ, then the distance from the origin and from every hyperplane scales as well (i.e., $d_1(\lambda\tilde{x}, 0) = \lambda d_1(\tilde{x}, 0)$ and $d_1\left(\lambda\tilde{x}, \tilde{H}_j\right) = \lambda d_1\left(\tilde{x}, \tilde{H}_j\right)$), and thus for every $\lambda > 0$, we have $\alpha(\lambda\tilde{x}) = \lambda\alpha(\tilde{x})$. Consequently, if we write $\tilde{x} = \|\tilde{x}\|_2 \tilde{s}$, where $\tilde{s} \in S^{k!-1}$, and $S^{k!-1}$ denotes the $(k!-1)$-sphere (not to be confused with S_k^n, the set of ranking profiles on n voters and k candidates), then we have $\alpha(\tilde{x}) = \|\tilde{x}\|_2 \alpha(\tilde{s})$.

The same scaling property holds for β as well, and hence we have

$$\overline{q}(c) = \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \alpha\left(\tilde{S}\right) \leq 2c\right), \tag{3.3}$$
$$\underline{q}(c) = \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \beta\left(\tilde{S}\right) \leq 2c\right). \tag{3.4}$$

Recall that our condition that for every $a \in [k]$, we have $\mathbb{P}(f(\sigma) = a) \geq \varepsilon$, implies that for every $\eta > 0$ and for every $a \in [k]$ there exists $\tilde{x} \in \mathbb{R}^{k!}|_{H_0}$ such that $\|\tilde{x}\|_2 \leq \eta$ and $\tilde{F}(\tilde{x}) = a$. Consequently for every $\tilde{x} \in \mathbb{R}^{k!}|_{H_0}$ we must have $\alpha(\tilde{x}) \leq d_1(\tilde{x}, 0)$ and $\beta(\tilde{x}) \leq d_1(\tilde{x}, 0)$. In particular, for $\tilde{s} \in S^{k!-1}$ we have $d_1(\tilde{s}, 0) \leq \sqrt{k!} d_2(\tilde{s}, 0) = \sqrt{k!}$ and so $\alpha(\tilde{s}), \beta(\tilde{s}) \leq \sqrt{k!}$. This immediately implies that for every $c > 0$ we have

$$\underline{q}(c) \geq \mu\left(\tilde{X} : \left\|\tilde{X}\right\|_2 \leq \frac{2c}{\sqrt{k!}}\right) > 0.$$

To show that $\overline{q}(c) < 1$, note that since the boundary is contained in the union of finitely many hyperplanes, there exists $\tilde{s}^* \in S^{k!-1}$ such that $\alpha(\tilde{s}^*) > 0$. By continuity of α, there exists a neighborhood $U \subseteq S^{k!-1}$ of $\tilde{s}^*$ such that for every $\tilde{s} \in U$, $\alpha(\tilde{s}) \geq \alpha(\tilde{s}^*)/2$. For any $\tilde{x}$ such that $\tilde{x}/\|\tilde{x}\|_2 \in U$ and $\|\tilde{x}\|_2 > \frac{4c}{\alpha(\tilde{s}^*)}$, we have

$$\alpha(\tilde{x}) = \|\tilde{x}\|_2 \alpha(\tilde{x}/\|\tilde{x}\|_2) > \frac{4c}{\alpha(\tilde{s}^*)} \frac{\alpha(\tilde{s}^*)}{2} = 2c.$$

So consequently

$$\overline{q}(c) \leq 1 - \mu\left(\tilde{X} : \tilde{X}/\left\|\tilde{X}\right\|_2 \in U, \left\|\tilde{X}\right\|_2 > \frac{4c}{\alpha(\tilde{s}^*)}\right) < 1.$$

Finally, the fact that $\underline{q}(c)$ and $\overline{q}(c)$ are continuously differentiable follows from the formulas (3.3) and (3.4), since $\underline{q}(c)$ and $\overline{q}(c)$ are both written as the Gaussian volume of a subset of $\mathbb{R}^{k!}|_{H_0}$, and in both cases this subset grows continuously as c increases. The derivative of both $\underline{q}(c)$ and $\overline{q}(c)$ is bounded at zero (by Corollary 3.9 and Claim 3.10), while as $c \to \infty$ the derivative approaches zero, and since the derivative is continuous, it must be bounded by a constant for the whole half-line.

3.3.3.2 The quantities $\overline{r}$ and $\underline{r}$

In the previous setup when the coalition of size $c\sqrt{n}$ was not specified, the ranking profile could be changed arbitrarily within a Hamming ball of radius $c\sqrt{n}$. On the probability simplex $\Delta^{k!}$ this corresponded to an L^1 ball of radius $2c/\sqrt{n}$, and in the rescaled limiting setting it corresponded to an L^1 ball in $\mathbb{R}^{k!}|_{H_0}$ of radius $2c$. When the coalition of size $c\sqrt{n}$ is specified, things are slightly different. In particular, when we look at the probability distribution on the probability simplex $\Delta^{k!}$ induced by the distribution on ranking profiles (or, in the limiting setting, the Gaussian distribution on $\mathbb{R}^{k!}|_{H_0}$), then we have lost track of the votes of any specific coalition. Nonetheless, the Gaussian limiting setting still provides formulas for the limiting probabilities $\underline{r}(c)$ and $\overline{r}(c)$.

We can first draw a random ranking profile for the other $n - c\sqrt{n}$ voters not in the coalition, $\sigma^{n-c\sqrt{n}}$, and then the voters in the coalition can set their votes arbitrarily. The question is, how can the coalition affect the outcome of the vote? In particular, (a) can they change the outcome of the election, and (b) can they elect any candidate?

The ranking profile $\sigma^{n-c\sqrt{n}}$ corresponds to a point $x\left(\sigma^{n-c\sqrt{n}}\right)$ on the probability simplex $\Delta^{k!}$, and by setting their votes the coalition can move this point on the probability simplex in some neighborhood of $x\left(\sigma^{n-c\sqrt{n}}\right)$. We omit the calculation for finite n and only present the result in the limiting setting.

Suppose that the limiting ranking profile of the voters other than the coalition corresponds to the point $\tilde{x} \in \mathbb{R}^{k!}|_{H_0}$. Then the set of points the coalition can reach is the following:

$$R_c(\tilde{x}) := \left\{\tilde{y} \in \mathbb{R}^{k!}|_{H_0} : \forall \pi \in S_k : \tilde{y}_\pi - \tilde{x}_\pi + cp(\pi) \geq 0\right\}.$$

We can then define

$$\varphi(\tilde{x}) := \inf\left\{\gamma : \exists \tilde{y} \in R_\gamma(\tilde{x}) \text{ such that } \tilde{F}(\tilde{y}) \neq \tilde{F}(\tilde{x})\right\},$$
$$\psi(\tilde{x}) := \inf\left\{\gamma : \forall a \in [k]\, \exists \tilde{y} \in R_\gamma(\tilde{x}) \text{ such that } \tilde{F}(\tilde{y}) = a\right\},$$

and it follows immediately that we can then write

$$\overline{r}\left(c\right)=\mu\left(\tilde{X}:\varphi\left(\tilde{X}\right)\leq c\right),$$
$$\underline{r}\left(c\right)=\mu\left(\tilde{X}:\psi\left(\tilde{X}\right)\leq c\right).$$

In the same way as in Section 3.3.3.1 one can argue that φ and ψ scale: if $\lambda>0$ then $\varphi\left(\lambda\tilde{x}\right)=\lambda\varphi\left(\tilde{x}\right)$ and $\psi\left(\lambda\tilde{x}\right)=\lambda\psi\left(\tilde{x}\right)$. Hence we have

$$\overline{r}\left(c\right)=\mu\left(\tilde{X}:\left\|\tilde{X}\right\|_{2}\varphi\left(\tilde{S}\right)\leq c\right),\tag{3.5}$$
$$\underline{r}\left(c\right)=\mu\left(\tilde{X}:\left\|\tilde{X}\right\|_{2}\psi\left(\tilde{S}\right)\leq c\right).\tag{3.6}$$

For every $0<c<\infty$ we have $\overline{r}\left(c\right)\leq\overline{q}\left(c\right)<1$ (using Section 3.3.3.1). Let us now show that also $\underline{r}\left(c\right)>0$. We claim that for all $\tilde{s}\in S^{k!-1}|_{H_0}$, we have $\psi\left(\tilde{s}\right)\leq\frac{2}{\delta}$. This follows from the fact that if $\tilde{s}\in S^{k!-1}|_{H_0}$ then $S^{k!-1}|_{H_0}\subseteq R_{\frac{2}{\delta}}\left(\tilde{s}\right)$, which is true because if $\tilde{y}\in S^{k!-1}|_{H_0}$ then for all $\pi\in S_k$, we have $\tilde{y}_\pi-\tilde{s}_\pi+\frac{2}{\delta}p\left(\pi\right)\geq-1-1+\frac{2}{\delta}\delta=0$. Thus we have

$$\underline{r}\left(c\right)\geq\mu\left(\tilde{X}:\left\|\tilde{X}\right\|_{2}\leq\frac{c\delta}{2}\right)>0$$

as claimed.

Finally, the fact that $\underline{r}\left(c\right)$ and $\overline{r}\left(c\right)$ are continuously differentiable follows from the formulas (3.5) and (3.6) using an argument given above: $\underline{r}\left(c\right)$ and $\overline{r}\left(c\right)$ are written as the Gaussian volume of subsets of $\mathbb{R}^{k!}|_{H_0}$, and these subsets grow continuously as c increases. The derivative of both $\underline{r}\left(c\right)$ and $\overline{r}\left(c\right)$ is bounded at zero (by Corollary 3.9 and Claim 3.10), while as $c\to\infty$ the derivative approaches zero, and since the derivative is continuous, it must be bounded by a constant for the whole half-line. □

3.4 Decomposing $\mathbb{R}^d$ as the disjoint union of finitely many convex cones: only via hyperplanes

For self-containment, we reproduce here the main definitions and results of Kemperman [45] that make precise the claim used in the proof of Lemma 3.3 that the only way to decompose $\mathbb{Q}^d_{\geq 0}\setminus\{0\}$ into the disjoint union of finitely many $\mathbb{Q}$-convex cones is via hyperplanes. Kemperman's paper deals with convex sets in general, but here we summarize the results about convex *cones* that are relevant to us. Kemperman's results pertain to finite dimensional linear spaces and we will state them in this form; in the end we show how results for $\mathbb{R}^d_{\geq 0}$ follow immediately from these, and as a consequence we also obtain the claim used in the proof of Lemma 3.3.

Let us start with the main definitions. In the following, all linear spaces are over the reals and are finite dimensional. Let X be a linear space. A convex cone is a subset $K\subseteq X$ such that $x, y \in K$ and $\lambda > 0$ imply $x + y \in K$ and $\lambda x \in K$. (We do not require that $0 \in K$.) For a set $A \subseteq X$, denote its affine hull by $\mathrm{aff}\,(A)$, its convex hull by $\mathrm{cvx}\,(A)$, and its closure by $\mathrm{cl}\,(A)$. Note that if $K \subseteq X$ is a convex cone, then $\mathrm{aff}\,(K)$ is a linear subspace of X.

We define two special types of convex cones: basic convex cones and elementary convex cones.

Definition 3.10 (Basic convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. We say that K is a* basic convex cone *(in X) if K is a member $K = K_0$ of some partition*

$$X = K_0 \dot\cup K_1 \dot\cup \ldots \dot\cup K_r$$

of X into finitely many disjoint convex cones $\{K_i\}_{i=0}^r$.

Note that any linear subspace Y of X is a basic convex cone, from which it immediately follows that K is a basic convex cone in X if and only if it is a basic convex cone in $\mathrm{aff}\,(K)$.

In order to define elementary convex cones, we need a few more definitions.

Definition 3.11 (Open polyhedral convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. We say that K is an* open polyhedral convex cone *relative to X if K can be expressed as the intersection of finitely many open halfspaces $H_1, \ldots, H_\ell$ of X, each of which has the origin on its boundary. The whole linear space X is an open polyhedral convex cone with $\ell = 0$.*

Definition 3.12 (Relatively open polyhedral convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. Then K is a* relatively open polyhedral convex cone *if either $K = \emptyset$ or K is an open polyhedral convex cone relative to* $\mathrm{aff}\,(K)$*.*

Definition 3.13 (Elementary convex cone)**.** *Let K be a convex cone in a finite dimensional linear space X. We say that K is an* elementary convex cone *if K can be represented as a disjoint union of finitely many relatively open polyhedral convex cones.*

To understand these definitions better, let us consider an example. We can write $\mathbb{R} = (-\infty, 0) \dot\cup [0, \infty)$, and so it is immediate that both $(-\infty, 0)$ and $[0, \infty)$ are basic convex cones. Clearly $(-\infty, 0)$ and $(0, \infty)$ are open polyhedral convex cones, and therefore elementary convex cones too. While $\{0\}$ is not an open polyhedral convex cone in $\mathbb{R}$, it is a *relatively open* polyhedral convex cone. Consequently $[0, \infty)$ is an elementary convex cone as well, since we can write $[0, \infty) = \{0\} \dot\cup (0, \infty)$. In particular, elementary convex cones are not necessarily open.

The main result of Kemperman concerning convex cones is the following [45, Theorem 2].

Theorem 3.14. *Let K be a convex cone in $\mathbb{R}^d$. Then K is a basic convex cone if and only if it is an elementary convex cone.*

In Lemma 3.3 we only use the "only if" direction, and we thus leave the proof of the "if" direction as an exercise for the reader.

Proof of "only if" direction. Let X be a finite dimensional linear space and let K be a basic convex cone in X of dimension $d = \dim(K) = \dim(Y)$, where $Y = \operatorname{aff}(K)$. We prove by induction on d the following:

(i) The relative interior of K, denoted by K^0, is a relatively open polyhedral convex cone.

(ii) If $K^0 \neq Y$, then denote by $F_1, \dots, F_\ell$ the $(d-1)$-dimensional hyperplanes in Y corresponding to the finitely many faces of the polyhedron $\operatorname{cl}(K) = \operatorname{cl}(K^0)$. Then the convex cones $F_i \cap K$, where $i = 1, \dots, \ell$, are elementary convex cones of dimension at most $d-1$ (but they need not be disjoint).

(iii) The convex cone K is also an elementary convex cone.

If $K = \emptyset$, then properties (i) - (iii) hold. If $d = 0$, then necessarily $K = \{0\}$, since K is a convex cone, and again K satisfies properties (i) - (iii) above.

So we may assume that $d \geq 1$ and that each basic convex cone of dimension at most $d-1$ satisfies properties (i) - (iii) above. Since K is a basic convex cone, there exists a partition

$$Y = K_0 \dot\cup K_1 \dot\cup \dots \dot\cup K_r \tag{3.7}$$

of Y into finitely many disjoint convex cones $\{K_j\}_{j=0}^r$, with $K_0 = K$. We may assume that $r \geq 0$ is minimal, and hence the K_j are non-empty. Note that K^0 is also non-empty since $\dim(K) = \dim(Y)$.

If $r = 0$ then $K = K_0 = Y$ and the properties (i) - (iii) above are immediately satisfied, so we may assume that $r \geq 1$. For $j = 1, \dots, r$, let H_j be a hyperplane in Y which separates the convex cone $K = K_0$ with non-empty interior K^0 from the non-empty convex cone K_j. (Such hyperplanes exist by the hyperplane separation theorem, and, moreover, each such hyperplane goes through the origin, because each K_j contains at least one point from every open ball around the origin, since each K_j is a cone.) Let H_j^0 be the associated open half space in Y which contains the interior K^0 of K. Let

$$L^0 = H_1^0 \cap \dots \cap H_r^0.$$

Then L^0 is a polyhedral convex cone, which is open relative to Y, and contains the interior K^0 of K.

We claim that $L^0 = K^0$. It is enough to show that $L^0 \subseteq K$, because then $L^0 \subseteq K^0$ follows from the definition of K^0. Suppose on the contrary that there exists $x \in L^0$ such that $x \notin K$. Then from the partition (3.7) there must exist an index $1 \leq j \leq r$ with $x \in K_j$. This implies that $x \notin H_j^0$ and thus $x \notin L^0$, which is a contradiction. This proves (i).

Now let us show (ii). By (3.7), we can write the linear space F_i as the disjoint union of the convex cones $F_i \cap K_j$, where $j = 0, \dots, r$, and thus $F_i \cap K$ is a basic convex cone and hence, by induction, an elementary convex cone.

Finally, let us show that K is an elementary convex cone. Since K^0 is a polyhedral convex cone which is open relative to Y, it only remains to show that $K \setminus K^0$ can be written as a finite disjoint union of relatively open polyhedral convex cones. By (ii), we can write $K \setminus K^0$ as the finite union of elementary convex cones:

$$K \setminus K^0 = \cup_{i=1}^{\ell} (F_i \cap K),$$

so what remains is to show that we can write this as a finite *disjoint* union of relatively open polyhedral convex cones. We may assume w.l.o.g. that $F_i \cap K \neq \emptyset$ for all i and that $(F_i \cap K) \nsubseteq (F_j \cap K)$ for all $i \neq j$ (otherwise we can leave out $F_i \cap K$ from the union).

We claim that then for every i,

$$\operatorname{rel\,int}(F_i \cap K) \subseteq (F_i \cap K) \setminus \bigcup_{j \neq i} (F_j \cap F_i \cap K), \tag{3.8}$$

from which it immediately follows that $\operatorname{rel\,int}(F_i \cap K) \cap \operatorname{rel\,int}(F_j \cap K) = \emptyset$ for $i \neq j$. To show (3.8), let the two open halfspaces on either side of the hyperplane F_j be denoted by F_j^+ and F_j^-. W.l.o.g. assume that $K \cap F_j^- = \emptyset$. Since $(F_i \cap K) \nsubseteq (F_j \cap K)$, we must have $(F_i \cap K) \cap F_j^+ \neq \emptyset$. Let $x \in (F_i \cap K) \cap F_j^+$ and let $y \in F_j \cap F_i \cap K$. Since $F_i \cap K$ is convex, the interval from x to y is contained in $F_i \cap K$, but because $(F_i \cap K) \cap F_j^- = \emptyset$, no points on this line past the point y can be in $F_i \cap K$; hence $y \notin \operatorname{rel\,int}(F_i \cap K)$.

Since $F_i \cap K$ is a basic convex cone, $\operatorname{rel\,int}(F_i \cap K)$ is a relatively open polyhedral convex cone by induction. If $F_i \cap K = \operatorname{aff}(F_i \cap K)$ then $\operatorname{rel\,int}(F_i \cap K) = F_i \cap K$. If not, then denote by $F_{i,1}, \ldots, F_{i,\ell_i}$ the hyperplanes in $\operatorname{aff}(F_i \cap K)$ corresponding to the finitely many faces of the polyhedron $\operatorname{cl}(F_i \cap K)$. By induction, the convex cones $F_{i,j} \cap F_i \cap K$, where $j = 1, \ldots, \ell_i$, are elementary convex cones, and we can write

$$K \setminus K^0 = \left(\dot\cup_{i=1}^{\ell} \operatorname{rel\,int}(F_i \cap K)\right) \dot{\bigcup} \left(\cup_{i=1}^{\ell} \cup_{j=1}^{\ell_i} (F_{i,j} \cap F_i \cap K)\right).$$

What remains to be shown is that $\cup_{i=1}^{\ell} \cup_{j=1}^{\ell_i} (F_{i,j} \cap F_i \cap K)$ can be written as a finite disjoint union of relatively open polyhedral convex cones; this follows by iterating the previous argument. □

Let us now show that $\mathbb{R}^d_{\geq 0}$ is a basic convex cone in $\mathbb{R}^d$. For $i = 1, \ldots, d$, define the closed halfspace $H_i^{\geq 0} = \left\{x \in \mathbb{R}^d : x_i \geq 0\right\}$ and its complement $H_i^{<0} = \left\{x \in \mathbb{R}^d : x_i < 0\right\}$, and from these define the convex cones

$$K_i = H_1^{\geq 0} \cap \cdots \cap H_{i-1}^{\geq 0} \cap H_i^{<0}, \qquad i = 1, \ldots, d.$$

Then we can write $\mathbb{R}^d$ as the disjoint union of the convex cones $\mathbb{R}^d_{\geq 0}$ and $K_1, \ldots, K_d$, showing that indeed $\mathbb{R}^d_{\geq 0}$ is a basic convex cone. This implies that if we can write $\mathbb{R}^d_{\geq 0}$ as the disjoint union of the convex cones $C_1, \ldots, C_r$, then each C_i is a basic convex cone, and hence, by Theorem 3.14, an elementary convex cone.

Now let us turn to the claim in the proof of Lemma 3.3. In Lemma 3.3, we write $\mathbb{Q}^{k!}_{\geq 0} \setminus \{0\}$ as the disjoint union of finitely many $\mathbb{Q}$-convex cones: $\mathbb{Q}^{k!}_{\geq 0} \setminus \{0\} = C_0 \dot\cup C_1 \dot\cup \ldots \dot\cup C_r$. For

$i = 0, \ldots, r$, let $\tilde{C}_i = \mathrm{cvx}\,(C_i)$. It is known (see, e.g., [79]) that $C_i = \mathbb{Q}^{k!} \cap \tilde{C}_i$. The $\tilde{C}_i$ are therefore disjoint convex cones which satisfy

$$\tilde{C}_0 \dot{\cup} \tilde{C}_1 \dot{\cup} \ldots \dot{\cup} \tilde{C}_r \subseteq \mathbb{R}^{k!}_{\geq 0} \tag{3.9}$$

and

$$\mathrm{cl}\left(\tilde{C}_0\right) \cup \mathrm{cl}\left(\tilde{C}_1\right) \cup \cdots \cup \mathrm{cl}\left(\tilde{C}_r\right) = \mathbb{R}^{k!}_{\geq 0}\,. \tag{3.10}$$

Our goal is to show that each $\tilde{C}_i$ is an elementary convex cone. Conditions (3.10) and (3.9) are very similar to the definition of a basic convex cone; in this spirit let us introduce the following definition.

Definition 3.14 (Basic convex cone up to closure)**.** *Let K_0 be a convex cone in a finite dimensional linear space X. We say that K_0 is a* basic convex cone up to closure *(in X) if there exist disjoint convex cones $K_1, \ldots, K_r$ such that*

$$K_0 \dot{\cup} K_1 \dot{\cup} \ldots \dot{\cup} K_r \subseteq X$$

and

$$\mathrm{cl}\,(K_0) \cup \mathrm{cl}\,(K_1) \cup \cdots \cup \mathrm{cl}\,(K_r) = X.$$

Since $\mathbb{R}^d_{\geq 0}$ is a basic convex cone, the $\tilde{C}_i$ above are basic convex cones up to closure.

In fact, every basic convex cone up to closure is an elementary convex cone; the proof is exactly the same as the one shown above for the "only if" direction of Theorem 3.14, one just needs to replace "basic convex cone" with "basic convex cone up to closure" everywhere in the proof, and make the appropriate changes. Moreover, the other direction of Theorem 3.14 implies that actually every basic convex cone up to closure is a basic convex cone.

Hence the $\tilde{C}_i$ are elementary convex cones, which is what we need in Lemma 3.3.

3.5 Most voting rules are hyperplane rules: examples

In the following we show that the following voting rules are all hyperplane rules: positional scoring rules, instant-runoff voting, Coombs' method, contingent vote, the Kemény-Young method, Bucklin voting, Nanson's method, Baldwin's method, and Copeland's method.

- **Positional scoring rules.** Let $w \in \mathbb{R}^k$ be a weight vector. Given a ranking profile vector σ, the (normalized) *score* of candidate $a \in [k]$ is $s_a = \frac{1}{n} \sum_{i=1}^{n} w\left(\sigma_i^{-1}(a)\right)$. The *positional scoring rule* associated to the weight vector w elects the candidate who has the highest score. (In case of a tie, there is some tie-breaking rule, but we do not care about this here.) We denote such a SCF on n voters by f_n^w. Examples include plurality (with weight vector $w = (1, 0, 0, \ldots, 0)$), Borda count (with weight vector $w = (k-1, k-2, \ldots, 0)$) and veto (with weight vector $w = (1, 1, \ldots, 1, 0)$).

To a sequence of SCFs $\{f_n^w\}_{n\geq 1}$ we can associate a function $F^w : \Delta^{k!} \to [k]$ in the following way. For a candidate $a \in [k]$ and $x \in \Delta^{k!}$, define the (normalized) score $s_a(x) = \sum_{\pi \in S_k} x_\pi w(\pi^{-1}(a))$, and let

$$F^w(x) := \arg\max_{a\in[k]} s_a(x),$$

if this arg max is unique, and if it is not unique, then there is some tie-breaking rule. This construction guarantees that $f_n^w = F^w|_{D_n}$. For candidates $a \neq b$, define

$$H_{a,b} := \left\{x \in \Delta^{k!} : s_a(x) = s_b(x)\right\},$$

which is an affine hyperplane of the probability simplex $\Delta^{k!}$. Clearly the boundary B^w is contained in the union of $\binom{k}{2}$ such affine hyperplanes:

$$B^w \subseteq \bigcup_{a\neq b\in[k]} H_{a,b}.$$

- **Instant-runoff voting.** If a candidate receives absolute majority of first preference votes, then that candidate wins. If no candidate receives an absolute majority, then the candidate with fewest top votes is eliminated. In the next round the votes are counted again, with each ballot counted as one vote for the advancing candidate who is ranked highest on that ballot. This is repeated until the winning candidate receives a majority of the vote against the remaining candidates.

  The boundary corresponds to two kinds of situations: either (1) there is a tie at the top at the end, when only two candidates remain; or (2) there is a tie for eliminating a candidate at the end of one of the rounds. Technically situation (1) is also contained in situation (2), since at the very end one can view choosing a winner as eliminating the second placed candidate. One can see that if candidates a and b are tied for elimination after candidates $C \subseteq [k] \setminus \{a,b\}$ (where $C = \emptyset$ is allowed) have been eliminated, then necessarily

$$\sum_{C'\subseteq C} \sum_{\substack{\{\pi(1),\dots,\pi(|C'|)\}=C',\\ \pi(|C'|+1)=a}} x_\pi = \sum_{C'\subseteq C} \sum_{\substack{\{\pi(1),\dots,\pi(|C'|)\}=C',\\ \pi(|C'|+1)=b}} x_\pi. \tag{3.11}$$

  Consequently, denoting by $s_{a,C}(x)$ the quantity on the left hand side of (3.11), the boundary B is contained in the union of at most $k^2 2^k$ affine hyperplanes:

$$B \subseteq \bigcup_{a\neq b} \bigcup_{C\subseteq[k]\setminus\{a,b\}} \left\{x \in \Delta^{k!} : s_{a,C}(x) = s_{b,C}(x)\right\}.$$

- **Coombs' method.** This is similar to IRV, but the elimination rule is different. If a candidate receives absolute majority of first preference votes, then that candidate wins. If no candidate receives an absolute majority, then the candidate who is ranked last by the most voters is eliminated. In the next round the votes are counted again, with each ballot counted as one vote for the advancing candidate who is ranked highest on that ballot. This is repeated until the winning candidate receives a majority of the vote against the remaining candidates.

The boundary corresponds to two kinds of situations: either (1) there is a tie at the top at the end, when only two candidates remain; or (2) there is a tie for eliminating a candidate at the end of one of the rounds. Technically situation (1) is also contained in situation (2), since at the very end one can view choosing a winner as eliminating the second placed candidate. One can see that if candidates a and b are tied for elimination after candidates $C \subseteq [k] \setminus \{a, b\}$ (where $C = \emptyset$ is allowed) have been eliminated, then necessarily

$$\sum_{C' \subseteq C} \sum_{\substack{\{\pi(k),\ldots,\pi(k-|C'|+1)\}=C',\\ \pi(k-|C'|)=a}} x_\pi = \sum_{C' \subseteq C} \sum_{\substack{\{\pi(k),\ldots,\pi(k-|C'|+1)\}=C',\\ \pi(k-|C'|)=b}} x_\pi. \tag{3.12}$$

Consequently, denoting by $s_{a,C}(x)$ the quantity on the left hand side of (3.12), the boundary B is contained in the union of at most $k^2 2^k$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \bigcup_{C \subseteq [k] \setminus \{a,b\}} \left\{ x \in \Delta^{k!} : s_{a,C}(x) = s_{b,C}(x) \right\}.$$

- **Contingent vote.** This is also similar to IRV, except here all but two candidates get eliminated after the first round. If a candidate receives absolute majority of first preference votes, then he/she wins. If no candidate receives an absolute majority, then all but the top two leading candidates are eliminated and there is a second count, where the votes of those who supported an eliminated candidate are redistributed among the two remaining candidates. The candidate who then achieves absolute majority wins.

