universität wien

MAGISTERARBEIT / MASTER'S THESIS

Titel der Magisterarbeit / Title of the Master's Thesis

"From Gerrymandering to Ballot Stuffing"

–

How statistical methods can detect election bias

verfasst von / submitted by

Stefan Ortner, BSc

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Magister der Sozial- und Wirtschaftswissenschaften (Mag.rer.soc.oec.)

Wien, 2020 / Vienna, 2020

Studienkennzahl lt. Studienblatt / degree programme code as it appears on the student record sheet:	A 066 951
Studienrichtung lt. Studienblatt / degree programme as it appears on the student record sheet:	Magisterstudium Statistik
Betreut von / Supervisor:	Univ.-Prof. i.R. Dr. Wilfried Grossmann

Abstract - English

This thesis analyses election bias from a new and broad perspective. Data-driven methods for the detection of election fraud and methods for the analysis of legal practices of influencing election results are presented. Legal practices include the conscious manipulation of political district boundaries which is also known as gerrymandering. Methods considered as election fraud contain among others so-called ballot stuffing which is the addition of votes never cast by a voter towards the ballot box. These phenomena are analysed and methods for detection are presented. The statistical framework, similarities and differences between the two approaches are presented in detail and applications on election data are discussed. The thesis concludes that the problem settings show many similarities in both cases and the difficulties in application are comparable as well. Differences were spotted regarding the complexity of the methods used for detection.

Deutsche Zusammenfassung

Die vorliegende Arbeit betrachtet das Thema der strukturellen Wahlbeeinflussung aus einer neuen, breiteren Perspektive. Es werden sowohl datengetriebene Methoden vorgestellt, die sich mit Wahlbetrug beschäftigen, aber auch Methoden, die legale Beeinflussungsmöglichkeiten analysieren. Darunter wird beispielsweise das bewusste Verschieben von Wahlbezirksgrenzen aus politischen Gründen - auch bekannt als Gerrymandering - verstanden. Unter dem Begriff Wahlbetrug sind unter anderem Praktiken wie das Hinzufügen von Stimmzetteln, die nicht von Wahlberechtigten abgegeben wurden - auch als Ballot Stuffing bekannt - gemeint. Die Phänomene werden analysiert und Methoden zu deren Aufdeckung vorgestellt. Die statistischen Rahmenbedingungen, Ähnlichkeiten und Unterschiede in der Herangehensweise werden ausgearbeitet und Anwendungsbeispiele der Methoden präsentiert. Die Arbeit kommt zu dem Ergebnis, dass die Problemstellungen in beiden Fällen ähnlich sind und auch die Schwierigkeiten in der praktischen Anwendung vergleichbar sind. Unterschiede wurden hinsichtlich der Komplexität der Verfahren zur Erfassung der Praktiken festgestellt.

Eidesstattliche Erklärung

Ich erkläre eidesstattlich, dass ich die vorliegende Arbeit selbst verfasst habe. Alle Hilfsmittel, die ich verwendet habe, habe ich angegeben. Alle übernommenen Formulierungen und Gedanken habe ich als solche gekennzeichnet. Ich habe diese Arbeit bisher an keiner anderen Stelle vorgelegt.

Stefan Ortner
Wien, 2020

Danksagung

Ich möchte mich an dieser Stelle besonders bei meinen Eltern für die moralische, finanzielle und oftmals auch kulinarische Unterstützung während meiner Studienzeit bedanken. Ich bedanke mich auch ganz herzlich bei meiner Freundin Natalie, die meine endlosen Monologe über Gerrymandering mit viel Geduld ertragen hat und mich während des Verfassens der Arbeit stets unterstützt hat.

Ich möchte mich bei meinen Studienkolleg*innen dafür bedanken, dass die letzten Jahre neben harter Prüfungen auch viele lustige und schöne Momente beinhaltet haben und wahrscheinlich vor allem deshalb sehr schnell vergangen sind.

Ebenfalls bedanken möchte ich mich bei Prof. Grossmann für die Betreuung der folgenden Arbeit.

Contents

1 Introduction

A candidate receiving 234,000 votes, although only 15,000 people were registered to vote - that is the story of the 1927 presidential election in Liberia. [Blundell, 1980]

Unfortunately election fraud is not always as obvious as it was back then in Liberia and one cannot decide on the presence of fraud by solely looking at the results. Numerous different methods have been invented to analyse elections and rate the credibility of their results. In recent years data-driven approaches became increasingly important as new methods were developed and computational costs are decreasing. This thesis provides an overview about statistical methods and models used to model different types of "election bias" by which legal or illegal practices are meant which result in an advantage for a party in the democratic election process.

The approach of this thesis may seem unorthodox. It does not exclusively treat election fraud practices, but also discusses questionable practices which are indeed legal. The goal is to compare statistical methods used for discovering these types of election bias to outline similarities and differences. The thesis is consequently split into two main parts where the first discusses legal practices of influencing election outcomes with a focus on the practice of gerrymandering by which the practice of manipulating voting district lines to achieve political advantages is meant. This phenomenon can take on different forms which are discussed in more detail. The tasks analysed in this section include the optimal gerrymandering strategy from the perspective of the party in charge, how a district ideally should be split from a neutral perspective, the effect of imposed constraints on gerrymandering strategies and measurements for detecting the presence of gerrymandering.

In the second part this thesis digs deeper into the topic of actual election fraud. Different types of election fraud are categorized to give an overview on existing practices and the challenges that methods trying to detect those face. A statistical framework for how these methods proceed is introduced and the variables most important for the analyses are discussed.

Within that framework different methods for detecting concrete forms of election fraud are examined. The first method is Benford's Law which is a simple approach for detecting anomalies in datasets. This section is followed by the discussion of the practices of ballot stuffing and voter rigging. To further illustrate the presented methods, papers in which they are applied on election data are discussed as well.

In a final section the two parts of this thesis are summarized with a focus on what the methods for detecting gerrymandering and election fraud have in common and in which areas differences between the two were detected. Future potential areas of research are outlined and the impact these methods have had so far is discussed.

1.1 Research Question

The research question of this thesis is
"How can statistical methods be used to detect different forms of election bias?"

The research question is intentionally stated in this general form as the novelty of this thesis is that it compares a broad range of different methods tackling diverse types of election bias.
However, this thesis does not "discover" illegal practices in any concrete country or claims fraud to be

present in any specific election. The goal of this thesis is to focus on the statistical methods and compare them across a wide range of potential challenges.

1.2 Methods and Literature

The open formulation of the research question has the effect that the methods presented also form a wide range. In the context of gerrymandering different ways of formulating this process as an optimization problem will be presented. Depending on the assumptions, the consequent optimization problem will differ in its complexity. Different ways of detecting gerrymandering are discussed, which are data-driven, non-parametric measurements with different underlying concepts of fair districtings.

In the second part about election fraud the focus is on different ways of modelling the relationship between turnout and the percentage of the winning party. Some of the methods can be classified as being part of "visual analytics", but have a theoretical foundation in that they determine so-called fraud parameters via goodness-of-fit procedures. In the section about voter rigging, the procedure discussed is based on outlier detection in that it compares new election data towards a reference set of trustworthy elections. The used models will be described in more detail in the respective sections.

Literature on the subject is closely related to the methods presented. Published articles and books from 1988 until 2019 stemming from a variety of research fields are included. These articles were published in different fields such as statistics, data analysis, complexity science, political science and law.

2 Legal Election Bias - Gerrymandering

2.1 Introduction

Gerrymandering - a word many have already heard, but only few actually know what it means. A common reason for that is its reasonably complex definition. In this thesis gerrymandering corresponds to the practice of redrawing district borders (lines) for political purposes. The term “Gerrymandering” dates back to former Massachusetts Governor and later US Vice president Elbridge Gerry (1744-1814) who was criticized for redrawing district lines in his party’s favour. One of those districts looked similar to a salamander and consequently, the word “Gerrymander” is a mixture between “Gerry” and “Salamander”.

Figure 1: Original “Gerrymander”

People knowing about the existence of gerrymandering associate that practice with the United States of America. That is no coincidence, as the political system is perfect for perofrming gerrymandering in two different ways:

1. In the US electoral boundaries are (in some states)drawn by political parties themselves instead of an independent commission. [Friedman and Holden, 2005]
2. The American electoral system is majority based.

It is not difficult at all to find prominent examples of where political parties used the redrawing of district lines in their favour. In 2005, the states Florida, Michigan and Pennsylvania which were almost evenly divided among voters, were on aggregate represented by 39 Republicans and 20 democrats. On the other hand the same phenomenon was observed in Texas. Texas consistently voted 2 to 1 Republican in the popular vote, but more than half of the members of the Texas congressional delegation were Democrats until a new redistricting plan was deployed in 2004. [Friedman and Holden, 2005]

A common question regarding gerrymandering is its legitimacy. It contradicts common intuition towards legislative processes that practices favouring one party in an election process are legitimate. The situation is indeed more complex. Gerrymandering using unequal district sizes or racial characteristics is unlawful, yet partisan gerrymandering is controversial, but legal. [Friedman and Holden, 2005]
The biggest practical problem is to determine where redistricting stops and where gerrymandering begins.

In the following section the problem of gerrymandering will be looked at from different perspectives, starting with the influence of a country’s voting system. Majority voting systems such as in the US, the United Kingdom and in France enable the possibility for gerrymandering while proportional representation systems rule out the possibility of common gerrymandering practices. In practice, voting systems can seldom be classified as purely majority based or proportional, but consist out of elements of both.

Gerrymandering is a general and broad term involving different strategies. In the US the phrase "pack or crack" is used to describe the two main strategies, but there are more approaches towards political redistricting.
The next step after lining out the strategies is to take the perspective of the party in charge of the districting process. Phrasing gerrymandering as an optimization problem trying to give the ruling party the biggest advantage possible will reveal interesting insights. By modifying assumptions and constraints of the problem, the results obtained are quite diverse.

The next perspective is the one of an independent institution trying to find the districting which optimally ensures proper representation of each voter group. That goal is summarized by the term "socially-optimal districting".

The last interesting perspective in this section is the one of a neutral election observer trying to reveal the presence of gerrymandering. There exist many different measurements in literature for that purpose which will be discussed.

2.2 Gerrymandering in different voting systems

The goal of any electoral system is to guarantee fair representation of the voting result in the elected political body. The question attached to that statement is what is meant by "fair representation". This section focuses on the two most widely spread voting systems:

1. Proportional Representation System: Suppose a US state elects its representatives in the House of Representatives. This state owns 5 seats (e.g. Connecticut) and party A receives 60% of the vote while party B wins 40% of the vote. Under this system party A gains 3 seats while party B gains 2 seats.

2. Majority based system: The actual system in the United States works differently. The state is split into 5 districts (number of seats) where each district gets to choose its "own" representative. The seat is won by the party winning most votes in the respective district. The result depends on how the state is split into districts and even after knowing the outcome of the state-wide result, the number of seats won by party cannot be determined without looking at voting district results.

In the first system any type of gerrymandering is difficult if not impossible. There is no similar phenomenon to gerrymandering if the delegates in the elected legislative body are split according to the nation-wide election result between the political parties. Yet, there exist minor tools with a potential impact on the seat ratio in parliament. One such tool is the establishment of a threshold for minor parties. Common thresholds in European countries are four or five percent of the nation-wide vote. If a party's national vote percentage is below that threshold, it is not represented in parliament, even though the party would gain seats if they were spread solely according to the proportional representation system. For example, suppose a legislative body consists of 100 seats and the voting system established a 5% threshold for minor parties. A party winning 3% of the nation-wide vote should be represented with 3 seats in parliament, but it is not as result of the threshold. If minor parties fail to reach the threshold, the strongest party in parliament profits as votes for parties in the parliament become more influential, of which the biggest party has the most. Even though the threshold can cause changes in the composition of parliament, it is not comparable to gerrymandering which can invert entire election results.

There is a second small effect in proportional representation systems worth discussing. In Austria, the size of a voting district is determined according to the so-called "Wahlzahl" ("voting number") which is the number of votes cast in a voting district, divided by the number of seats that voting district has in parliament. In the nineteenfifties and sixties, Austrian parties discussed the potential difference of the Wahlzahl referring to the population of a voting district or to the number of eligible voters, as the claim was that people in rural areas had more children who were no eligible voters. It was claimed back then that a Wahlzahl referring to the total population would give more weight to rural areas. That effect is negligible today as there is no impact in the current electoral system.

Conversely, the second voting system opens up numerous possibilities for performing gerrymandering. The drawing of district lines is essential to the voting result, but redrawing is in general not something negative. It enables to adjust to latest demographic developments which is why districts in urban areas are in general smaller than districts in rural areas to ensure that delegates roughly represent the same number of voters. Another advantage is that by changing political district borders, proper representation of minorities can be ensured more easily.
In practice gerrymandering has often been used to do the opposite. By splitting up African American neighbourhoods into various districts, proper political representation of minorities can be thwarted. Each US state has its own procedure for redrawing district boundaries which caused some states handing the task to independent commissions. [Friedman and Holden, 2005]

Still states face only few constraints when setting their districts. Congressional districts must have roughly the same population and must be contiguous which is a tolerant constraint in practice. [Gul and Pesendorfer, 2010]

2.3 Strategies of Gerrymandering

Not all types of gerrymandering are equal as there is no guidebook of performing it. In fact there exist different strategies for reaching political goals via the redrawing of district lines which are listed below. [Owen and Grofman, 1988]

- Concentration gerrymander (also called "packing"): concentrating one's opposite voters in one district which the opposition wins by a large margin. Consequently, the opposition loses influence in all other districts.
  This tactic is meaningful when the proper representation of a minority group shall be ensured. Yet there is a fine line between concentrating voters to secure their representation and over-concentrating a group to remove their ability to influence the outcome of surrounding districts. [Payne-Riley, 2017]
  This strategy turns out to be of better use for minority parties, as majority parties discovered the cracking strategy as more efficient. However, it makes sense for the minority party to concentrate large numbers of the opponent's voters into a single district as it is thereby equalizing the possibilities of winning other districts.