  Here the boundary B corresponds to two kinds of situations: either (1) there are two distinct top candidates, and when the votes of the voters who voted for other candidates are redistributed, then the two top candidates are in a dead heat; or (2) there are two or more candidates who receive an equal number of votes in the first round. Both of these situations can be described as subsets of affine hyperplanes, and so B is contained in the union of at most $k(k-1)$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \left\{ x \in \Delta^{k!} : \sum_{\pi:\pi(1)=a} x_\pi + \sum_{\pi:\pi(1)\notin\{a,b\}, a \overset{\pi}{>} b} x_\pi = \sum_{\pi:\pi(1)=b} x_\pi + \sum_{\pi:\pi(1)\notin\{a,b\}, b \overset{\pi}{>} a} x_\pi \right\}$$
$$\cup \bigcup_{a \neq b} \left\{ x \in \Delta^{k!} : \sum_{\pi:\pi(1)=a} x_\pi = \sum_{\pi:\pi(1)=b} x_\pi \right\}.$$

- **Kemény-Young method.** Denote by K the Kendall tau distance, which is a metric on permutations which counts the number of pairwise disagreements between the two permutations, i.e.,

$$K(\tau_1, \tau_2) = \sum_{\{a,b\}} \mathbf{1}\left[a \text{ and } b \text{ are in the opposite order in } \tau_1 \text{ and } \tau_2\right],$$

where the sum is over all unordered pairs of distinct candidates. Given a ranking profile σ^n, the Kemény-Young method selects the ranking which minimizes the sum of Kendall tau distances from the votes:

$$\tau^* = \arg\min_{\tau} \sum_{i=1}^{n} K(\sigma_i, \tau),$$

and then the winner of the election is declared to be $\tau^*(1)$. For us it will be convenient to write τ^* as

$$\tau^* = \arg\min_{\tau} \sum_{\pi} x_\pi(\sigma^n) K(\pi, \tau).$$

Here if we are on the boundary B then there must exist two rankings τ_1 and τ_2 such that $\tau_1(1) \neq \tau_2(1)$ and $\sum_\pi x_\pi K(\pi, \tau_1) = \sum_\pi x_\pi K(\pi, \tau_2)$. Thus B is contained in the union of at most $(k!)^2$ affine hyperplanes:

$$B \subseteq \bigcup_{\tau_1 \neq \tau_2} \left\{ x \in \Delta^{k!} : \sum_\pi x_\pi K(\pi, \tau_1) = \sum_\pi x_\pi K(\pi, \tau_2) \right\}.$$

- **Bucklin voting.** First, every candidate gets a point from all the voters who ranked them at the top. If there is a candidate who has a majority (i.e., more than $n/2$ points), then that candidate wins. If not, then every candidate gets a point from all the voters who ranked them second. If there is a candidate who has more than $n/2$ points after this, then the candidate with the most points wins (there might be multiple candidates with more than $n/2$ points after a given round). This process is iterated until there is a candidate with more than $n/2$ points.

  Here a point on the boundary B corresponds to a situation where some pair of candidates have the same number of points after some number of rounds. Therefore B is contained in the union of at most $k^2(k-1)/2$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \bigcup_{m=1}^{k} \left\{ x \in \Delta^{k!} : \sum_{i=1}^{m} \sum_{\pi : \pi(i) = a} x_\pi = \sum_{i=1}^{m} \sum_{\pi : \pi(i) = b} x_\pi \right\}.$$

- **Nanson's method.** This is Borda count combined with a variation of the instant-runoff voting procedure. First, the Borda scores of all candidates are computed, and then those candidates with Borda score no greater than the average Borda score are eliminated. Then the Borda scores of each remaining candidate are recomputed, as if the eliminated candidates were not on the ballot. This is repeated until there is a final candidate left.

The boundary corresponds to situations when a candidate's Borda score exactly equals the average score after some candidates have been eliminated. For $C \subseteq [k]$, denote by $s_{a,C}(x)$ the score of candidate a after exactly the candidates in C have been eliminated ($s_{a,C}(x)$ is a linear function of $\{x_\pi\}_{\pi \in S_k}$), and denote by $\bar{s}_C(x)$ the average score of remaining candidates after exactly the candidates in C have been eliminated. The boundary B is contained in the union of at most $k2^k$ affine hyperplanes:

$$B \subseteq \bigcup_{a \in [k]} \bigcup_{C \subseteq [k] \setminus \{a\}} \left\{ x \in \Delta^{k!} : s_{a,C}(x) = \bar{s}_C(x) \right\}.$$

- **Baldwin's method.** This is essentially Borda count combined with the instant-runoff voting procedure. First, the Borda scores of all candidates are computed, and then the candidate with the lowest score is eliminated. Then the Borda scores of each remaining candidate are recomputed, as if the eliminated candidate were not on the ballot. This is repeated until there is a final candidate left.

  The boundary corresponds to ties for eliminating a candidate at the end of one of the rounds. Borrow the notation $s_{a,C}(x)$ from the previous example. The boundary B is thus contained in the union of at most $k^2 2^k$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \bigcup_{C \subseteq [k] \setminus \{a,b\}} \left\{ x \in \Delta^{k!} : s_{a,C}(x) = s_{b,C}(x) \right\}.$$

- **Copeland's method.** This is a pairwise aggregation method: every candidate gets 1 point for each other candidate it beats in a pairwise majority election, and 1/2 a point for each candidate it ties with in a pairwise majority election. The winner is the candidate who receives the most points. This method corresponds to cutting the simplex $\Delta^{k!}$ up into finitely many regions via $\binom{k}{2}$ affine hyperplanes, and in each region the winner is the candidate with the most points.

  While in the previous examples tie-breaking rules were not an issue, here it does become important. We do not care about tie-breaking rules when we are on an affine hyperplane where two candidates tie each other in a pairwise majority election. However, there are open regions in the intersection of halfspaces defined by the affine hyperplanes where candidates are tied at the top with having the same scores. In this case, in order for Copeland to be a hyperplane rule, we need to break ties in favor of the same candidate for the whole region. (This is also how Xia and Conitzer break ties for Copeland's method in [77].)

Using this tie-breaking rule Copeland's method is indeed a hyperplane rule, since the boundary is contained in the union of at most $\binom{k}{2}$ affine hyperplanes:

$$B \subseteq \bigcup_{a \neq b} \left\{ x \in \Delta^{k!} : \sum_{\pi : a \overset{\pi}{>} b} x_\pi = \sum_{\pi : b \overset{\pi}{>} a} x_\pi \right\}.$$

Part II

Influences in Growing Networks

Chapter 4

The influence of the seed in growing networks

4.1 Introduction

In Part II of the thesis we study statistical inference questions in growing networks, and in particular we are interested in the following question. Suppose we generate a large graph according to some model of randomly growing graphs—can we say anything about the initial (seed) graph? A precise answer to this question could lead to new insights for the diverse applications of randomly growing graphs, for instance in the area of control. There have been several heuristic approaches to this question, and experimental evidence of the seed's influence already exists in the literature, see, e.g., [39, 57, 68]. In this chapter we initiate the theoretical study of the seed's influence. For sake of simplicity we focus on *trees*, in particular those grown according to linear preferential attachment and uniform attachment, though the questions we study are of interest more broadly.

In general, one can define a sequence of randomly growing graphs as follows. For $n \geq k \geq 2$ and a graph T on k vertices, define the random graph $G(n,T)$ by induction. First, set $G(k,T) = T$. Then, given $G(n,T)$, $G(n+1,T)$ is formed from $G(n,T)$ by adding a new vertex and some new edges according to some adaptive rule. We focus on randomly growing trees, where at each time step a single edge is added connecting the new vertex to an existing one. In particular, we study preferential attachment and uniform attachment trees, which we now define.

For a tree T denote by $d_T(u)$ the degree of vertex u in T, $\Delta(T)$ the maximum degree in T, and $\vec{d}(T) \in \mathbb{N}^{\mathbb{N}}$ the vector of degrees arranged in decreasing order.[1] We refer to $\vec{d}(T)$ as the degree profile of T. For $n \geq k \geq 2$ and a tree T on k vertices we define the random tree $\mathrm{PA}(n,T)$ by induction. First, let $\mathrm{PA}(k,T) = T$. Then, given $\mathrm{PA}(n,T)$, $\mathrm{PA}(n+1,T)$ is formed from $\mathrm{PA}(n,T)$ by adding a new vertex u and a new edge uv where v is selected at

[1] We artificially continue the vector of degrees with zeros after the $|T|^{\text{th}}$ coordinate to put all degree profiles on the same space.

random among vertices in $\mathrm{PA}(n,T)$ according to the following probability distribution:

$$\mathbb{P}\left(v=i \,|\, \mathrm{PA}(n,T)\right) = \frac{d_{\mathrm{PA}(n,T)}(i)}{2\,(n-1)}.$$

This model was introduced in [48] under the name *Random Plane-Oriented Recursive Trees* but we use here the modern terminology of Preferential Attachment graphs, see [6, 11].

Uniform attachment trees are even more simply defined. For $n \geq k \geq 2$ and a tree T on k vertices, we define the random tree $\mathrm{UA}\,(n,T)$ by induction as follows. First, let $\mathrm{UA}\,(k,T)=T$. Then, given $\mathrm{UA}\,(n,T)$, $\mathrm{UA}\,(n+1,T)$ is formed from $\mathrm{UA}\,(n,T)$ by adding a new vertex u and adding a new edge uv where the vertex v is chosen uniformly at random among vertices of $\mathrm{UA}\,(n,T)$, independently of all past choices.

We want to understand whether there is a relation between T and $G(n,T)$ when n becomes very large. We investigate three ways to make this question more formal, which correspond to three different points of view on the limiting graph obtained by letting n go to infinity.

4.2 Notions of influence and main results

The least refined point of view is to consider the graph $G(\infty,T)$ defined on a countable set of vertices that one obtains by continuing the graph growing process indefinitely. In the case of preferential attachment trees, Kleinberg and Kleinberg [46] observed that the seed does not have any influence in this sense: indeed, for any tree T, almost surely, $\mathrm{PA}(\infty,T)$ will be the unique isomorphism type of tree with countably many vertices and in which each vertex has infinite degree. In fact, this statement holds for *any* model where the degree of each fixed vertex diverges to infinity as the tree grows; this is the case for uniform attachment trees as well.

Next we consider the much more subtle and fine-grained notion of a weak local limit introduced in [8]. This notion of graph limits contains information about local neighborhoods of a typical vertex (see Section 5.4 for a precise definition), and is more powerful than the one considered in the previous paragraph as it can, e.g., distinguish between models having different limiting degree distributions. The question is thus whether we can say anything about T using only local statistics of $G(n,T)$.

The weak local limit of the preferential attachment graph was first studied in the case of trees in [66] using branching process techniques, and then later in general in [9] using Pólya urn representations. These papers show that $\mathrm{PA}(n,S_2)$ tends to the so-called Pólya-point graph in the weak local limit sense, where S_2 is the 2-vertex star, i.e., two vertices connected by a single edge. Our first theorem utilizes this result to obtain the same for an arbitrary seed:

Theorem 4.1. *For any tree T the weak local limit of* $\mathrm{PA}(n,T)$ *is the Pólya-point graph described in [9] with $m=1$.*

This result says that “locally” (in the Benjamini-Schramm sense) the seed has no effect. The intuitive reason for this result is that in the preferential attachment model most nodes are far from the seed graph and therefore it is expected that their neighborhoods will not reveal any information about it.

In the case of uniform attachment trees it is even simpler to see that the seed has no effect locally. This follows from the fact that with high probability all vertices in $\mathrm{UA}(n,T)$ have degree at most logarithmic in n, so any finite neighborhood of a vertex has size at most polylogarithmic in n.

Finally, we consider the most refined point of view, which we believe to be the most natural one for this problem as well as the richest one (both mathematically and in terms of insights for potential applications). First, we rephrase our main question in the terminology of hypothesis testing. Given two potential seed graphs S and T, and an observation R which is a graph on n vertices, one wishes to test whether $R \sim G(n,S)$ or $R \sim G(n,T)$. Our original question then boils down to whether one can design a test with asymptotically (in n) nonnegligible power. This is equivalent to studying the total variation distance between $G(n,S)$ and $G(n,T)$, where recall that the total variation distance between two random variables X and Y taking values in a finite space $\mathcal{X}$ with laws μ and ν is defined as $\mathrm{TV}\left(X,Y\right) = \frac{1}{2}\sum_{x \in \mathcal{X}} \left|\mu\left(x\right) - \nu\left(x\right)\right|$. Thus we naturally define

$$\delta(S,T) := \lim_{n\to\infty} \mathrm{TV}(G(n,S),G(n,T)), \tag{4.1}$$

where $G(n,S)$ and $G(n,T)$ are random elements in the finite space of unlabeled graphs with n vertices. This limit is well-defined because $\mathrm{TV}(G(n,S),G(n,T))$ is nonincreasing in n (since if $G(n,S) = G(n,T)$, then the evolution of the random graphs can be coupled such that $G(n',S) = G(n',T)$ for all $n' \geq n$) and always nonnegative.

One can propose a test with asymptotically nonnegligible power (i.e., a nontrivial test) if and only if $\delta(S,T) > 0$. We believe that in fact this is the case for natural models of randomly growing graphs (except in trivial situations). In particular, due to the work presented in this thesis, and also the paper by Curien et al. [21], we now know that this is true for preferential attachment and uniform attachment trees. Let δ_{PA} and δ_{UA}, respectively, denote the limiting total variation distance as in (4.1) in the case of preferential attachment and uniform attachment trees, respectively.

Theorem 4.2. *For any trees S and T that are nonisomorphic and have at least 3 vertices, we have that $\delta_{\mathrm{PA}}(S,T) > 0$.*

Theorem 4.3. *For any trees S and T that are nonisomorphic and have at least 3 vertices, we have that $\delta_{\mathrm{UA}}(S,T) > 0$.*

In some cases our methods can say more. As a proof of concept we show the following results, which state that the limiting total variation distance between a fixed tree and a star can be arbitrarily close to 1 if the star is large enough. Let S_k denote the k-vertex star, i.e., the tree where a central vertex is connected to all $k-1$ other vertices.

Theorem 4.4. *For any fixed tree T one has*

$$\lim_{k \to \infty} \delta_{\mathrm{PA}}(S_k, T) = 1.$$

Theorem 4.5. *For any fixed tree T one has*

$$\lim_{k \to \infty} \delta_{\mathrm{UA}}(S_k, T) = 1.$$

4.3 Discussion

4.3.1 Chronological history of main results

In joint work with Sébastien Bubeck and Elchanan Mossel [14], which forms Chapter 5 of this thesis, we initiated the theoretical study of the influence of the seed in preferential attachment trees. We conjectured that Theorem 4.2 holds, and showed that it indeed does provided that the seed trees S and T have different degree profiles.

Theorem 4.6. *Let S and T be two finite trees on at least 3 vertices. If $\vec{d}(S) \neq \vec{d}(T)$, then $\delta_{\mathrm{PA}}(S,T) > 0$.*

We prove this theorem in Chapter 5, and in fact our proof shows a stronger statement, namely that different degree profiles lead to different limiting distributions for the (appropriately normalized) maximum degree. The smallest pair of trees that our method cannot as of yet distinguish is depicted in Figure 4.1.

Figure 4.1: Two trees with six vertices and $\vec{d}(S) = \vec{d}(T)$.

Following the posting of our results and conjectures on the preprint server arXiv, in a beautiful work, Curien et al. [21] proved that our conjecture is indeed true, i.e., that Theorem 4.2 holds. Their proof utilizes some of the ideas presented in Chapter 5, in particular by using statistics which are very similar to those we consider in Section 5.3.2. The proof approach of [21] is much more abstract than ours. By constructing and analyzing a large family of martingales, they are able to show that the limiting distribution of these martingales must differ when starting from two different trees. One of the advantages of the more computational proof presented in Chapter 5 is that it allows to more easily derive quantitative bounds for the limiting total variation distance in cases where our results show that the distance is nonzero.

Building on both of these previous works, we then studied the influence of the seed in uniform attachment trees, showing that the seed does indeed have an influence, as described in Theorem 4.3 (see also Theorem 4.5). This is joint work with Sébastien Bubeck, Ronen Eldan, and Elchanan Mossel [15], and forms Chapter 6 of this thesis. We next describe how preferential attachment and uniform attachment are related.

4.3.2 Comparison of preferential attachment and uniform attachment

The main idea in both works on preferential attachment trees (Chapter 5 and [21]), motivated by the rich-get-richer property of the preferential attachment model, is to consider various statistics based on large degree nodes, and to show that the initial seed influences the distribution of these statistics.

Consider, for instance, the problem of differentiating between the two seed trees in Figure 4.2. On the one hand, in S the degree of v_ℓ is greater than that of v_r, and this unbalancedness in the degrees likely remains as the tree grows according to preferential attachment. On the other hand, in T the degrees of v_ℓ and v_r are the same, so they will have the same distribution at larger times as well. This difference in the balancedness vs. unbalancedness of the degrees of v_ℓ and v_r is at the root of why the seed trees S and T are distinguishable in the preferential attachment model. A precise understanding of the balancedness properties of the degrees relies on the classical theory of Pólya urns.

In the uniform attachment model the degrees of vertices do not play a special role. In particular, in the example of Figure 4.2, v_ℓ and v_r will have approximately similar degrees in a large tree grown according to the uniform attachment model, irrespective of whether the seed tree is S or T. Nonetheless, we are able to distinguish the seed trees S and T, but the statistics we use to do this are based on more global balancedness properties of these trees.

An edge of a tree partitions the tree into two parts on either side of the edge. For most edges in a tree, this partition has very unbalanced sizes; for instance, if an edge is adjacent to a leaf, then one part contains only a single vertex. On the other hand, for edges that are in some sense "central" the partition is more balanced, in the sense that the sizes of the two parts are comparable. Intuitively, the edges of the seed tree will be among the "most central" edges of the uniform attachment tree at large times, and so we expect that the seed should influence the global balancedness properties of such trees.

Consider again the example of the two seed trees S and T in Figure 4.2. The edge e_0 partitions the tree into two parts: a subtree under v_ℓ and a subtree under v_r. In S these subtree sizes are unbalanced, and this likely remains the case as the tree grows according to uniform attachment. On the other hand, in T the subtree sizes are equal, and they will likely remain balanced as the tree grows. Again, Pólya urns play an important role, since the subtree sizes evolve according to a classical Pólya urn initialized by the subtree sizes in the seed tree. The difference in the balancedness vs. unbalancedness of the subtree sizes is at the root of why S and T are distinguishable in the uniform attachment model. To prove Theorem 4.3 we need to analyze statistics based on more general global balancedness properties of such trees, but the underlying intuition is what is described in the preceding

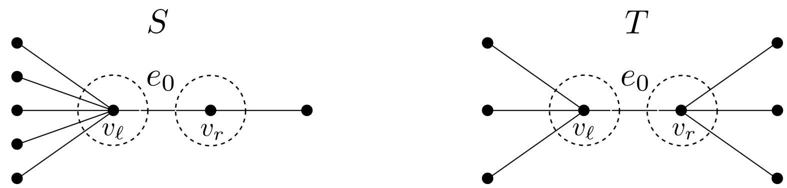

Preferential attachment. The degrees of v_ℓ and v_r are unbalanced in S but balanced in T, and this likely remains the case as the trees grow according to preferential attachment. This is at the root of why S and T are distinguishable as seed trees in the preferential attachment model.

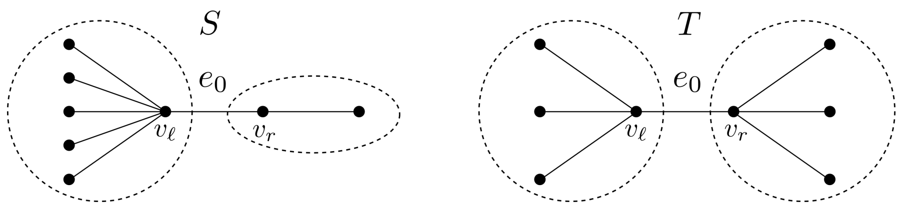

Uniform attachment. The sizes of the subtrees under v_ℓ and v_r are unbalanced in S but balanced in T, and this likely remains the case as the trees grow according to uniform attachment. This is at the root of why S and T are distinguishable as seed trees in the uniform attachment model.

Figure 4.2: Distinguishing between two trees requires different approaches for the uniform and the preferential attachment models.

paragraphs.

To formalize this intuition we essentially follow the proof scheme developed in [21]. However, the devil is in the details: since the underlying statistics are markedly different—in particular, statistics based on degrees are local, whereas those based on balancedness properties of subtree sizes are global—some of the essential steps of the proof become different. We provide a more detailed comparison to the work of [21] in Section 6.4, after we present our proof.

4.3.3 Further related work

A tree with node set $[n] := \{1, \dots, n\}$ is called recursive if the node numbers along the unique path from 1 to j increase for every $j \in \{2, \dots, n\}$. A stochastic process of random growing trees where nodes are labeled according to the time they are born is thus a sequence of random recursive trees. If we choose a recursive tree with node set $[n]$ uniformly at random, then the resulting tree has the same distribution as $\mathrm{UA}(n)$. A random tree grown according to the preferential attachment process starting from a single node is also known as a random plane-oriented recursive tree, see [48].

There is a large literature on random recursive trees and their various statistics; we refer the reader to the book by Drmota [24]. Of particular interest to the question we study here are recent works on a boundary theory approach to the convergence of random recursive trees in the limit as the tree size goes to infinity, see [26] and [36]. The main difference between these and the current work is that they consider labeled and rooted trees, whereas we are interested in what can be said about the seed given an unlabeled and unrooted copy of the tree.

4.4 Future directions and open problems

We conclude this chapter with suggestions for future directions and a collection of open problems.

1. Our results so far are essentially about the *testing* version of the problem. Can anything be said about the *estimation* version? Perhaps a first step would be to understand the multiple hypothesis testing problem where one is interested in testing whether the seed belongs to the family of trees $\mathcal{T}_1$ or to the family $\mathcal{T}_2$.

2. Starting from two seeds S and T with different spectrum, is it always possible to distinguish (with nontrivial probability) between $\mathrm{PA}(n, S)$ and $\mathrm{PA}(n, T)$ with spectral techniques? What about for uniform attachment and other models? More generally, it would be interesting to understand what properties are invariant under modifications of the seed.

3. Is it possible to give a combinatorial description of the metrics δ_{PA} and δ_{UA}?

4. Under what conditions on two tree sequences (T_k), (R_k) do we have $\lim_{k\to\infty} \delta(T_k, R_k) = 1$? In Theorems 4.4 and 4.5 we showed that a sufficient condition for δ_{PA} and δ_{UA} is to have $T_k = T$ and $R_k = S_k$. If T_k and R_k are independent (uniformly) random trees on k vertices, do we have $\lim_{k\to\infty} \mathbb{E}\delta(T_k, R_k) = 1$ for $\delta = \delta_{\mathrm{PA}}$ and $\delta = \delta_{\mathrm{UA}}$?

5. What can be said about the general preferential attachment and uniform attachment models, when multiple edges are added at each step? What about other models of randomly growing graphs?

6. A simple variant on the model studied in this paper is to consider probabilities of connection proportional to the degree of the vertex raised to some power α. For $\alpha = 1$ and $\alpha = 0$ the results of this thesis and those of [21] show that different seeds are distinguishable. What about for other α? Is $\delta_\alpha(S, T) > 0$ whenever S and T are nonisomorphic and have at least three vertices? (Here δ_α is defined analogously to δ_{PA} and δ_{UA} for general α.) What can be said about $\delta_\alpha(S, T)$ as a function of α? Is it monotone in α? Is it convex?

When $\alpha > 1$, i.e., in the case of superlinear preferential attachment, we expect the seed to have an influence in the strongest sense, i.e., that if S and T are nonisomorphic trees on at least three vertices, then

$$\mathrm{TV}\left(\mathrm{PA}_\alpha\left(\infty, S\right), \mathrm{PA}_\alpha\left(\infty, T\right)\right) > 0. \tag{4.2}$$

When $\alpha > 1$, Oliveira and Spencer [58] give a precise description of the infinite tree $\mathrm{PA}_\alpha\left(\infty, S_2\right)$, which contains exactly one vertex of infinite degree, with all other vertices having finite degree. From this it is possible to give a similar description of the infinite tree $\mathrm{PA}_\alpha\left(\infty, S\right)$ for any seed tree S. We believe that from this description it is possible to deduce that (4.2) holds for $\alpha > 1$, but have not pursued this question further.

Chapter 5

Preferential attachment trees

5.1 Overview

This chapter is devoted to proving the results described in the previous chapter on the influence of the seed in preferential attachment trees. To simplify notation, in this chapter we let $\delta \equiv \delta_{\mathrm{PA}}$.

In the next section we derive results on the limiting distribution of the maximum degree $\Delta(\mathrm{PA}(n,T))$ that are useful in proving Theorems 4.6 and 4.4, which we then prove in Section 5.3.1. In Section 5.3.2 we describe a particular way of generalizing the notion of maximum degree which we believe should provide an alternative way to prove Theorem 4.2. At present we are missing a technical result which we state separately as Conjecture 5.1 in the same section. The proof of Theorem 4.1 is in Section 5.4, while the proof of a key lemma described in Section 5.2 is presented in Section 5.5.

5.2 Useful results on the maximum degree

We first recall several results that describe the limiting degree distributions of preferential attachment graphs (Section 5.2.1), and from these we determine the tail behavior of the maximum degree in Section 5.2.2, which we then use in the proofs of Theorems 4.6 and 4.4. Throughout the chapter we label the vertices of $\mathrm{PA}(n,T)$ by $\{1,2,\ldots,n\}$ in the order in which they are added to the graph, with the vertices of the initial tree labeled in decreasing order of degree, i.e., satisfying $d_T(1) \geq d_T(2) \geq \cdots \geq d_T(|T|)$ (with ties broken arbitrarily). We also define the constant

$$c(a,b) = \frac{\Gamma(2a-2)}{2^{b-1}\Gamma(a-1/2)\Gamma(b)}, \tag{5.1}$$

which will occur multiple times.

5.2.1 Previous results

5.2.1.1 Starting from an edge

Móri [49] used martingale techniques to study the maximum degree of the preferential attachment tree starting from an edge, and showed that $\Delta(\mathrm{PA}(n, S_2))/\sqrt{n}$ converges almost surely to a random variable which we denote by $D_{\max}(S_2)$. He also showed that for each fixed $i \geq 1$, $d_{\mathrm{PA}(n,S_2)}(i)/\sqrt{n}$ converges almost surely to a random variable which we denote by $D_i(S_2)$, and furthermore that $D_{\max}(S_2) = \max_{i\geq 1} D_i(S_2)$ almost surely. In light of this, in order to understand $D_{\max}(S_2)$, it is useful to study $\{D_i(S_2)\}_{i\geq 1}$. Móri [49] computes the joint moments of $\{D_i(S_2)\}_{i\geq 1}$; in particular, we have (see [49, eq. (2.4)]) that for $i \geq 2$,

$$\mathbb{E}D_i(S_2)^r = \frac{\Gamma(i-1)\Gamma(1+r)}{\Gamma\left(i-1+\frac{r}{2}\right)}. \tag{5.2}$$

Using different methods and slightly different normalization, Peköz et al. [59] also study the limiting distribution of $d_{\mathrm{PA}(n,S_2)}(i)$; in particular, they give an explicit expression for the limiting density. Fix $s \geq 1/2$ and define

$$\kappa_s(x) = \Gamma(s)\sqrt{\frac{2}{s\pi}}\exp\left(-\frac{x^2}{2s}\right)U\left(s-1, \frac{1}{2}, \frac{x^2}{2s}\right)\mathbf{1}_{\{x>0\}},$$

where $U(a,b,z)$ denotes the confluent hypergeometric function of the second kind, also known as the Kummer U function (see [1, Chapter 13]); it can be shown that this is a density function. Peköz et al. [59] show that for $i \geq 2$ the distributional limit of

$$d_{\mathrm{PA}(n,S_2)}(i) \,/\, \left(\mathbb{E}d_{\mathrm{PA}(n,S_2)}(i)^2\right)^{1/2}$$

has density κ_{i-1} (they also give rates of convergence to this limit in the Kolmogorov metric). Let W_s denote a random variable with density κ_s. The moments of W_s (see [59, Section 2]) are given by

$$\mathbb{E}W_s^r = \left(\frac{s}{2}\right)^{r/2}\frac{\Gamma(s)\Gamma(1+r)}{\Gamma\left(s+\frac{r}{2}\right)}, \tag{5.3}$$

and thus comparing (5.2) and (5.3) we see that $D_i(S_2) \stackrel{d}{=} \sqrt{2/(i-1)}W_{i-1}$ for $i \geq 2$.

5.2.1.2 Starting from an arbitrary seed graph

Since we are interested in the effect of the seed graph, we desire similar results for $\mathrm{PA}(n,T)$ for an arbitrary tree T. One way of viewing $\mathrm{PA}(n,T)$ is to start growing a preferential attachment tree from a single edge and condition on it being T after reaching $|T|$ vertices; $\mathrm{PA}(n,T)$ has the same distribution as $\mathrm{PA}(n,S_2)$ conditioned on $\mathrm{PA}(|T|, S_2) = T$. Due to this, the almost sure convergence results of [49] carry over to the setting of an arbitrary seed tree. Thus for every fixed $i \geq 1$, $d_{\mathrm{PA}(n,T)}(i)/\sqrt{n}$ converges almost surely to a random variable which we denote by $D_i(T)$, $\Delta(\mathrm{PA}(n,T))/\sqrt{n}$ converges almost surely to a random variable which we denote by $D_{\max}(T)$, and furthermore $D_{\max}(T) = \max_{i\geq 1} D_i(T)$ almost surely.

In order to understand these limiting distributions, the basic observation is that for any i such that $1 \leq i \leq |T|$, the pair $\left(2(n-1) - d_{\mathrm{PA}(n,T)}(i), d_{\mathrm{PA}(n,T)}(i)\right)$ evolves according to a Pólya urn with replacement matrix $\left(\begin{smallmatrix} 2 & 0 \\ 1 & 1 \end{smallmatrix}\right)$ starting from $(2(|T|-1) - d_T(i), d_T(i))$. Indeed, when a new vertex is added to the tree, either it attaches to vertex i, with probability $d_{\mathrm{PA}(n,T)}(i)/(2n-2)$, in which case both $d_{\mathrm{PA}(n,T)}(i)$ and $2(n-1) - d_{\mathrm{PA}(n,T)}(i)$ increase by one (and hence why the second row of the replacement matrix is (1 1)), or otherwise it attaches to some other vertex, in which case $d_{\mathrm{PA}(n,T)}(i)$ does not increase but $2(n-1) - d_{\mathrm{PA}(n,T)}(i)$ increases by two (and hence why the first row of the replacement matrix is (2 0)). Janson [42] gives limit theorems for triangular Pólya urns, and also provides information about the limiting distributions; for instance [42, Theorem 1.7] gives a formula for the moments of $D_i(T)$, extending (5.2) for arbitrary trees T: for every i such that $1 \leq i \leq |T|$, we have

$$\mathbb{E}D_i(T)^r = \frac{\Gamma(|T|-1)\,\Gamma(d_T(i)+r)}{\Gamma(d_T(i))\,\Gamma\left(|T|-1+\frac{r}{2}\right)}, \tag{5.4}$$

and for $i > |T|$ we have $\mathbb{E}D_i(T)^r = \Gamma(i-1)\,\Gamma(1+r)/\Gamma(i-1+r/2)$, just like in (5.2).

The joint distribution of the limiting degrees in the seed graph, $\left(D_1(T), \ldots, D_{|T|}(T)\right)$, can be understood by viewing the evolution of $\left(d_{\mathrm{PA}(n,T)}(1), \ldots, d_{\mathrm{PA}(n,T)}(|T|)\right)$ in the following way. When adding a new vertex, first decide whether it attaches to one of the initial $|T|$ vertices (with probability $\sum_{i=1}^{|T|} d_{\mathrm{PA}(n,T)}(i)/(2n-2)$) or not (with the remaining probability); if it does, then independently pick one of them to attach to with probability proportional to their degrees. In other words, if viewed at times when a new vertex attaches to one of the initial $|T|$ vertices, the joint degree counts of the initial vertices evolve like a standard Pólya urn with $|T|$ colors and identity replacement matrix.

Let $\mathrm{Beta}(a,b)$ denote the beta distribution with parameters a and b (with density proportional to $x^{a-1}(1-x)^{b-1}\mathbf{1}_{\{x\in[0,1]\}}$), let $\mathrm{Dir}(\alpha_1, \ldots \alpha_s)$ denote the Dirichlet distribution with density proportional to $x_1^{\alpha_1-1}\cdots x_s^{\alpha_s-1}\mathbf{1}_{\left\{x\in[0,1]^s, \sum_{i=1}^s x_i = 1\right\}}$, and write $X \sim \mathrm{GGa}(a,b)$ for a random variable X having the generalized gamma distribution with density proportional to $x^{a-1}e^{-x^b}\mathbf{1}_{\{x>0\}}$. On the one hand, the pair $\left(2(n-1) - \sum_{i=1}^{|T|} d_{\mathrm{PA}(n,T)}(i), \sum_{i=1}^{|T|} d_{\mathrm{PA}(n,T)}(i)\right)$ evolves according to a Pólya urn with replacement matrix $\left(\begin{smallmatrix} 2 & 0 \\ 1 & 1 \end{smallmatrix}\right)$ starting from $(0, 2(|T|-1))$. Janson [42] gives the limiting distribution of $\sum_{i=1}^{|T|} d_{\mathrm{PA}(n,T)}(i)/\sqrt{n}$ (see Theorem 1.8 and Example 3.1): $\sum_{i=1}^{|T|} D_i(T) \overset{d}{=} 2Z_{|T|}$, where $Z_{|T|} \sim \mathrm{GGa}(2|T|-1, 2)$. On the other hand, it is known that in a standard Pólya urn with identity replacement matrix the vector of proportions of each color converges almost surely to a random variable with a Dirichlet distribution with parameters given by the initial counts. These facts, together with the observation in the previous paragraph, lead to the following representation: if X and $Z_{|T|}$ are independent,

$X \sim \mathrm{Dir}\left(d_T(1), \ldots, d_T(|T|)\right)$, and $Z_{|T|} \sim \mathrm{GGa}\left(2|T|-1, 2\right)$, then

$$\left(D_1(T), \ldots, D_{|T|}(T)\right) \stackrel{d}{=} 2Z_{|T|}X. \tag{5.5}$$

Recently, Peköz et al. [60] gave useful representations for $\left(D_1(T), \ldots, D_r(T)\right)$ for general r, and the representation above appears as a special case (see [60, Remark 1.9]).

5.2.2 Tail behavior

In order to prove Theorem 4.6 our main tool is to study the tail of the limiting degree distributions. In particular, we use the following key lemma.

Lemma 5.1. *Let T be a finite tree.*

(a) Let $U \subseteq \{1, 2, \ldots, |T|\}$ be a nonempty subset of the vertices of T, and let $d = \sum_{i \in U} d_T(i)$. Then

$$\mathbb{P}\left(\sum_{i \in U} D_i(T) > t\right) \sim c\left(|T|, d\right) t^{1-2|T|+2d} \exp\left(-t^2/4\right) \tag{5.6}$$

as $t \to \infty$, where the constant c is as in (5.1).[1]

(b) For every $L > |T|$ there exists a constant $C(L) < \infty$ such that for every $t \geq 1$ we have

$$\sum_{i=L}^{\infty} \mathbb{P}\left(D_i(T) > t\right) \leq C(L)\, t^{3-2L} \exp\left(-t^2/4\right). \tag{5.7}$$

We postpone the proof of Lemma 5.1 to Section 5.5, as it results from a lengthy computation. As an immediate corollary we get the asymptotic tail behavior of $D_{\max}(T)$.

Corollary 5.2. *Let T be a finite tree and let $m := \left|\{i \in \{1, \ldots, |T|\} : d_T(i) = \Delta(T)\}\right|$, where recall that $\Delta(T)$ is the maximum degree in T. Then*

$$\mathbb{P}\left(D_{\max}(T) > t\right) \sim m \times c\left(|T|, \Delta(T)\right) t^{1-2|T|+2\Delta(T)} \exp\left(-t^2/4\right) \tag{5.8}$$

as $t \to \infty$, where the constant c is as in (5.1).

Proof. Recall the fact that $D_{\max}(T) = \max_{i \geq 1} D_i(T)$ almost surely. First, a union bound gives us that

$$\mathbb{P}\left(D_{\max}(T) > t\right) \leq \sum_{i=1}^{m} \mathbb{P}\left(D_i(T) > t\right) + \sum_{i=m+1}^{|T|} \mathbb{P}\left(D_i(T) > t\right) + \sum_{i=|T|+1}^{\infty} \mathbb{P}\left(D_i(T) > t\right).$$

[1] Throughout the paper we use standard asymptotic notation; for instance, $f(t) \sim g(t)$ as $t \to \infty$ if $\lim_{t \to \infty} f(t)/g(t) = 1$.

Then using Lemma 5.1 we get the upper bound required for (5.8): the first sum gives the right hand side of (5.8), while the other two sums are of smaller order. For the lower bound we first have that

$$\mathbb{P}\left(D_{\max}\left(T\right) > t\right) \geq \sum_{i=1}^{m} \mathbb{P}\left(D_i\left(T\right) > t\right) - \sum_{i=1}^{m} \sum_{j=i+1}^{m} \mathbb{P}\left(D_i\left(T\right) > t, D_j\left(T\right) > t\right). \tag{5.9}$$

Lemma 5.1(a) with $U = \{i, j\}$ implies that for any $1 \leq i < j \leq m$, we have

$$\mathbb{P}\left(D_i\left(T\right) > t, D_j\left(T\right) > t\right) \leq \mathbb{P}\left(D_i\left(T\right) + D_j\left(T\right) > 2t\right) \leq C_{i,j}\left(T\right) t^{1-2|T|+4\Delta(T)} \exp\left(-t^2\right) \tag{5.10}$$

for some constant $C_{i,j}\left(T\right)$ and all t large enough. The exponent $-t^2$, appearing on the right hand side of (5.10), is smaller by a constant factor than the exponent $-t^2/4$, appearing in the asymptotic expression for $\mathbb{P}\left(D_i\left(T\right) > t\right)$ (see (5.6)). Consequently the second sum on the right hand side of (5.9) is of smaller order than the first sum, and so we have that $\mathbb{P}\left(D_{\max}\left(T\right) > t\right) \geq \left(1 - o\left(1\right)\right) \sum_{i=1}^{m} \mathbb{P}\left(D_i\left(T\right) > t\right)$ as $t \to \infty$. We conclude using Lemma 5.1. $\square$

5.3 Distinguishing trees using the maximum degree

In this section we first prove Theorems 4.6 and 4.4, both using Corollary 5.2 (see Section 5.3.1). Then in Section 5.3.2 we describe a particular way of generalizing the notion of maximum degree which we believe should provide an alternative way to prove Theorem 4.2. At present we are missing a technical result, see Conjecture 5.1 below, and we prove Theorem 4.2 assuming that this holds. Although Curien et al. [21] have now proven Theorem 4.2, we believe this alternative approach could be of interest by itself due to its simplicity, and it may also lead to better bounds. Moreover, as described at the end of the section, the statistics used by [21] are very similar to the ones we considered, and it would be interesting to understand this connection better.

5.3.1 Proofs

Proof of Theorem 4.6. We first provide a simple proof of distinguishing two trees of the same size but with different maximum degree, and then show how to extend this argument to the other cases.

Case 1: $|S| - \Delta\left(S\right) \neq |T| - \Delta\left(T\right)$**.** W.l.o.g. suppose that $|S| - \Delta\left(S\right) < |T| - \Delta\left(T\right)$. Clearly for any $t > 0$ and $n \geq \max\left\{|S|, |T|\right\}$ one has

$$\begin{aligned}\mathrm{TV}\left(\mathrm{PA}\left(n, S\right), \mathrm{PA}\left(n, T\right)\right) &\geq \mathrm{TV}\left(\Delta\left(\mathrm{PA}\left(n, S\right)\right), \Delta\left(\mathrm{PA}\left(n, T\right)\right)\right) \\ &\geq \mathbb{P}\left(\Delta\left(\mathrm{PA}\left(n, S\right)\right) > t\sqrt{n}\right) - \mathbb{P}\left(\Delta\left(\mathrm{PA}\left(n, T\right)\right) > t\sqrt{n}\right).\end{aligned}$$

Taking the limit as $n \to \infty$ this implies that

$$\delta\left(S, T\right) \geq \sup_{t>0} \left[\mathbb{P}\left(D_{\max}\left(S\right) > t\right) - \mathbb{P}\left(D_{\max}\left(T\right) > t\right)\right]. \tag{5.11}$$

By Corollary 5.2 and the fact that $|S|-\Delta(S)<|T|-\Delta(T)$ we have that $\mathbb{P}(D_{\max}(S)>t)>\mathbb{P}(D_{\max}(T)>t)$ for large enough t, which concludes the proof in this case.

Case 2: $|S|\neq|T|$**.** W.l.o.g. suppose that $|S|<|T|$. If $|S|-\Delta(S)\neq|T|-\Delta(T)$ then by Case 1 we have that $\delta(S,T)>0$, so we may assume that $|S|-\Delta(S)=|T|-\Delta(T)$. Just as in the proof of Case 1 we have that

$$\delta(S,T)\geq\sup_{t>0}\left[\mathbb{P}(D_{\max}(T)>t)-\mathbb{P}(D_{\max}(S)>t)\right]. \tag{5.12}$$

Corollary 5.2 provides the asymptotic behavior for $\mathbb{P}(D_{\max}(T)>t)$ in the form of (5.8), where $m\geq1$.

To find an upper bound for $\mathbb{P}(D_{\max}(S)>t)$, first notice that $\Delta(\mathrm{PA}(|T|,S))\leq\Delta(T)$, with equality holding if and only if all of the $|T|-|S|$ vertices of $\mathrm{PA}(|T|,S)$ that were added to S connect to the same vertex $i\in\{1,2,\dots,|S|\}$ and $d_S(i)=\Delta(S)$. Consequently, if $\Delta(\mathrm{PA}(|T|,S))=\Delta(T)$, then there is exactly one vertex $j\in\{1,2,\dots,|T|\}$ such that $d_{\mathrm{PA}(|T|,S)}(j)=\Delta(T)$. This, together with Corollary 5.2, shows that on the one hand

$$\mathbb{P}\left(D_{\max}(S)>t\,|\,\Delta(\mathrm{PA}(|T|,S))<\Delta(T)\right)=o\left(t^{1-2|T|+2\Delta(T)}\exp\left(-t^2/4\right)\right),$$

as $t\to\infty$, and on the other hand

$$\begin{aligned}&\mathbb{P}\left(D_{\max}(S)>t\,|\,\Delta(\mathrm{PA}(|T|,S))=\Delta(T)\right)\\&\qquad\leq(1+o(1))\,c(|T|,\Delta(T))\,t^{1-2|T|+2\Delta(T)}\exp\left(-t^2/4\right)\end{aligned}$$

as $t\to\infty$. Consequently we have that

$$\begin{aligned}&\mathbb{P}(D_{\max}(S)>t)\\&\qquad\leq(1+o(1))\,\mathbb{P}(\Delta(\mathrm{PA}(|T|,S))=\Delta(T))\,c(|T|,\Delta(T))\,t^{1-2|T|+2\Delta(T)}\exp\left(-t^2/4\right)\end{aligned}$$

as $t\to\infty$, which combined with the tail behavior of $D_{\max}(T)$ gives that

$$\begin{aligned}&\mathbb{P}(D_{\max}(T)>t)-\mathbb{P}(D_{\max}(S)>t)\\&\qquad\geq(1-o(1))\,\mathbb{P}(\Delta(\mathrm{PA}(|T|,S))<\Delta(T))\,c(|T|,\Delta(T))\,t^{1-2|T|+2\Delta(T)}\exp\left(-t^2/4\right)\end{aligned}$$

as $t\to\infty$. To conclude the proof, notice that $\mathbb{P}(\Delta(\mathrm{PA}(|T|,S))<\Delta(T))$ is at least as great as the probability that vertex $|S|+1$ connects to a leaf of S, which has probability at least $1/(2|S|-2)$.

Case 3: $|S|=|T|$**, different degree profiles.** Let $z\in\{1,\dots,|T|\}$ be the first index such that $d_S(z)\neq d_T(z)$ and assume w.l.o.g. that $d_S(z)<d_T(z)$. First we have that

$$\begin{aligned}\mathbb{P}(D_{\max}(T)>t)\geq\;&\mathbb{P}(\exists i\in[z-1]:D_i(T)>t)+\mathbb{P}(D_z(T)>t)\\&-\sum_{i=1}^{z-1}\mathbb{P}(D_z(T)>t,D_i(T)>t)\end{aligned}$$

and

$$\mathbb{P}\left(D_{\max}\left(S\right) > t\right) \leq \mathbb{P}\left(\exists i \in [z-1] : D_i\left(S\right) > t\right) + \sum_{i=z}^{\infty} \mathbb{P}\left(D_i\left(S\right) > t\right).$$

Now observe that one can couple the evolution of $\mathrm{PA}(n, T)$ and $\mathrm{PA}(n, S)$ in such a way that the degrees of vertices $1, \ldots, z-1$ stay the same in both trees. Thus one clearly has

$$\mathbb{P}\left(\exists i \in [z-1] : D_i\left(T\right) > t\right) = \mathbb{P}\left(\exists i \in [z-1] : D_i\left(S\right) > t\right).$$

Putting the three above displays together one obtains

$$\begin{aligned}
&\mathbb{P}\left(D_{\max}\left(T\right) > t\right) - \mathbb{P}\left(D_{\max}\left(S\right) > t\right) \\
&\qquad \geq \mathbb{P}\left(D_z\left(T\right) > t\right) - \sum_{i=1}^{z-1} \mathbb{P}\left(D_z\left(T\right) > t, D_i\left(T\right) > t\right) - \sum_{i=z}^{\infty} \mathbb{P}\left(D_i\left(S\right) > t\right).
\end{aligned}$$

Now using Lemma 5.1 one easily gets (for some constant $C > 0$) that

$$\begin{aligned}
&\mathbb{P}\left(D_z\left(T\right) > t\right) \sim c\left(\left|T\right|, d_T(z)\right) t^{1-2|T|+2d_T(z)} \exp\left(-t^2/4\right), \\
&\sum_{i=1}^{z-1} \mathbb{P}\left(D_z\left(T\right) > t, D_i\left(T\right) > t\right) \leq \sum_{i=1}^{z-1} \mathbb{P}\left(D_z\left(T\right) + D_i\left(T\right) > 2t\right) \\
&\qquad \leq \sum_{i=1}^{z-1} \left(1 + o\left(1\right)\right) c\left(\left|T\right|, d_T(z) + d_T(i)\right) \left(2t\right)^{1-2|T|+2(d_T(z)+d_T(i))} \exp\left(-t^2\right), \\
&\sum_{i=z}^{\infty} \mathbb{P}\left(D_i\left(S\right) > t\right) \leq C t^{1-2|T|+2d_S(z)} \exp\left(-t^2/4\right).
\end{aligned}$$

In particular, since $d_S(z) < d_T(z)$ and $t^\alpha \exp(-t^2) = o(\exp(-t^2/4))$ for any α, this shows that

$$\mathbb{P}\left(D_{\max}\left(T\right) > t\right) - \mathbb{P}\left(D_{\max}\left(S\right) > t\right) \geq \left(1 - o\left(1\right)\right) c\left(\left|T\right|, d_T(z)\right) t^{1-2|T|+2d_T(z)} \exp\left(-t^2/4\right),$$

which, together with (5.12), concludes the proof. $\square$

Proof of Theorem 4.4. As before we have that

$$\begin{aligned}
\delta\left(S_k, T\right) &\geq \sup_{t \geq 0} \left[\mathbb{P}\left(D_{\max}\left(S_k\right) > t\right) - \mathbb{P}\left(D_{\max}\left(T\right) > t\right)\right] \\
&\geq \mathbb{P}\left(D_{\max}\left(S_k\right) > \sqrt{k}/2\right) - \mathbb{P}\left(D_{\max}\left(T\right) > \sqrt{k}/2\right). \qquad (5.13)
\end{aligned}$$

By Corollary 5.2, we know that the second term in (5.13) goes to zero as $k \to \infty$ for any fixed T. We can lower bound the first term in (5.13) by $\mathbb{P}\left(D_1\left(S_k\right) > \sqrt{k}/2\right) = 1 - \mathbb{P}\left(D_1\left(S_k\right) \leq \sqrt{k}/2\right)$. From (5.4) we have that the first two moments of $D_1\left(S_k\right)$ are

$\mathbb{E}D_1(S_k) = \Gamma(k)/\Gamma(k-1/2)$ and $\mathbb{E}D_1(S_k)^2 = \Gamma(k+1)/\Gamma(k) = k$. From standard facts about the Γ function and Stirling series one has that $0 \leq \mathbb{E}D_1(S_k) - \sqrt{k-1} \leq \left(6\sqrt{k-1}\right)^{-1}$ and then also

$$\mathrm{Var}(D_1(S_k)) = \mathbb{E}D_1(S_k)^2 - (\mathbb{E}D_1(S_k))^2 \leq k - (k-1) = 1.$$

Therefore Chebyshev's inequality implies that $\lim_{k\to\infty} \mathbb{P}\left(D_1(S_k) \leq \sqrt{k}/2\right) = 0$. □

5.3.2 Towards an alternative proof of Theorem 4.2

Our proof of Theorem 4.6 above relied on the precise asymptotic tail behavior of $D_{\max}(T)$, as described in Corollary 5.2. In order to distinguish two trees with the same degree profile (such as the pair of trees in Figure 4.1), it is necessary to incorporate information about the graph structure. Indeed, if S and T have the same degree profiles, then it is possible to couple $\mathrm{PA}(n,S)$ and $\mathrm{PA}(n,T)$ such that they have the same degree profiles for every n.

Thus a possible way to prove Theorem 4.2 is to generalize the notion of maximum degree in a way that incorporates information about the graph structure, and then use similar arguments as in the proofs above. A candidate is the following.

Definition 5.1. *Given a tree U, define the U-*maximum degree *of a tree T, denoted by $\Delta_U(T)$, as*

$$\Delta_U(T) = \max_{\varphi} \sum_{u \in V(U)} d_T(\varphi(u)),$$

where $V(U)$ denotes the vertex set of U, and the maximum is taken over all injective graph homomorphisms from U to T. That is, φ ranges over all injective maps from $V(U)$ to $V(T)$ such that $\{u,v\} \in E(U)$ implies that $\{\varphi(u), \varphi(v)\} \in E(T)$, where $E(U)$ denotes the edge set of U, and $E(T)$ is defined similarly.

When U is a single vertex, then $\Delta_U \equiv \Delta$, so this indeed generalizes the notion of maximum degree.

Intuitively, the main contributor to the tail of $\Delta_T(\mathrm{PA}(n,T))$ should be the homomorphism that maps T to the vertices making up the initial seed. In other words, the tail should behave like the tail of the sum of the degrees of the initial vertices. On the other hand, if S is not isomorphic to T (and assume for simplicity that $|S| = |T|$), then any homomorphism from T to $\mathrm{PA}(n,S)$ must use a vertex that is not part of the seed. Because of this, one expects that the tail of $\Delta_T(\mathrm{PA}(n,S))$ is lighter than the tail of $\Delta_T(\mathrm{PA}(n,T))$. In particular, we conjecture the following.

Conjecture 5.1. *Suppose that S and T are two nonisomorphic trees of the same size. Then*

$$\limsup_{n\to\infty} \mathbb{P}\left(\Delta_T(\mathrm{PA}(n,S)) > t\sqrt{n}\right) = o\left(t^{2|T|-3}\exp\left(-t^2/4\right)\right)$$

as $t \to \infty$.

If this conjecture were true, then Theorem 4.2 also follows, as we now show.

Proof of Theorem 4.2 assuming Conjecture 5.1 holds. Assume that $|S| = |T|$; if $|S| \neq |T|$ then we already know from Theorem 4.6 that $\delta(S,T) > 0$. As in the proof of Theorem 4.6, for any $t > 0$ and $n \geq \max\{|S|, |T|\}$ we have that

$$\begin{aligned}\mathrm{TV}(\mathrm{PA}(n,S), \mathrm{PA}(n,T)) &\geq \mathrm{TV}(\Delta_T(\mathrm{PA}(n,S)), \Delta_T(\mathrm{PA}(n,T))) \\ &\geq \mathbb{P}\left(\Delta_T(\mathrm{PA}(n,T)) > t\sqrt{n}\right) - \mathbb{P}\left(\Delta_T(\mathrm{PA}(n,S)) > t\sqrt{n}\right),\end{aligned}$$

and consequently

$$\delta(S,T) \geq \sup_{t>0}\left\{\liminf_{n\to\infty} \mathbb{P}\left(\Delta_T(\mathrm{PA}(n,T)) > t\sqrt{n}\right) - \limsup_{n\to\infty} \mathbb{P}\left(\Delta_T(\mathrm{PA}(n,S)) > t\sqrt{n}\right)\right\}. \tag{5.14}$$

Since $\varphi(i) = i$ for $1 \leq i \leq |T|$ is an injective graph homomorphism from T to $\mathrm{PA}(n,T)$, we have that

$$\begin{aligned}\liminf_{n\to\infty} \mathbb{P}\left(\Delta_T(\mathrm{PA}(n,T)) > t\sqrt{n}\right) &\geq \liminf_{n\to\infty} \mathbb{P}\left(\sum_{i=1}^{|T|} d_{\mathrm{PA}(n,T)}(i) > t\sqrt{n}\right) \\ &= \mathbb{P}\left(\sum_{i=1}^{|T|} D_i(T) > t\right).\end{aligned}$$

By Lemma 5.1 we know that

$$\mathbb{P}\left(\sum_{i=1}^{|T|} D_i(T) > t\right) \sim c(|T|, 2|T|-2)\, t^{2|T|-3} \exp\left(-t^2/4\right)$$

as $t \to \infty$, which together with (5.14) and Conjecture 5.1 shows that $\delta(S,T) > 0$. □

We note that the statistics considered by [21] are very similar to the ones considered above based on the U-maximum degree. More precisely, instead of taking a maximum over homomorphisms, they take a sum over them, and instead of taking a sum over vertices, they take a product over them. (They also consider decorated trees, which essentially means raising the degrees appearing in the statistic to appropriate powers.) Furthermore, while we considered the tail behavior of statistics based on the U-maximum degree, they constructed appropriate martingales, for which they needed to estimate the first two moments of these statistics. Understanding the connection between these two related approaches would be interesting.

5.4 The weak limit of $\mathrm{PA}(n, T)$

In this section we prove Theorem 4.1. For two graphs G and H we write $G = H$ if G and H are isomorphic, and we use the same notation for rooted graphs. Recalling the definition of the Benjamini-Schramm limit and the Pólya-point graph (see [8] and [9, Section 2.3 and Definition 2.1]), we want to prove that

$$\lim_{n\to\infty} \mathbb{P}\left(B_r\left(\mathrm{PA}\left(n,T\right),k_n\left(T\right)\right)=\left(H,y\right)\right)=\mathbb{P}\left(B_r\left(\mathcal{T},\left(0\right)\right)=\left(H,y\right)\right),$$

where $B_r(G, v)$ is the rooted ball of radius r around vertex v in the graph G, $k_n(T)$ is a uniformly random vertex in $\mathrm{PA}(n, T)$, (H, y) is a finite rooted tree, and $(\mathcal{T}, (0))$ is the Pólya-point graph (with $m = 1$).

We construct a forest F based on T as follows. To each vertex v in T we associate $d_T(v)$ isolated nodes with self loops, that is, F consists of $2(|T| - 1)$ isolated vertices with self loops. Our convention here is that a node with k regular edges and one self loop has degree $k + 1$. The graph evolution process $\mathrm{PA}(n, F)$ for forests is defined in the same way as for trees, and we couple the processes $\mathrm{PA}(n, T)$ and $\mathrm{PA}(n+|T|-2, F)$ in the natural way: when an edge is added to vertex v of T in $\mathrm{PA}(n, T)$ then an edge is also added to one of the $d_T(v)$ corresponding vertices of F in $\mathrm{PA}(n+|T|-2, F)$, and furthermore newly added vertices are always coupled. We first observe that, clearly, the weak limit of $\mathrm{PA}(n + |T| - 2, F)$ is the Pólya-point graph, that is,

$$\lim_{n\to\infty} \mathbb{P}\left(B_r\left(\mathrm{PA}\left(n+|T|-2,F\right),k_n\left(F\right)\right)=\left(H,y\right)\right)=\mathbb{P}\left(B_r\left(\mathcal{T},\left(0\right)\right)=\left(H,y\right)\right),$$

where $k_n(F)$ is a uniformly random vertex in $\mathrm{PA}(n+|T|-2, F)$. We couple $k_n(F)$ and $k_n(T)$ in the natural way, that is, if $k_n(F)$ is the t^{th} newly created vertex in $\mathrm{PA}(n+|T|-2, F)$, then $k_n(T)$ is the t^{th} newly created vertex in $\mathrm{PA}(n, T)$. To conclude the proof it is now sufficient to show that

$$\lim_{n\to\infty} \mathbb{P}\left(B_r\left(\mathrm{PA}\left(n+|T|-2,F\right),k_n\left(F\right)\right)\neq B_r\left(\mathrm{PA}\left(n,T\right),k_n\left(T\right)\right)\right)=0.$$

The following inequalities hold true (with a slight—but clear—abuse of notation when we write $v \in F$) for any $u > 0$,

$$\begin{aligned}
\mathbb{P}&\left(B_r\left(\mathrm{PA}\left(n+|T|-2,F\right),k_n\left(F\right)\right)\neq B_r\left(\mathrm{PA}\left(n,T\right),k_n\left(T\right)\right)\right)\\
&\leq \mathbb{P}\left(\exists v\in F \text{ s.t. } v\in B_r\left(\mathrm{PA}\left(n+|T|-2,F\right),k_n\left(F\right)\right)\right)\\
&\leq \mathbb{P}\left(\exists v\in F, d_{\mathrm{PA}(n+|T|-2,F)}\left(v\right)<u\right)\\
&\quad+\mathbb{P}\left(\exists v\in B_r\left(\mathrm{PA}\left(n+|T|-2,F\right),k_n\left(F\right)\right) \text{ s.t. } d_{\mathrm{PA}(n+|T|-2,F)}\left(v\right)\geq u\right).
\end{aligned}$$

It is easy to verify that for any $u > 0$, we have

$$\lim_{n\to\infty} \mathbb{P}\left(\exists v\in F, d_{\mathrm{PA}(n+|T|-2,F)}\left(v\right)<u\right)=0.$$

Furthermore since $B_r\left(\mathrm{PA}\left(n+|T|-2,F\right),k_n\left(F\right)\right)$ tends to the Pólya-point graph we also have

$$\lim_{n\to\infty}\mathbb{P}\left(\exists v\in B_r\left(\mathrm{PA}\left(n+|T|-2,F\right),k_n\left(F\right)\right)\text{ s.t. }d_{\mathrm{PA}(n+|T|-2,F)}\left(v\right)\geq u\right)$$
$$=\mathbb{P}\left(\exists v\in B_r\left(\mathcal{T},(0)\right)\text{ s.t. }d_{\mathcal{T}}\left(v\right)\geq u\right).$$

By looking at the definition of $(\mathcal{T},(0))$ given in [9] one can easily show that

$$\lim_{u\to\infty}\mathbb{P}\left(\exists v\in B_r\left(\mathcal{T},(0)\right)\text{ s.t. }d_{\mathcal{T}}\left(v\right)\geq u\right)=0,$$

which concludes the proof.