- Dispersal gerrymander (also called "cracking"): spreading the opposite party's voters across as many districts as possible to deny majorities in as many districts as possible. That strategy is the most common gerrymandering strategy. By preventing voting blocs of the opposite party, the arithmetic in the majority voting system enables one party to win a large majority of the districts while only having a slight edge in the popular vote.

- Incumbent displacement gerrymander: Eliminating seats held by members of the opposing party by moving the homes of two incumbents into the same district such that only one of them can win the district in the next election. However, incumbents can also cooperate to raise the chances of both incumbents being re-elected. [Owen and Grofman, 1988]
- Prison-based gerrymandering: In most US states prisoners are not allowed to vote, but count when calculating the population size. As the majority of prisoners usually stems from urban areas, their weight is shifted to rural areas where prisons are usually located. [Wagner, 2012]

The following analysis only discusses the first two types of gerrymandering as the other two types require different analyses.

2.4 Optimal Gerrymandering

The next step is to present different approaches of writing gerrymandering as an optimization problem. The following first two approaches are extended descriptions of the two cases discussed by Owen and Grofman who take the perspective of someone being in charge of the redistricting process trying to maximize the chances of their party.

Maximizing the expected number of seats

Assume a two party model in which one party (Party 1) is in charge of the redistricting process. The electoral outcomes are assumed to have a probabilistic component. A random variable Z is introduced and a number characterizing the j-th district which is named α_j. Party 1 wins all districts for which $\alpha_j > Z$ and loses districts in which $\alpha_j < Z$. The case where $\alpha_j = Z$ is neglected as it occurs with probability 0.
Z is a random variable with a long-term distribution $H(z)$. The numbers α_j are interpreted as general favour or disfavour of Party 1 in a particular district. These can be changed by gerrymandering.

Define $y(x)$ as the number of districts j such that $\alpha_j = x$. x is a run variable with the same scale as α_j while $y(x)$ is the number of districts in which α_j is exactly at x. This will be used for the formulation of the constraints of the optimization problem: [Owen and Grofman, 1988]

$$\sum_x y(x) = N$$

$$\sum_x xy(x) = A$$

$$y(x) \geq 0 \ \forall x$$

N is the total number of districts. By summing up $y(x)$ for all potential values of x each district is included in the calculation and the result of the sum is the total number of districts.
A is a measure for the overall support of Party 1 and Party 2. A is large if the values of α_j (support in districts) are large for many districts and is low if many α_j yield low values.

In addition to the constraints above, Owen and Grofman fix x such that $-1 \leq x \leq 1$ for simplicity and interpretability.

Of particular interest is the long-term probability of winning a district. For the district j, this probability is given by the distribution function introduced above, $H(\alpha_j)$. Thus, the expected number of seats won $S(y)$ is obtained by multiplying the number of districts with characteristic α_j with the probability of winning that district $H(\alpha_j)$ and summing up over all potential values for α_j (or x respectively).

$$\mathbb{E}[S(y)] = \sum_x y(x)H(x) \to \max!$$

This expected value is the objective function which shall be maximized under the constraints stated above. As the expected value can be written as above, the optimization problem is a linear program and by formulating the dual problem one obtains

$$Nu + Av \to \min!$$

subject to

$$u + xv \geq H(x)\ \forall x$$

The most notable difference is that the dual problem only has two variables and therefore only two solutions. It follows that the primal problem has a solution in which at most two of the variables $y(x)$ are different from zero which in optimization theory corresponds to the solution being at a corner. An optimal gerrymandering strategy generally consists of setting one set of districts at one level of support for one's party and the remaining districts at another level of support for the party. [Owen and Grofman, 1988]
In practice this means that the opposition party is given a large majority in a small number of districts. The number of such districts depends on the specific values of N, A and H i.e. the total number of districts, the overall support for the party and the distribution of the random variable Z.

Owen and Grofman simplify the analysis by setting $N = 1$ and $A = 0$. This can be done without loss of generality. Setting $N = 1$ causes the result to be interpreted differently i.e. it does not represent a number of districts, but a fraction of districts. Setting $A = 0$ can always be accomplished by shifting the value of x by A/N. The modified problem therefore looks different:

$$\sum_x y(x)H(x) \to \max!$$

s.t.

$$\sum_x y(x) = 1$$

$$\sum_x xy(x) = 0$$

$$y(x) \geq 0\ \forall x$$

The dual problem changes to

$$u \to \min!$$

s.t.

$$u + xv \geq H(x) \Leftrightarrow u \geq H(x) - vx\ \forall x$$

It follows that for a given value v the minimal u satisfies

$$u = \max_x \{H(x) - vx\}$$

Assume $H(x)$ to be differentiable with derivative $h(x)$ and under the restriction $-1 \leq x \leq 1$ the maximum can only be obtained at a point x at which one of the following three conditions holds:

1. $h(x) = v$
2. $x = -1, h(-1) \leq v$
3. $x = 1, h(-1) \geq v$

Three possible results for u as a function of v are obtained from the conditions above:

1. $u(v) = H(h^{-1}(v)) - vh^{-1}(v)$
2. $u(v) = H(-1) + v$
3. $u(v) = H(1) - v$

Both extreme values, $h(-1)$ and $h(1)$ are expected to be very small which corresponds to the idea that large swings are infrequent and therefore the third case, in which the condition $h(1) \geq 1$ has to be satisfied, is seldom applicable. [Owen and Grofman, 1988]
The focus is therefore on the case where the first two conditions hold. As stated above, not all of the conditions need to hold, it suffices that one condition holds. The optimum is obtained when the two corresponding values of $u(v)$ coincide, so that (from above): [Owen and Grofman, 1988]

$$H(h^{-1}(v)) - vh^{-1}(v) = H(-1) + v$$

equivalently if $v = h(x)$

$$H(x) - xh(x) = H(-1) + h(x)$$

which can be reduced to

$$H(x) - (x+1)h(x) = H(-1)$$

This equation has a root $\tilde{x}$ which is usually greater than 0. The optimal gerrymandering strategy consequently requires

$$y(-1) = \frac{\tilde{x}}{1+\tilde{x}}$$
$$y(\tilde{x}) = \frac{1}{1+\tilde{x}}$$

where $y(x) = 0$ for all other x. That means that a fraction of all districts $\left(\frac{\tilde{x}}{1+\tilde{x}}\right)$ will vote overwhelming in favour of the opposition ($\alpha_j = -1$) under the optimal gerrymandering plan, while the remaining districts $\left(\frac{1}{1+\tilde{x}}\right)$ vote solidly in favour of Party 1. A situation in which both parties win seats with large majorities is constructed, but more seats are won by the party in control of the gerrymandering process. Since $\tilde{x} < 1$ usually, Party 1 wins a solid majority of seats in most cases. The expected number of seats according to Owen and Grofman is

$$\mathbb{E}[S(y)] = \frac{\tilde{x}H(-1) + H(\tilde{x})}{1+\tilde{x}}$$

Maximizing the probability of controlling a working majority of seats

The task can be modified by assuming a different goal of the gerrymandering party. Instead of trying to maximize the expected number of seats, the party wants to ensure to win the majority of seats with the highest possible probability. This majority does not necessarily have to be a bare majority of 50%, but can be defined as whatever the party in charge of the gerrymandering considers to be its goal. To specify that, a number k out of n districts is fixed which is the desired number of districts the party needs to win for its working majority.

The number α_j again characterizes the district j and the party carries the district if $Z < \alpha_j$. Assuming restrictions from above

$$y(\hat{x}) = \frac{k}{n}$$

$$y(-1) = 1 - \frac{k}{n}$$

$y(x)$ is 0 for all other x and $\hat{x} = \frac{n-k}{k}$. The probability of of gaining a working majority of seats is then given by $H(\hat{x})$ where H is as above. This results in a bimodal redistricting scheme which would be expected under a bipartisan districting scheme. [Owen and Grofman, 1988].

For this analysis to be valid a strong positive correlation in the year-to year swings in votes between legislative districts is assumed. That can also be interpreted as having a factor explaining year to year swings which is present in all districts. The alternative would be to assume the swings to be totally uncorrelated. In this case it can be shown that the optimal gerrymander calls for strengthening the party equally in a specified number of districts. [Owen and Grofman, 1988]
In other words, the party would be given strong majorities in r of the districts. However, it is difficult to specify r analytically.

Under the above analysis the districting scheme by which the probability for winning the working majority is maximized is by maximizing α_j uniformly in exactly k districts while making it small in the remaining $n - k$ districts. [Owen and Grofman, 1988]

An example for comparing the objectives

Suppose the random variable Z follows a normal distribution with mean 0 and standard deviation 1/2. Then $H(x) = \Phi(2x)$ where Φ is the cumulative standard normal distribution function, while $h(x) = 2\phi(2x)$ with ϕ being the normal density function. To assume Z to have 0 mean can be interpreted as having two parties of even strength trying to win a majority or maximize the number of seats. With the formulas from above one obtains that

$$H(x) - (x+1)h(x) = H(-1)$$

$$\Leftrightarrow \Phi(2x) - 2(x+1)\phi(2x) = \Phi(-2)$$

$$\Leftrightarrow \Phi(2x) - 2(x+1)\phi(2x) = 0.02$$

This equation can't be solved analytically, but numerically. Owen and Grofman determined $\tilde{x} = 0.42$ to be the optimal solution. It immediately follows that

$$y(-1) = \frac{\tilde{x}}{1+\tilde{x}} = \frac{0.42}{1.42} = 0.30$$

$$y(x^*) = \frac{1}{1+\tilde{x}} = \frac{1}{1.42} = 0.70$$

The optimal gerrymandering strategy when maximizing the expected number of seats is therefore to give the party in charge of the redistricting process a strong majority ($x = 0.42$) in 70% of the districts while the second party wins the other 30% of the districts overwhelmingly ($x = -1$).

In case Party 1 wants to ensure having a working majority (here arbitrarily fixed at 52% of the seats), $\tilde{x}$ is given by

$$\tilde{x} = \frac{n-k}{k} = \frac{0.48}{0.52} = 0.92$$

The same calculations yield

$$y(\tilde{x}) = 0.52$$

$$y(-1) = 0.48$$

This result is different from the result obtained above. The dominant party should construct itself an overwhelming majority in 52% of the districts while handing the other party an even stronger majority in the other 48%. This is due to the fact that the parties start on an even level.

In the first case the dominant party usually controls around 70 percent of the seats, but once in a while the opposition wins by large numbers (probability 0.2) and then even wins all of the seats. On the other hand there is a slight chance of the dominant party controlling all of the seats. The party will control 0, 70 or 100 percent of the seats with probabilities 0.20, 0.78 and 0.02. The expected percentage of seats won is 56.6%. [Owen and Grofman, 1988]

In the second case the party usually controls 52% of the seats. There is a small chance that the opposition wins all of the seats (probability 0.03) and an even smaller chance of the dominant party winning all seats (probability 0.02). Therefore Party 1 controls 0, 52 or 100 percent of the seats with probabilities 0.03, 0.95 and 0.02 which results in an expected percentage of seats won of 51.5%.

The first plan gives Party 1 a substantially higher expected value of seats won, with the disadvantage being that the majority is safer in strategy 2 as it controls the legislature 97% of the time compared to 80% under strategy 1. [Owen and Grofman, 1988]

New model: Throwing-away districts?

A different approach towards modelling the gerrymander process was introduced by Friedman and Holden and focusses on the question whether it is ideal to concentrate extreme opposite voters in throwing-away districts while spreading one's supporters over so-called winnable districts which is the solution obtained above.

Their model assumes two parties D and R (Democrats and Republicans) and R is assumed to be the gerrymandering party which creates district profiles. There is a continuum of voters which is characterized by a policy preference parameter which can't be observed. Instead the gerrymander receives a noisy signal of it. The party also observes the posterior distribution of the policy preference parameter conditional on the signal. The marginal distribution of the signal is referred to as "signal distribution". The problem the party faces is creating N voting districts by allocating voters from the signal distribution by which the expected number of seats shall be maximized. The probability for each party to win a district is determined by the median voter in that district.[Friedman and Holden, 2005]

The constraints are:

1. each voter must be allocated to one and only one district
2. all districts must contain an identical number of voters

The political preference is given by the unobserved parameter $\beta \in \mathbb{R}$ for which a noisy signal, $s \in \mathbb{R}$ is known. The joint distribution is given by $F(\beta, s)$ on support $\mathbb{R}^2$. R is the gerrymanderer and R's Bayesian posterior distribution of preferences given an observed signal is $G(\beta|s)$. This distribution is referred to as the "conditional preference distribution". Both F and G are assumed to be absolutely continuous, while the marginal distribution of s is given as:

$$H(s) = \int F(\beta, s) d\beta$$

This function represents not only a characterization of a single draw from the population, but also the mass of voters within the population. H is the "signal distribution". The gerrymandering party R allocates mass from this distribution to form districts. [Friedman and Holden, 2005]

In an election, voters choose the candidate closest to them on the ideological spectrum. This is represented in the model by assuming that there is a "stochastic break point" in each election with distribution function B. Voters positioned above the realization of the breakpoint vote Republican, those below vote for the Democratic candidate.

For example, Friedman and Holden choose B to be standard normal and some voter to have political preference $\beta = 1.96$. This voter elects the Republican candidate with probability 0.975. Once the above described breakpoint is calculated, the uncertainty is resolved and the position of voter and breakpoint determines the voting behaviour with probability 1. The only uncertainty involved is that the breakpoint itself is stochastic. To simplify the problem, the breakpoint is assumed to be equal for all districts, so that no more local effect is in the model. The gerrymanderer R divides the population into N equally sized districts such that the expected number of seats won is maximized.

The optimal strategy is characterized in four steps which describe how the distribution H should be chopped up optimally. The authors first show that all other shapes than vertical slices or so-called "parfaits" of the signal distribution h are not optimal and therefore the basic shape of districts must be vertical slices of $h(s)$. If one district n has a different mean than the other districts, then it must be the case that this district consists of vertical slices of $h(s)$. On the other hand if districts n and m share a median, they must together comprise vertical slices of $h(s)$. These slices could be split between the two districts in many different ways, but the median and the density of district preferences must remain the same. One such way to split vertical slices between the districts would be an equal split of $h(s)$ for all s in the districts. [Friedman and Holden, 2005]

In the second step the authors show that the mass of higher-median districts must lie outside that of lower value districts. It is inefficient to have mass for a district with a higher median between voters from a lower-median district because that mass could be reallocated above or below the lower district for a profitable deviation. [Friedman and Holden, 2005]

In step three the authors rule out the shape of a parfait as it is not optimal and from only focus on the shape of vertical slices of the signal distribution. In the fourth and last step the authors show that voters in a high-median district cannot lie within the set of all voters in a lower-median district. The intuition is similar to step 2: If there was such a group of voters, a profitable deviation would exist by moving this group to an outer slice.

The result by Owen and Grofman can be seen as a special case of this model. But in a general setting the two models give different explanations. In the previous model, the authors recommend to create "throw-away" districts which are intentionally sacrificed to ensure majorities in other districts.

The authors of the second model conclude that existing models make simplifying assumptions with drastic implications for the conclusions they draw. The optimal strategy they derive from vertical slicing involves creating districts which match extreme Republicans and extreme Democrats rather than throwing away districts and then smoothing over the rest of the districts.

To illustrate how the optimal strategy works a numerical example of Friedman and Holden is presented. Assume that there are five districts and the gerrymanderer is Republican. The joint distribution of preferences and signals, $F(\beta, s)$ is multivariate normal with parameters $\mu_\beta = \mu_s = 0$ and covariance matrix

$$\Sigma = \begin{pmatrix} \sigma_\beta^2 & \rho\sigma_\beta\sigma_s \\ \rho\sigma_\beta\sigma_s & \sigma_s^2 \end{pmatrix}$$

In this case a distribution of $F(\beta, s)$ is assumed such that $\beta \sim N(0, 5)$ and $\rho = 0.5$. For simplicity, $\sigma_s = \rho\sigma_\beta$ which causes an uncomplicated form of $G(\beta|s)$. The distribution function of the breakpoint (B) is a standard normal distribution and $N = 5$. The assumptions imply that without gerrymandering both parties expect to win 2.5 seats as the voters are half Republicans and half Democrats.

Applying the procedure of Friedman and Holden the highest-median district (district 1) consists of 62% from a slice from the right tail of the distribution and 38% from the left tail. The upper slices become larger in the lower-median districts. Obviously the fifth and last district can only consist out of voters which are not part of any other district. The win probabilities for the 5 districts (from the perspective of the gerrymanderer) are 87.5% , 74.5%, 65.7%, 41.7% and 13.7%. We see that the probability of winning the highest median district is reasonably high and therefore the voters at the very left of the distribution (extreme distribution) are unlikely to be represented. It is important to note that no district is "thrown away", the gerrymandering party even has a 13% chance of winning the least favourable district (district 5). The number of districts won in expectation is now at 2.8 which is a 12% increase to the 2.5 expected seats without the gerrymandering strategy. [Friedman and Holden, 2005]

Another key result is that the expected number of districts won depends on the quality of the signal. As $\sigma_{\beta|s}^2$ increases, the probability of winning each district is lower. By plugging in $\sigma_{\beta|s}^2 = 4.5$ one can observe that the expected number of seats won drops to 2.54 which is barely more than 2.5. The expected number of districts won is a monotonic function in $\sigma_{\beta|s}^2$, the probability of winning each district is not. As the signal becomes more informative the gerrymanderer can cut the districts finer, but the probability of winning votes of those with a low signal (preference) decreases. These effects work in opposite directions which leads to the potential non-monotonicity of winning districts with low medians. [Friedman and Holden, 2005]

The authors also looked at the effects of varying σ_β^2 which causes the voter preferences to be more spread. That results in a monotonic increase of the probability of winning districts 1-4 as voter preference becomes more spread out, because fewer extreme voters are necessary to provide a solid margin of victory in expectation.

The last analysed aspect is how a change in the mean affects the expected number of districts won under this strategy i.e. if the number of Democrats/Republicans in the population changes. As the mean increases, the share of nominal Republicans also increases and vice versa. The expected number of seats won increases as the percentage of nominal Republicans increases. The value of being the gerrymandering party is characterized as the difference in expected seats won compared to proportional representation. This value decreases as the number of nominal Republicans increases. [Friedman and Holden, 2005]

Socially Optimal Districting: The seat-vote curve

One of the goals of this thesis is to establish methods for detecting partisan gerrymandering. A different look at the problem is yielded by looking at how an optimal districting could look like. This section is based in large parts on the concepts of [Coate and Knight, 2005].

In previous sections the consequences of different districtings were already outlined. By using seat-vote curves, which relate the fraction of seats parties obtain to their share of the popular vote, districting plans can be evaluated. This section focusses on their properties.

Seat-vote curves are denoted as $S(V)$ where V is the aggregate function of votes received by (for example) the Democratic Party. S is the fraction of seats the Democrats hold in the legislature. The two key properties of a seat-vote curve are its partisan bias and its responsiveness. Partisan bias is defined as $S(1/2) - 1/2$. Responsiveness is defined as $\Delta S/\Delta V$. A seat-vote curve is unbiased if a party reaches exactly half of the seats with half of the popular vote. Responsiveness describes the effect of small changes in the popular vote on the seat distribution. The consequences of bias of a seat-vote curve and the optimal level of responsiveness will be subject of further analysis.

Coate and Knight start by developing a micro-founded model providing a framework for the problem which is consistent with concerns in literature. Under the model districting determines the relationship between seats and votes and social welfare consequences depend on that. The ideal relationship between seats and votes shall be characterized as the optimal seat-vote curve. Another connected question is whether an optimal seat-vote curve is implementable in the sense that there exist districtings that generate this curve. Such districtings are called socially optimal districtings.

Assume now that there are three different types of voters: Democrats, Republicans and Independents where only Independents can change their voting behaviour from one election to another. They are represented by two parties: Democrats and Republicans. The relationship between seats and votes determines how responsive the legislature's policy choices are to swings in the popular vote share created by changes in the voting behaviour of Independents. The following model ignores geographical constraints, i.e. the planner can allocate citizens to districts in any desired way.

The model assumes a community divided into n equally sized districts where each district chooses its representative. The outcome depends on the average ideology of the district which is measured on a 0 to 1 scale. Democrats and Republicans have ideologies 0 and 1, Independents have ideologies uniformly distributed on the interval $[m - \tau, m + \tau]$ with $\tau > 0$. m is the realization of a random variable on the interval $[1/2 - \epsilon, 1/2 + \epsilon]$ where $\epsilon \in (0, \tau)$ and $\epsilon + \tau \leq 1/2$. This guarantees that the ideologies of Independents lie between those of Democrats and Republicans. Still some Independents lean Democrat and some Republican. The fraction of Democratic voters in district i is given by $\pi_D(i)$. $\pi_R(i)$ and $\pi_I(i)$ are defined analogously. Thus π_D, π_R and π_I give the fraction of voters in the entire community.

Elections are held simultaneously in all districts. It is assumed that the average ideology of a party representative (no matter if Democrat or Republican) is α' and a randomly chosen citizen has ideology α. The citizen has a quadratic loss function with payoff given by $-(\alpha - \alpha')^2$. Every citizen votes for the representative whose ideology is closest to his own. The fraction of Democratic voters in district i, in case the median independent has ideology m, can be specified as: [Coate and Knight, 2005]

$$V(i, m) = \pi_D(i) + \pi_I(i) \left[\frac{1/2 - (m - \tau)}{2\tau} \right]$$

The aggregate vote share of Democrats is

$$V(m) = \pi_D + \pi_I \left[\frac{1/2 - (m - \tau)}{2\tau} \right]$$

The maximum and minimum aggregate Democrat vote shares are denoted as $\overline{V}$ and $\underline{V}$ which are used to derive the relationship between seats and aggregate Democratic vote share. The group of Democratic voters consists of people who are Democrats and Independents with ideologies less than 1/2. The ideology of the median Independent that would generate the vote share V is denoted as $m(V)$. The equations from above enable us to specify

$$m(V) = 1/2 + \tau \left[\frac{\pi_I + 2\pi_D - 2V}{\pi_I} \right]$$

By substituting into the first equation one obtains

$$V(i; m(V)) = \pi_D(i) + \pi_I(i) \left[\frac{V - \pi_D}{\pi_I} \right]$$

Therefore district i elects a Democrat if $V(i; m(V)) \geq 1/2$ or, with $V^*(i)$ being the critical aggregate vote threshold above which the district i elects a Democrat, if

$$V \geq V^*(i) = \pi_D(i) + \pi_I(i) \left[\frac{1/2 - \pi_D(i)}{\pi_I(i)} \right]$$

A seat is called safe Democrat if $V^*(i) \leq \underline{V}$, otherwise it is competitive (or analogously safe Republican). The districts are reordered according to the Democratic vote share such that $V^*(1) \leq V^*(2) \leq \ldots V^*(n)$. The fraction of seats that Democrats win with aggregate vote share V is given by the equilibrium seat-vote curve

$$S(V) = \frac{\max\{i : V^*(i) \leq V\}}{n}$$

This seat-vote curve depends on the allocation of citizens across districts which determines the critical thresholds for each district. [Coate and Knight, 2005]

A neutral planner wants to allocate voters in a way such that the curve looks similar to the equilibrium seat-vote curve. This seat vote curve is determined by the pattern of critical vote thresholds across districts and can therefore be achieved by many different allocations of districts. Coate and Knight begin their analysis by looking at what the optimal relationship between seats and votes looks like if the constraint that there has to be an equilibrium for some districting is ignored.

The planner has to decide on the number of seats S that Democrats should receive after gaining vote share V. Aggregate utility when the median Independent has ideology m is given by:

$$W(S;m) = -\left[\pi_D(1-S)^2 + \pi_R S^2 + \pi_I \int_{m-\tau}^{m+\tau} (1-S-x)^2 \frac{dx}{2\tau}\right]$$

The optimal seat share is given by

$$S_{opt}(V) = \underset{S\in\{\frac{i}{n}\}}{\arg\max}\, W(S;m(V))$$

After solving the first order condition, Coate and Knight state the optimal seat-vote-curve as

$$S_{opt}(V) = 1/2 + (\pi_D - \pi_R)(1/2-\tau) + 2\tau(V-1/2)$$

This equation shows that the curve is linear with bias $(\pi_D - \pi_R)(1/2-\tau)$ and responsiveness 2τ.[Coate and Knight, 2005]

Since $\tau < 1/2$ from above, the slope of the optimal curve is less than 1. Therefore the fraction of Democrat seats increases at a constant, but less than proportional rate as the aggregate Democratic vote increases. The curve intersects with the 45 degrees line when exactly half of Independents leans Democratic and the vote share is $\pi_D + \pi_I/2$. Note that $S_{opt}(\underline{V}) > 0$ and $S_{opt}(\overline{V}) < 1$ which results in having a few safe seats for both parties.

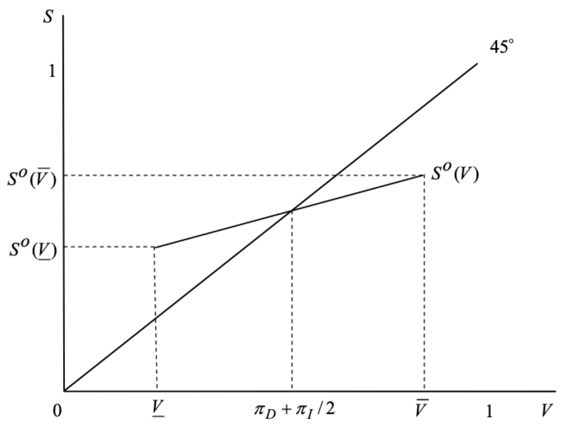

Figure 2: Optimal seat-vote curve [Coate and Knight, 2005]

The optimal responsiveness is at 2τ. One key aspect of responsiveness is that the welfare maximizing Democratic seat share must be such that the social gains from increasing it marginally just equal the social losses.