5.5 Proof of Lemma 5.1

In this section we prove Lemma 5.1. In light of the representation (5.5) in Section 5.2.1.2, part (a) of Lemma 5.1 follows from a lengthy computation, the result of which we state separately.

Lemma 5.3. *Fix positive integers a and b. Let B and Z be independent random variables such that $B\sim\mathrm{Beta}\left(a,b\right)$ and $Z\sim\mathrm{GGa}\left(a+b+1,2\right)$, and let $V=2BZ$. Then*

$$\mathbb{P}\left(V>t\right)\sim c\left(\frac{a+b+2}{2},a\right)t^{-1+a-b}\exp\left(-t^2/4\right)\tag{5.15}$$

as $t\to\infty$, where the constant c is as in (5.1).

Proof. By definition we have for $t>0$ that

$$\begin{aligned}\mathbb{P}\left(V>t\right)=\mathbb{P}\left(2BZ>t\right)&=\int_{t/2}^{\infty}\int_{t/(2z)}^{1}\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}x^{a-1}\left(1-x\right)^{b-1}dx\frac{2}{\Gamma\left(\frac{a+b+1}{2}\right)}z^{a+b}e^{-z^2}dz\\&=\int_{t/2}^{\infty}\left[1-I_{t/(2z)}(a,b)\right]\frac{2}{\Gamma\left(\frac{a+b+1}{2}\right)}z^{a+b}e^{-z^2}dz,\end{aligned}$$

where $I_x(a,b)=\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}\int_0^x y^{a-1}(1-y)^{b-1}dy$ is the regularized incomplete Beta function. For positive integers a and b, integration by parts and induction gives that

$$I_x(a,b)=1-\sum_{j=0}^{a-1}\binom{a+b-1}{j}x^j\left(1-x\right)^{a+b-1-j}.$$

Plugging this back in to the integral and doing a change of variables $y=2z$, we get that

$$\mathbb{P}(V>t)=\frac{2^{-(a+b)}}{\Gamma\left(\frac{a+b+1}{2}\right)}\sum_{j=0}^{a-1}\binom{a+b-1}{j}\int_t^{\infty}t^j\left(y-t\right)^{a+b-1-j}y\exp\left(-y^2/4\right)dy.$$

Expanding $(y-t)^{a+b-1-j}$ we arrive at the alternating sum formula

$$\mathbb{P}(V>t)=\frac{2^{-(a+b)}}{\Gamma\left(\frac{a+b+1}{2}\right)}\sum_{j=0}^{a-1}\sum_{k=0}^{a+b-1-j}\binom{a+b-1}{j}\binom{a+b-1-j}{k}(-1)^{a+b-1-j-k}\,t^{a+b-1-k}A_{k+1}, \tag{5.16}$$

where for $m\geq 0$ let

$$A_m:=\int_t^\infty y^m\exp\left(-y^2/4\right)dy.$$

Thus in order to show (5.15) it is enough to show that for every j such that $0\leq j\leq a-1$ we have

$$\sum_{k=0}^{a+b-1-j}\binom{a+b-1-j}{k}(-1)^{a+b-1-j-k}\,t^{a+b-1-k}A_{k+1}\sim\frac{2^{a+b-j}(a+b-1-j)!}{t^{a+b-1-2j}}\exp\left(-t^2/4\right). \tag{5.17}$$

To do this, we need to evaluate the integrals $\{A_m\}_{m\geq 0}$. Recall that the complementary error function is defined as $\operatorname{erfc}(z)=1-\operatorname{erf}(z)=(2/\sqrt{\pi})\int_z^\infty\exp(-u^2)\,du$, and thus $A_0=\sqrt{\pi}\operatorname{erfc}(t/2)$; also $A_1=2\exp(-t^2/4)$. Integration by parts gives that for $m\geq 2$ we have $A_m=2t^{m-1}\exp(-t^2/4)+2(m-1)A_{m-2}$. Iterating this, and using the values for A_0 and A_1, gives us that for m odd we have

$$A_m=2t^{m-1}\exp\left(-t^2/4\right)\sum_{\ell=0}^{\frac{m-1}{2}}\frac{(m-1)!!}{(m-2\ell-1)!!}\left(\frac{2}{t^2}\right)^\ell, \tag{5.18}$$

and for m even we have

$$A_m=2t^{m-1}\exp\left(-t^2/4\right)\sum_{\ell=0}^{\frac{m}{2}-1}\frac{(m-1)!!}{(m-2\ell-1)!!}\left(\frac{2}{t^2}\right)^\ell+2^{\frac{m}{2}}\times(m-1)!!\times\sqrt{\pi}\operatorname{erfc}(t/2). \tag{5.19}$$

In the following we fix j such that $0\leq j\leq a-1$ and $a+b-1-j$ is odd—showing (5.17) when $a+b-1-j$ is even can be done in the same way. In order to abbreviate notation we let $r=(a+b-2-j)/2$. Plugging in the formulas (5.18) and (5.19) into the left hand side of (5.17) we get that

$$\begin{aligned}&\sum_{k=0}^{a+b-1-j}\binom{a+b-1-j}{k}(-1)^{a+b-1-j-k}\,t^{a+b-1-k}A_{k+1}\\&\quad=\sum_{k=0}^{2r+1}\binom{2r+1}{k}(-1)^{2r+1-k}\,t^{2r+1+j-k}A_{k+1}\\&\quad=-\sum_{\ell=0}^{r}\binom{2r+1}{2\ell}t^{2r+1+j-2\ell}A_{2\ell+1}+\sum_{\ell=0}^{r}\binom{2r+1}{2\ell+1}t^{2r+1+j-(2\ell+1)}A_{2\ell+2}\end{aligned}$$

$$
\begin{aligned}
&= -\sum_{\ell=0}^{r}\binom{2r+1}{2\ell}t^{2r+1+j-2\ell}2\exp\left(-t^2/4\right)\sum_{u=0}^{\ell}2^u\frac{(2\ell)!!}{(2\ell-2u)!!}t^{2\ell-2u}\\
&\quad+\sum_{\ell=0}^{r}\binom{2r+1}{2\ell+1}t^{2r+1+j-(2\ell+1)}2\exp\left(-t^2/4\right)\sum_{u=0}^{\ell}2^u\frac{(2\ell+1)!!}{(2\ell+1-2u)!!}t^{2\ell+1-2u}\\
&\quad+\sum_{\ell=0}^{r}\binom{2r+1}{2\ell+1}t^{2r+1+j-(2\ell+1)}2^{\ell+1}\left(2\ell+1\right)!!\sqrt{\pi}\,\mathrm{erfc}\left(t/2\right)\\
&= 2\exp\left(-t^2/4\right)\sum_{u=0}^{r}t^{2r+1+j-2u}2^u\sum_{k=2u}^{2r+1}\binom{2r+1}{k}\left(-1\right)^{k+1}\frac{k!!}{(k-2u)!!} \qquad (5.20)\\
&\quad+\sqrt{\pi}\,\mathrm{erfc}\left(t/2\right)\sum_{\ell=0}^{r}\binom{2r+1}{2\ell+1}t^{2r+1+j-(2\ell+1)}2^{\ell+1}\left(2\ell+1\right)!!. \qquad (5.21)
\end{aligned}
$$

An important fact that we will use is that for every polynomial P with degree less than n we have

$$\sum_{k=0}^{n}\binom{n}{k}\left(-1\right)^k P(k)=0. \tag{5.22}$$

Consequently, applying this to the polynomial $P(k)=k\left(k-2\right)\cdots\left(k-2\left(u-1\right)\right)$ we get that

$$
\begin{aligned}
&\sum_{k=2u}^{2r+1}\binom{2r+1}{k}\left(-1\right)^{k+1}k\left(k-2\right)\cdots\left(k-2\left(u-1\right)\right)\\
&\qquad= \sum_{k=0}^{2u-1}\binom{2r+1}{k}\left(-1\right)^{k}k\left(k-2\right)\cdots\left(k-2\left(u-1\right)\right)\\
&\qquad= -\sum_{\ell=0}^{u-1}\binom{2r+1}{2\ell+1}\left(2\ell+1\right)\left(2\ell-1\right)\cdots\left(2\ell+1-2\left(u-1\right)\right)\\
&\qquad= -\sum_{\ell=0}^{u-1}\binom{2r+1}{2\ell+1}\left(2\ell+1\right)!!\left(2\left(u-1-\ell\right)-1\right)!!\left(-1\right)^{u-1-\ell}. \qquad (5.23)
\end{aligned}
$$

Thus we see that in the sum (5.20) the cofficient of the term involving t^{2r+1+j} is zero, while the coefficient of the term involving $t^{2r+1+j-2u}$ for $1\leq u\leq r$ is $2^{u+1}\exp\left(-t^2/4\right)$ times the expression in (5.23). These are cancelled by terms coming from the sum in (5.21) as we will see shortly; to see this we need the asymptotic expansion of erfc to high enough order. In particular we have (see [1, equations 7.1.13 and 7.1.24]) that

$$\sqrt{\pi}\,\mathrm{erfc}\left(t/2\right)=2\exp\left(-t^2/4\right)\sum_{n=0}^{2r}\left(-1\right)^n2^n\left(2n-1\right)!!t^{-2n-1}+R\left(t\right), \tag{5.24}$$

where the approximation error $R(t)$ satisfies

$$|R(t)| \leq 2^{2r+2}(4r+1)!!t^{-(4r+3)}\exp\left(-t^2/4\right).$$

Plugging (5.24) back into (5.21), we first see that the error term satisfies

$$|R(t)|\sum_{\ell=0}^{r}\binom{2r+1}{2\ell+1}t^{2r+1+j-(2\ell+1)}2^{\ell+1}(2\ell+1)!! = O\left(t^{2j-1-(a+b)}\exp\left(-t^2/4\right)\right) \tag{5.25}$$

as $t\to\infty$. The main term of (5.21) becomes the sum

$$2\exp\left(-t^2/4\right)\sum_{\ell=0}^{r}\sum_{n=0}^{2r}\binom{2r+1}{2\ell+1}2^{\ell+n+1}(2\ell+1)!!(2n-1)!!(-1)^n t^{2r+1+j-2(\ell+n+1)}.$$

For u such that $1\leq u\leq r$, the coefficient of the term involving $t^{2r+1+j-2u}$ is $2^{u+1}\exp(-t^2/4)$ times

$$\sum_{\ell=0}^{u-1}\binom{2r+1}{2\ell+1}(2\ell+1)!!(2(u-1-\ell)-1)!!(-1)^{u-1-\ell},$$

which cancels out the coefficient of the same term coming from the other sum (5.20), see (5.23). For u such that $r<u\leq 2r$, the coefficient of the term involving $t^{2r+1+j-2u}$ is $2^{u+1}\exp(-t^2/4)$ times

$$\begin{aligned}\sum_{\ell=0}^{r}&\binom{2r+1}{2\ell+1}(2\ell+1)!!(2(u-1-\ell)-1)!!(-1)^{u-1-\ell}\\ &=\sum_{\ell=0}^{r}\binom{2r+1}{2\ell+1}(2\ell+1)(2\ell-1)\dots((2\ell+1)-2(u-1))\\ &=-\sum_{k=0}^{2r+1}\binom{2r+1}{k}(-1)^k k(k-2)\dots(k-2(u-1))=0,\end{aligned}$$

where we again used (5.22), together with the fact that $u\leq 2r$. Finally, the coefficient of the term involving $t^{2j+1-(a+b)}$ is $2^{2r+2}\exp(-t^2/4)$ times

$$\begin{aligned}\sum_{\ell=0}^{r}&\binom{2r+1}{2\ell+1}(2\ell+1)!!(2(2r-\ell)-1)!!(-1)^{2r-\ell}\\ &=-\sum_{k=0}^{2r+1}\binom{2r+1}{k}(-1)^k k(k-2)\dots(k-4r)\\ &=-\sum_{k=0}^{2r+1}\binom{2r+1}{k}(-1)^k k^{2r+1}=-(-1)^{2r+1}(2r+1)!=(2r+1)!,\end{aligned}$$

where we used (5.22) in the second equality. Since all other terms are of lower order (see (5.25)), this concludes the proof. □

Proof of Lemma 5.1. (a) If $U \neq T$, then $d = \sum_{i \in U} d_T(i) \in \{1, \ldots, 2|T| - 3\}$. Similarly to the third paragraph in Section 5.2.1.2, we can view the evolution of $\sum_{i \in U} d_{\mathrm{PA}(n,T)}(i)$ in the following way. When adding a new vertex, first decide whether it attaches to one of the initial $|T|$ vertices (with probability $\sum_{i=1}^{|T|} d_{\mathrm{PA}(n,T)}(i) / (2n-2)$) or not (with the remaining probability); if it does, then independently pick one of them to attach to with probability proportional to their degree—a vertex in U is chosen with probability $\sum_{i \in U} d_{\mathrm{PA}(n,T)}(i) / \sum_{i=1}^{|T|} d_{\mathrm{PA}(n,T)}(i)$. This implies the following representation: we have that $\sum_{i \in U} D_i(T) \stackrel{d}{=} 2BZ$, where B and Z are independent, $B \sim \mathrm{Beta}(d, 2|T| - 2 - d)$, and $Z \sim \mathrm{GGa}(2|T| - 1, 2)$. This also follows directly from the representation (5.5). Thus (5.6) is a direct consequence of Lemma 5.3.

If $U = T$, then $\sum_{i \in U} D_i(T) \stackrel{d}{=} 2Z$, where $Z \sim \mathrm{GGa}(2|T| - 1, 2)$ (see Section 5.2.1.2), and then (5.6) follows from a calculation that is contained in the proof of Lemma 5.3.

(b) To show (5.7) we use the results of [59] as described in Section 5.2.1.1. In addition we use the following tail bound of [59, Lemma 2.6], which says that for $x > 0$ and $s \geq 1$ we have $\int_x^\infty \kappa_s(y)\, dy \leq \frac{s}{x} \kappa_s(x)$. Consequently, for any $i > |T|$ we have the following tail bound:

$$\begin{aligned}
\mathbb{P}(D_i(T) > t) &= \mathbb{P}\left(W_{i-1} > \sqrt{\frac{i-1}{2}} t\right) = \int_{\sqrt{\frac{i-1}{2}} t}^\infty \kappa_{i-1}(y)\, dy \\
&\leq \frac{\sqrt{2i-2}}{t} \kappa_{i-1}\left(\sqrt{\frac{i-1}{2}} t\right) = \frac{2}{\sqrt{\pi} t} \exp\left(-t^2/4\right)(i-2)! U\left(i-2, \frac{1}{2}, \frac{t^2}{4}\right).
\end{aligned}$$

The following integral representation is useful for us [1, eq. 13.2.5]:

$$\Gamma(a) U(a, b, z) = \int_0^\infty e^{-zw} w^{a-1} (1+w)^{b-a-1}\, dw.$$

Consequently, we have

$$\begin{aligned}
\sum_{i=3}^\infty (i-2)! U\left(i-2, \frac{1}{2}, \frac{t^2}{4}\right) &= \sum_{i=3}^\infty (i-2) \int_0^\infty e^{-\frac{t^2}{4} w} \frac{1}{w\sqrt{1+w}} \left(\frac{w}{1+w}\right)^{i-2} dw \\
&= \int_0^\infty e^{-\frac{t^2}{4} w} \frac{1}{w\sqrt{1+w}} \sum_{i=3}^\infty (i-2) \left(\frac{w}{1+w}\right)^{i-2} dw \\
&= \int_0^\infty e^{-\frac{t^2}{4} w} \frac{1}{w\sqrt{1+w}} w(1+w)\, dw \\
&\leq \int_0^\infty e^{-\frac{t^2}{4} w} (1+w)\, dw = \frac{4}{t^2} + \frac{16}{t^4},
\end{aligned}$$

which shows (5.7) for $L = 3$. Similarly, for $L \geq 4$ we have

$$\begin{aligned}
\sum_{i=L}^{\infty} (i-2)! U\left(i-2, \frac{1}{2}, \frac{t^2}{4}\right) &= \int_0^{\infty} e^{-\frac{t^2}{4}w} \frac{1}{w\sqrt{1+w}} \sum_{i=L}^{\infty} (i-2) \left(\frac{w}{1+w}\right)^{i-2} dw \\
&= \int_0^{\infty} e^{-\frac{t^2}{4}w} \frac{1}{w\sqrt{1+w}} \frac{(L-2)\left(\frac{w}{1+w}\right)^{L-2} + (3-L)\left(\frac{w}{1+w}\right)^{L-1}}{1/(1+w)^2} dw \\
&\leq \int_0^{\infty} e^{-\frac{t^2}{4}w} (L-2) \left(\frac{w}{1+w}\right)^{L-3} \sqrt{1+w} dw \\
&\leq \int_0^{\infty} e^{-\frac{t^2}{4}w} (L-2) w^{L-3} dw = \frac{4^{L-2} \times (L-2)!}{t^{2L-4}},
\end{aligned}$$

where the first inequality follows from dropping the nonpositive term $(3-L)\left(\frac{w}{1+w}\right)^{L-1}$, and the second one follows because $L \geq 4$. This shows (5.7) for $L \geq 4$ and thus concludes the proof. $\square$

Chapter 6

Uniform attachment trees

6.1 Overview

This chapter is devoted to proving the results described in Chapter 4 on the influence of the seed in uniform attachment trees. We begin in Section 6.2 by providing some simple examples that formalize the intuition described in Section 4.3.2. We also prove Theorem 4.5 along the way. Section 6.3 is entirely devoted to the proof of Theorem 4.3, although a few technical estimates are deferred until Section 6.5. We conclude with a detailed comparison of our proof to the work of Curien et al. [21] in Section 6.4. To simplify notation, in this chapter we let $\delta \equiv \delta_{\mathrm{UA}}$.

6.2 Partitions and their balancedness: simple examples

In this section we show on a simple example how to formalize the intuition described in Section 4.3.2. We define a simple statistic based on this intuition, and after collecting some preliminary facts in Section 6.2.1, we show in Section 6.2.2 that $\delta(P_4, S_4) > 0$, where P_4 and S_4 are the path and the star on four vertices, respectively. We conclude the section by proving Theorem 4.5 in Section 6.2.3. The goal of this section is thus to provide a gentle introduction into the methods and statistics used to distinguish different seed trees, before analyzing more general statistics in the proof of Theorem 4.3 in Section 6.3.

For a tree T and an edge $e \in E(T)$, let T_1 and T_2 be the two connected components of $T \setminus \{e\}$. Define

$$g(T, e) = |T_1|^2 |T_2|^2 / |T|^4,$$

where $|T|$ denotes the number of vertices of T. Clearly, $0 \leq g(T, e) \leq 1/16$, and for "peripheral" edges e, $g(T, e)$ is closer to 0, while for more "central" edges e, $g(T, e)$ is closer to 1/16. Define the following statistic:

$$G(T) = \sum_{e \in E(T)} g(T, e).$$

The statistic $G(T)$ thus measures in a particular way the global balancedness properties of the tree T, and "central" edges contribute the most to this statistic.

6.2.1 Preliminary facts

For all $\alpha, \beta, n \in \mathbb{N}$, let $B_{\alpha,\beta,n}$ be a random variable such that $B_{\alpha,\beta,n} - \alpha$ has the beta-binomial distribution with parameters (α, β, n), i.e., it is a random variable satisfying

$$\mathbb{P}\left(B_{\alpha,\beta,n} = \alpha + k\right) = \frac{(k+\alpha-1)!(n-k+\beta-1)!(\alpha+\beta-1)!}{(n+\alpha+\beta-1)!(\alpha-1)!(\beta-1)!}\binom{n}{k}, \quad \forall k \in \{0, 1, \ldots, n\}.$$

The key to understanding the statistic G is the following distributional identity:

$$g\left(\mathrm{UA}(n, S), e\right) \stackrel{d}{=} \frac{1}{n^4} B^2_{|T_1|,|T_2|,n-|S|}\left(n - B_{|T_1|,|T_2|,n-|S|}\right)^2, \quad \forall e \in E\left(S\right), \tag{6.1}$$

where T_1 and T_2 are defined, given e, as above, and $\stackrel{d}{=}$ denotes equality in distribution. This is an immediate consequence of the characterization of $(B_{\alpha,\beta,n}, n + (\alpha+\beta) - B_{\alpha,\beta,n})$ as the distribution after n draws of a classical Pólya urn with replacement matrix $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ and starting state (α, β). Similarly, for edges not in the seed S we have

$$g\left(\mathrm{UA}(n, S), e_j\right) \stackrel{d}{=} \frac{1}{n^4} B^2_{1,j,n-j-1}\left(n - B_{1,j,n-j-1}\right)^2, \tag{6.2}$$

where $e_j \in E\left(\mathrm{UA}\left(j+1, S\right)\right) \setminus E\left(\mathrm{UA}\left(j, S\right)\right)$ and $j \in \{|S|, \ldots, n-1\}$.

We use the following elementary facts about the beta-binomial distribution, which we prove in Section 6.5.1:

Fact 6.1. *For every $p \geq 1$ there exists a constant $C(p)$ such that for all α, β, and n such that $n \geq \alpha + \beta$, we have*

$$\left(\mathbb{E}[B^p_{\alpha,\beta,n-\alpha-\beta}]\right)^{1/p} \leq C(p) n \frac{\alpha}{\alpha+\beta}. \tag{6.3}$$

Fact 6.2. *There exists a universal constant $C > 0$ such that whenever $\alpha, \beta \geq 1$, $n \geq \alpha + \beta$, and $t \geq 0$, we have*

$$\mathbb{P}\left(B_{\alpha,\beta,n-\alpha-\beta} < tn\frac{\alpha}{\alpha+\beta}\right) \leq Ct. \tag{6.4}$$

6.2.2 A simple example

After these preliminaries we are now ready to show that $\delta(P_4, S_4) > 0$. To abbreviate notation, in the following we write simply $P \equiv P_4$ and $S \equiv S_4$. In order to show that $\delta(P, S) > 0$, it is enough to show two things:

$$\liminf_{n\to\infty} |\mathbb{E}[G(\mathrm{UA}(n,P))] - \mathbb{E}[G(\mathrm{UA}(n,S))]| > 0, \tag{6.5}$$

and

$$\limsup_{n\to\infty} (\mathrm{Var}[G(\mathrm{UA}(n,P))] + \mathrm{Var}[G(\mathrm{UA}(n,S))]) < \infty. \tag{6.6}$$

The proof can then be concluded using the Cauchy-Schwarz inequality (for more detail on this point, see the proof of Theorem 4.3 in Section 6.3.2).

For $j \geq 4$, let e_j^P denote the edge in $\mathrm{UA}(j+1,P) \setminus \mathrm{UA}(j,P)$, and define e_j^S similarly. Towards (6.5), we first observe that

$$g\left(\mathrm{UA}(n,P), e_j^P\right) \stackrel{d}{=} g\left(\mathrm{UA}(n,S), e_j^S\right), \quad \forall j \in \{4, \ldots, n-1\}.$$

Consequently, we have

$$\mathbb{E}[G(\mathrm{UA}(n,P))] - \mathbb{E}[G(\mathrm{UA}(n,S))] = \sum_{e\in P} \mathbb{E}[g(\mathrm{UA}(n,P), e)] - \sum_{e\in S} \mathbb{E}[g(\mathrm{UA}(n,S), e)].$$

Moreover, note that P has two edges, e_1 and e_2, such that $P \setminus \{e_i\}$ has two connected components of sizes 1 and 3 for $i = 1, 2$. Since this is true for all edges of the star S, we conclude that

$$\mathbb{E}[G(\mathrm{UA}(n,P))] - \mathbb{E}[G(\mathrm{UA}(n,S))] = \mathbb{E}[g(\mathrm{UA}(n,P), e_3)] - \mathbb{E}[g(\mathrm{UA}(n,S), e_1)],$$

where e_3 is the remaining edge of P (i.e., the middle edge of the path). Using (6.1), we thus have

$$\begin{aligned}
&\mathbb{E}[G(\mathrm{UA}(n,P))] - \mathbb{E}[G(\mathrm{UA}(n,S))] \\
&\qquad = \frac{1}{n^4}\left(\mathbb{E}\left[B_{2,2,n-4}^2 (n - B_{2,2,n-4})^2\right] - \mathbb{E}\left[B_{1,3,n-4}^2 (n - B_{1,3,n-4})^2\right]\right) \\
&\qquad = \frac{2n^3 + 5n^2 + 8n + 5}{140n^3},
\end{aligned}$$

where the last equality is attained via a straightforward calculation using explicit formulae for the first four moments of the beta-binomial distribution. We see that

$$\lim_{n\to\infty} (\mathbb{E}[G(\mathrm{UA}(n,P))] - \mathbb{E}[G(\mathrm{UA}(n,S))]) = \frac{1}{70} \neq 0,$$

which establishes (6.5).

It remains to prove (6.6). We show now that $\limsup_{n\to\infty} \mathrm{Var}\left[G\left(\mathrm{UA}\left(n,P\right)\right)\right] < \infty$; the proof that $\limsup_{n\to\infty} \mathrm{Var}\left[G\left(\mathrm{UA}\left(n,S\right)\right)\right] < \infty$ is identical. To abbreviate notation, write T_n for $\mathrm{UA}\left(n,P\right)$. Similarly as above, for $j \geq 4$ let e_j be the edge in $T_{j+1} \setminus T_j$, and let e_1, e_2, and e_3 be the edges of P in some arbitrary order. Using Cauchy-Schwarz we have that

$$\mathrm{Var}\left[G\left(T_n\right)\right] \leq \left(\sum_{j=1}^{n-1} \sqrt{\mathrm{Var}\left[g\left(T_n, e_j\right)\right]}\right)^2. \tag{6.7}$$

For any edge e_i we clearly have $0 \leq g(T_n, e_i) \leq 1$, and so

$$\sum_{j=1}^{3} \sqrt{\mathrm{Var}\left[g\left(T_n, e_j\right)\right]} \leq 3. \tag{6.8}$$

Next, fix j such that $4 \leq j \leq n-1$. Using formula (6.2) we know that

$$g\left(T_n, e_j\right) \stackrel{d}{=} \frac{1}{n^4} B_{1,j,n-j-1}^2 (n - B_{1,j,n-j-1})^2.$$

The estimate (6.3) yields that $\mathbb{E}\left[B_{1,j,n-j-1}^4\right] \leq Cn^4/j^4$, where $C > 0$ is a universal constant. Consequently, we have

$$\mathbb{E}\left[g\left(T_n, e_j\right)^2\right] \leq C/j^4, \quad \forall j \in \left\{4, \ldots, n-1\right\},$$

which, in turn, implies that

$$\sqrt{\mathrm{Var}\left[g\left(T_n, e_j\right)\right]} \leq C/j^2.$$

Plugging this inequality and (6.8) into (6.7) establishes (6.6). This completes the proof of $\delta\left(P_4, S_4\right) > 0$.

The statistic $G\left(\cdot\right)$ cannot distinguish between all pairs of nonisomorphic trees; however, appropriate generalizations of it can. An alternative description of $G\left(\cdot\right)$ is as follows. Let τ be a tree consisting of two vertices connected by a single edge. Up to normalization, the quantity $G\left(T\right)$ is equal to the sum over all embeddings $\phi : \tau \to T$ of the product of the squares of the sizes of the connected components of $T \setminus \phi\left(\tau\right)$. A natural generalization of this definition is to take τ to be an arbitrary finite tree. Moreover, we can assign natural numbers to each vertex of τ, which determines the power to which we raise the size of the respective connected components. In this way we obtain a family of statistics associated with so-called *decorated trees*. It turns out that this generalized family of statistics can indeed distinguish between any pair of nonisomorphic trees; for details see Section 6.3.

6.2.3 Distinguishing large stars: a proof of Theorem 4.5

In the following we first give an upper bound on the probability that $G\left(\mathrm{UA}\left(n,T\right)\right)$ is small, and then we give an upper bound on the probability that $G\left(\mathrm{UA}\left(n,S_k\right)\right)$ is not too small. The two together will prove Theorem 4.5.