When the mean Independent has ideology m, the optimal Democratic seat share should be $\pi_D+\pi_I(1-m)$ because this would make the average ideology in the legislature equal to the population average, which is $\pi_R + \pi_I m$. The higher m is, the "more Republican" it is. When the Democrat vote share increases marginally, the change in the mean Independent's ideology is $dm/dV = -2\tau/\pi_I$ and hence the increase in the optimal Democrat seat share is just 2τ. [Coate and Knight, 2005].

Consider the case in which Democrats receive exactly half of the popular vote i.e. $V = 1/2$. If the curve was unbiased Democrats should receive exactly half of the seats i.e. $S_{opt}(1/2) = 1/2$. But that is not the case under the optimal seat-vote curve which is therefore biased. It is unbiased only if the median voter has ideology 1/2 and the average ideology equals 1/2 which is the case only when the fractions of Democrats and Republicans are equal.

When $V = 1/2$, the median Independent's ideology is $m(1/2) = 1/2 + \tau(\pi_D - \pi_R)/\pi_I$ which implies the average ideology in the population to be $1/2 + (\pi_R - \pi_D)(1/2 - \tau)$. The bias in the optimal seat-vote curve arises from the fact that the ideology of the median voter differs from the ideology of the average voter. So as the average legislator's ideology should be equal to the population average, the Democratic seat share is needed to be greater than 1/2 in case $\pi_D > \pi_R$. [Coate and Knight, 2005]

The important question attached to the optimal seat-vote curve is its implementability. The idea of creating districts as microcosms of the entire population is rejected immediately as in this case all districts tend to vote the same and therefore an all Democrat or all Republican legislature is the consequence. The fraction of Independents is an important factor for the implementability of the optimal seat-vote curve. If there were no Independents, the seat-vote curve would only be a single point and could be implemented by creating a fraction π_R districts with Republican majorities and the rest Democratic. The opposite would be to have Independents only without any Democratic or Republican voters. Then all districts would be identical and the optimal seat-vote is not implementable as $S_{opt}(V) = 1/2 + 2\tau(V - 1/2)$ while the equilibrium vote curve is $S(V) = 0$ if $V < 1/2$ and $S(V) = 1$ if $V > 1/2$. It is therefore necessary to find conditions under which the optimal-seat vote curve is implementable.

For that purpose Coate and Knight look at the inverse seat-vote curve which is described by a triple $(\underline{i}, \bar{i}, V^*(.))$, where $V^*(i)$ is defined as the critical value in the competitive district i and $\underline{i}$ and $\bar{i}$ are scalars representing the fraction of districts which are safe Democrat, or respectively safe Republican. The inverse seat-vote curve is formed by setting $\underline{i}$ as $S(\underline{V})$, while $\bar{i}$ is $S(\overline{V})$ and for all $i \in [\underline{i}, \bar{i}]$ it is such that $S(V) = i$.

Checking the implementability is based upon a minimization problem for finding the districting with the minimal fraction of Democratic voters and a maximization problem choosing the districting with the maximal fraction of Democratic voters (hence, minimal fraction of Republicans). These values are defined as $\underline{\Omega}(\underline{i}, \bar{i}, V^*(i))$ and $\overline{\Omega}(\underline{i}, \bar{i}, V^*(i))$. To see whether the optimal seat-vote curve $S_{opt}(V)$ is implementable one starts by determining the associated inverse seat-vote curve. The values of the associated minimization and maximization problems $\overline{\Omega}$ and $\underline{\Omega}$ are computed to compare them with the actual fraction of Democrats π_D. The following result is established: [Coate and Knight, 2005]

The optimal seat-vote curve is implementable if and only if

$$\pi_I(\frac{\epsilon}{2\tau} + \epsilon - (\tau + \epsilon)\ln(1 + \frac{\epsilon}{\tau})) \leq \pi_D$$

and

$$\pi_I(\frac{\epsilon}{2\tau} + \epsilon - (\tau + \epsilon)\ln(1 + \frac{\epsilon}{\tau})) \leq 1 - \pi_D - \pi_I = \pi_R$$

It states that enough Democrats and Republicans are needed compared to the number of Independent voters. Note that the coefficient with which π_I is multiplied on the left side of both inequalities converges to 0 as ϵ converges to zero for all values of τ. The optimal seat-vote curve is implementable when the degree of uncertainty in the identity of the median Independent is sufficiently small. [Coate and Knight, 2005]

For a given value of ϵ, the entire coefficient with which π_I is multiplied on the left side is decreasing in τ which the authors interpret as the curve being more likely to be implementable when there is more diversity in the ideologies of Independents.
For any values of ϵ and τ satisfying the assumptions, the coefficient is less than 1/2 and therefore a useful sufficient condition for the optimal seat-vote curve to be implementable can be introduced:

$$\pi_I \leq 2\min\{\pi_D, \pi_R\}$$

In practice it is hard to classify voters as pure Democrats, Independents or Republicans beforehand. Therefore another component is added by classifying "Democrat-leaning" districts which are not classified as safe seats, but are likely to vote Democratic. In the model the competitive districts are divided into Democrat-leaning districts and Republican-leaning districts. Democrat-leaning districts are only populated by Democrats and Independents, where the fraction of Independents is varying from $\tau/(\tau+\epsilon)$ to 1. Consequently the Democrats do not necessarily win this district, but they win it in case the majority of Independents prefers the Democrats ($V \geq \pi_D + \frac{\pi_I}{2}$). These districts differ in their critical vote thresholds as they contain varying fractions of Independent voters and the fraction of these districts electing Democrats varies smoothly as the aggregate Democrat vote share increases from $\underline{V}$ to $\pi_D + \frac{\pi_I}{2}$ and vice versa in Republican leaning districts. This property is exploited for the optimal districting. [Coate and Knight, 2005]

To summarize there is no particular interest in districting safe seats, but whenever one of the two above conditions is satisfied, there is a unique districting that generates the optimal seat-vote curve. One straightforward case is when the fraction of Independents is constant across all districts and only the fraction of Democrats and Republicans varies. A lack of Democrats or Republicans in a district causes problems for the implementability of the optimal seat-vote curve. The authors provide a different solution which is called the constrained optimal seat-vote curve for cases in which the fraction of Independents is too large. While in some states in New England the fraction of Republicans is too small, these restrictions are reversed in other states. The constrained optimal seat-vote curve is found by "solving for the implementable inverse seat-vote curve which maximizes aggregate welfare where implementable means that there exists a feasible districting that generates such an inverse seat-vote curve." [Coate and Knight, 2005]

Denote F^{-1} as the set of all inverse seat-vote curves $\{\underline{i}, \bar{i}, V^*()\}$ with the property that $V^*()$ is piecewise continuously differentiable. Let $EW(\{\underline{i}, \bar{i}, V^*()\})$ denote expected aggregate utility under the specified inverse seat-vote curve. The constrained optimal seat-vote curve is corresponding to the solution of the problem

$$\max_{\{\underline{i},\bar{i},V^*()\}\in F^{-1}} EW(\{\underline{i}, \bar{i}, V^*()\})$$

s.t.

$$\overline{\Omega}(\{\underline{i}, \bar{i}, V^*()\}) \geq \pi_D \geq \underline{\Omega}(\{\underline{i}, \bar{i}, V^*()\})$$

The optimal seat-vote curve solves this problem if the two conditions stated above are satisfied. In case there are not enough Democrats or Republicans the constrained optimal seat-vote curve can still be

implemented. When either Democrats or Republicans are in short supply at least some fraction of them are optimally concentrated to construct safe seats for their party. [Coate and Knight, 2005]
Whenever there is a shortage of only one group of partisan voters, the constrained optimal seat-vote curve is biased towards the party with the larger partisan voter base, but this is not the case when there is a shortage of both partisan voters. A big difference between the constrained curve and the optimal curve is the responsiveness. In case of the constrained optimal seat-vote curve it can vary from 0 to infinity as the notion of responsiveness is not meaningful for the constrained optimal seat-vote curve which serves a different purpose.

Majority-Minority Districts

Another discussion in the context of gerrymandering is a controversial, but important aspect of redistricting: the purpose and effects of majority-minority districts (MMDs). It is of enormous interest to analyse whether the creation of such districts leads to more minority candidates being elected to congress. The section is based on the paper by Grigg and Katz.

It is by construction that in the system of single member districts, it is rare that a minority group is able to select a candidate of its choice without a large concentration of minority voters in a district. The principle of majority-minority districts is to concentrate minority voters in a single district to increase the probability of a minority candidate being elected. The question related is whether the creation of MMDs leads to pro-Republican gerrymandering as minorities tend to vote for Democratic candidates. As MMDs increase the possibility of a minority candidate being elected in a particular district, they might limit the influence of minorities in other districts. This is referred to as "perverse-effect claim". Originally, most of MMDs were created because of constraints imposed by the US Department of Justice to increase the opportunity for minorities of electing a candidate of their choice. The model which Grigg and Katz construct assumes being the party in charge of the redistricting process with the constraint of a Federal court that MMDs shall be created. Their model is compared to the model constructed by Shotts.

The seat share a party wins is denoted as $s(v, \lambda, \rho)$ and is a function of the partisan bias λ, the responsiveness ρ and the average vote share of the party v. Another assumption is that the way partisan bias and responsiveness influence the translation of votes to seats is given by

$$\left(\frac{s}{1-s}\right) = e^{\lambda}\left(\frac{v}{1-v}\right)^{\rho}$$

The responsiveness of the districting plan is measured relative to $\rho = 1$. When $\rho = 1$, both parties receive seats equal to their vote shares (as long as partisan bias $\lambda = 0$). In contrast, if $\rho < 1$, the corresponding seat share is greater than the vote share for the party winning the smaller share of votes. [Grigg and Katz, 2005]

The gerrymandering party is assumed to have some belief about the voting behaviour which is given by a cumulative probability distribution F where $F(v)$ represents the probability that the stronger party receives a vote share less than or equal to v. A positive value for λ represents a bias in favour of the stronger party which is why the stronger party usually profits from bias. An interesting aspect in this model is to which extent a party tolerates bias compared to responsiveness which depends on the expected vote share of the two parties. The more confident the stronger party is that it wins a larger part of the vote share, the more willing it is to give up bias in favour of more responsiveness. Therefore the party faces a trade-off between the level of bias and the level of responsiveness.

Constraints for the construction of MMDs are added to the model. The strong party selects the parameters in a way such that $\mathbb{E}[u(s(v;\rho,\lambda))]$ is maximized, where $s(x;\rho,\lambda)$ can be derived from above as

$$s(x;\rho,\lambda) = \frac{\exp[\rho x]}{\exp[-\lambda] + \exp[\rho x]}$$

The strong party's maximization problem can be expressed as [Grigg and Katz, 2005]

$$\max_{(\rho,\lambda)\in A} \int_{-\infty}^{\infty} u\left[\frac{\exp[\rho x]}{\exp[-\lambda] + \exp[\rho x]}\right] f(x)dx$$

where f is the density representing x whereas $x = \ln(\frac{v}{1-v})$.

The restrictions are imposed in form of the constraint set A. Suppose the legislature has to create at least some mandated level of MMDs and is in control of the Democrats. Democrats might be forced to select a suboptimal plan and the optimal constrained plan from their perspective necessarily contains higher bias and lower responsiveness levels than the original optimal plan.[Grigg and Katz, 2005]
If the perverse-effect claim was true, Grigg and Katz conjecture that the MMD constraint is binding for Democratic plans.

The constraint enters the maximization problem in two ways: it imposes an upper bound on the level of responsiveness and a lower bound on the level of democratic bias the plan can exhibit. Therefore the stronger party wants to maximize

$$\mathbb{E}[u(s(x;\rho,\lambda))]$$

s.t.

$$(\rho,\lambda) \in D = A \cap \{(\rho,\lambda) : \rho^{MMD} - \rho \geq 0; \lambda - \lambda^{MMD} \geq 0\}$$

The problem can be solved by forming the Lagrangian dual function and applying the Kuhn-Tucker theorem. The set M_{MMD} is obtained which denotes the set of all critical points under the restrictions. It is compared to the set of solutions of the unconstrained maximization problem to examine the impact of the constraints. This comparison is done empirically by estimating $\hat{\rho}$ and $\hat{\lambda}$ for plans with MMDs and without MMDs and testing for significant differences.

Grigg and Katz analysed election data to complete the empirical analysis. By plotting the average Democratic vote share against the seat share for each state-year pair they obtain an empirical seat-vote curve which is shown below.

If the perverse-effect claim was true, a rightward shift in the curve for observations with MMDs indicating a pro-Republican bias is expected. [Grigg and Katz, 2005]
In general the author's note that in their analysis there was no significant difference in the estimates for the bias. They claim this finding to be in contrast to the claim of MMDs generating a significant pro-Republican bias in the redistricting plan.