First, fix a tree T and choose an arbitrary edge $e_1 \in E(T)$. Let T'_n and T''_n be the two connected components of $\mathrm{UA}(n,T) \setminus \{e_1\}$, defined consistently such that $T'_j \subset T'_{j+1}$ and $T''_j \subset T''_{j+1}$ (and otherwise the order is chosen arbitrarily). We have

$$g\left(\mathrm{UA}\left(n,T\right),e_1\right) = \frac{1}{n^4}\left|T'_n\right|^2\left|T''_n\right|^2 = \frac{1}{n^4}\left|T'_n\right|^2\left(n-\left|T'_n\right|\right)^2.$$

By equation (6.1) we have

$$|T'_n| \stackrel{d}{=} B^2_{a,|T|-a,n-|T|},$$

where $a := \left|T'_{|T|}\right| \geq 1$. Using Fact 6.2 we then have that for all $t>0$,

$$\mathbb{P}\left(\frac{|T'_n|^2}{n^2} < t\frac{1}{|T|^2}\right) \leq \mathbb{P}\left(\frac{|T'_n|^2}{n^2} < t\frac{a^2}{|T|^2}\right) = \mathbb{P}\left(\frac{B^2_{a,|T|-a,n-|T|}}{n^2} < t\frac{a^2}{|T|^2}\right) \leq C\sqrt{t}. \tag{6.9}$$

Consider the event $E_n = \{|T'_n| \leq n/2\}$ and note that if E_n holds then $g\left(\mathrm{UA}\left(n,T\right),e_1\right) \geq \left|T'_n\right|^2/\left(4n^2\right)$. This, together with (6.9), gives

$$\mathbb{P}\left(E_n \cap \left\{g\left(\mathrm{UA}\left(n,T\right),e_1\right) < t\frac{1}{4\left|T\right|^2}\right\}\right) \leq C\sqrt{t}, \quad \forall t>0.$$

By repeating the above argument with T''_n instead of T'_n, we also have

$$\mathbb{P}\left(E_n^C \cap \left\{g\left(\mathrm{UA}\left(n,T\right),e_1\right) < t\frac{1}{4\left|T\right|^2}\right\}\right) \leq C\sqrt{t}, \quad \forall t>0,$$

and thus we conclude that

$$\mathbb{P}\left(g\left(\mathrm{UA}\left(n,T\right),e_1\right) < t\frac{1}{4\left|T\right|^2}\right) \leq 2C\sqrt{t}, \quad \forall t>0.$$

Now since $G\left(\mathrm{UA}\left(n,T\right)\right) \geq g\left(\mathrm{UA}\left(n,T\right),e_1\right)$, and by setting $z = t/\left(4\left|T\right|^2\right)$, we finally get that

$$\mathbb{P}\left(G\left(\mathrm{UA}\left(n,T\right)\right) < z\right) \leq 4C\sqrt{z}\left|T\right|, \quad \forall z>0. \tag{6.10}$$

In order to understand the distribution of $G\left(\mathrm{UA}\left(n,S_k\right)\right)$ we first estimate its mean:

$$\begin{aligned}
&\mathbb{E}\left[G\left(\mathrm{UA}\left(n,S_k\right)\right)\right] \\
&\quad= \frac{k-1}{n^4}\mathbb{E}\left[B^2_{1,k-1,n-k}\left(n-B_{1,k-1,n-k}\right)^2\right] + \sum_{j=k}^{n-1}\frac{1}{n^4}\mathbb{E}\left[B^2_{1,j,n-j-1}\left(n-B_{1,j,n-j-1}\right)^2\right] \\
&\quad\leq \frac{k-1}{n^2}\mathbb{E}\left[B^2_{1,k-1,n-k}\right] + \frac{1}{n^2}\sum_{j=k}^{n-1}\mathbb{E}\left[B^2_{1,j,n-j-1}\right] \stackrel{(6.3)}{\leq} \frac{C'}{k} + \sum_{j=k}^{n-1}\frac{C'}{j^2} \leq \frac{3C'}{k}
\end{aligned}$$

for some absolute constant C'. Now using Markov's inequality with this estimate, and also taking $z = 3C'/\sqrt{k}$ in the inequality (6.10), we get that

$$\mathbb{P}\left(G\left(\mathrm{UA}\left(n, S_k\right)\right) \geq 3C'/\sqrt{k}\right) \leq \frac{1}{\sqrt{k}} \qquad \text{and} \qquad \mathbb{P}\left(G\left(\mathrm{UA}\left(n, T\right)\right) < 3C'/\sqrt{k}\right) \leq \frac{C''}{k^{1/4}}$$

for some absolute constant C''. This then immediately implies that $\delta\left(S_k, T\right) \to 1$ as $k \to \infty$.

6.3 Proof of Theorem 4.3

After the intuition and simple examples provided in Section 6.2, in this section we fully prove Theorem 4.3. As mentioned in Section 4.3.2, the proof shares some features with the proof in [21] for preferential attachment, but is different in several ways. These differences are discussed in detail after the proof, in Section 6.4.

Notation. For a graph G, denote by $V(G)$ the set of its vertices, by $E(G)$ the set of its edges, and by $\mathrm{diam}(G)$ its diameter. For brevity, we often write $v \in G$ instead of $v \in V(G)$. For integers $k, j \geq 1$, define the descending factorial $[k]_j = k(k-1)\ldots(k-j+1)$, and also let $[k]_0 = 1$. For two sequences of real numbers $\{a_n\}_{n\geq 0}$ and $\{b_n\}_{n\geq 0}$, we write $a_n \precsim b_n$ (to be read as a_n is less than b_n up to log factors) if there exist constants $c > 0$, $\gamma \in \mathbb{R}$, and n_0 such that $|a_n| \leq c\left(\log(n)\right)^{\gamma}|b_n|$ for all $n \geq n_0$. For a sequence $\{a_n\}_{n\geq 0}$ of real numbers, define $\Delta_n a = a_{n+1} - a_n$ for $n \geq 0$.

6.3.1 Decorated trees

A *decorated tree* is a pair $\underline{\tau} = (\tau, \ell)$ consisting of a tree τ and a family of nonnegative integers $(\ell(v); v \in \tau)$, called labels, associated with its vertices; see Figure 6.1 for an illustration. Let $\mathcal{D}$ denote the set of all decorated trees, $\mathcal{D}_+$ the set of all decorated trees where every label is positive, $\mathcal{D}_0$ the set of all decorated trees where there exists a zero label, and finally let $\mathcal{D}_0^*$ denote the set of all decorated trees where there exists a leaf which has label zero.

Define $|\underline{\tau}|$ to be the number of vertices of τ, and let $w(\underline{\tau}) := \sum_{v\in\tau} \ell(v)$ denote the total weight of $\underline{\tau}$. For $\underline{\tau}, \underline{\tau}' \in \mathcal{D}$, let $\underline{\tau} \prec \underline{\tau}'$ if $|\underline{\tau}| < |\underline{\tau}'|$ and $w(\underline{\tau}) \leq w(\underline{\tau}')$ or $|\underline{\tau}| \leq |\underline{\tau}'|$ and $w(\underline{\tau}) < w(\underline{\tau}')$. This defines a strict partial order $\prec$, and let $\preccurlyeq$ denote the associated partial order, i.e., $\underline{\tau} \preccurlyeq \underline{\tau}'$ if and only if $\underline{\tau} \prec \underline{\tau}'$ or $\underline{\tau} = \underline{\tau}'$.

For $\underline{\tau} \in \mathcal{D}$, let $L(\underline{\tau})$ denote the set of leaves of τ, let $L_0(\underline{\tau}) := \{v \in L(\underline{\tau}) : \ell(v) = 0\}$, $L_1(\underline{\tau}) := \{v \in L(\underline{\tau}) : \ell(v) = 1\}$, and $L_{0,1}(\underline{\tau}) := L_0(\underline{\tau}) \cup L_1(\underline{\tau})$. For $\underline{\tau} \in \mathcal{D}$ and $v \in L(\underline{\tau})$, define $\underline{\tau}_v$ to be the same as $\underline{\tau}$ except the leaf v and its label are removed. For $\underline{\tau} \in \mathcal{D}$ and a vertex $v \in \underline{\tau}$ such that $\ell(v) \geq 2$, define $\underline{\tau}'_v$ to be the same as $\underline{\tau}$ except the label of v is decreased by one, i.e., $\ell_{\underline{\tau}'_v}(v) = \ell_{\underline{\tau}}(v) - 1$.

6.3.2 Statistics and distinguishing martingales

Given two trees τ and T, a map $\phi : \tau \to T$ is called an *embedding* if ϕ is an injective graph homomorphism. That is, ϕ is an injective map from $V(\tau)$ to $V(T)$ such that $\{u, v\} \in E(\tau)$

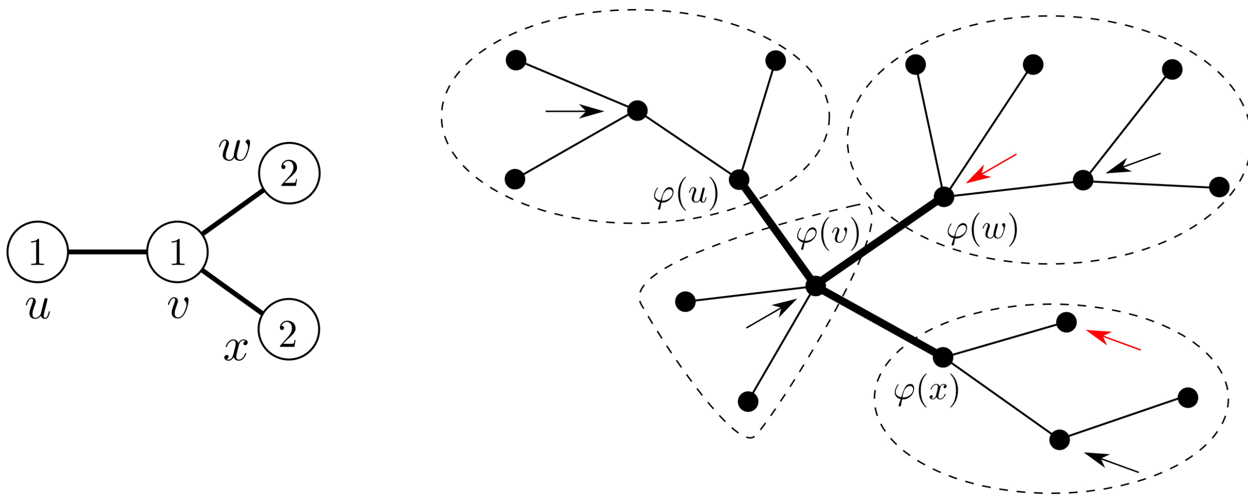

Figure 6.1: **A decorated tree and a decorated embedding.** On the left is a decorated tree $\underline{\tau} = (\tau, \ell)$ with four vertices, two of them having label 1, and two of them having label 2. On the right is a larger tree T, and an embedding $\phi : \tau \to T$ depicted in bold. The connected components of the forest $\widehat{T}(T, \tau, \phi)$ are circled with dashed lines, with the component sizes being $f_{\phi(u)}(T) = 5$, $f_{\phi(v)}(T) = 3$, $f_{\phi(w)}(T) = 6$, and $f_{\phi(x)}(T) = 4$. A decorated embedding $\underline{\phi}$ is also depicted, which consists of the embedding ϕ together with the mapping of $w(\underline{\tau}) = 6$ arrows to vertices of T. The arrows in each subtree are distinguishable, which is why they are depicted using different colors.

implies that $\{\phi(u), \phi(v)\} \in E(T)$.

For two trees τ and T, and an embedding $\phi : \tau \to T$, denote by $\widehat{T} = \widehat{T}(T, \tau, \phi)$ the forest obtained from T by removing the images of the edges of τ under the embedding ϕ; see Figure 6.1 for an illustration. Note that the forest $\widehat{T}$ consists of exactly $|\tau|$ trees, and each tree contains exactly one vertex which is the image of a vertex of τ under ϕ. For $v \in \tau$, denote by $f_{\phi(v)}(T)$ the number of vertices of the tree in $\widehat{T}$ which contains the vertex $\phi(v)$. Using this notation, for a decorated tree $\underline{\tau} \in \mathcal{D}$ define

$$F_{\underline{\tau}}(T) = \sum_{\phi} \prod_{v \in \tau} \left[f_{\phi(v)}(T) \right]_{\ell(v)}, \tag{6.11}$$

where the sum is over all embeddings $\phi : \tau \to T$. If there are no such embeddings then $F_{\underline{\tau}}(T) = 0$ by definition. Note that if $\underline{\tau}$ consists of a single vertex with label $k \geq 0$, then $F_{\underline{\tau}}(T) = |T| \times [|T|]_k$.

The quantity $F_{\underline{\tau}}(T)$ has a combinatorial interpretation which is useful: it is the number of *decorated embeddings* of $\underline{\tau}$ in T, defined as follows. Imagine that for each vertex $v \in \tau$ there are $\ell(v)$ distinguishable (i.e., ordered) arrows pointing to v. A decorated embedding $\underline{\phi}$ is an embedding ϕ of τ in T, together with a mapping of the arrows to vertices of T in such a way that each arrow pointing to $v \in \tau$ is mapped to a vertex in the tree of $\widehat{T}$ that contains $\phi(v)$, with distinct arrows mapped to distinct vertices. See Figure 6.1 for an illustration.

The quantities $F_{\underline{\tau}}$ are also more amenable to analysis than other statistics, because their expectations satisfy recurrence relations, as described in Section 6.3.3. Using the statistics $F_{\underline{\tau}}$ it is possible to create martingales that distinguish between different seeds.

Proposition 6.3. *Let $\underline{\tau} \in \mathcal{D}_+$. There exists a family of constants*

$$\{c_n(\underline{\tau}, \underline{\tau}') : \underline{\tau}' \in \mathcal{D}_+, \underline{\tau}' \preccurlyeq \underline{\tau}, n \geq 2\}$$

with $c_n(\underline{\tau}, \underline{\tau}) > 0$ such that for every seed tree S, the process $\left\{M_{\underline{\tau}}^{(S)}(n)\right\}_{n \geq |S|}$ defined by

$$M_{\underline{\tau}}^{(S)}(n) = \sum_{\underline{\tau}' \in \mathcal{D}_+ : \underline{\tau}' \preccurlyeq \underline{\tau}} c_n(\underline{\tau}, \underline{\tau}') F_{\underline{\tau}'}(\mathrm{UA}(n, S))$$

is a martingale with respect to the natural filtration $\mathcal{F}_n = \sigma\{\mathrm{UA}(|S|, S), \ldots, \mathrm{UA}(n, S)\}$, and is bounded in L^2.

Note that in the construction of these martingales we only use decorated trees where every label is positive. As we shall see, we analyze decorated trees having a zero label in order to show that the martingales above are bounded in L^2. See Sections 6.3.4 and 6.4 for more details and discussion on this point.

We now prove Theorem 4.3 using Proposition 6.3, which we then prove in the following subsections.

Proof of Theorem 4.3. Let S and T be two nonisomorphic trees with at least three vertices, and let $n_0 := \max\{|S|, |T|\}$. First we show that there exists $\underline{\tau} \in \mathcal{D}_+$ such that

$$\mathbb{E}[F_{\underline{\tau}}(\mathrm{UA}(n_0, S))] \neq \mathbb{E}[F_{\underline{\tau}}(\mathrm{UA}(n_0, T))]. \tag{6.12}$$

Assume without loss of generality that $|S| \leq |T|$, and let $\underline{\tau}$ be equal to T with labels $\ell(v) = 1$ for all $v \in T$. Then for every tree T' with $|T'| = |T|$ we have $F_{\underline{\tau}}(T') = F_{\underline{\tau}}(T) \times \mathbf{1}_{\{T' = T\}}$, since if T' and T are nonisomorphic then there is no embedding of T in T'. Note also that $F_{\underline{\tau}}(T)$ is the number of automorphisms of T, which is positive. Consequently we have

$$\mathbb{E}[F_{\underline{\tau}}(\mathrm{UA}(n_0, S))] = F_{\underline{\tau}}(T) \times \mathbb{P}[\mathrm{UA}(n_0, S) = T].$$

When $|S| = |T|$, we have $\mathbb{P}[\mathrm{UA}(n_0, S) = T] = 0$. When $|S| < |T|$, it is easy to see that the isomorphism class of $\mathrm{UA}(n_0, S)$ is nondeterministic (here we use the fact that $|S| \geq 3$), and so $\mathbb{P}[\mathrm{UA}(n_0, S) = T] < 1$. In both cases we have that (6.12) holds.

Now let $\underline{\tau} \in \mathcal{D}_+$ be a minimal (for the partial order $\preccurlyeq$ on $\mathcal{D}_+$) decorated tree for which (6.12) holds. By definition we then have that $\mathbb{E}[F_{\underline{\tau}'}(\mathrm{UA}(n_0, S))] = \mathbb{E}[F_{\underline{\tau}'}(\mathrm{UA}(n_0, T))]$ for every $\underline{\tau}' \in \mathcal{D}_+$ such that $\underline{\tau}' \prec \underline{\tau}$. By the construction of the martingales in Proposition 6.3 we then have that

$$\mathbb{E}\left[M_{\underline{\tau}}^{(S)}(n_0)\right] \neq \mathbb{E}\left[M_{\underline{\tau}}^{(T)}(n_0)\right].$$

Clearly for any $n \geq n_0$ we have that $\mathrm{TV}(\mathrm{UA}(n,S), \mathrm{UA}(n,T)) \geq \mathrm{TV}\left(M_{\underline{\tau}}^{(S)}(n), M_{\underline{\tau}}^{(T)}(n)\right)$. Now let (X,Y) be a coupling of $\left(M_{\underline{\tau}}^{(S)}(n), M_{\underline{\tau}}^{(T)}(n)\right)$. We need to bound from below $\mathbb{P}(X \neq Y)$ in order to obtain a lower bound on $\mathrm{TV}\left(M_{\underline{\tau}}^{(S)}(n), M_{\underline{\tau}}^{(T)}(n)\right)$. Using the Cauchy-Schwarz inequality we have

$$\mathbb{P}(X \neq Y) \geq \frac{(\mathbb{E}[|X-Y|])^2}{\mathbb{E}\left[(X-Y)^2\right]}.$$

By Jensen's inequality we have that $(\mathbb{E}[|X-Y|])^2 \geq (\mathbb{E}[X] - \mathbb{E}[Y])^2$, and furthermore

$$\begin{aligned}\mathbb{E}\left[(X-Y)^2\right] &= \mathbb{E}\left[(X - \mathbb{E}[X] + \mathbb{E}[X] - \mathbb{E}[Y] + \mathbb{E}[Y] - Y)^2\right] \\ &= \mathrm{Var}(X) + \mathrm{Var}(Y) + (\mathbb{E}[X] - \mathbb{E}[Y])^2 + 2\mathbb{E}[(X - \mathbb{E}[X])(\mathbb{E}[Y] - Y)] \\ &\leq 2\,\mathrm{Var}(X) + 2\,\mathrm{Var}(Y) + (\mathbb{E}[X] - \mathbb{E}[Y])^2.\end{aligned}$$

Thus we have shown that

$$\begin{aligned}&\mathrm{TV}(\mathrm{UA}(n,S), \mathrm{UA}(n,T)) \\ &\qquad \geq \frac{\left(\mathbb{E}\left[M_{\underline{\tau}}^{(S)}(n)\right] - \mathbb{E}\left[M_{\underline{\tau}}^{(T)}(n)\right]\right)^2}{2\,\mathrm{Var}\left(M_{\underline{\tau}}^{(S)}(n)\right) + 2\,\mathrm{Var}\left(M_{\underline{\tau}}^{(T)}(n)\right) + \left(\mathbb{E}\left[M_{\underline{\tau}}^{(S)}(n)\right] - \mathbb{E}\left[M_{\underline{\tau}}^{(T)}(n)\right]\right)^2}.\end{aligned}$$

Since $M_{\underline{\tau}}^{(S)}$ and $M_{\underline{\tau}}^{(T)}$ are martingales, for every $n \geq n_0$ we have that $\mathbb{E}\left[M_{\underline{\tau}}^{(S)}(n)\right] - \mathbb{E}\left[M_{\underline{\tau}}^{(T)}(n)\right] = \mathbb{E}\left[M_{\underline{\tau}}^{(S)}(n_0)\right] - \mathbb{E}\left[M_{\underline{\tau}}^{(T)}(n_0)\right] \neq 0$. Also, since the two martingales are bounded in L^2, we have that $\mathrm{Var}\left(M_{\underline{\tau}}^{(S)}(n)\right) + \mathrm{Var}\left(M_{\underline{\tau}}^{(T)}(n)\right)$ is bounded as $n \to \infty$. We conclude that $\delta(S,T) > 0$. □

6.3.3 Recurrence relation

The following recurrence relation for the conditional expectations of $F_{\underline{\tau}}(\mathrm{UA}(n,S))$ is key to estimating the moments of $F_{\underline{\tau}}(\mathrm{UA}(n,S))$.

Lemma 6.4. *Let $\underline{\tau} \in \mathcal{D}$ be such that $|\underline{\tau}| \geq 2$. Then for every seed tree S and for every $n \geq |S|$ we have*

$$\begin{aligned}\mathbb{E}\left[F_{\underline{\tau}}(\mathrm{UA}(n+1,S)) \mid \mathcal{F}_n\right] &= \left(1 + \frac{w(\underline{\tau})}{n}\right) F_{\underline{\tau}}(\mathrm{UA}(n,S)) \\ &\quad + \frac{1}{n}\left\{\sum_{v \in \tau : \ell(v) \geq 2} \ell(v)(\ell(v)-1) F_{\underline{\tau}_v'}(\mathrm{UA}(n,S)) + \sum_{v \in L_{0,1}(\underline{\tau})} F_{\underline{\tau}_v}(\mathrm{UA}(n,S))\right\}. \quad (6.13)\end{aligned}$$

Proof. Fix $\underline{\tau} \in \mathcal{D}$ with $|\underline{\tau}| \geq 2$, fix a seed tree S, and let $n \geq |S|$. To simplify notation we omit the dependence on S and write T_n instead of $\mathrm{UA}(n, S)$. When evaluating $\mathbb{E}\left[F_{\underline{\tau}}(T_{n+1}) \,|\, \mathcal{F}_n\right]$ we work conditionally on $\mathcal{F}_n$, so we may consider T_n as being fixed.

Let u_{n+1} denote the vertex present in T_{n+1} but not in T_n, and let u_n denote its neighbor in T_{n+1}. Let $\mathcal{E}_{n+1}$ denote the set of all embeddings $\phi : \tau \to T_{n+1}$; we can write $\mathcal{E}_{n+1}$ as the disjoint union of the set of those using only vertices of T_n, denoted by $\mathcal{E}_n$, and the set of those using the new vertex u_{n+1}, denoted by $\mathcal{E}_{n+1} \setminus \mathcal{E}_n$. To simplify notation, if $\underline{\tau} \in \mathcal{D}$, T is a tree, and $\phi : \tau \to T$ is an embedding, write $\mathcal{W}_\phi(T) = \prod_{v \in \tau} \left[f_{\phi(v)}(T)\right]_{\ell(v)}$ for the number of decorated embeddings of $\underline{\tau}$ in T that use the embedding ϕ. We then have $F_{\underline{\tau}}(T_{n+1}) = \sum_{\phi \in \mathcal{E}_n} \mathcal{W}_\phi(T_{n+1}) + \sum_{\phi \in \mathcal{E}_{n+1} \setminus \mathcal{E}_n} \mathcal{W}_\phi(T_{n+1})$ and we deal with the two sums separately.

First let $\phi \in \mathcal{E}_n$. For $v \in \tau$, denote by E_v the event that u_n is in the same tree of $\widehat{T_n}$ as $\phi(v)$ (recall the definition of $\widehat{T_n}$ from Section 6.3.2). Clearly $\mathbb{P}(E_v \,|\, \mathcal{F}_n) = f_{\phi(v)}(T_n)/n$. Under the event E_v we have that $f_{\phi(v)}(T_{n+1}) = f_{\phi(v)}(T_n) + 1$, while for every $v' \in \tau \setminus \{v\}$ we have $f_{\phi(v')}(T_{n+1}) = f_{\phi(v')}(T_n)$. Now using the identities $[d+1]_\ell = [d]_\ell + \ell \times [d]_{\ell-1}$ and $d \times [d]_{\ell-1} = [d]_\ell + (\ell-1) \times [d]_{\ell-1}$, which hold for every $d, \ell \geq 1$, and also using $[d+1]_0 = [d]_0$, we have that

$$
\begin{aligned}
&\mathbb{E}\left[\mathcal{W}_\phi(T_{n+1}) \,|\, \mathcal{F}_n\right] \\
&\quad = \sum_{v \in \tau} \frac{f_{\phi(v)}(T_n)}{n} \left[f_{\phi(v)}(T_n) + 1\right]_{\ell(v)} \prod_{v' \in \tau \setminus \{v\}} \left[f_{\phi(v')}(T_n)\right]_{\ell(v')} \\
&\quad = \mathcal{W}_\phi(T_n) + \frac{1}{n} \sum_{v \in \tau : \ell(v) \geq 1} \ell(v) \, f_{\phi(v)}(T_n) \left[f_{\phi(v)}(T_n)\right]_{\ell(v)-1} \prod_{v' \in \tau \setminus \{v\}} \left[f_{\phi(v')}(T_n)\right]_{\ell(v')} \\
&\quad = \left(1 + \frac{w(\underline{\tau})}{n}\right) \mathcal{W}_\phi(T_n) \\
&\qquad + \frac{1}{n} \sum_{v \in \tau : \ell(v) \geq 2} \ell(v)\left(\ell(v) - 1\right) \left[f_{\phi(v)}(T_n)\right]_{\ell(v)-1} \prod_{v' \in \tau \setminus \{v\}} \left[f_{\phi(v')}(T_n)\right]_{\ell(v')} \\
&\quad = \left(1 + \frac{w(\underline{\tau})}{n}\right) \mathcal{W}_\phi(T_n) + \frac{1}{n} \sum_{v \in \tau : \ell(v) \geq 2} \ell(v)\left(\ell(v) - 1\right) \mathcal{W}_{\phi'_v}(T_n),
\end{aligned}
$$

where ϕ'_v is the embedding equal to ϕ of the decorated tree $\underline{\tau}'_v$. Now as ϕ runs through the embeddings of $\underline{\tau}$ in T_n, ϕ'_v runs exactly through the embeddings of $\underline{\tau}'_v$. So we have that

$$
\sum_{\phi \in \mathcal{E}_n} \mathbb{E}\left[\mathcal{W}_\phi(T_{n+1}) \,|\, \mathcal{F}_n\right] = \left(1 + \frac{w(\underline{\tau})}{n}\right) F_{\underline{\tau}}(T_n) + \frac{1}{n} \sum_{v \in \tau : \ell(v) \geq 2} \ell(v)\left(\ell(v) - 1\right) F_{\underline{\tau}'_v}(T_n). \quad (6.14)
$$

Now fix T_{n+1} and consider $\phi \in \mathcal{E}_{n+1} \setminus \mathcal{E}_n$. Let $w \in \tau$ be such that $\phi(w) = u_{n+1}$. Since ϕ is an embedding, we must have $w \in L(\underline{\tau})$. Note that if $w \notin L_{0,1}(\underline{\tau})$ then $\mathcal{W}_\phi(T_{n+1}) = 0$. If $w \in L_{0,1}(\underline{\tau})$ then denote by $\mathcal{E}_w$ the set of all embeddings $\phi \in \mathcal{E}_{n+1} \setminus \mathcal{E}_n$ such that $\phi(w) = u_{n+1}$.

Now fix $w \in L_{0,1}(\underline{\tau})$ and $\phi \in \mathcal{E}_w$. Note that ϕ restricted to $\tau \setminus \{w\}$ is an embedding of $\underline{\tau}_w$ in T_n; call this ϕ_w. Let x be the neighbor of w in τ. We then must have $\phi(x) = u_n$, and also $\left[f_{\phi(w)}(T_{n+1})\right]_{\ell(w)} = 1$ (irrespective of whether $w \in L_0(\underline{\tau})$ or $w \in L_1(\underline{\tau})$). Furthermore, for every $w' \in \tau \setminus \{w\}$ we have $f_{\phi(w')}(T_{n+1}) = f_{\phi(w')}(T_n)$. Thus we have

$$\mathcal{W}_\phi(T_{n+1}) = \mathcal{W}_{\phi_w}(T_n)\,\mathbf{1}_{\{\phi(w)=u_{n+1}\}}.$$

For fixed $w \in L_{0,1}(\underline{\tau})$, as ϕ runs through $\mathcal{E}_w$, ϕ_w runs through all the embeddings of $\underline{\tau}_w$ in T_n. So summing over $w \in L$ we obtain

$$\begin{aligned}\sum_{\phi\in\mathcal{E}_{n+1}\setminus\mathcal{E}_n} \mathcal{W}_\phi(T_{n+1}) &= \sum_{w\in L_{0,1}(\underline{\tau})}\ \sum_{\phi_w:\underline{\tau}_w\to T_n} \mathcal{W}_{\phi_w}(T_n)\,\mathbf{1}_{\{\phi(w)=u_{n+1}\}}\\ &= \sum_{w\in L_{0,1}(\underline{\tau})} F_{\underline{\tau}_w}(T_n)\,\mathbf{1}_{\{\phi(w)=u_{n+1}\}}.\end{aligned}$$

Now taking conditional expectation given $\mathcal{F}_n$, we get that

$$\mathbb{E}\left[\sum_{\phi\in\mathcal{E}_{n+1}\setminus\mathcal{E}_n} \mathcal{W}_\phi(T_{n+1}) \,\middle|\, \mathcal{F}_n\right] = \frac{1}{n}\sum_{w\in L_{0,1}(\underline{\tau})} F_{\underline{\tau}_w}(T_n). \tag{6.15}$$

Summing (6.14) and (6.15) we obtain (6.13). □

6.3.4 Moment estimates

Using the recurrence relation of Lemma 6.4 proved in the previous subsection, we now establish moment estimates on the number of decorated embeddings $F_{\underline{\tau}}(\mathrm{UA}(n,S))$. These are then used in the next subsection to show that the martingales of Proposition 6.3 are bounded in L^2.

The first moment estimates are a direct corollary of Lemma 6.4.

Corollary 6.5. *Let $\underline{\tau} \in \mathcal{D}$ be a decorated tree and let S be a seed tree.*

(a) We have that $n^{w(\underline{\tau})} \gtrsim \mathbb{E}[F_{\underline{\tau}}(\mathrm{UA}(n,S))]$.

(b) If $|\underline{\tau}| \geq 2$ and $\underline{\tau} \in \mathcal{D} \setminus \mathcal{D}_0^*$, then we have that $\mathbb{E}[F_{\underline{\tau}}(\mathrm{UA}(n,S))] \gtrsim n^{w(\underline{\tau})}$.

(c) If $|\underline{\tau}| = 1$ or $\underline{\tau} \in \mathcal{D}_0^*$, then we have that $\mathbb{E}[F_{\underline{\tau}}(\mathrm{UA}(n,S))] \gtrsim n^{w(\underline{\tau})+1}$.

Proof. Fix a seed tree S and, as before, write T_n instead of $\mathrm{UA}(n,S)$ in order to simplify notation. First, recall that if $|\underline{\tau}| = 1$, then $F_{\underline{\tau}}(T_n) = n \times [n]_{w(\underline{\tau})}$, so the statements of part (a) and (c) hold in this case. In the following we can therefore assume that $|\underline{\tau}| \geq 2$.