The estimates of responsiveness were all lower for the plans including MMDs. This result is explained by the more general effect that Democratic responsiveness is raised by plans which generate districts that are little microcosms of the state which is the exact opposite of what occurs when creating MMDs. This

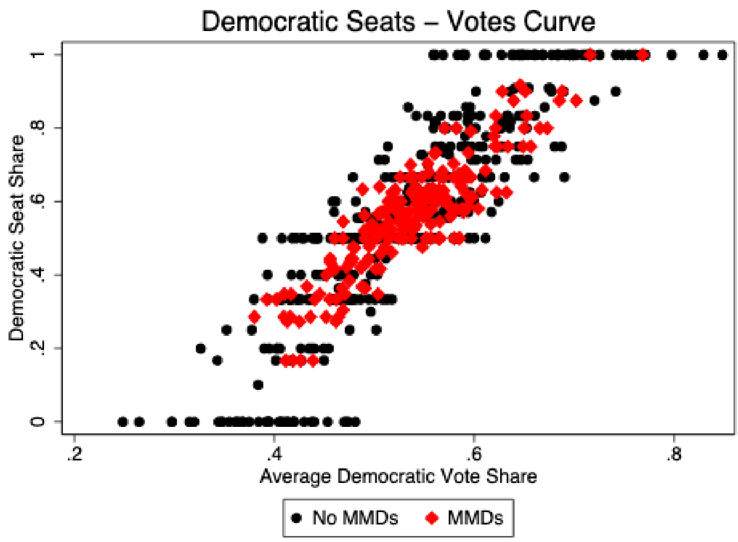

Figure 3: Empirical seat vote curve [Grigg and Katz, 2005]

impact is not significant as well.

A game-theoretical approach towards MMDs

Shotts also analyses the perverse-effect which he characterizes as the concentration of liberal voters into a few districts causing the remaining districts to elect conservatives who do not represent political issues which matter to minority voters. A new model with the goal to predict the influence of gerrymandering on policy consequences is introduced. [Shotts, 2002]
It is based on three steps: gerrymandering, election and policy choices.

Gerrymandering is controlled by a single actor in each state. Liberal gerrymanders elect as many liberals as possible while conservative gerrymanders do the same for conservative candidates. The gerrymanderer faces two constraints: he does not know the policy preferences of individual voters and he has geographical constraints i.e. there are neighbourhoods containing voters with a variety of different policy preferences which cannot be split into different districts. Hence, it is impossible to concentrate all voters with a particular preference into a single district or draw a district without voters with a particular preference. Those constraints are represented by a parameter limiting the actor's ability to split voters with different preferences into different districts.

The model of Shotts assumes S states with D_s districts in state s. The total number of legislators is $N = \sum_{s=1}^{S} D_s$ which is assumed to be odd. Voters in state s have policy preferences given by f_s. The S acting gerrymanderers also have policy preferences on a scale between 0 and 1 which is denoted G_s. A strategy in state s is given by a vector of functions allocating voters to districts: $\gamma_s = (\gamma_{s_1}, \ldots \gamma_{s_{D_s}})$. This function (allocation) has to fulfil three constraints: [Shotts, 2002]

1. An equal number of voters has to be allocated to each of the D_s districts in state s, i.e.

$$\int_0^1 \gamma_{s_d}(x)dx = \frac{1}{D_s} \forall s_d$$

2. Each voter has to be part of exactly one district, i.e. $\sum_{d=1}^{D_s} \gamma_{s_d}(x) = f_s(x) \forall$ ideal points $x \in [0, 1]$;

x is a preference parameter.

3. The minimum density constraint which prevents a gerrymanderer from drawing a district that contains no voters with a particular policy preference, i.e.
$\gamma_{s_d}(x) \geq \delta f_s(x), \forall x \in [0,1]$ where the minimum density parameter $\delta \in \left(0, \frac{1}{2D_s}\right)$

The policy outcome is determined in three steps. First, the gerrymanderer decides to pick an allocation γ_s. Each district elects a legislator with an "ideal point" L_{s_d} at the district median. The national policy outcome is the median of the elected legislators ideal points i.e. $x = median\{L_{s_d}\}$.

Shotts states that the "game" has a unique strong Nash equilibrium, but although it is unique there are various gerrymandering plans consistent with that equilibrium. Voters to the left of some point $x_M \in (0,1)$ are assumed to be of type M(minority voters) and they are assumed to vote for Democrats in this model. To ensure their proper representation each gerrymanderer must create a certain number of districts in which at least half of the voters are to the left of x_M. M_s is the number of MMDs a gerrymanderer has to create in state s.

Shotts states that even under this constrained setting there exists a unique equilibrium policy outcome. In addition to that the policy outcome with the constraint is at least as liberal (left) as the unconstrained outcome, as long as geographical and informational constraints are minimal. This statement contradicts the perverse-effect claim. The model predicts that the addition of MMD constraints moves policy outcomes to the left. According to Shotts this follows from the fact that MMD-mandates have an asymmetric effect on the gerrymandering. They do not hinder a liberal gerrymanderer as much as they do with conservative gerrymanderers. A few examples will illustrate this effect.
Assume a state with three districts is assumed of which one has to be a MMD. Voter's preferences are uniformly distributed, the cutpoint is $x_M = 0.3$ and the geographical constraint $\delta = 0.01$ [Shotts, 2002].

- **Liberal Gerrymanderer without constraints:** The gerrymanderer wants to elect as many representatives as possible to the left of 0.45 (assumed to be the national legislative median). It is obvious that liberals cannot win all three seats (median vote 0.5), but by gerrymandering they can elect two representatives to the left of 0.45 and for that scenario we have many different allocations. One such scenario is to create two identical liberal districts which elect the liberal candidate at 0.417, while the third district is conservative and elects a representative at 0.623. That represents the classical gerrymandering strategy from above, where large numbers of opposite voters are moved into district 3 while the rest of the voters is divided among the remaining districts.

- **Liberal Gerrymanderer with constraints:** The strategy from above has to be adapted as there was no district with a median left to $x_M = 0.3$. Nevertheless, the strategy can be adopted easily. District 1 is the MMD that elects a representative at 0.278, but the gerrymanderer constructs it in a way that it contains only few non-minority liberal voters, because they have to be used in the second district. District 2 elects its representative with a median of 0.41 while district 3 is more conservative than before and elects a representative at 0.681. The gerrymanderer's goal to elect liberal representatives is more or less unaffected by the constraints.

- **Conservative Gerrymanderer without constraints:** In this scenario the gerrymanderer creates as many districts with a median greater than 0.45 as he is able to. It is possible to create all three districts as such, by mirroring the composition as a whole in each district. Representatives would therefore be elected at 0.5 in all three districts.

- **Conservative Gerrymanderer with constraints:** The strategy from above does not work here. The gerrymanderer is in a worse situation as there is no scenario in which all three districts elect their representative at a higher number than 0.45 as one district has to be the MMD in which the median is to the left of 0.3. The districting is adopted such that two conservatives are elected with ideal points at 0.654, but one liberal is elected at 0.239. Because of the constraint, the conservatives lose one seat.

Another special case to illustrate how the statement from above (policy under constraint stays the same or moves to the left) affects the process is presented. Assume $S = 1$, which corresponds to redistricting for a single state's legislature and consider a liberal gerrymanderer. With the statement from above it follows that as $\delta \to 0$ the policy stays the same or moves to the left.
Now consider a conservative gerrymanderer. The author could show that policy only moves strictly to the left if more than half of the districts have to be MMDs which is unrealistic in practice. The creation of a few liberal districts does not cause any difference as the gerrymanderer can still ensure a majority of districts to be conservative. But in this scenario it makes a difference on the national level as the loss of a few districts can have national implications. [Shotts, 2002]

The presented model therefore does not provide any evidence for the perverse-effect claim. Conversely, it even states that a majority-minority mandate moves policy outcomes to the left, because it does not constrain liberal gerrymanderers. There are still additional scenarios which are not covered by this model which do not allow us to categorically rule out any perverse-effect claim. It could happen that minority policy views are more complex and cannot be mapped on a scale from liberal to conservative. The constraint may also be formulated differently. If gerrymanderers are required to create districts with a clear majority of minority voters, then democrats couldn't avoid "wasting" votes in these districts any more. Despite such limitations, the model provides strong reasons to question the theoretical logic behind the perverse-effect claim. [Shotts, 2002]

2.5 Limitations towards Gerrymandering

So far most of the models assumed the gerrymanderer being able to divide districts in any way possible. In this section some limitations of gerrymandering are discussed.

The first and biggest constraint is the geographical constraint which was already mentioned briefly at the end of the previous section. As Owen and Grofman state that it would theoretically be optimal to concentrate opposite voters in certain districts, this is not possible in practice. The authors state that "it may require drawing district lines that pass through the middle of a bedroom to separate a Republican husband from his Democratic wife". [Owen and Grofman, 1988]
Nevertheless they argue that districtings close to the optimal one are possible.
On the other hand Seabrook argues that the effects of gerrymandering are over-exaggerated as there are too many constraints to fulfil in practice. They mention legal and political constraints as in many US states the redistricting process does not lie in a single hand as assumed in previous models.[Seabrook, 2010]
Another constraint is the forced creation of MMDs which also limits the flexibility of the gerrymanderer.

According to Owen and Grofman another limitation is that there might be conflicts between the individual legislator's motivation and overall partisan advantage. The argument is that the group of people in charge of the redistricting process is rarely guided by a pure desire to improve the party's political situation. Hardy argues that districts are created with only four different interests in mind: individual preservation, mutual preservation between incumbents, preservation of political power by the majority party and the preservation of blocs such as an rural bloc. Therefore the usual view on gerrymandering might be a “gross simplification” as it is only seen as an action by one party to bolster the majority. Purely implementing the party’s strategy requires a strong political party organization as legislators are acting out of individual interest. [Hardy, 1977]

The motive of mutual agreement between incumbents is of particular interest. It is easy to imagine a situation in which legislators of both parties profit from cooperating and decide to create safe districts for incumbents. That claim is supported by the empirical analysis of Friedman and Holden who notice an upward trend in the reelection rate, as for example in the 2004 congressional elections more than 97.9% of incumbents were re-elected. The authors link the rise in the reelection rate to better techniques for imposing a gerrymander. [Friedman and Holden, 2009]
As noted in the papers by Owen and Grofman and Hardy the strategy of spreading own voters across districts while concentrating voters of the opposite party would be optimal for party interests, incumbents favour concentration gerrymanders which is due to the fact that their reelection chances are augmented.

There exist even more extreme versions than gerrymandering in favour of the incumbent which Owen and Grofman call “personalized gerrymandering” where district lines are drawn to help or hinder the reelection chances of a single personality. Those side effects of partisan redistricting cannot be denied, but still they are not a phenomenon of high interest for research and this thesis. The main constraints to keep in mind are the geographical constraint and that partisan goals are not as easy to impose on individual state legislators.

2.6 Measuring Gerrymandering

Detecting intentional gerrymandering in favour of one political party is not a trivial task. Owen and Grofman argue that the best way of performing the gerrymander is by concentrating opponent’s voters while splitting the voters of the gerrymanderer’s party among the remaining districts. One approach to detect such strategies is to look for districts with a large gap between the percentages of the two parties. But that is not a good indicator as such districts are often constructed “naturally” as rural districts tend to vote for Republicans in large numbers while urban areas vote for Democratic candidates by large margins. Therefore measurements of the translation from votes to seats are needed. The following analysis is based on the paper by Best et al. and considers five different ways of detecting partisan or racial gerrymandering.

Efficiency gap

The first measure is the so-called “efficiency gap” which was introduced by Stephanopoulos and McGhee. The idea behind is to measure “wasted votes” which are votes not necessarily required to win a certain district. A vote is wasted if it is cast for a loosing candidate or for a winning candidate in excess of the votes the candidate needed to win the district. The measurement works such that it checks the results of an election and then gives an indication of whether gerrymandering had been present beforehand. [Stephanopoulos and McGhee, 2015]

Stephanopoulos and McGhee give an example for how the efficiency gap works. Suppose a state has 10 districts where each district contains 100 voters. Party A wins 55 % of the total vote in the state while it wins around 70 votes in districts 1-3, 54 votes in districts 4-8 and 35 votes in districts 9-10. Party A wastes 20 votes in districts 1-3, 4 votes in districts 4-8 and 35 votes in districts 9 and 10. On the other hand Party B wastes 30 votes in the first three districts, 46 in districts 4-8 and 15 in districts 9 and 10. In total Party A wastes 150 votes and Party B wastes 350 votes.

The difference between the two quantities is 200 which is divided by the total number of votes cast and that results in an efficiency gap of 20%. A simple interpretation of the efficiency gap is that Party A wins 20 percent more seats than it would have won if both parties had wasted the same number of votes. The case in which the efficiency gap is zero is seen as ideal.

The definition is criticized by Best et al.. They spot three key problems when talking about the efficiency gap. The first is that it does not work in one district states. Obviously there can't be any gerrymandering, but still the desirable case of the efficiency gap being zero is only reached in case the result is 75-25% . One may conclude that it only works in multi-district states, but it also does not fulfil its purpose in a three district state with the majority party receiving 48, 52 and 56% as there is a 8.3% efficiency gap in favour of the majority party. Stephanopoulos and McGhee suggested an efficiency gap of 8% as being the barrier after which redistricting can be considered as gerrymandering. Suppose there is an overall 2% shift in the voting behaviour of all three districts such that they are at 46, 50 and 54% then. Out of a sudden there is no more efficiency gap at all although the districting stayed the same. Another uniform shift of 2% in all districts such that the results are at 44, 48 and 52% causes the efficiency gap to indicate the presence of gerrymandering once again, but in favour of the other party. Best et al. state that "reading a gerrymander from the efficiency gap can and often will vary depending on the underlying percentage level of the votes a party receives". [Best et al., 2018]

The third problem they see with the efficiency gap is related to the translation of votes to seats. Suppose a party wins 60% of the seats while winning 55% of the votes. While the first value is at 10% above 50%, the second value is only 5% above 50. This situation is called a two-to-one seat-vote ratio. The efficiency gap indicates that there is no gerrymandering taking place in such a situation while the authors argue that this is not necessarily the case.