Lemma 6.4 then implies that for every $n \geq |S|$ we have

$$\mathbb{E}[F_{\underline{\tau}}(T_{n+1})] \geq \left(1 + \frac{w(\underline{\tau})}{n}\right)\mathbb{E}[F_{\underline{\tau}}(T_n)].$$

Since there exists n_0 such that $\mathbb{E}\left[F_{\underline{\tau}}\left(T_{n_0}\right)\right] > 0$ (one can take, e.g., $n_0 = |S| + |\underline{\tau}|$), this immediately implies part (a) (see Section 6.5.2 for further details).

We prove parts (b) and (c) by induction on $\underline{\tau}$ for the partial order $\preccurlyeq$. We have already checked that the statement holds when $|\underline{\tau}| = 1$, so the base case of the induction holds. Now fix $\underline{\tau}$ such that $|\underline{\tau}| \geq 2$, and assume that (b) and (c) hold for all $\underline{\tau}'$ such that $\underline{\tau}' \prec \underline{\tau}$. There are two cases to consider: either $\underline{\tau} \in \mathcal{D} \setminus \mathcal{D}_0^*$ or $\underline{\tau} \in \mathcal{D}_0^*$.

If $\underline{\tau} \in \mathcal{D} \setminus \mathcal{D}_0^*$ then $L_0\left(\underline{\tau}\right) = \emptyset$. Note that for every $v \in L_1\left(\underline{\tau}\right)$ we have $w\left(\underline{\tau}_v\right) = w\left(\underline{\tau}\right) - 1$, and also for every $v \in \tau$ such that $\ell\left(v\right) \geq 2$ we have $w\left(\underline{\tau}_v'\right) = w\left(\underline{\tau}\right) - 1$. Therefore by induction for every $v \in L_1\left(\underline{\tau}\right)$ we have $\mathbb{E}\left[F_{\underline{\tau}_v}\left(T_n\right)\right] \lesssim n^{w(\underline{\tau})}$ and also for every $v \in \tau$ such that $\ell\left(v\right) \geq 2$ we have $\mathbb{E}\left[F_{\underline{\tau}_v'}\left(T_n\right)\right] \lesssim n^{w(\underline{\tau})}$. So by Lemma 6.4 there exist constants $C, \gamma > 0$ such that

$$\mathbb{E}\left[F_{\underline{\tau}}\left(T_{n+1}\right)\right] \leq \left(1 + \frac{w\left(\underline{\tau}\right)}{n}\right) \mathbb{E}\left[F_{\underline{\tau}}\left(T_n\right)\right] + C\left(\log\left(n\right)\right)^{\gamma} n^{w(\underline{\tau})-1}. \tag{6.16}$$

This then implies that $\mathbb{E}\left[F_{\underline{\tau}}\left(T_n\right)\right] \lesssim n^{w(\underline{\tau})}$; see Section 6.5.2 for details.

If $\underline{\tau} \in \mathcal{D}_0^*$ then $L_0\left(\underline{\tau}\right) \neq \emptyset$ and note that for every $v \in L_0\left(\underline{\tau}\right)$ we have $w\left(\underline{\tau}_v\right) = w\left(\underline{\tau}\right)$. If for every $v \in L_0\left(\underline{\tau}\right)$ we have $\underline{\tau}_v \in \mathcal{D} \setminus \mathcal{D}_0^*$, then the same argument as in the previous paragraph goes through, and we have that $\mathbb{E}\left[F_{\underline{\tau}}\left(T_n\right)\right] \lesssim n^{w(\underline{\tau})}$. However, if there exists $v \in L_0\left(\underline{\tau}\right)$ such that $\underline{\tau}_v \in \mathcal{D}_0^*$, then (6.16) does not hold in this case; instead we have from Lemma 6.4 that there exist constants $C, \gamma > 0$ such that

$$\mathbb{E}\left[F_{\underline{\tau}}\left(T_{n+1}\right)\right] \leq \left(1 + \frac{w\left(\underline{\tau}\right)}{n}\right) \mathbb{E}\left[F_{\underline{\tau}}\left(T_n\right)\right] + C\left(\log\left(n\right)\right)^{\gamma} n^{w(\underline{\tau})}.$$

Similarly as before, this then implies that $\mathbb{E}\left[F_{\underline{\tau}}\left(T_n\right)\right] \lesssim n^{w(\underline{\tau})+1}$; see Section 6.5.2 for details. $\square$

The second moment estimates require additional work. First, recall again that if $|\underline{\tau}| = 1$, then $F_{\underline{\tau}}\left(\mathrm{UA}\left(n, S\right)\right) = n \times [n]_{w(\underline{\tau})}$. Consequently, we have that $F_{\underline{\tau}}\left(\mathrm{UA}\left(n, S\right)\right)^2 \lesssim n^{2w(\underline{\tau})+2}$ and also that $\left(F_{\underline{\tau}}\left(\mathrm{UA}\left(n+1, S\right)\right) - F_{\underline{\tau}}\left(\mathrm{UA}\left(n, S\right)\right)\right)^2 \lesssim n^{2w(\underline{\tau})}$.

Lemma 6.6. *Let $\underline{\tau} \in \mathcal{D}_+$ with $|\underline{\tau}| \geq 2$ and let S be a seed tree.*

(a) We have that $\mathbb{E}\left[F_{\underline{\tau}}\left(\mathrm{UA}\left(n, S\right)\right)^2\right] \lesssim n^{2w(\underline{\tau})}$.

(b) We have that $\mathbb{E}\left[\left(F_{\underline{\tau}}\left(\mathrm{UA}\left(n+1, S\right)\right) - F_{\underline{\tau}}\left(\mathrm{UA}\left(n, S\right)\right)\right)^2\right] \lesssim n^{2w(\underline{\tau})-2}$.

We note that part (a) of the lemma follows in a short and simple way once part (b) is proven. However, for expository reasons, we first prove part (a) directly, and then prove part (b), whose proof is similar to, and builds upon, the proof of part (a).

Proof. Fix $\underline{\tau} \in \mathcal{D}_+$ with $|\underline{\tau}| \geq 2$ and a seed tree S. Define

$$K \equiv K\left(\underline{\tau}\right) := \max\left\{4\left(|\underline{\tau}| + w\left(\underline{\tau}\right)\right), 20\right\}.$$

We always have $F_{\underline{\tau}}(\mathrm{UA}(n,S)) \leq n^{K/4}$, since the number of embeddings of τ in $\mathrm{UA}(n,S)$ is at most $n^{|\underline{\tau}|}$, and the product of the subtree sizes raised to appropriate powers is at most $n^{w(\underline{\tau})}$. By Lemma 6.8 in Section 6.5.3 there exists a constant $C(S)$ such that

$$\mathbb{P}(\operatorname{diam}(\mathrm{UA}(n,S)) > K\log(n)) \leq C(S)\, n^{-K/2}.$$

Therefore we have

$$\mathbb{E}\left[F_{\underline{\tau}}(\mathrm{UA}(n,S))^2\, \mathbf{1}_{\{\operatorname{diam}(\mathrm{UA}(n,S))>K\log(n)\}}\right] \leq C(S),$$

and similarly

$$\mathbb{E}\left[\left(F_{\underline{\tau}}(\mathrm{UA}(n+1,S)) - F_{\underline{\tau}}(\mathrm{UA}(n,S))\right)^2 \mathbf{1}_{\{\operatorname{diam}(\mathrm{UA}(n,S))>K\log(n)\}}\right] \precsim 1.$$

Therefore in the remainder of the proof we may, roughly speaking, assume that

$$\operatorname{diam}(\mathrm{UA}(n,S)) \leq K\log(n);$$

this will be made precise later.

To simplify notation, write simply T_n instead of $\mathrm{UA}(n,S)$. Our proof is combinatorial and uses the notion of decorated embeddings as described in Section 6.3.2. We start with the proof of (a) which is simpler. We say that $\underline{\phi} = \underline{\phi}_1 \times \underline{\phi}_2$ is a *decorated map* if it is a map such that both $\underline{\phi}_1$ and $\underline{\phi}_2$ are decorated embeddings from $\underline{\tau}$ to T_n. Note that $\underline{\phi}$ is not necessarily injective. If $\underline{\phi}$ is a decorated embedding or a decorated map, we denote by ϕ the map of the tree without the choices of vertices associated with the arrows.

Now observe that $F_{\underline{\tau}}(T_n)^2$ is exactly the number of decorated maps $\underline{\phi} = \underline{\phi}_1 \times \underline{\phi}_2$. We partition the set of such decorated maps into two parts: let $\mathcal{E}^1_{\underline{\tau}}(T_n)$ denote the set of all such decorated maps where $\phi_1(\tau)\cap\phi_2(\tau) \neq \emptyset$, and let $\mathcal{E}^2_{\underline{\tau}}(T_n)$ denote the set of all such decorated maps where $\phi_1(\tau)\cap\phi_2(\tau) = \emptyset$. Clearly $F_{\underline{\tau}}(T_n)^2 = \left|\mathcal{E}^1_{\underline{\tau}}(T_n)\right| + \left|\mathcal{E}^2_{\underline{\tau}}(T_n)\right|$. This partition is not necessary for the proof, but it is helpful for exposition.

We first estimate $\left|\mathcal{E}^1_{\underline{\tau}}(T_n)\right|$. To do this, we associate to each decorated map $\underline{\phi} \in \mathcal{E}^1_{\underline{\tau}}(T_n)$ a decorated tree $\underline{\sigma}$ and a decorated embedding $\underline{\psi}$ of it in T_n, in the following way; see also Figure 6.2 for an illustration. We take simply the union of the images of the decorated embeddings $\underline{\phi}_1$ and $\underline{\phi}_2$, and if these share any vertices, edges, or arrows, then we identify them (i.e., we take only a single copy). The resulting union is the image of a decorated tree $\underline{\sigma}$ under a decorated embedding $\underline{\psi}$; note that $\underline{\sigma}$ is uniquely defined, and $\underline{\psi}$ is uniquely defined up to the ordering of the arrows associated with $\underline{\sigma}$. To define $\underline{\psi}$ uniquely, we arbitrarily define the ordering of the arrows associated with $\underline{\sigma}$ to be the concatenation of the orderings associated with $\underline{\phi}_1$ and $\underline{\phi}_2$. Here we used the fact that $\underline{\phi} \in \mathcal{E}^1_{\underline{\tau}}(T_n)$, since when $\phi_1(\tau)\cap\phi_2(\tau) = \emptyset$, the union of the two decorated embeddings cannot be the image of a single decorated tree under a decorated embedding.

Note that when taking the union of the decorated embeddings we do not introduce any new arrows, so we must have $w(\underline{\sigma}) \leq 2w(\underline{\tau})$. Note also that, due to the nonlocality of

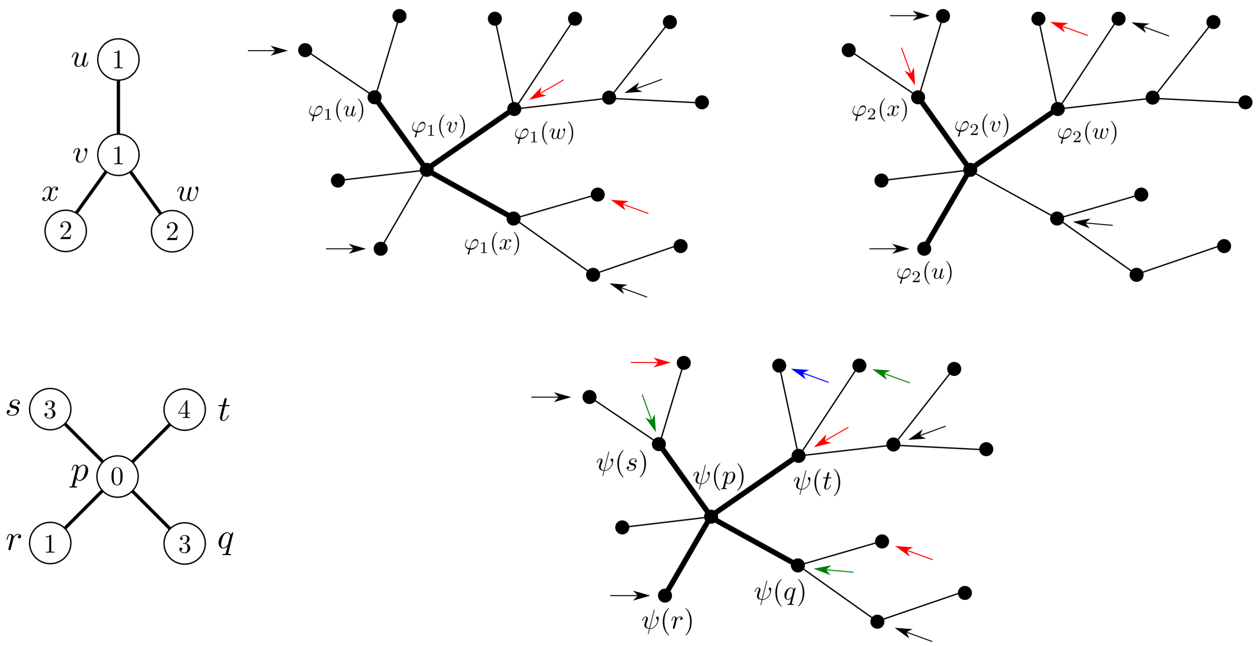

Figure 6.2: **A decorated map and an associated decorated embedding.** The top row depicts a decorated tree $\underline{\tau}$ and two decorated embeddings, $\underline{\phi}_1$ and $\underline{\phi}_2$, of it into a larger tree T. The bottom row depicts the associated decorated tree $\underline{\sigma}$, and the decorated embedding $\underline{\psi}$ of it into T.

the decorations, $\underline{\sigma}$ might have vertices having a zero label, see, e.g., Figure 6.2. However, importantly, the construction implies that all leaves of $\underline{\sigma}$ have positive labels, i.e., $\underline{\sigma} \in \mathcal{D} \setminus \mathcal{D}_0^*$.

Let $\mathcal{U}(\underline{\tau})$ denote the set of all decorated trees $\underline{\sigma}$ that can be obtained in this way. The cardinality of $\mathcal{U}(\underline{\tau})$ is bounded above by a constant depending only on $\underline{\tau}$, as we now argue. The number of ways that two copies of τ can be overlapped clearly depends only on τ. Once the union σ of the two copies of τ is fixed, only the arrows need to be associated with vertices of σ. There are at most $2w(\underline{\tau})$ arrows, and σ has at most $2|\underline{\tau}|$ vertices, so there are at most $(2|\underline{\tau}|)^{2w(\underline{\tau})}$ ways to associate arrows to vertices.

The function $\underline{\phi} \mapsto (\underline{\sigma}, \underline{\psi})$ is not necessarily one-to-one. However, there exists a constant $c(\underline{\tau})$ depending only on $\underline{\tau}$ such that any pair $(\underline{\sigma}, \underline{\psi})$ is associated with at most $c(\underline{\tau})$ decorated maps $\underline{\phi}$. To see this, note that given $(\underline{\sigma}, \underline{\psi})$, in order to recover $\underline{\phi}$, it is sufficient to know the following: (i) for every edge of $\psi(\sigma)$, whether it is a part of $\phi_1(\tau)$, a part of $\phi_2(\tau)$, or a part of both, (ii) for every arrow of $\underline{\psi}(\underline{\sigma})$, whether it is a part of $\underline{\phi}_1(\underline{\tau})$, a part of $\underline{\phi}_2(\underline{\tau})$, or a part of both, and (iii) the ordering of the arrows for $\underline{\phi}_1$ and $\underline{\phi}_2$. Since $|\underline{\sigma}| \leq 2|\underline{\tau}|$ and $w(\underline{\sigma}) \leq 2w(\underline{\tau})$, we can take $c(\underline{\tau}) = 3^{2|\underline{\tau}|+2w(\underline{\tau})}(w(\underline{\tau})!)^2$.

We have thus shown that

$$\left|\mathcal{E}^1_{\underline{\tau}}\left(T_n\right)\right| \leq c\left(\underline{\tau}\right) \sum_{\underline{\sigma} \in \mathcal{U}(\underline{\tau})} F_{\underline{\sigma}}\left(T_n\right).$$

For every $\underline{\sigma} \in \mathcal{U}\left(\underline{\tau}\right)$ we have that $\underline{\sigma} \in \mathcal{D} \setminus \mathcal{D}^*_0$, $|\underline{\sigma}| \geq |\underline{\tau}| \geq 2$, and $w\left(\underline{\sigma}\right) \leq 2w\left(\underline{\tau}\right)$, and so by Corollary 6.5 we have that $\mathbb{E}\left[F_{\underline{\sigma}}\left(T_n\right)\right] \precsim n^{2w(\underline{\tau})}$. Since the cardinality of $\mathcal{U}\left(\underline{\tau}\right)$ depends only on $\underline{\tau}$, this implies that $\mathbb{E}\left[\left|\mathcal{E}^1_{\underline{\tau}}\left(T_n\right)\right|\right] \precsim n^{2w(\underline{\tau})}$.

Now we turn to estimating $\left|\mathcal{E}^2_{\underline{\tau}}\left(T_n\right)\right|$. Again, we associate to each decorated map $\underline{\phi} \in \mathcal{E}^2_{\underline{\tau}}\left(T_n\right)$ a decorated tree $\underline{\sigma}$ and a decorated embedding $\underline{\psi}$ of it in T_n. This is done by first, just as before, taking the union of the images of the decorated embeddings $\underline{\phi}_1$ and $\underline{\phi}_2$, and if these share any arrows, identifying them (i.e., we take only a single copy). (Note that, since $\phi_1\left(\tau\right) \cap \phi_2\left(\tau\right) = \emptyset$, the decorated embeddings do not share any vertices or edges; however, due to the nonlocality of decorations they might share arrows.) We then take the union of this with the unique path in T_n that connects $\phi_1\left(\tau\right)$ and $\phi_2\left(\tau\right)$. The result of this is a tree in T_n, together with a set of at most $2w\left(\underline{\tau}\right)$ arrows associated with vertices of T_n; this is thus the image of a decorated tree $\underline{\sigma}$ under a decorated embedding $\underline{\psi}$, and this is how we define $\left(\underline{\sigma}, \underline{\psi}\right)$. See Figure 6.3 for an illustration.

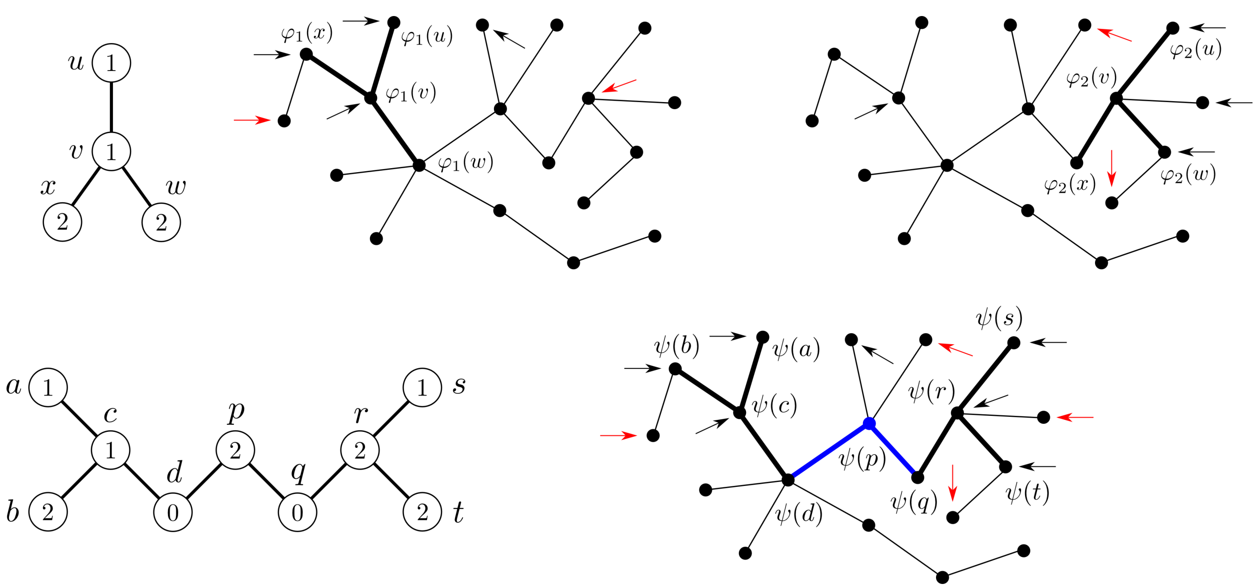

Figure 6.3: **Another decorated map and an associated decorated embedding.** The top row depicts a decorated tree $\underline{\tau}$ and two decorated embeddings, $\underline{\phi}_1$ and $\underline{\phi}_2$, of it into a larger tree T, where now $\phi_1\left(\tau\right) \cap \phi_2\left(\tau\right) = \emptyset$. The bottom row depicts the associated decorated tree $\underline{\sigma}$, and the decorated embedding $\underline{\psi}$ of it into T. The path connecting $\phi_1\left(\tau\right)$ and $\phi_2\left(\tau\right)$ is depicted in blue.

Again we have that any such $\underline{\sigma}$ must satisfy $w(\underline{\sigma}) \leq 2w(\underline{\tau})$ and $\underline{\sigma} \in \mathcal{D} \setminus \mathcal{D}_0^*$. The important difference now is that a priori we have no bound on $|\underline{\sigma}|$. This is where we use that $\operatorname{diam}(T_n) \leq K \log(n)$ with high probability.

Let $\mathcal{U}_2^{(n)}(\underline{\tau})$ denote the set of all decorated trees $\underline{\sigma}$ of diameter at most $K \log(n)$ that can be obtained in this way. The cardinality of $\mathcal{U}_2^{(n)}(\underline{\tau})$ cannot be bounded above by a constant depending only on $\underline{\tau}$, but it is at most polylogarithmic in n, as we now argue. There are at most $|\underline{\tau}|^2$ ways to choose which vertices of $\phi_1(\tau)$ and $\phi_2(\tau)$ are closest to each other, and the path connecting them has length at most $K \log(n)$. So the number of trees σ that can be obtained is at most $|\underline{\tau}|^2 K \log(n)$. Once the tree σ is fixed, only the arrows need to be associated with vertices of σ. There are at most $2w(\underline{\tau})$ arrows, and σ has at most $K \log(n) + 2|\underline{\tau}|$ vertices, which shows that

$$\left|\mathcal{U}_2^{(n)}(\underline{\tau})\right| \leq |\underline{\tau}|^2 K \log(n) \left(K \log(n) + 2|\underline{\tau}|\right)^{2w(\underline{\tau})} \gtrsim 1.$$

The function $\underline{\phi} \mapsto \left(\underline{\sigma}, \underline{\psi}\right)$ is not one-to-one. However, there exists a constant $c_2(\underline{\tau})$ depending only on $\underline{\tau}$ such that any pair $\left(\underline{\sigma}, \underline{\psi}\right)$ is associated with at most $c_2(\underline{\tau})$ decorated maps $\underline{\phi}$, as we now show. First, given $\left(\underline{\sigma}, \underline{\psi}\right)$, we know that $\phi_1(\tau)$ and $\phi_2(\tau)$ are at the two "ends" of $\psi(\sigma)$. The two "ends" of $\psi(\sigma)$ are well-defined: an edge e of $\psi(\sigma)$ is part of the path connecting $\phi_1(\tau)$ and $\phi_2(\tau)$ (and hence not part of an "end") if and only if there are at least $|\underline{\tau}|$ vertices on both sides of the cut defined by e. In order to recover $\underline{\phi}$, we also need to know for each arrow of $\underline{\psi}(\underline{\sigma})$, whether it is a part of $\underline{\phi}_1(\underline{\tau})$, a part of $\underline{\phi}_2(\underline{\tau})$, or a part of both. Finally, we need to know the ordering of the arrows for $\underline{\phi}_1$ and $\underline{\phi}_2$. Since $w(\underline{\sigma}) \leq 2w(\underline{\tau})$, we can thus take $c_2(\underline{\tau}) = 2 \times 3^{2w(\underline{\tau})} \left(w(\underline{\tau})!\right)^2$.

We have thus shown that

$$\left|\mathcal{E}_{\underline{\tau}}^2(T_n)\right| \mathbf{1}_{\{\operatorname{diam}(T_n) \leq K \log(n)\}} \leq c_2(\underline{\tau}) \sum_{\underline{\sigma} \in \mathcal{U}_2^{(n)}(\underline{\tau})} F_{\underline{\sigma}}(T_n).$$

For every $\underline{\sigma} \in \mathcal{U}_2^{(n)}(\underline{\tau})$ we have that $\underline{\sigma} \in \mathcal{D} \setminus \mathcal{D}_0^*$, $|\underline{\sigma}| \geq |\underline{\tau}| \geq 2$, and $w(\underline{\sigma}) \leq 2w(\underline{\tau})$, and so by Corollary 6.5 we have that $\mathbb{E}\left[F_{\underline{\sigma}}(T_n)\right] \gtrsim n^{2w(\underline{\tau})}$. Since we have $\left|\mathcal{U}_2^{(n)}(\underline{\tau})\right| \gtrsim 1$, we thus have

$$\mathbb{E}\left[\left|\mathcal{E}_{\underline{\tau}}^2(T_n)\right| \mathbf{1}_{\{\operatorname{diam}(T_n) \leq K \log(n)\}}\right] \gtrsim n^{2w(\underline{\tau})}.$$

This concludes the proof of (a).

For the proof of (b) we work conditionally on $\mathcal{F}_n$. As in the proof of Lemma 6.4, let u_{n+1} denote the vertex present in T_{n+1} but not in T_n, and let u_n denote its neighbor in T_{n+1}. Observe that $F_{\underline{\tau}}(T_{n+1}) - F_{\underline{\tau}}(T_n)$ is equal to the number of decorated embeddings of $\underline{\tau}$ in T_{n+1} that use the new vertex u_{n+1}. There are two ways that this may happen, and we call such decorated embeddings "type A" and "type B" accordingly (see Figure 6.4 for an illustration):

- **Type A.** The decorated embedding maps a vertex $v \in \tau$ to u_{n+1}. Since $\underline{\tau} \in \mathcal{D}_+$ and the arrows are mapped to different vertices, we must then have $\ell(v) = 1$, and the arrow pointing to v in $\underline{\tau}$ must be mapped to u_{n+1}.

- **Type B.** The decorated embedding maps τ in T_n, but there exists an arrow of $\underline{\tau}$ which it maps to u_{n+1}.

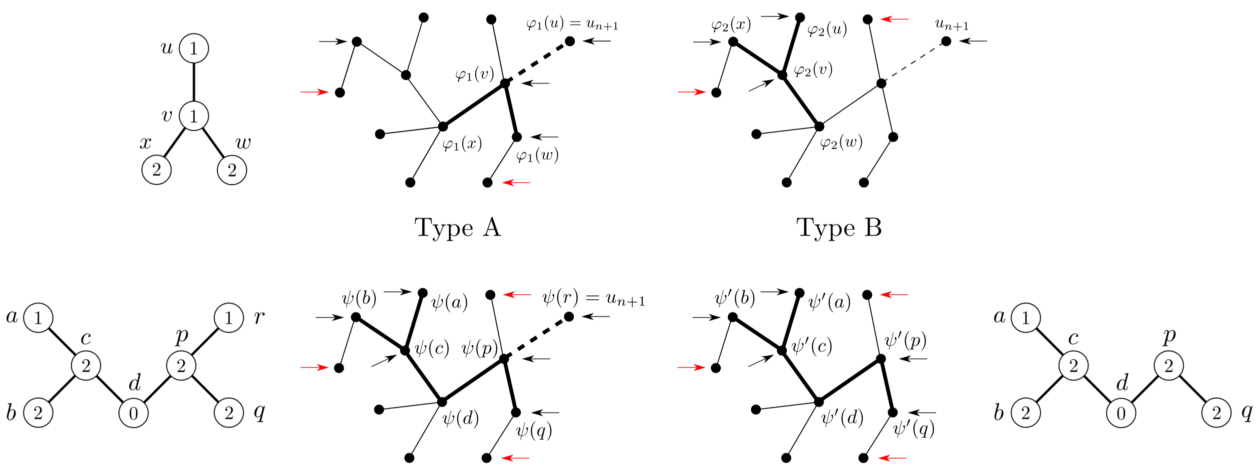

Figure 6.4: **Type A and type B decorated embeddings.** The top row depicts a decorated tree $\underline{\tau}$ and two decorated embeddings, $\underline{\phi}_1$ and $\underline{\phi}_2$, of it into a larger tree T_{n+1}. Here $\underline{\phi}_1$ is of type A, and $\underline{\phi}_2$ is of type B. In the bottom left is the associated decorated tree $\underline{\sigma}$, together with the decorated embedding $\underline{\psi}$ of it into T_{n+1}. In the bottom right is the pair $\left(\underline{\sigma}', \underline{\psi}'\right)$.

Consequently $\left(F_{\underline{\tau}}(T_{n+1}) - F_{\underline{\tau}}(T_n)\right)^2$ is equal to the number of decorated maps $\underline{\phi} = \underline{\phi}_1 \times \underline{\phi}_2$ such that $\underline{\phi}_1$ is either of type A or of type B, and the same holds for $\underline{\phi}_2$. We denote by $\tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1})$ the set of all such decorated maps where $\phi_1(\tau) \cap \phi_2(\tau) \neq \emptyset$, and let $\tilde{\mathcal{E}}^2_{\underline{\tau}}(T_{n+1})$ denote the set of all such decorated maps where $\phi_1(\tau) \cap \phi_2(\tau) = \emptyset$. Thus we have $\left(F_{\underline{\tau}}(T_{n+1}) - F_{\underline{\tau}}(T_n)\right)^2 = \left|\tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1})\right| + \left|\tilde{\mathcal{E}}^2_{\underline{\tau}}(T_{n+1})\right|$. Again, this partition is not necessary for the proof, but it helps the exposition.