Consider the situation of a 40-40-60-65-70 voting result which is asymmetrical as the mean does not equal the median (if equal turnout in all districts is assumed), yet the efficiency gap is 0 as the amount of wasted votes of the two parties is equal. So the efficiency gap does not indicate gerrymandering even though the majoritarian ratio is at two to one and the vote distribution is asymmetrical. Best et al. claim that " despite its proponents' claims to the contrary, the efficiency gap standard does not comport with nor arise from the idea of partisan symmetry." [Best et al., 2018]

They argue that the efficiency gap is mainly attractive because it is intuitive and easy to explain, but they criticize the definition of wasted votes.
Another disadvantage of the efficiency gap is that one has to use historical data to determine until which threshold efficiency gaps can be considered as normal. There is no concrete definition as the considered threshold of an efficiency gap of 8% is hard to argue for although Stephanopoulos and McGhee make their case using historical data.

The efficiency gap is a simple to calculate and easy to interpret measure. But by performing closer analyses it is easy to construct examples in which the efficiency gap does not fulfil its purpose. It is therefore not the ideal measure for the detection of gerrymandering.

Comparing number of districts carried

One alternative is to look at a test comparing seats won to neutral expectations which was introduced by Chen et al.. The idea behind is to perform various simulations on how district lines could be drawn. By doing so, an estimate for how many seats a party is expected to win is obtained and the so-called "measure of comparing wins" indicates that gerrymandering may be present if one party wins fewer seats than expected in district plans produced through partisan blind line-drawings. [Best et al., 2018]
These line drawings are viewed as the null-set and large deviations from it are considered as gerrymandering. Still, a variety of problems arise with that method.

It is hard to trust the result of the simulations done by a black-box algorithm in that the results are not explainable. That would not necessarily be a problem if the procedure was used for research only, but in reality the method should also be applicable by courts to check districtings on whether partisan or racial gerrymandering had been present and that seems unrealistic if the analysis is based on a black box simulation. Best et al. argue that the procedure is not yet established and authors still use different evaluation methods in determining whether gerrymandering is present or not. One idea is to analyse the number of competitive districts while others evaluate the number of districts in which each racial group or political party holds a majority.

Another factor causing trouble when comparing the number of districts won is that the number of expected districts won depends on the overall vote percentage a party obtains, making no difference if competitive districts or majority held districts are measured. It can also be stated that a match or mismatch between the observed and expected number of districts carried is not a robust and structural feature of a district plan. [Best et al., 2018]
Suppose a party wins 40% of the seats while a packing gerrymander is in place and the overall vote for that party is in the first case at 40%, in the second case at 50% and in the third case at 60%. The result of the gerrymandering sometimes matches the expected number of districts and in other scenarios it does not.
As a consequence of that the number of seats won does not depend solely on the districting plan, but on the system of party support as well. This interaction complicates the definition of what gerrymandering is under this approach. Best et al. conclude that the approach of forming a null set which represents a neutral districting is promising, but there is no good way of setting the benchmark. It therefore has to be figured out how this method can be made more transparent and efficient.

Equal vote weight test

Another approach is the equal vote weight principle which was introduced by McDonald and Best. The principle focuses on detecting unequal vote weights as a result of gerrymandering and tries to distinguish between gerrymandering as a political fact or as a "legally significant constitutional offence." [McDonald and Best, 2015]

The procedure relies on two observable facts: it compares the median district vote percentage of a party to the mean district vote percentage of the party and checks whether the majority rule (defined below) is violated. The principle of equal vote weights can be violated in case one group of partisan supporters is more packed than the other. The procedure relies on three ideas: [Best et al., 2018]

1. Leading indicator: Asymmetrical packing is present when the median vote percentage for one party is persistently lower than the mean district vote percentage.

2. Objectionable harm: The clearest way of identifying unequal vote weight is in case a majority casts for one party, but the same party wins less than the majority of the districts (majority rule violation)
3. Cause: As district line drawings are the main reason for unequal vote weight, the votes counted in the entire system (state) are compared to the votes counted after the division into districts. System-wide counted votes are by definition of equal weight, but if the two forms of counting yield different results, those results may be caused by the district line placements.

Best et al. consider this method to have many advantages, but also some disadvantages. Comparing the median and the mean district percentages and checking for violations of the majority rule is a simple underlying principle. But this method is not as aggressive as one might wish for the detection of systemic gerrymandering and consequently its effectiveness is questioned. Notice that equal median and mean of percentages only indicate average symmetry, but not full-scale symmetry which will be required when talking about partisan symmetry. [Best et al., 2018]

An example by Best et al. considers a five district plan where the Republican Party wins 44, 46, 51, 52 and 62% of the vote within the districts. Mean and median percentage are both at 51% which is symmetric according to the principles defined above. The majority rule is preserved as the Republican Party wins more votes and 3 out of 5 districts. A problem arises in case vote swings occur. A uniform upward shift for the Republican Party results in a 45, 47, 52, 53, 63% distribution where Republicans win 3 seats with 52% of the vote. But a downward shift of three points results in a 41, 43, 48, 49, 59% distribution in which the Democrats win 52% of the overall vote, but with that result they are able to win 4 seats.
Obviously this example is constructed to illustrate this weakness, but the underlying problem is obvious. The majority rule is not violated in any of these examples and therefore the districting plan does not raise any suspicion regarding unequal vote weights. On the contrary this "idea of vote inequality is not as aggressive as it might be in the sense that different rewards (seats) can be acquired from the same resources (votes)." [Best et al., 2018]
The procedure does not account for situations in which vote shifts produce different seat outcomes while both parties are winning the same vote percentage.

Partisan symmetry

Another approach is to analyse the degree of partisan symmetry. The formulation used in this section goes back to Grofman and King. The idea is to fix the problem of the equal votes approach above. The goal is to construct partisan symmetry on the basis of a fair seat-vote translation. [Grofman and King, 2007]
The approach is called "seat-denominated symmetry standard" by Best et al. by which they mean that each party is expected to win the same seat percentage for the same vote percentage. To make it more concrete, suppose that Democrats win 35 of 50 seats (70%) with 55% of the vote. Under this definition of partisan symmetry Republicans must be able to win 70% of the seats by winning 55% of the vote as well.[Best et al., 2018]

This ensures that the majority rule is satisfied, but it adds another component to the measurement. If both parties gain half of the overall vote, they also split the seats 50:50 and even if Democrats win 5 more seats by just winning 53% of the vote, then the same has to be true for Republicans winning 53%.

The procedure shows similarities with the wasted votes approach.

In the example where the equal vote principle failed, the mean and median of the district percentages were equal, but a uniform 3% swing in favour of the Republicans left them with three districts won while a swing in favour of the Democratic Party left them with four additional seats. This situation is resolved with this approach as partisan symmetry is violated in that case. However, swings may also be non-uniform in a way that some districts swing more than others and in that case the procedure doesn't help any longer as more information is needed to determine whether this is due to partisan gerrymandering or not. Best et al. conclude that this approach is more comprehensive than the equal vote standard, but can "under reach in practice by requiring supporting analysis that makes some decision makers wary of relying on it".[Best et al., 2018]
The reason therefore is that the assumption of uniform swings is not easy to evaluate and also computationally intensive. The approach also shows weaknesses in its application.

Three prongs

The last method presented in this section is the so-called "Three prongs" method. It was introduced by Wang and uses the idea of combining multiple criteria. The original paper proposes three tests for assessing asymmetry in districting schemes. The first test assesses whether there is an unrepresentative distortion in the number of seats won compared to an expected number from nationwide districts. The second test checks for discrepancies in the winning vote between the two parties while the third test checks for the construction of reliable wins for the party in charge of the gerrymandering process which is measured either by comparing median and mean district vote share or by checking for an unusually even distribution of votes across districts. [Wang, 2016]

Best et al. summarise the framework by formulating that an excess seat test, a lopsided outcome test and a reliable wins test is yielded: [Best et al., 2018]

1. Excess seat test: Checks whether the seat to vote responsiveness is within a range between its proportionality and what would be expected under a seat-vote relationship which was estimated from other states.
2. Lopsided outcomes test: Compare each party's average winning margin above 50% to check for unequal average lopsidedness.
3. Reliable wins test: Check whether median and mean average district percentage are equal if the district is in a competitive jurisdiction. If that is not the case the dominant party's standard deviation of the vote percentages shall equal the standard deviation obtained by simulations for the party's vote based on other jurisdictions.

This set-up appears to be a more comprehensive approach which combines the measurements used earlier. However, it cannot be neglected that this approach inherits many disadvantages of other measurements. Similar to the problem of the efficiency gap, the three prongs method relies on past election results which are used for comparison. While Wang argues that it even gives a good indication to perform only one of the three tests (instead of all three which is the original idea), this raises the question of how to interpret cases in which 2 tests indicate no gerrymander, but one does so or vice versa.

As Best et al. point out, the three prongs can contradict each other and they construct a simple example to demonstrate this problem. Assume a five district state with a party's vote percentage distribution given by 40, 40, 60, 60 and 60% of the vote. This districting plan satisfies the first two prongs (proportionality and equal average lopsidedness), but the median vote percentage is not equal to the mean percentage (mean is at 52, median at 60) which is required for the symmetry standard of prong three. They also mention that a swing ratio could possible lie within the bounds of proportionality, but fail when it comes to lopsidedness and symmetry.

If only one prong needs to be fulfilled for neglecting the claim of gerrymandering, then a gerrymanderer could modify its districting plan such that it satisfies one of the three criteria. And even though the evaluation fo gerrymandering using three different tests seems to be strong and powerful, Best et al. see a lack of a coherent framework and state that no coordinating principle supplies clarity about whether a gerrymander exists according to any or all three prongs. [Best et al., 2018]

Other methods

Another approach of determining whether gerrymandering is present in a districting is by looking at the geographical design of the map. According to Cervas and Grofman the number of county splits can be used to check for gerrymandering as many state constitutions already incorporate to avoid county splits which is one way of preventing gerrymandering. While some county splits are necessary to bring electoral districts closer to their desired population sizes, other county splits may be due to gerrymandering. The selection of counties getting split and the way they are split can have drastic impacts on the election outcome. The Pennsylvania Supreme Court found that an unusually large number of counties were split in the redistricting constructing the 2011 congressional map and found this fact to be deeply troubling. [Cervas and Grofman, 2019]

The second type of geographical indicator is a measure of compactness. By looking at the geometrical shape of legislative districts and comparing it to geographical constraints, one can detect anomalies as well. It was one of the first measures for trying to detect gerrymandering, but districtings appearing "compact" at first glance, turned out not to be compact from another perspective. One measure used to calculate compactness is the so called Polsby-Popper measure which examines the area of the district by comparing it to the area of a circle with the same perimeter. [Polsby and Popper, 1991]
An alternative, elder way of calculating compactness is the Reock measure which compares the area of the district with that of a circle that is circumscribing the district. [Reock, 1961]

Another measure for asymmetry is partisan bias which is closely connected to the idea of partisan symmetry discussed above. It indicates bias in favour of one party in the seat-vote translation and is measured at the point where both parties reach 50% of the popular vote and therefore should win half of the seats. A partisan bias of zero still does not imply proportional representation. [Cervas and Grofman, 2019]
It only implies that both parties are treated equally. In case Democrats can win 70% of the seats with 52% of the vote, this doesn't indicate partisan bias as long as the same is possible for Republicans.

A further measurement is the so-called declination. It is a more recent measure and was introduced to a larger audience by Warrington. It uses angles which are created by ordering all districts by vote share and computing the mean vote-share for each party separately for the seats they won.[Warrington, 2018]
After doing so one compares for differences in the distributions.

Intuitively a line is drawn from the mean vote share to the 0.5 line for both parties. The angle formed by the different slopes of the line is the declination. If no gerrymandering is present, it is assumed that vote-shares are distributed uniformly across districts, but the angle is greater if a packing gerrymandering strategy is present. [Cervas and Grofman, 2019]

2.7 Discussion

In this section different aspects of the broad topic of gerrymandering were analysed. Different strategies of performing gerrymandering were presented with the result that in case only obvious assumptions are made, the optimal strategy is to concentrate voters of the opposite party in a number of districts which are "sacrificed". Supporters of the party in charge of the gerrymandering are split among the remaining districts. There is a difference in whether the number of seats won shall be maximized or the probability of obtaining a majority. [Owen and Grofman, 1988]

By looking at a different model by Friedman and Holden which works with a signal distribution, one obtains a different optimal strategy as the assumptions were different. The signal distribution is chopped into vertical slices and it turns out to be optimal to match extreme Democratic voters with extreme Republican voters. The result of this model differs in one key point from the first one in that there are no more throwing-away districts. It is assumed that the voting preference cannot be observed directly, but some signal for party identification is observable. The quality of the gerrymandering depends on the quality of this signal.

By switching towards a neutral perspective of being a non-partisan official trying to find the perfect districting, one can analyse how a theoretical, perfect districting looks like. One therefore splits voters into three groups (Democrats, Republicans and Independents) with the result that there exists an optimal seat-vote curve which is implementable in districts when there are enough partisan voters in relation to Independent voters. [Coate and Knight, 2005]
The important message of this section is that there exists a unique districting under which the optimal seat-vote curve desired (although it is biased) is implementable.