We first estimate $\left|\tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1})\right|$. In the same way as in part (a), we associate to each decorated map $\underline{\phi} \in \tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1})$ a pair $\left(\underline{\sigma}, \underline{\psi}\right)$. Note that both $\underline{\phi}_1$ and $\underline{\phi}_2$ map an arrow to u_{n+1}, so $w(\underline{\sigma}) \leq 2w(\underline{\tau}) - 1$, and also there exists an arrow $\underline{a}^* \in \underline{\sigma}$ that is mapped to u_{n+1}, denoted by $\underline{\psi}(\underline{a}^*) = u_{n+1}$. We again have $\underline{\sigma} \in \mathcal{D} \setminus \mathcal{D}_0^*$. As before, the set $\tilde{\mathcal{U}}(\underline{\tau})$ of all decorated trees $\underline{\sigma}$ that can be obtained in this way has cardinality bounded above by a constant depending only on $\underline{\tau}$. Furthermore, there exists a constant $\tilde{c}(\underline{\tau})$ depending only on $\underline{\tau}$ such that any pair $\left(\underline{\sigma}, \underline{\psi}\right)$ is associated with at most $\tilde{c}(\underline{\tau})$ decorated maps $\underline{\phi}$.

We partition $\tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1})$ further into two parts. Let $\tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})$ denote the set of decorated maps $\underline{\phi} \in \tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1})$ such that at least one of $\underline{\phi}_1$ and $\underline{\phi}_2$ is of type A, and let $\tilde{\mathcal{E}}^{1,B}_{\underline{\tau}}(T_{n+1}) := \tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1}) \setminus \tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})$. That is, $\tilde{\mathcal{E}}^{1,B}_{\underline{\tau}}(T_{n+1})$ consists of those decorated maps $\underline{\phi} \in \tilde{\mathcal{E}}^1_{\underline{\tau}}(T_{n+1})$ such that both $\underline{\phi}_1$ and $\underline{\phi}_2$ is of type B.

We first estimate $\left|\tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})\right|$. We associate to each $\underline{\phi} \in \tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})$ a pair $\left(\underline{\sigma}, \underline{\psi}\right)$ as above. Let $v \in \sigma$ denote the vertex such that $\psi(v) = u_{n+1}$, and let $v' \in \sigma$ denote the vertex such that $\psi(v') = u_n$ (these vertices exist because $\underline{\phi} \in \tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})$). Define the decorated tree $\underline{\sigma}'$ from $\underline{\sigma}$ by removing the vertex v from $\underline{\sigma}$, as well as the arrow $\underline{a}^*$ pointing to it. Define also the decorated embedding $\underline{\psi}' : \underline{\sigma}' \to T_n$ to be equal to $\underline{\psi}$ on $\underline{\sigma}'$, i.e., $\underline{\psi}' = \underline{\psi}|_{\underline{\sigma}'}$; see Figure 6.4 for an illustration. We have that $w(\underline{\sigma}') = w(\underline{\sigma}) - 1 \leq 2w(\underline{\tau}) - 2$, it can be checked that $\underline{\sigma}' \in \mathcal{D} \setminus \mathcal{D}^*_0$, and we also have $\psi'(v') = u_n$. Let $\tilde{\mathcal{U}}'(\underline{\tau})$ denote the set of all decorated trees $\underline{\sigma}'$ that can be obtained in this way, and note that the cardinality of $\tilde{\mathcal{U}}'(\underline{\tau})$ is bounded from above by a constant depending only on $\underline{\tau}$. Since the map $\left(\underline{\sigma}, \underline{\psi}\right) \mapsto \left(\underline{\sigma}', \underline{\psi}', v'\right)$ is one-to-one, we have obtained that

$$\left|\tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})\right| \leq \sum_{\underline{\sigma}' \in \tilde{\mathcal{U}}'(\underline{\tau})} \sum_{v' \in \sigma'} \sum_{\underline{\psi}' : \underline{\sigma}' \to T_n} \tilde{c}(\underline{\tau}) \mathbf{1}_{\{\psi'(v') = u_n\}}.$$

Since u_n is uniform, we obtain

$$\mathbb{E}\left[\left|\tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})\right|\right] = \mathbb{E}\left[\mathbb{E}\left[\left|\tilde{\mathcal{E}}^{1,A}_{\underline{\tau}}(T_{n+1})\right| \,\middle|\, \mathcal{F}_n\right]\right] \leq \sum_{\underline{\sigma}' \in \tilde{\mathcal{U}}'(\underline{\tau})} \sum_{v' \in \sigma'} \frac{\tilde{c}(\underline{\tau})}{n} \mathbb{E}\left[F_{\underline{\sigma}'}(T_n)\right] \lesssim n^{2w(\underline{\tau})-3},$$

where in the last inequality we used that for every $\underline{\sigma}' \in \tilde{\mathcal{U}}'(\underline{\tau})$ we have $w(\underline{\sigma}') \leq 2w(\underline{\tau}) - 2$ and $\underline{\sigma}' \in \mathcal{D} \setminus \mathcal{D}^*_0$, and so by Corollary 6.5 we have that $\mathbb{E}\left[F_{\underline{\sigma}'}(T_n)\right] \lesssim n^{2w(\underline{\tau})-2}$.

We now turn to estimating $\left|\tilde{\mathcal{E}}^{1,B}_{\underline{\tau}}(T_{n+1})\right|$. Let

$$v^* := \underset{v \in \sigma}{\arg\min}\, \mathrm{dist}_{T_{n+1}}(\psi(v), u_{n+1}) = \underset{v \in \sigma}{\arg\min}\, \mathrm{dist}_{T_n}(\psi(v), u_n),$$

where dist_G denotes graph distance in a graph G. Note that the arrow $\underline{a}^* \in \underline{\sigma}$ is associated with v^* in $\underline{\sigma}$. Define the decorated map $\underline{\sigma}^*$ from $\underline{\sigma}$ by removing the arrow $\underline{a}^*$ from $\underline{\sigma}$. We have that $w(\underline{\sigma}^*) = w(\underline{\sigma}) - 1 \leq 2w(\underline{\tau}) - 2$. Furthermore, either $\underline{\sigma}^* \in \mathcal{D} \setminus \mathcal{D}^*_0$ or v^* is the only leaf of $\underline{\sigma}$ that has label zero. Let $\tilde{\mathcal{U}}^*(\underline{\tau})$ denote the set of all decorated trees $\underline{\sigma}^*$ that can be obtained in this way, and note that the cardinality of $\tilde{\mathcal{U}}^*(\underline{\tau})$ is bounded from above by a constant depending only on $\underline{\tau}$. Define also the decorated embedding $\underline{\psi}^* : \underline{\sigma}^* \to T_n$ to be equal to $\underline{\psi}$ on $\underline{\sigma}^*$, i.e., $\underline{\psi}^* = \underline{\psi}|_{\underline{\sigma}^*}$. Define furthermore z^* to be the neighbor of $\underline{\psi}(\underline{a}^*)$ in T_{n+1}; we thus have $z^* = u_n$. Due to the ordering of the arrows, the map $\left(\underline{\sigma}, \underline{\psi}\right) \mapsto \left(\underline{\sigma}^*, \underline{\psi}^*, z^*\right)$ is not necessarily one-to-one, but any triple $\left(\underline{\sigma}^*, \underline{\psi}^*, z^*\right)$ is associated with at most $w(\underline{\tau})$ pairs

$\left(\underline{\sigma}, \underline{\psi}\right)$. Thus, defining $\tilde{c}'\left(\underline{\tau}\right) := \tilde{c}\left(\underline{\tau}\right) w\left(\underline{\tau}\right)$, we have that

$$\begin{aligned}
&\left|\tilde{\mathcal{E}}^{1,B}_{\underline{\tau}}\left(T_{n+1}\right)\right| \\
&\quad \leq \sum_{\underline{\sigma}^* \in \tilde{\mathcal{U}}^*\left(\underline{\tau}\right)} \sum_{\underline{\psi}^* : \underline{\sigma}^* \to T_n} \sum_{z^* \in T_n} \tilde{c}'\left(\underline{\tau}\right) \mathbf{1}_{\{z^* = u_n\}} \mathbf{1}_{\left\{\underline{\sigma}^* \in \mathcal{D} \setminus \mathcal{D}_0^*\right\} \cup \left\{\underline{\sigma}^* \in \mathcal{D}_0^*, \arg\min_{v \in \sigma^*} \mathrm{dist}_{T_n}(\psi(v), z^*) \in L_0(\underline{\sigma}^*)\right\}} \\
&\quad \leq \tilde{c}'\left(\underline{\tau}\right) \sum_{\underline{\sigma}^* \in \tilde{\mathcal{U}}^*\left(\underline{\tau}\right)} \sum_{\underline{\psi}^* : \underline{\sigma}^* \to T_n} \mathbf{1}_{\left\{\underline{\sigma}^* \in \mathcal{D} \setminus \mathcal{D}_0^*\right\} \cup \left\{\underline{\sigma}^* \in \mathcal{D}_0^*, \arg\min_{v \in \sigma^*} \mathrm{dist}_{T_n}(\psi(v), u_n) \in L_0(\underline{\sigma}^*)\right\}}.
\end{aligned}$$

Now if $\underline{\sigma}^* \in \tilde{\mathcal{U}}^*\left(\underline{\tau}\right) \cap \left(\mathcal{D} \setminus \mathcal{D}_0^*\right)$, then the sum over embeddings $\underline{\psi}^* : \underline{\sigma}^* \to T_n$ becomes $F_{\underline{\sigma}^*}\left(T_n\right)$, and by Corollary 6.5 we have that $\mathbb{E}\left[F_{\underline{\sigma}^*}\left(T_n\right)\right] \lesssim n^{2w(\underline{\tau})-2}$. If $\underline{\sigma}^* \in \tilde{\mathcal{U}}^*\left(\underline{\tau}\right) \cap \mathcal{D}_0^*$, then, as mentioned above, $L_0\left(\underline{\sigma}^*\right) = \{v^*\}$, and we have

$$\mathbb{P}\left(\arg\min_{v \in \sigma^*} \mathrm{dist}_{T_n}\left(\psi\left(v\right), u_n\right) = v^* \,\middle|\, \mathcal{F}_n \right) = \frac{f_{\psi(v^*)}\left(T_n\right)}{n}.$$

So by summing over $\underline{\psi}^* : \underline{\sigma}^* \to T_n$, if $\underline{\sigma}^* \in \tilde{\mathcal{U}}^*\left(\underline{\tau}\right) \cap \mathcal{D}_0^*$, then

$$\mathbb{E}\left[\sum_{\underline{\psi}^* : \underline{\sigma}^* \to T_n} \mathbf{1}_{\left\{\underline{\sigma}^* \in \mathcal{D}_0^*, \arg\min_{v \in \sigma^*} \mathrm{dist}_{T_n}(\psi(v), u_n) \in L_0(\underline{\sigma}^*)\right\}} \,\middle|\, \mathcal{F}_n \right] = \frac{1}{n} F_{\underline{\sigma}}\left(T_n\right).$$

Since $\underline{\sigma} \in \mathcal{D} \setminus \mathcal{D}_0^*$ and $w\left(\underline{\sigma}\right) \leq 2w\left(\underline{\tau}\right) - 1$, by Corollary 6.5 we have that $\mathbb{E}\left[F_{\underline{\sigma}}\left(T_n\right)\right] \lesssim n^{2w(\underline{\tau})-1}$ and thus $\mathbb{E}\left[n^{-1} F_{\underline{\sigma}}\left(T_n\right)\right] \lesssim n^{2w(\underline{\tau})-2}$. Putting everything together we thus obtain that

$$\mathbb{E}\left[\left|\tilde{\mathcal{E}}^{1,B}_{\underline{\tau}}\left(T_{n+1}\right)\right|\right] \lesssim n^{2w(\underline{\tau})-2}.$$

To estimate $\left|\tilde{\mathcal{E}}^{2}_{\underline{\tau}}\left(T_{n+1}\right)\right|$ we can do the same thing as in part (a), and we obtain the same bound as for $\left|\tilde{\mathcal{E}}^{1}_{\underline{\tau}}\left(T_{n+1}\right)\right|$ up to polylogarithmic factors in n. We omit the details. This concludes the proof of part (b). $\square$

6.3.5 Constructing the martingales

We now construct the martingales of Proposition 6.3 with the help of the recurrence relation of Lemma 6.4. In order to show that these martingales are bounded in L^2, we use the moment estimates of Section 6.3.4.

Proof of Proposition 6.3. Fix a seed tree S with $|S| = n_0 \geq 2$. For a decorated tree $\underline{\tau} \in \mathcal{D}_+$ and $n \geq 2$, define

$$\beta_n\left(\underline{\tau}\right) := \prod_{j=2}^{n-1} \left(1 + \frac{w\left(\underline{\tau}\right)}{j}\right)^{-1}, \qquad \text{when } |\underline{\tau}| \geq 2;$$

$$\text{and} \qquad \beta_n\left(\underline{\tau}\right) := \left(n \times [n]_{w(\underline{\tau})}\right)^{-1}, \qquad \text{when } |\underline{\tau}| = 1.$$

Note that when $|\underline{\tau}| \geq 2$, we have $n^{-w(\underline{\tau})} \lesssim \beta_n(\underline{\tau}) \lesssim n^{-w(\underline{\tau})}$.

We now construct, by induction on the order $\preccurlyeq$ on decorated trees, coefficients

$$\{a_n(\underline{\tau}, \underline{\tau}') : \underline{\tau}, \underline{\tau}' \in \mathcal{D}_+, \underline{\tau}' \prec \underline{\tau}, n \geq n_0\}$$

such that

$$a_n(\underline{\tau}, \underline{\tau}') \lesssim 1, \qquad \Delta_n a(\underline{\tau}, \underline{\tau}') \lesssim 1/n, \tag{6.17}$$

and

$$M_{\underline{\tau}}^{(S)}(n) = \beta_n(\underline{\tau}) \left(F_{\underline{\tau}}(\mathrm{UA}(n,S)) - \sum_{\underline{\tau}' \in \mathcal{D}_+ : \underline{\tau}' \prec \underline{\tau}} a_n(\underline{\tau}, \underline{\tau}') F_{\underline{\tau}'}(\mathrm{UA}(n,S)) \right) \tag{6.18}$$

is a martingale. Importantly, we shall see that the coefficients $a_n(\underline{\tau}, \underline{\tau}')$ do not depend on S. To simplify notation, in the following we omit dependence on S and write $M_{\underline{\tau}}(n)$ for $M_{\underline{\tau}}^{(S)}(n)$. Also, as before, we write T_n for $\mathrm{UA}(n,S)$.

First, when $|\underline{\tau}| = 1$, we have $M_{\underline{\tau}}(n) = \beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) = 1$, which is a martingale. Now fix $\underline{\tau} \in \mathcal{D}_+$ with $|\underline{\tau}| \geq 2$. Assume that the coefficients $a_n(\underline{\sigma}, \underline{\sigma}')$ have been constructed for every $\underline{\sigma}, \underline{\sigma}' \in \mathcal{D}_+$ such that $\underline{\sigma}' \prec \underline{\sigma} \prec \underline{\tau}$ and every $n \geq n_0$, and that they have the desired properties. We first claim that there exist constants $\{b_n(\underline{\sigma}, \underline{\sigma}') : \underline{\sigma}' \prec \underline{\sigma} \prec \underline{\tau}, n \geq n_0\}$ such that $b_n(\underline{\sigma}, \underline{\sigma}') \lesssim 1$ and

$$F_{\underline{\sigma}}(T_n) = \frac{1}{\beta_n(\underline{\sigma})} M_{\underline{\sigma}}(n) + \sum_{\underline{\sigma}' \in \mathcal{D}_+ : \underline{\sigma}' \prec \underline{\sigma}} \frac{b_n(\underline{\sigma}, \underline{\sigma}')}{\beta_n(\underline{\sigma}')} M_{\underline{\sigma}'}(n) \tag{6.19}$$

for $n \geq n_0$. To see this, define the matrix $A_n = (A_n(\underline{\sigma}, \underline{\sigma}'))_{\underline{\sigma}, \underline{\sigma}' \prec \underline{\tau}}$ by $A_n(\underline{\sigma}, \underline{\sigma}') = -a_n(\underline{\sigma}, \underline{\sigma}')$ if $\underline{\sigma}' \prec \underline{\sigma}$, $A_n(\underline{\sigma}, \underline{\sigma}') = 1$ if $\underline{\sigma} = \underline{\sigma}'$, and $A_n(\underline{\sigma}, \underline{\sigma}') = 0$ otherwise. Then, using (6.18), we have for every $n \geq n_0$ the following equality of vectors indexed by $\underline{\sigma} \in \mathcal{D}_+$ such that $\underline{\sigma} \prec \underline{\tau}$:

$$\left(\frac{1}{\beta_n(\underline{\sigma})} M_{\sigma}(n) \right)_{\underline{\sigma} \prec \underline{\tau}} = A_n \cdot (F_{\underline{\sigma}}(T_n))_{\underline{\sigma} \prec \underline{\tau}}. \tag{6.20}$$

We can write $\{\underline{\sigma} \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}\} = \{\underline{\sigma}_1, \ldots, \underline{\sigma}_K\}$ in such a way that $\underline{\sigma}_i \prec \underline{\sigma}_j$ implies $i < j$. With this convention, A_n is a lower triangular matrix with all diagonal entries equal to 1 and all entries satisfying $A_n(\underline{\sigma}, \underline{\sigma}') \lesssim 1$. Therefore A_n is invertible, and its inverse also satisfies these properties. That is, if we write $A_n^{-1} = (b_n(\underline{\sigma}, \underline{\sigma}'))_{\underline{\sigma}, \underline{\sigma}' \prec \underline{\tau}}$, then A_n^{-1} is a lower triangular matrix that satisfies $b_n(\underline{\sigma}, \underline{\sigma}) = 1$ for all $\underline{\sigma} \prec \underline{\tau}$, and $b_n(\underline{\sigma}, \underline{\sigma}') \lesssim 1$ for all $\underline{\sigma}, \underline{\sigma}' \prec \underline{\tau}$. So (6.19) follows directly from (6.20).

Note that we can write equation (6.13) of Lemma 6.4 more compactly as follows:

$$\begin{aligned} \mathbb{E}\left[F_{\underline{\tau}}(\mathrm{UA}(n+1,S)) \,|\, \mathcal{F}_n\right] = &\left(1 + \frac{w(\underline{\tau})}{n}\right) F_{\underline{\tau}}(\mathrm{UA}(n,S)) \\ &+ \frac{1}{n} \sum_{\underline{\tau}' \in \mathcal{D} : \underline{\tau}' \prec \underline{\tau}} c(\underline{\tau}, \underline{\tau}') F_{\underline{\tau}'}(\mathrm{UA}(n,S)), \end{aligned} \tag{6.21}$$

for appropriately defined constants $\{c(\underline{\tau},\underline{\tau}') : \underline{\tau},\underline{\tau}' \in \mathcal{D}, \underline{\tau}' \prec \underline{\tau}\}$, and note that since $\underline{\tau} \in \mathcal{D}_+$, we have $c(\underline{\tau},\underline{\tau}') = 0$ if $\underline{\tau}' \notin \mathcal{D}_+$. Therefore, using (6.21) and (6.19), together with the identities $\beta_{n+1}(\underline{\tau})(1 + w(\underline{\tau})/n) = \beta_n(\underline{\tau})$ and $\beta_{n+1}(\underline{\tau}) n^{-1} = \beta_n(\underline{\tau})(n + w(\underline{\tau}))^{-1}$, we have for $n \geq n_0$ that

$$\begin{aligned}
&\mathbb{E}\left[\beta_{n+1}(\underline{\tau}) F_{\underline{\tau}}(T_{n+1}) \,|\, \mathcal{F}_n\right] \\
&\qquad = \beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) + \frac{\beta_n(\underline{\tau})}{n + w(\underline{\tau})} \sum_{\underline{\tau}' \in \mathcal{D}_+ : \underline{\tau}' \prec \underline{\tau}} c(\underline{\tau},\underline{\tau}') F_{\underline{\tau}'}(T_n) \\
&\qquad = \beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) + \frac{1}{n + w(\underline{\tau})} \times \\
&\qquad\qquad \times \sum_{\underline{\sigma} \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}} \left(c(\underline{\tau},\underline{\sigma}) + \sum_{\underline{\tau}' \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}' \prec \underline{\tau}} c(\underline{\tau},\underline{\tau}') b_n(\underline{\tau}',\underline{\sigma}) \right) \frac{\beta_n(\underline{\tau})}{\beta_n(\underline{\sigma})} M_{\underline{\sigma}}(n).
\end{aligned}$$

For $n \geq n_0$ define

$$\overline{a}_n(\underline{\tau},\underline{\sigma}) = \sum_{j=n_0}^{n-1} \frac{1}{j + w(\underline{\tau})} \left(c(\underline{\tau},\underline{\sigma}) + \sum_{\underline{\tau}' \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}' \prec \underline{\tau}} c(\underline{\tau},\underline{\tau}') b_j(\underline{\tau}',\underline{\sigma}) \right) \frac{\beta_j(\underline{\tau})}{\beta_j(\underline{\sigma})}.$$

We thus have

$$\mathbb{E}\left[\beta_{n+1}(\underline{\tau}) F_{\underline{\tau}}(T_{n+1}) \,|\, \mathcal{F}_n\right] = \beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) + \sum_{\underline{\sigma} \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}} \left(\overline{a}_{n+1}(\underline{\tau},\underline{\sigma}) - \overline{a}_n(\underline{\tau},\underline{\sigma})\right) M_{\underline{\sigma}}(n).$$

By our induction hypothesis, $\{M_{\underline{\sigma}}(n)\}_{n \geq n_0}$ is an $(\mathcal{F}_n)$-martingale for every $\underline{\sigma} \prec \underline{\tau}$, and consequently

$$\beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) - \sum_{\underline{\sigma} \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}} \overline{a}_n(\underline{\tau},\underline{\sigma}) M_{\underline{\sigma}}(n)$$

is also an $(\mathcal{F}_n)$-martingale. By (6.18) we have

$$\begin{aligned}
&\beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) - \sum_{\underline{\sigma} \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}} \overline{a}_n(\underline{\tau},\underline{\sigma}) M_{\underline{\sigma}}(n) \\
&\quad = \beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) - \sum_{\underline{\sigma} \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}} \overline{a}_n(\underline{\tau},\underline{\sigma}) \beta_n(\underline{\sigma}) \left(F_{\underline{\sigma}}(T_n) - \sum_{\underline{\tau}' \in \mathcal{D}_+ : \underline{\tau}' \prec \underline{\sigma}} a_n(\underline{\sigma},\underline{\tau}') F_{\underline{\tau}'}(T_n) \right) \\
&\quad = \beta_n(\underline{\tau}) F_{\underline{\tau}}(T_n) \\
&\qquad - \beta_n(\underline{\tau}) \sum_{\underline{\sigma} \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}} \left[\overline{a}_n(\underline{\tau},\underline{\sigma}) \frac{\beta_n(\underline{\sigma})}{\beta_n(\underline{\tau})} - \sum_{\underline{\tau}' \in \mathcal{D}_+ : \underline{\sigma} \prec \underline{\tau}' \prec \underline{\tau}} \overline{a}_n(\underline{\tau},\underline{\tau}') a_n(\underline{\tau}',\underline{\sigma}) \frac{\beta_n(\underline{\tau}')}{\beta_n(\underline{\tau})} \right] F_{\underline{\sigma}}(T_n).
\end{aligned}$$

So if we set

$$a_n(\underline{\tau},\underline{\sigma}) := \overline{a}_n(\underline{\tau},\underline{\sigma})\frac{\beta_n(\underline{\sigma})}{\beta_n(\underline{\tau})} - \sum_{\underline{\tau}'\in\mathcal{D}_+:\underline{\sigma}\prec\underline{\tau}'\prec\underline{\tau}} \overline{a}_n(\underline{\tau},\underline{\tau}')\, a_n(\underline{\tau}',\underline{\sigma})\frac{\beta_n(\underline{\tau}')}{\beta_n(\underline{\tau})},$$

then it is clear that $\{M_{\underline{\tau}}(n)\}_{n\geq n_0}$ defined as in (6.18) is a martingale.

Now let us establish that the coefficients are of the correct order, i.e., let us show (6.17). First note that $(n+w(\underline{\tau}))^{-1} \lessapprox 1/n$, and that when $|\underline{\tau}|\geq 2$, $\beta_n(\underline{\tau})\, n^{w(\underline{\tau})}$ has a positive and finite limit as $n\to\infty$. Therefore a simple computation shows that for $\underline{\sigma},\underline{\sigma}'\in\mathcal{D}_+$ with $|\underline{\sigma}|,|\underline{\sigma}'|\geq 2$, we have

$$\frac{\beta_n(\underline{\sigma})}{\beta_n(\underline{\sigma}')} \lessapprox n^{w(\underline{\sigma}')-w(\underline{\sigma})} \qquad \text{and} \qquad \Delta_n\frac{\beta(\underline{\sigma})}{\beta(\underline{\sigma}')} \lessapprox n^{w(\underline{\sigma}')-w(\underline{\sigma})-1}.$$

Furthermore, by the induction hypothesis we have that $b_n(\underline{\sigma},\underline{\sigma}') \lessapprox 1$ for every $\underline{\sigma},\underline{\sigma}'\prec\underline{\tau}$. From the definition of $\overline{a}_n(\underline{\tau},\underline{\sigma})$ we then immediately get that $\Delta_n\overline{a}(\underline{\tau},\underline{\sigma}) \lessapprox n^{w(\underline{\sigma})-w(\underline{\tau})-1}$, and consequently also $\overline{a}_n(\underline{\tau},\underline{\sigma}) \lessapprox n^{w(\underline{\sigma})-w(\underline{\tau})}$, for every $\underline{\sigma}\in\mathcal{D}_+$ such that $\underline{\sigma}\prec\underline{\tau}$ and $|\underline{\sigma}|\geq 2$. So for every $\underline{\sigma}\in\mathcal{D}_+$ such that $\underline{\sigma}\prec\underline{\tau}$ and $|\underline{\sigma}|\geq 2$, we have that

$$\overline{a}_n(\underline{\tau},\underline{\sigma})\frac{\beta_n(\underline{\sigma})}{\beta_n(\underline{\tau})} \lessapprox 1 \qquad \text{and} \qquad \Delta_n\left(\overline{a}(\underline{\tau},\underline{\sigma})\frac{\beta(\underline{\sigma})}{\beta(\underline{\tau})}\right) \lessapprox \frac{1}{n}. \tag{6.22}$$

One can easily check that (6.22) holds also when $|\underline{\sigma}|=1$. Now combining all of these estimates with the definition of $a_n(\underline{\tau},\underline{\sigma})$, we get that (6.17) holds. This completes the induction.

Finally, what remains to show is that the martingales $M_{\underline{\tau}}$ are bounded in L^2. Since $M_{\underline{\tau}}$ is a martingale, its increments are orthogonal in L^2, and so

$$\mathbb{E}\left[M_{\underline{\tau}}(n)^2\right] = \sum_{j=n_0}^{n-1}\mathbb{E}\left[(M_{\underline{\tau}}(j+1)-M_{\underline{\tau}}(j))^2\right] + \mathbb{E}\left[M_{\underline{\tau}}(n_0)^2\right].$$

Clearly $\mathbb{E}\left[M_{\underline{\tau}}(n_0)^2\right]<\infty$ and so it suffices to show that

$$\sum_{n=n_0}^{\infty}\mathbb{E}\left[(M_{\underline{\tau}}(n+1)-M_{\underline{\tau}}(n))^2\right]<\infty.$$

Recalling the definition of $M_{\underline{\tau}}$ from (6.18) we have

$$\mathbb{E}\left[(\Delta_n(M_{\underline{\tau}}))^2\right] = \mathbb{E}\left[\left(\Delta_n(\beta_\cdot(\underline{\tau})\,F_{\underline{\tau}}(T_\cdot)) - \sum_{\underline{\tau}\in\mathcal{D}_+:\underline{\tau}'\prec\underline{\tau}}\Delta_n(\beta_\cdot(\underline{\tau})\,a_\cdot(\underline{\tau},\underline{\tau}')\,F_{\underline{\tau}'}(T_\cdot))\right)^2\right],$$

where the dots in the subscripts denote dependence on n, on which the difference operator Δ_n acts. By the Cauchy-Schwarz inequality, there exists a positive and finite constant c that depends only on $\underline{\tau}$ such that for every $n \geq n_0$, the quantity $c \times \mathbb{E}\left[(\Delta_n (M_{\underline{\tau}}))^2\right]$ is bounded from above by

$$\mathbb{E}\left[(\Delta_n (\beta_{\cdot}(\underline{\tau}) F_{\underline{\tau}}(T_{\cdot})))^2\right] + \sum_{\underline{\tau} \in \mathcal{D}_+ : \underline{\tau}' \prec \underline{\tau}} \mathbb{E}\left[(\Delta_n (\beta_{\cdot}(\underline{\tau}) a_{\cdot}(\underline{\tau}, \underline{\tau}') F_{\underline{\tau}'}(T_{\cdot})))^2\right]. \tag{6.23}$$

Since

$$\Delta_n (\beta_{\cdot}(\underline{\tau}) F_{\underline{\tau}}(T_{\cdot})) = \beta_{n+1}(\underline{\tau}) \Delta_n (F_{\underline{\tau}}(T_{\cdot})) + (\Delta_n (\beta_{\cdot}(\underline{\tau}))) F_{\underline{\tau}}(T_n),$$

we have that

$$\mathbb{E}\left[(\Delta_n (\beta_{\cdot}(\underline{\tau}) F_{\underline{\tau}}(T_{\cdot})))^2\right] \leq 2 (\beta_{n+1}(\underline{\tau}))^2 \mathbb{E}\left[(\Delta_n (F_{\underline{\tau}}(T_{\cdot})))^2\right] + 2 (\Delta_n (\beta_{\cdot}(\underline{\tau})))^2 \mathbb{E}\left[F_{\underline{\tau}}(T_n)^2\right].$$

We have seen that $(\beta_{n+1}(\underline{\tau}))^2 \lesssim n^{-2w(\underline{\tau})}$ and $(\Delta_n (\beta_{\cdot}(\underline{\tau})))^2 \lesssim n^{-2w(\underline{\tau})-2}$, and by Lemma 6.6 we have that $\mathbb{E}\left[F_{\underline{\tau}}(T_n)^2\right] \lesssim n^{2w(\underline{\tau})}$ and $\mathbb{E}\left[(\Delta_n (F_{\underline{\tau}}(T_{\cdot})))^2\right] \lesssim n^{2w(\underline{\tau})-2}$. Putting these together we thus have that $\mathbb{E}\left[(\Delta_n (\beta_{\cdot}(\underline{\tau}) F_{\underline{\tau}}(T_{\cdot})))^2\right] \lesssim n^{-2}$. For the other terms in (6.23) we similarly have

$$\begin{aligned}\mathbb{E}&\left[(\Delta_n (\beta_{\cdot}(\underline{\tau}) a_{\cdot}(\underline{\tau}, \underline{\tau}') F_{\underline{\tau}'}(T_{\cdot})))^2\right] \\ &\leq 2 (a_{n+1}(\underline{\tau}, \underline{\tau}'))^2 \mathbb{E}\left[(\Delta_n (\beta_{\cdot}(\underline{\tau}) F_{\underline{\tau}'}(T_{\cdot})))^2\right] \\ &\quad + 2 (\Delta_n (a_{\cdot}(\underline{\tau}, \underline{\tau}')))^2 \mathbb{E}\left[(\beta_n (\underline{\tau}) F_{\underline{\tau}'}(T_n))^2\right].\end{aligned}$$

We have seen that $(a_{n+1}(\underline{\tau}, \underline{\tau}'))^2 \lesssim 1$ and $(\Delta_n (a_{\cdot}(\underline{\tau}, \underline{\tau}')))^2 \lesssim n^{-2}$. Furthermore, by Lemma 6.6 we have that $\mathbb{E}\left[(\beta_n (\underline{\tau}) F_{\underline{\tau}'}(T_n))^2\right] \lesssim n^{2w(\underline{\tau}')-2w(\underline{\tau})} \leq 1$, and similarly to the computation above we have that $\mathbb{E}\left[(\Delta_n (\beta_{\cdot}(\underline{\tau}) F_{\underline{\tau}'}(T_{\cdot})))^2\right] \lesssim n^{2w(\underline{\tau}')-2w(\underline{\tau})-2} \leq n^{-2}$. Putting everything together we get that

$$\mathbb{E}\left[(M_{\underline{\tau}}(n+1) - M_{\underline{\tau}}(n))^2\right] \lesssim n^{-2},$$

which is summable, so M_τ is indeed bounded in L^2. □

6.4 Comparison to the work of Curien et al.

As discussed in Section 4.3.2, the key difference in our proof for uniform attachment compared to the proof of [21] for preferential attachment is the underlying family of statistics. For preferential attachment these are based on the degrees of the nodes, whereas for uniform attachment they are based on partition sizes when embedding a given tree, i.e., they are based on global balancedness properties of the tree.

The statistics $F_{\underline{\tau}}(T)$ are defined in this specific way in order to make the analysis simpler. In particular, it is useful that $F_{\underline{\tau}}(T)$ has a combinatorial interpretation as the number of decorated embeddings of $\underline{\tau}$ in T, similarly to the statistics of [21]. However, the notion of a decorated embedding is different in the two settings. In [21], arrows associated with the decorated tree $\underline{\tau}$ are mapped by $\underline{\phi}$ to corners around the vertices of $\phi(\tau)$, or in other words, the decorations are local. In contrast, in the notion of a decorated embedding as defined in this chapter, arrows associated with a decorated tree $\underline{\tau}$ can be mapped to any vertex in the graph T, or in other words, the decorations are nonlocal/global.

While the general structure of our proof is identical to that of [21], this local vs. global difference in the underlying statistics manifests itself in the details. In particular, the main challenge is the second moment estimate provided in Lemma 6.6. Here, we associate to each decorated map $\underline{\phi} = \underline{\phi}_1 \times \underline{\phi}_2$ a decorated tree $\underline{\sigma}$ and a decorated embedding $\underline{\psi}$ of it in UA (n, S). In the case of preferential attachment, the decorated tree $\underline{\sigma}$ necessarily has all labels positive, due to the decorations being local. However, in the case of uniform attachment, it might happen that a vertex of $\underline{\sigma}$ has a zero label, due to the global nature of decorations. This is the reason why we need to deal with decorated trees having zero labels, in contrast with the preferential attachment model, where it suffices to consider decorated trees with positive labels. The recurrence relation and the subsequent moment estimates show that there is a subtlety in dealing with decorated trees having zero labels, as it matters whether the vertices with label zero are leaves or not.

Finally, in our proof of the second moment estimate we also use the fact that the diameter of UA (n, S) is on the order of $\log(n)$ with high probability (see Lemma 6.8). This is again due to the global nature of decorations, and such an estimate is not necessary in the case of preferential attachment.

6.5 Technical results used in the proofs

6.5.1 Facts about the beta-binomial distribution

We prove here Facts 6.1 and 6.2 stated in Section 6.2.1. Let $\{M_k\}$, where $k = \alpha+\beta, \alpha+\beta+1, \ldots$, be the martingale associated with the standard Pólya urn process with starting state (α, β). In other words, the martingale M_k is defined by $M_{\alpha+\beta} = \frac{\alpha}{\alpha+\beta}$ and

$$(k+1)M_{k+1} = \begin{cases} kM_k + 1 & \text{with probability } M_k \\ kM_k & \text{with probability } 1 - M_k \end{cases}$$

independently for different values of k. Note that for $n \geq \alpha + \beta$, $nM_n \stackrel{d}{=} B_{\alpha,\beta,n-\alpha-\beta}$, so all results for the martingale M_n transfer to results for $B_{\alpha,\beta,n-\alpha-\beta}$. Define $M_\infty = \lim_{k\to\infty} M_k$, and note that this limit exists almost surely by the martingale convergence theorem. It is a well-known fact about Pólya urns that M_∞ has a beta distribution with parameters α and β, i.e., the density of M_∞ with respect to the Lebesgue measure is

$$h(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}\mathbf{1}_{\{x\in[0,1]\}}.$$

By the formula for the moments of M_∞ (see, e.g., [43, Chapter 21]), we have

$$\mathbb{E}\left[M_\infty^p\right] = \prod_{j=0}^{p-1} \frac{\alpha+j}{\alpha+\beta+j} \leq \left(\frac{\alpha+p}{\alpha+\beta}\right)^p \leq (p+1)^p \left(\frac{\alpha}{\alpha+\beta}\right)^p, \quad \forall p \in \mathbb{N}.$$

Moreover, since M_n is a martingale, M_n^p is a submartingale for all $p \geq 1$, and thus $\mathbb{E}\left[M_n^p\right] \leq \mathbb{E}\left[M_\infty^p\right]$. So we have that

$$\mathbb{E}\left[(nM_n)^p\right] \leq n^p (p+1)^p \left(\frac{\alpha}{\alpha+\beta}\right)^p,$$

which establishes Fact 6.1 with $C\left(p\right) = p+1$.

Next, in order to prove Fact 6.2, we first use the formula for the negative first moment of M_∞ (see, e.g., [43, Chapter 21]): for every $\alpha > 1$ we have $\mathbb{E}\left[M_\infty^{-1}\right] = \left(\alpha+\beta-1\right) / \left(\alpha-1\right)$. Thus by Markov's inequality we have that $\mathbb{P}\left(M_\infty < z\right) \leq z\left(\alpha+\beta-1\right) / \left(\alpha-1\right)$ for every $z > 0$, and thus

$$\mathbb{P}\left(M_\infty < t\frac{\alpha}{\alpha+\beta}\right) \leq 2t. \tag{6.24}$$

In the case that $\alpha = 1$, we have $h\left(x\right) \leq \beta$ which implies that $\int_0^{\frac{t}{\beta+1}} h(x)dx \leq t$. We conclude that (6.24) holds for any $\alpha, \beta \geq 1$. Since M_k is a positive martingale, we have

$$\mathbb{P}\left(M_\infty \leq 2z \,|\, M_n \leq z\right) \geq 1/2, \quad \forall z \in \left(0,1\right).$$

Combining this inequality with (6.24) gives

$$\mathbb{P}\left(M_n \leq t\frac{\alpha}{\alpha+\beta}\right) \leq 8t,$$

and formula (6.4) then follows with $C = 8$.

6.5.2 Estimates on sequences

Lemma 6.7. *Suppose that $\left\{a_n\right\}_{n\geq 1}$ is a sequence of nonnegative real numbers and that there exists n_0 such that $a_{n_0} > 0$. Let α be a positive integer.*

(a) If there exists N such that $a_{n+1} \geq \left(1+\alpha/n\right) a_n$ for every $n \geq N$, then

$$\liminf_{n\to\infty} a_n / n^\alpha > 0.$$

(b) If there exist constants c, γ, and N such that for every $n \geq N$,

$$a_{n+1} \leq \left(1+\alpha/n\right) a_n + c\left(\log\left(n\right)\right)^\gamma n^{\alpha-1},$$

then $a_n \gtrsim n^\alpha$.

(c) If there exist constants c, γ, and N such that for every $n \geq N$,

$$a_{n+1} \leq (1 + \alpha/n)\, a_n + c\, (\log(n))^{\gamma}\, n^{\alpha},$$

then $a_n \precsim n^{\alpha+1}$.

Proof. (a) By the assumption we have that $a_n \geq a_{n_0} \exp\left(\sum_{j=n_0}^{n-1} \log(1+\alpha/j)\right)$, where $a_{n_0} > 0$. For $0 \leq x \leq 1$ we have that $\log(1+x) \geq x - x^2$, and so using the fact that $\sum_{j=1}^{\infty} 1/j^2 < \infty$, we have that there exists $c > 0$ such that $a_n \geq c \exp\left(\alpha \sum_{j=n_0 \vee \alpha}^{n-1} 1/j\right)$. To conclude, recall that $\sum_{j=1}^{n-1} 1/j > \log(n)$.

(b) Let $b_n := a_n / n^{\alpha}$. We then have that $b_{n+1} \leq (1+\alpha/n)\,(n/(n+1))^{\alpha}\, b_n + c\,(\log(n))^{\gamma}/n$. There exists a constant $c'' = c''(\alpha)$ such that $(n/(n+1))^{\alpha} \leq 1 - \alpha/n + c''/n^2$ for every $n \geq 1$. Therefore there exists a constant $c' = c'(\alpha)$ such that $(1+\alpha/n)\,(n/(n+1))^{\alpha} \leq 1 + c'/n^2$ for every $n \geq 1$. Thus we have that $b_{n+1} \leq (1 + c'/n^2)\, b_n + c\,(\log(n))^{\gamma}/n$, and iterating this we get that

$$b_n \leq b_1 \prod_{j=1}^{n-1} \left(1 + c'/j^2\right) + \sum_{j=1}^{n-1} \left(\prod_{i=j+1}^{n-1} \left(1 + c'/i^2\right) \right) c\,(\log(j))^{\gamma}/j.$$

Since $\prod_{j=1}^{\infty} (1 + c'/j^2) < \infty$, we immediately get that $b_n \precsim 1$, and so $a_n \precsim n^{\alpha}$.

(c) This is similar to (b) so we do not repeat the argument. □

6.5.3 Tail behavior of the diameter

We reproduce a simple argument of [22] to obtain a tail bound for the diameter of a uniform attachment tree.

Lemma 6.8. *For every seed tree S there exists a constant $C = C(S)$ such that for every $K > 20$ we have*

$$\mathbb{P}(\mathrm{diam}(\mathrm{UA}(n,S)) > K \log(n)) \leq \frac{C(S)}{n^{K/2}}.$$

Proof. First, if we set $C(S) := (\mathbb{P}(\mathrm{UA}(|S|, S_2) = S))^{-1}$, then we have

$$\begin{aligned}\mathbb{P}(\mathrm{diam}(\mathrm{UA}(n,S)) > K\log(n)) &= \mathbb{P}(\mathrm{diam}(\mathrm{UA}(n,S_2)) > K\log(n) \mid \mathrm{UA}(|S|, S_2) = S) \\ &\leq C(S)\, \mathbb{P}(\mathrm{diam}(\mathrm{UA}(n,S_2)) > K\log(n)),\end{aligned}$$

so it remains to bound the tail of $\mathrm{diam}(\mathrm{UA}(n, S_2))$.

For notational convenience, shift the names of the vertices so that they consist of the set $\{0, 1, \ldots, n-1\}$ (instead of $\{1, 2, \ldots, n\}$), and call vertex 0 the root. With this convention, the label of the parent of vertex j is distributed as $\lfloor jU \rfloor$ where U is uniform on $[0,1]$. Similarly, an ancestor ℓ generations back has a label distributed like $\lfloor \ldots \lfloor \lfloor jU_1 \rfloor U_2 \rfloor \ldots U_{\ell} \rfloor$, where the U_i's are i.i.d. uniform on $[0,1]$.

Define R_j to be the distance from vertex j to the root. By the triangle inequality we have that $\operatorname{diam}(\mathrm{UA}(n,S)) \leq 2\max_{1\leq j\leq n-1} R_j$, so it suffices to bound the tail of this latter quantity. Using a union bound, it then suffices to bound the tail of R_j for each j. Now notice that $R_j \leq \min\{t : jU_1 \dots U_t < 1\}$. Consequently, for any $\lambda > 0$ we have

$$\mathbb{P}(R_j > t) \leq \mathbb{P}(jU_1 \dots U_t \geq 1) \leq \mathbb{E}\left[(jU_1 \dots U_t)^\lambda\right] = j^\lambda (\lambda+1)^{-t}.$$

This is optimized by choosing $\lambda = t/\log(j) - 1$ (provided $t > \log(j)$) to obtain

$$\mathbb{P}(R_j > t) \leq \exp\left(t - \log(j) - t\log\left(\frac{t}{\log(j)}\right)\right) \leq \exp\left(t - t\log\left(\frac{t}{\log(n)}\right)\right),$$

when $j \leq n$. Putting everything together we get that

$$\begin{aligned}\mathbb{P}(\operatorname{diam}(\mathrm{UA}(n,S_2)) > K\log(n)) &\leq \mathbb{P}\left(\max_{1\leq j\leq n-1} R_j > \frac{K}{2}\log(n)\right) \\ &\leq \sum_{j=1}^{n-1}\mathbb{P}\left(R_j > \frac{K}{2}\log(n)\right) \\ &\leq n^{-\frac{K}{2}\log\left(\frac{K}{2}\right)+\frac{K}{2}+1} \leq n^{-K/2},\end{aligned}$$

where the last inequality holds when $K > 20$. $\square$

Bibliography

[1] Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions*. Dover, 1964.

[2] Dimitris Achlioptas. "Threshold Phenomena in Random Graph Colouring and Satisfiability". PhD thesis. Department of Computer Science, University of Toronto, 1999.

[3] Dimitris Achlioptas, Assaf Naor, and Yuval Peres. "Rigorous location of phase transitions in hard optimization problems". In: *Nature* 435 (2005), pp. 759–764.

[4] Kenneth J. Arrow. "A Difficulty in the Concept of Social Welfare". In: *Journal of Political Economy* 58.4 (1950), pp. 328–346.

[5] Kenneth J. Arrow. *Social Choice and Individual Values*. 2nd. Yale University Press, 1963.

[6] Albert-László Barabási and Réka Albert. "Emergence of Scaling in Random Networks". In: *Science* 286.5439 (1999), pp. 509–512.

[7] J.J. Bartholdi III, C.A. Tovey, and M.A. Trick. "The Computational Difficulty of Manipulating an Election". In: *Social Choice and Welfare* 6.3 (1989), pp. 227–241.

[8] Itai Benjamini and Oded Schramm. "Recurrence of distributional limits of finite planar graphs". In: *Electronic Journal of Probability* 6.23 (2001), pp. 1–13.

[9] Noam Berger et al. "Asymptotic behavior and distributional limits of preferential attachment graphs". In: *The Annals of Probability* 42.1 (2014), pp. 1–40.

[10] Nadja Betzler, Rolf Niedermeier, and Gerhard J. Woeginger. "Unweighted Coalitional Manipulation Under the Borda Rule Is NP-Hard". In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 2011, pp. 55–60.

[11] Béla Bollobás et al. "The Degree Sequence of a Scale-Free Random Graph Process". In: *Random Structures & Algorithms* 18.3 (2001), pp. 279–290.

[12] Christer Borell. "Positivity improving operators and hypercontractivity". In: *Mathematische Zeitschrift* 180.3 (1982), pp. 225–234.

[13] Felix Brandt. "Some Remarks on Dodgson's Voting Rule". In: *Mathematical Logic Quarterly* 55.4 (2009), pp. 460–463.

[14] Sébastien Bubeck, Elchanan Mossel, and Miklós Z. Rácz. "On the influence of the seed graph in the preferential attachment model". In: *IEEE Transactions on Network Science and Engineering* 2.1 (2015), pp. 30–39.

[15] Sébastien Bubeck et al. "From trees to seeds: on the inference of the seed from large trees in the uniform attachment model". Preprint available at `http://arxiv.org/abs/1409.7685`. 2014.

[16] Gary Chamberlain and Michael Rothschild. "A Note on the Probability of Casting a Decisive Vote". In: *Journal of Economic Theory* 25.1 (1981), pp. 152–162.

[17] Peter Cheeseman, Bob Kanefsky, and William M. Taylor. "Where the Really Hard Problems Are". In: *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI)*. 1991, pp. 331–337.

[18] Nicolas de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. L'Imprimerie Royale, 1785.

[19] Vincent Conitzer and Tuomas Sandholm. "Nonexistence of Voting Rules That Are Usually Hard to Manipulate". In: *Proceedings of the 21st National Conference on Artificial Intelligence*. Vol. 21. 2006, pp. 627–634.

[20] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. "When Are Elections with Few Candidates Hard to Manipulate?" In: *Journal of the ACM* 54.3 (2007), pp. 1–33.

[21] Nicolas Curien et al. "Scaling limits and influence of the seed graph in preferential attachment trees". In: *Journal de l'École polytechnique — Mathématiques* 2 (2015), pp. 1–34.

[22] Luc Devroye and Svante Janson. "Long and short paths in uniform random recursive dags". In: *Arkiv för Matematik* 49.1 (2011), pp. 61–77.

[23] Shahar Dobzinski and Ariel D. Procaccia. "Frequent Manipulability of Elections: The Case of Two Voters". In: *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE)*. Springer, 2008, 653–664.

[24] Michael Drmota. *Random Trees: An Interplay between Combinatorics and Probability*. Springer, 2009.

[25] Cynthia Dwork et al. "Rank Aggregation Methods for the Web". In: *Proceedings of the 10th International Conference on World Wide Web (WWW)*. ACM. 2001, pp. 613–622.

[26] Steven N. Evans, Rudolf Grübel, and Anton Wakolbinger. "Trickle-down processes and their boundaries". In: *Electronic Journal of Probability* 17.1 (2012), pp. 1–58.

[27] Piotr Faliszewski and Ariel D. Procaccia. "AI's War on Manipulation: Are We Winning?" In: *AI Magazine* 31.4 (2010), pp. 53–64.

[28] Piotr Faliszewski et al. "Llull and Copeland Voting Computationally Resist Bribery and Constructive Control". In: *Journal of Artificial Intelligence Research* 35 (2009), pp. 275–341.

[29] Ehud Friedgut, Gil Kalai, and Noam Nisan. "Elections Can be Manipulated Often". In: *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2008, pp. 243–249.

[30] Ehud Friedgut et al. "A Quantitative Version of the Gibbard-Satterthwaite Theorem for Three Alternatives". In: *SIAM J. Comput.* 40.3 (2011), pp. 934–952.

[31] Yaotian Fu and P.W. Anderson. "Application of statistical mechanics to NP-complete problems in combinatorial optimisation". In: *Journal of Physics A: Mathematical and General* 19.9 (1986), pp. 1605–1620.

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

[33] Oded Goldreich et al. "Testing Monotonicity". In: *Combinatorica* 20.3 (2000), pp. 301–337.

[34] Carla Gomes and Toby Walsh. "Randomness and Structure". In: *Handbook of Constraint Programming*. Ed. by F. Rossi, P. van Beek, and T. Walsh. Foundations of Artificial Intelligence. Elsevier, 2006, pp. 637–662.

[35] I.J. Good and Lawrence S. Mayer. "Estimating the efficacy of a vote". In: *Behavioral Science* 20.1 (1975), pp. 25–33.

[36] Rudolf Grübel and Igor Michailow. "Random recursive trees: A boundary theory approach". In: *Electronic Journal of Probability* 20.37 (2015), pp. 1–22.

[37] L.H. Harper. *Global Methods for Combinatorial Isoperimetric Problems*. Cambridge University Press, 2004. ISBN: 0521832683.

[38] Javier C. Hernandez. *Dead Man on Ballot: An Election Mystery*. `http://www.nytimes.com/2011/01/30/nyregion/30dead.html`. January 28, 2011 (retrieved online October 24, 2011).

[39] Fereydoun Hormozdiari et al. "Not All Scale-Free Networks Are Born Equal: The Role of the Seed Graph in PPI Network Evolution". In: *PLoS Computational Biology* 3.7 (2007), e118.

[40] Marcus Isaksson, Guy Kindler, and Elchanan Mossel. "The Geometry of Manipulation: A Quantitative Proof of the Gibbard-Satterthwaite Theorem". In: *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2010, pp. 319–328.

[41] Marcus Isaksson, Guy Kindler, and Elchanan Mossel. "The Geometry of Manipulation: A Quantitative Proof of the Gibbard-Satterthwaite Theorem". In: *Combinatorica* 32.2 (2012), pp. 221–250.

[42] Svante Janson. "Limit theorems for triangular urn schemes". In: *Probability Theory and Related Fields* 134.3 (2006), pp. 417–452.

[43] Norman Lloyd Johnson, Samuel Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*. Second. Vol. 2. Wiley, 1995.

[44] Jerry S. Kelly. "Almost all social choice rules are highly manipulable, but a few aren't". In: *Social Choice and Welfare* 10.2 (1993), pp. 161–175.

[45] J.H.B. Kemperman. "Decomposing $\mathbb{R}^d$ into finitely many semigroups". In: *Indagationes Mathematicae (Proceedings)*. Vol. 89. 1. Elsevier. 1986, pp. 71–78.

[46] Robert D. Kleinberg and Jon M. Kleinberg. "Isomorphism and Embedding Problems for Infinite Limits of Scale-Free Graphs". In: *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2005, pp. 277–286.

[47] John H. Lindsey II. "Assignment of Numbers to Vertices". In: *The American Mathematical Monthly* 71.5 (1964), pp. 508–516.

[48] Hosam M. Mahmoud. "Distances in random plane-oriented recursive trees". In: *Journal of Computational and Applied Mathematics* 41.1–2 (1992), pp. 237–245.

[49] Tamás F. Móri. "The Maximum Degree of the Barabási–Albert Random Tree". In: *Combinatorics, Probability and Computing* 14.03 (2005), pp. 339–348.

[50] Elchanan Mossel. "A quantitative Arrow theorem". In: *Probab. Theory Relat. Fields* 154.1–2 (2012), pp. 49–88.

[51] Elchanan Mossel, Ariel D. Procaccia, and Miklós Z. Rácz. "A Smooth Transition from Powerlessness to Absolute Power". In: *Journal of Artificial Intelligence Research* 48 (2013), pp. 923–951.

[52] Elchanan Mossel and Miklós Z. Rácz. "A quantitative Gibbard-Satterthwaite theorem without neutrality". In: *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*. ACM. 2012, pp. 1041–1060.

[53] Elchanan Mossel and Miklós Z. Rácz. "A quantitative Gibbard-Satterthwaite theorem without neutrality". In: *Combinatorica, to appear* (2015+).

[54] Elchanan Mossel et al. "Non-interactive correlation distillation, inhomogeneous Markov chains, and the reverse Bonami-Beckner inequality". In: *Israel Journal of Mathematics* 154.1 (2006), pp. 299–336.

[55] Hervé Moulin. *The Strategy of Social Choice*. North-Holland, 1983.

[56] David P. Myatt. "On the Theory of Strategic Voting". In: *The Review of Economic Studies* 74.1 (2007), pp. 255–281.

[57] Saket Navlakha and Carl Kingsford. "Network Archaeology: Uncovering Ancient Networks from Present-Day Interactions". In: *PLoS Computational Biology* 7.4 (2011), e1001119.

[58] Roberto Oliveira and Joel Spencer. "Connectivity Transitions in Networks with Super-Linear Preferential Attachment". In: *Internet Mathematics* 2.2 (2005), pp. 121–163.

[59] Erol A. Peköz, Adrian Röllin, and Nathan Ross. "Degree asymptotics with rates for preferential attachment random graphs". In: *The Annals of Applied Probability* 23.3 (2013), pp. 1188–1218.

[60] Erol A. Peköz, Adrian Röllin, and Nathan Ross. "Joint degree distributions of preferential attachment random graphs". Preprint available at `http://arxiv.org/abs/1402.4686`. 2014.

[61] Bezalel Peleg. "A note on manipulability of large voting schemes". In: *Theory and Decision* 11.4 (1979), pp. 401–412.

[62] Geoffrey Pritchard and Arkadii Slinko. "On the average minimum size of a manipulating coalition". In: *Social Choice and Welfare* 27.2 (2006), pp. 263–277.

[63] Geoffrey Pritchard and Mark C. Wilson. "Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour". In: *Mathematical Social Sciences* 58.1 (2009), pp. 35–57.

[64] Ariel D. Procaccia and Jeffrey S. Rosenschein. "Average-Case Tractability of Manipulation in Voting via the Fraction of Manipulators". In: *Proceedings of the 6th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2007, 718–720.

[65] Ariel D. Procaccia and Jeffrey S. Rosenschein. "Junta Distributions and the Average-Case Complexity of Manipulating Elections". In: *Journal of Artificial Intelligence Research* 28 (2007), pp. 157–181.

[66] Anna Rudas, Bálint Tóth, and Benedek Valkó. "Random Trees and General Branching Processes". In: *Random Structures & Algorithms* 31.2 (2007), pp. 186–202.

[67] Mark Allen Satterthwaite. "Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions". In: *Journal of Economic Theory* 10.2 (1975), pp. 187–217.

[68] Regev Schweiger, Michal Linial, and Nathan Linial. "Generative probabilistic models for protein-protein interaction networks—the biclique perspective". In: *Bioinformatics* 27.13 (2011), pp. i142–i148.

[69] Arkadii Slinko. "How large should a coalition be to manipulate an election?" In: *Mathematical Social Sciences* 47.3 (2004), pp. 289–293.

[70] Steve Tawa. *Philadelphia Voting Machines Pass Inspection, Even If All Ballot Names Don't.* `http://philadelphia.cbslocal.com/2011/05/06/philadelphia-voting-machines-pass-inspection-even-if-all-ballot-names-dont/`. May 6, 2011 (retrieved online October 24, 2011).

[71] Toby Walsh. "The Interface between P and NP: COL, XOR, NAE, 1-in-k, and Horn SAT". In: *Proceedings of the 17th National Conference on AI (AAAI)*. 2002, pp. 695–700.

[72] Toby Walsh. "Where Are the Hard Manipulation Problems?" In: *Journal of Artifical Intelligence Research* 42 (2011), 1–29.

[73] Lirong Xia. "Computing the Margin of Victory for Various Voting Rules". In: *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. ACM. 2012, pp. 982–999.

[74] Lirong Xia. "How Many Vote Operations Are Needed to Manipulate A Voting System?" In: *Proceedings of the 4th International Workshop on Computational Social Choice*. 2012.

[75] Lirong Xia and Vincent Conitzer. "A Sufficient Condition for Voting Rules to be Frequently Manipulable". In: *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*. ACM. 2008, pp. 99–108.

[76] Lirong Xia and Vincent Conitzer. "Finite Local Consistency Characterizes Generalized Scoring Rules". In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*. 2009, pp. 336–341.

[77] Lirong Xia and Vincent Conitzer. "Generalized Scoring Rules and the Frequency of Coalitional Manipulability". In: *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*. ACM. 2008, pp. 109–118.

[78] Lirong Xia et al. "Complexity of Unweighted Coalitional Manipulation Under Some Common Voting Rules". In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*. 2009, pp. 348–353.

[79] H.P. Young. "Social Choice Scoring Functions". In: *SIAM Journal on Applied Mathematics* 28.4 (1975), pp. 824–838.

[80] Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. "Algorithms for the coalitional manipulation problem". In: *Artificial Intelligence* 173.2 (2009), pp. 392–412.