Another interesting phenomenon connected to redistricting is that the setting changes in case a MMD constraint is added. That constraint forces the gerrymanderer to add majority-minority districts such that minorities are properly represented in the legislative process. The perverse-effect claim was analysed which states that the creation of majority-minority districts leads to more Republicans being elected to the legislative who in general do not represent minority political views. By looking at publications by Shotts and Grigg and Katz the claim was assessed with the result that none of their models supports the claim, yet both state a need for more research in that area.

Another possible perspective on the phenomenon of gerrymandering is yielded by taking the role of a neutral court trying to detect whether gerrymandering was present under a certain districting plan. Possible measurements were presented, but no single measurement turned out to suffice for answering the question on the presence of gerrymandering. Most methods can be faced with concrete constructed examples in which they do not indicate gerrymandering although it could be present. However, the methods could be used combined as some (but not all) measurements analyse different aspects and are therefore almost independent from another. [Cervas and Grofman, 2019]

This field of research is still developing as also recently published articles could be included in the analysis. To summarise, one can state that the collective of possible statistical measurements gives a good indication of whether gerrymandering is present in a redistricting plan. Especially the growing amount of district plans available for analysis will further improve the power of the procedures presented in this section.

3 Illegal Election Bias - Ballot Stuffing and Voter rigging

3.1 Introduction

In this section the thesis discusses practices which, in contrast to gerrymandering, are considered as election fraud. The goal is to give a comprehensive overview on different types of election fraud and how they can be detected using statistical models, referring back to the research question stated in the introduction.

The section starts by categorising different types of election fraud, but focusses on two particular types. Benford's Law [Benford, 1938] is discussed which is an early method of how authorities tried to check for anomalies in election data. This method has considerable disadvantages which is why talking about alternatives is essential.

The first specific type of fraud discussed is ballot stuffing. By that term the presence of votes in ballot boxes is meant, with the speciality that these votes were never cast by a voter i.e. before the election starts, one party has already put votes in their favour into the ballot boxes. This phenomenon is expected to influence the distribution of vote and turnout according to Klimek et al., which is exploited for visualizing election fraud. For that purpose advantages and disadvantages of various models trying to detect ballot stuffing are presented.

The second type of fraud closer looked at is so-called voter rigging by which the action of the government party of preventing the opponent's voters from casting their vote or the more drastic version, where people are forced to vote for a specific party, is meant. [Lehoucq, 2003]
This phenomenon is particularly present in smaller polling stations where only few voters cast their votes. This does not necessarily have to be the case in small districts. Therefore the effects of these small polling stations on the outcome of an election are of particular interest. Note that a "rigged vote" is hard to distinguish from a regularly cast vote. Jimenez et al. differentiate between "voter rigging" and "clientelism" which is the practice of distributing benefits to citizens in exchange for electoral support. "Voter rigging entails coercion, affecting the free choice of the voters, and is part of the menu of manipulation that goes beyond the limits of democratic politics." [Jimenez et al., 2017b]

The discussed methods are tested by applying them on real election data. The results and potential explanations are reviewed in this section. An important factor is that election data is hard to access as it is usually not published as detailed as researchers wished for the purpose of the analyses. Therefore data availability is going to be a major constraint for the upcoming analyses.

3.2 Types of election fraud

Before talking about modelling and detecting election fraud, it needs to be clarified which actions are in the centre of the discussion. The next section on different fraud types is based on the paper by Lehoucq. He defines election fraud as "clandestine and illegal efforts to shape election results" [Lehoucq, 2003]
This definition is broad in the sense that no specific actions are described, but it already outlines that someone has to be "active" when election fraud is present and that there is no such thing as unintentional election fraud. That is a first difference towards gerrymandering as unintentional gerrymandering is a major concern. [Chen et al., 2013]

In addition to that election fraud by definition has to break some law. To be of interest for analysis the actions mentioned also need to have some effect on the election outcome as they are irrelevant and extremely hard to detect otherwise. The general setting is that there is a party in an election which feels the need to perform illegal actions to win an election of which the party believes it could not be won otherwise. [Lehoucq, 2003]

Numerous ways of performing election fraud according to that definition exist. They are categorized according to the approach they take:

- *structural manipulation:* The first type of electoral fraud are intentional structural mistakes in the organization of an election. An example is the phenomenon that a single voter could vote more than once by visiting different polling stations. In 19th century Peru elections, each party organised its own polling station and while the party in power took charge of the polling station at the main square, the opposition party had to position itself at a less important square. [Mücke, 2001]
  As there was no independent commission at that time, the probability of vote counts being correct is low.
  The scenario of "invented elections" is also added to this category. Assume an election is held, but the results publicized afterwards are made up and in favour of the party in charge.

- *ballot manipulation:* This form of election fraud is of interest for data-driven methods. It includes activities such as the addition of ballots which were never cast by a voter (ballot stuffing). Lehoucq mentions practices in 19th century Argentina where entire ballots of the opposition were destroyed and parallel polling stations were set up in which the votes did not count. If a party wants to perform this type of electoral fraud, it has to rely on a polling station official who could theoretically prevent these manipulations from happening. This category also includes simple number manipulations where the number of votes counted is manually changed before being sent to a central electoral agency which is known as the misrecording of votes. This is closely linked to the invalidation of valid votes for political purposes.

- *turnout manipulation:* This form involves a wide range of different activities. Parties try to prevent certain groups of people from voting or ensure that others are voting by performing illegal activities, such as paying them money for casting their vote. A more common version is to intimidate opposite voters or even use violence in more extreme cases. Even the threat of violence can cause groups of voters not to make their way to a polling station. This category includes the practice of raising turnout by using illegal actions, but also the case in which the government party tries to prevent voters from participating.
  Another practice linked is the practice of making it harder for a specific group to register for an election. Among such scenarios polling stations could be placed far away from certain regions such that the barrier to vote for people living in these regions is substantially higher than for others.

- *technologic manipulation:* This category includes cyberattacks on voting systems, but also the conscious and intentional manipulation of voting machines. Hardware and software manipulations of voting machines are included although it might also be appropriate to list them as structural manipulation practice and therefore place it in the first category.

This categorization is not complete as there are some minor fraud practices not mentioned here. The intention behind the grouping is to give a brief overview of what the following chapter is discussing.

3.3 Statistical Framework

The methods presented in this section all have in common that they are applied within a statistical framework and analyse similar variables. Assume that election data for fine aggregation levels is published and one tries to construct a procedure to determine the presence of fraudulent mechanisms.

All methods formulate a scenario of free and fair elections. This scenario is seen as the null set and statistically speaking, this null set is what is expected under the null hypothesis that no fraud is present. The way the analyses proceed differs depending on the type of election fraud potentially being present under the alternative hypothesis that "fraudulent mechanisms are present at the election".

Benford's Law solely focusses on the voting numbers by district and checks whether they follow the Benford distribution.
The methods targeting explicit types of election fraud make use of more variables:

1. turnout
2. voting district sizes
3. percentage of the winning party

Each of the methods on ballot stuffing and voter rigging exploits the correlation structures between these three variables. For modelling the variables of interest, the log-normal distribution is often chosen because of the advantage that it is not generating negative numbers. Therefore these methods assume that the number of potential voters per district is known, as well as the voting numbers by party out of which the turnout per district and the percentage of the winning party can be calculated.

In a fair election no noteworthy correlation structure between turnout and the winning party's percentage is assumed. Some methods assume a bivariate normal distribution. That changes in case fraudulent mechanisms are present. If the winning party's percentage rises for districts in which the turnout is considerably high, an indication for ballot stuffing can be deduced. Nevertheless, other reasons also cause correlation between turnout and the winning party's percentage, but their presence is often known beforehand (e.g. strong voter mobilization in certain regions). These analyses can therefore be seen as correlation structure analyses.

The correlation structure is visualized by looking at the two-dimensional histogram (later called the electoral fingerprint introduced by Klimek et al.). By taking special properties of each election into account, a fingerprint expected under the null hypothesis of fair elections can be generated. The methods presented in this section differ in their way of measuring the deviations of the null model and the actual election data.

In the ballot stuffing section, two parameters are fitted to the data via a goodness of fit procedure. As randomness is involved in constructing the model (normal distribution assumption), a considerable amount of simulations has to be perforned to draw conclusions on the goodness of fit parameters. The model is formulated in a way such that the parameters being zero indicates no fraud being present, while large values indicate deviations between the fair election model and the data. By performing a large number of simulations, a confidence interval for the parameters can be constructed.

In the voter rigging example the model is based on the assumption that voter rigging causes the result of the winning party to be different in small polling stations as these are particularly associated with voter rigging. A model under the null hypothesis is created as well and the deviations between the data and the model under the null hypothesis are calculated. The test approach is different in this case as a set of trustworthy elections is used as reference set. The magnitude of the deviations is interpreted in relation towards the set of trustworthy elections and the significance is measured relative to this set.

One model will be analysing the effect of a subset of non-regular votes on the overall election result. The assumption made in this model is that the non-regular votes follow the same behaviour as regular votes and can therefore be estimated by simple weighted regression.

All of these models are presented in more detail and with applications in the following section.

3.4 Benford's law and alternatives

The first type of fraud addressed is the publication of made-up numbers which do not coincide with the actual election outcome. Such elections are considered to be invented elections (or show-elections), but also the practice of submitting wrong results from the polling station to the central election agency is among those types treated in this section.

To statistically detect human intervention on the election results, one tries to exploit that the last digit of vote percentages or even of the absolute number of votes cast per party, follow a mathematical distribution. Humans trying to manipulate such "natural datasets", naturally overcompensate by making the numbers look even more random on first sight.

A powerful tool to detect such irregularities is Benford's law. [Benford, 1938]
After analysing diverse sources he stated that in naturally collected datasets used numbers started more often with the digit 1 than with the digit 9. Benford's law is a statement about the frequency distribution of the leading significant digit. As Berger et al. state, Benford's Law (BL) "is the observation that in many collections of numbers, be they e.g. mathematical tables, real-life data, or combinations thereof, the leading significant digits are not uniformly distributed,as might be expected, but are heavily skewed toward the smaller digits. More specifically, BL says that the significant digits in many datasets follow a very particular logarithmic distribution." [Berger et al., 2011]

Many theoretical explanations for Benford's Law were already stated, but none succeeded in convincing a broader academic community. [Berger and Hill, 2011]

Mathematically speaking, a set of numbers satisfies Benford's Law if the leading significant digit D_1 satisfies for all $d \in \{1, \dots 9\}$ that [Berger et al., 2011]

$$\mathbb{P}(D_1 = d) = \log_{10}\left(1 + \frac{1}{d}\right)$$

For a concrete number $d = 3$ this means that

$$\mathbb{P}(D_1 = 3) = \log_{10}\left(1 + \frac{1}{3}\right) = \log_{10}\left(\frac{4}{3}\right) = 0.1249$$

Note that the higher the digit is, the lower is its probability of being the leading significant digit. What is particularly noteworthy is the weight on the first two digits which is almost at 0.5.

The concrete distribution for all digits is given in the table below:

Digit	Probability
1	0.301
2	0.176
3	0.125
4	0.097
5	0.079
6	0.067
7	0.058
8	0.051
9	0.046

Graphically this corresponds to the following plot:

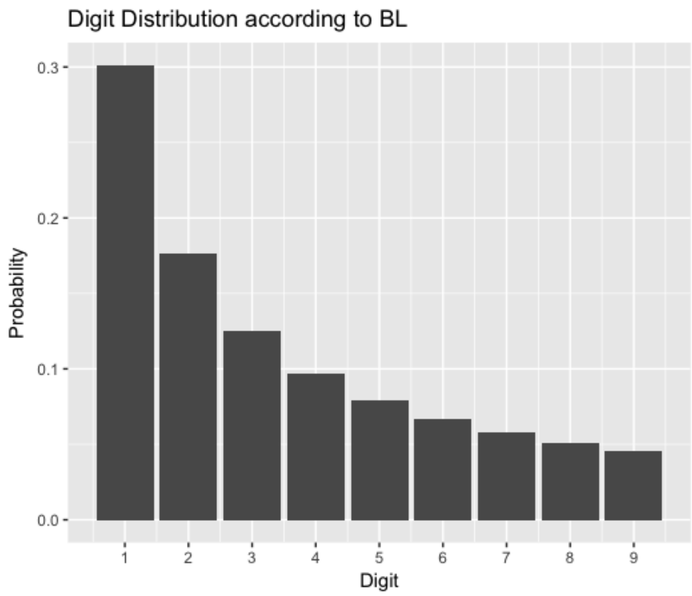

Figure 4: Benford's Law

This formulation can be rewritten in that the logarithm of the first significant digit follows a uniform distribution. [Klimek et al., 2012]
Benford's Law is used to detect manipulated data in many areas such as tax investigation. One of the main difficulties with Benford's Law is the definition of a "natural dataset". If a set of numbers does not follow Benford's Law, it is according to the law not generated "naturally" which yields an indication for potential manipulation. On the other hand with this argument the law itself can't be challenged which is subject of controversy.

If election results deviate massively from Benford's Law it may indicate manipulative tendencies. Klimek et al. assume that reported vote counts are invented by a person preferring to choose numbers that are multiple of ten. That would immediately raise the proportion of zeros as last digits compared to uncorrupted numbers.

In practice researchers have found many distributions obeying the law, but also distributions not doing so. Among those obeying it are the Fibonacci numbers, among those disobeying the law are square roots and reciprocals. [Washington, 1981] [Raimi, 1976]

The law can be stated in a more complete form according to Berger et al. as it is a statement about the joint distribution of all decimal digits and not only about the first. This is of interest in election outcomes as the first digit strongly depends on the size of the polling entity. The formulation for all digits is

$$\mathbb{P}(D_1 = d_1, D_2 = d_2, \dots D_m = d_m) = \log_{10}\left(1 + \left(\sum_{j=1}^{m} 10^{m-j} d_j\right)^{-1}\right)$$

where $d_1 \in \{1, \dots 9\}$ and for $j \geq 2$, $d_j \in \{0, \dots 9\}$. By simply plugging into the formula in a concrete example this means [Berger et al., 2011]

$$\mathbb{P}(D_1 = 3, D_2 = 1, D_3 = 4) = \log_{10}\left(1 + \left(3 \cdot 10^2 + 1 \cdot 10^1 + 4 \cdot 10^0\right)^{-1}\right) = \log_{10}\frac{315}{314} = 0.00138$$

In addition to that it is also possible to formulate unconditional probabilities for the second digit. [Deckert et al., 2011]

Digit	Probability
0	0.120
1	0.114
2	0.109
3	0.104
4	0.100
5	0.097
6	0.093
7	0.090
8	0.088
9	0.085

One way to determine whether a set of numbers is distributed according to Benford's Law is by applying a χ^2-Test. To have more statistical power one could apply the Kolmogorov-Smirnov or the Kuiper test instead. A minor problem in the common formulation of these tests arises which makes them no more applicable for testing on a violation of Benford's Law. These tests are based on the null hypothesis of a continuous distribution which causes them to be conservative for testing discrete distributions. So these tests should reject more often than they actually do. [Morrow, 2014]

This problem was solved by Morrow who derived new test statistics guaranteeing compliance with Benford's Law under the null hypothesis. He also derived methods for characterizing which distributions of a particular family follow Benford's Law and derived an upper bound for the rate of convergence.

One critique on the use of Benford's Law for the detection of election fraud was established by Deckert et al.. They argue that the focus on the first digit is misleading when trying to classify election results as fraudulent or trustworthy. Brady observes that in a two party race with district sizes varying between 100 and 1000, the first digit of each party's vote is most likely to be 4 or 5 and not 1 or 2 as Benford's Law requires. Therefore if the first digit is not following the law, this occurrence cannot be interpreted as a higher probability for the presence of fraudulent mechanisms. [Brady, 2005]

Therefore it is more common to look at the second digit in election analysis. However, Deckert et al. state there is no sophisticated reason to believe that fair and free elections are consistent with the second digit Benford Law. Another argument against the second-digit Benford test was originally raised by Beber and Scacco which discusses the relevance for looking at the last and next to last digit. Their key argument is that officials performing election fraud are afraid of being prosecuted and as a consequence they avoid zeros and fives as last digits to ensure randomness on first sight. In addition to that humans try to avoid double digits when inventing random numbers and therefore those repeated numbers occur less frequently in made up datasets.

The key question in discussing the properties of Benford's law for detecting election fraud is whether election data of fair elections satisfies the law. Mebane Jr argues that as in election numbers a lot of randomness is involved, it can be classified as a "natural" dataset which should fulfil the law. The randomness involved stems from a variety of stochastic choices involved in the voting process - the decision of whether a voter participates in an election, the party he votes for and even unintentional counting errors add random noise towards the numbers. [Mebane Jr, 2006]

Deckert et al. argue that this argument is problematic in that fraudulent election counts can also be generated while there are stochastic choices in-between and according to the argument the law will also hold in fraudulent elections. The simplest one is that fraud is performed by local election officers each of which is using their own procedure and thereby adding another stochastic component. That is extremely hard to model, but does not yield any reason on why fraudulent election data should not follow Benford's Law, while the results of fair elections should. The authors argue that there is "no reason for believing that fraud free data itself will correspond to the second digit Benford Law." [Deckert et al., 2011]

The authors assessed the value of Benford's Law as an indicator for election fraud by simulation. They simulated data stemming from fraud-free elections to check whether the data satisfied Benford's Law. The simulation was set up by simplifying the model to have as few assumptions as possible and using the so-called spatial model in which voters are identified by ideal points in an Euclidean "issue" space.
The only assumption is that the final vote count being higher for one candidate or the other is explained by the relative electoral strategies, or mathematically speaking, the election result is explained by the spatial position relative to the electorate.

Assume a model of an electorate in which all voters occupy positions in the two-dimensional space described above. As a consequence every voter elects the candidate closest to his own position. The positions of the candidates are subject of simulation such that the degree of competitiveness in elections varies. To introduce random and homogeneous preferences and turnout distribution, the structure of the simulations is formalized by letting for each voter i the following be true: [Deckert et al., 2011]

$$X_i = \beta_X V_i + u_{Xi}$$

$$Y_i = \beta_Y V_i + u_{Yi}$$

$$T_i = \beta_T V_i + u_{Ti}$$

The parameters X_i and Y_i represent a voter's position in the two-dimensional space. T_i is the "voting" coefficient and is essential for the turnout in an election as voter i only takes part in the election whenever T_i is greater than some fixed threshold $T^\star$.
$V_i \sim N(g, 2)$ where g itself follows a normal distribution such that $g \sim N(G, 0.15)$ where $G = 2$ for X_i and Y_i holds. The parameters β_X and β_Y vary between the districts and are given by two normal distributions; $\beta_X \sim N(2, 0.15)$ and $\beta_Y \sim N(-1, 0.15)$. Conversely, the parameter G for turnout is fixed at 4 and the critical threshold for turnout is fixed at $T^\star = 15$ which results in turnout varying between 40% and 60% across the different districts. $u \sim N(0, 2)$ is a noise term. [Deckert et al., 2011]

This model comes close to a fraud-free election. Suppose that G denotes the mean personal income of the entire voting population (of course G could also represent some other characteristic). The national income of the voting population is known and therefore can be fixed at some value. The mean income of every district depends on the national income and is denoted by g which is randomly drawn from the distribution with expected value G described above.
In this hierarchical model one can switch to the level of a specific voter. The income of voter i depends on the average income in the district and this income (V_i) is randomly drawn from the distribution with mean g. The model accounts for the possibility that income has a different influence on the voting behaviour in different districts. The income's impact on the voting behaviour is modelled by the random variable β_X which is multiplied by the voters income V_i. β_X and β_Y have different means which allows for different importance of the social parameters affecting the election. [Deckert et al., 2011]

To summarize, X_i and Y_i represent the voter's position on two different social factors each of which depends on the impact of that factor modelled with the respective coefficient of β. T_i determines whether voter i takes part in an election. The parameters X and Y determine together with the candidate's position in the "issue" space which party the voter is preferring.

The parameters of interest for the simulation are the district size and the margin of victory. At first Deckert et al. simulate an election by creating 1000 districts with each district having the same number of possible voters. They performed simulations in which each district contains 1,000 voters, 10,000 voters and some simulations with district size 20,000. The district level is interesting as that is usually the level at which election results are published, but the smaller the level of aggregation is, the better analyses can be performed. One problem with district sizes being the lowest aggregation level in the data is that the second digit could possibly be the last or at least next to last digit which is problematic as discussed above. These arguments are used to justify the selection of these three levels for the district sizes as they are reasonable, but avoid the problem mentioned. [Deckert et al., 2011]

The analysis was also performed in a different way where districts are not equally sized. Therefore the specific district sizes are randomly distributed around the means of 1000, 10 000 and 20 000 which are the district sizes in the first simulation. The distribution is such that smaller districts are more common than larger districts which is due to empirical observation. The size of each district is given by

$$S_d = 0.75m + e_d,$$

where m is the mean district size and e_d is an exponentially distributed random variable with mean 0.25.
The second parameter of interest is the share of the winning party's candidate in a two-candidate race (the share therefore yields the margin of victory as well). Deckert et al. assume three different values

for that margin which are at 52%, 57% and 66%. The set of simulations that the authors generated is divided into two equally sized parts of which only the first one is used to asses the hypothesis that the second digit Benford test can be used to detect a free and fair vote count.

Whether the data followed Benford's Law was determined by testing for the mean of the second digit of candidates absolute vote numbers to be significantly different from 4.187 which is the mean under Benford's Law. The result was that this mean was different from 4.187 in 80 of 102 simulations in which the district size was held constant. This simulation supports the claim that fair elections do not follow Benford's Law in general and it therefore is not a reliable indicator for a free and fair election. [Deckert et al., 2011]

To see and evaluate the performance of the second digit Benford test on fair elections was one goal of the analysis, while the second goal was to see the performance under fraudulent data. Deckert et al. inserted "fraud" in their simulation. The second half of the simulations was performed by taking the first fraud-free half and transferring some percentage of votes from the losing to the winning party. The percentage of votes transferred in each district is distributed according to a uniform distribution between 0% and 30%. This type of election fraud falls in the category of structural fraud.

Note that under the fraudulent data the number of simulations in which the mean is significantly different from 4.187 (still constant district size) drops from 80 to 72. Hence the test considered more fair elections to be fraudulent than fraudulent elections. The only positive aspect of these simulations is that the performance is solid in small districts (1,000 or 10,000 voters) and the number of means significantly different from 4.187 rises when moving from fraud-free to fraudulent data in this subgroup of simulations, but in larger districts with population of 20 000 the second digit Benford test does not fulfil its purpose at all. The number of significantly different means in the fraud-free version is higher than in the fraudulent dataset. [Deckert et al., 2011]

The calculation was repeated after dropping the assumption on equal district sizes. The result in the fraud-free dataset drastically improves to 57 of 170 simulations in which the mean was significantly different from 4.187 although no fraud was present. One explanation is that the added noise in choosing the district sizes moves the data closer towards a natural dataset and the assumption of constant district sizes biases the numbers in a way such that they cannot be distributed according to Benford's Law. But around 35% of fraud-free elections being labelled as potential fraudulent elections is still not a good result.

This result is compared to the performance under fraudulent data. Here 68 of 170 simulations yield a mean significantly different from the mean under Benford's Law. Therefore under exponentially distributed district sizes, it can be stated that the Type 1 error is at 34% by which is meant that fair elections are classified as fraudulent, while the Type 2 error is at 60% by which is meant that fraudulent elections are classified as being fair in that proportion of simulations.[Deckert et al., 2011]

The authors complain about the absence of developed theory which links Benford's Law towards the problem setting in the detection of election fraud. Without such a theory it is reasonably hard to determine whenever deviations from the mean under Benford's Law are concerning. Other topics not covered by this paper but possibly essential factors in finding a procedure to detect election fraud, are the impact of the number of parties competing in the election or the presence of strategical voting. But as these factors are unknown here, Deckert et al. conclude that "Benford's Law is problematic at best as a forensic tool when applied to elections". [Deckert et al., 2011]

Note that there is some harsh critique on the paper of Deckert et al. by Mebane. He criticizes the procedure used in the paper and argues that a test for a significant difference of the mean from the mean under Benford's Law cannot be associated with Benford's Law in general. [Mebane, 2011]

He questions the usefulness of Benford's Law for the detection of election fraud, but this research question can't be answered by the approach taken by Deckert et al.. He argues that the simulations do not contribute towards the analysis of the procedure as the test should be validated on real election data. Mebane is not surprised by the poor results of the simulations as he claims that the use of normally distributed random components might be the reason for that as the second digit of normally distributed numbers in general do not follow Benford's Law and therefore also have a different mean.
He doubts the meaningfulness of using the 2-dimensional spatial model for modelling a voter's electoral behaviour, as no actual voting behaviour is described by such a model and Deckert et al. have no evidence that the digits generated by such a simulation match the digits of any real world election. [Mebane, 2011]

The actual challenge according to Mebane is to start with real election data in which the second digits are distributed according to a second-digit Benford test and then develop simulation mechanisms matching the real world data and to check the patterns in the simulations for wider applicability. According to Mebane the work by Deckert et al. is a "purely negative contribution". [Mebane, 2011]

Both analyses have in common that they share scepticism towards the possible applicability of Benford's Law in detecting election fraud. The lack of mathematical theory is a constant problem in interpreting the results whether a set of numbers is matching Benford's Law.
A simple test on the second digit of voting results does not provide enough evidence for stating that an election was held free and fair, but on the other hand deviations of the second digits of election data not fulfilling Benford's Law is not evidence enough for stating that election fraud was present.

3.5 Modelling Ballot Stuffing - Electoral Fingerprints

The next type of election fraud analysed is ballot stuffing. The addition of votes never cast by a voter into ballot boxes is a phenomenon that is reasonably complex to detect. The strategy discussed in the next section is to have a closer look at the distribution of vote and turnout. [Klimek et al., 2012]
The developed procedure exploits that unusually high vote counts for the winning party coincide with unusually high turnout rates when additional ballots are added in favour of the winning party. Ballot Stuffing changes the shape of vote and turnout distribution and induces correlation between them. The goal is to find a statistical technique which does not depend on the size of the sample or the aggregation level of the data. Therefore ballot stuffing must not disappear if the same election dataset is aggregated differently. [Klimek et al., 2012]

Klimek et al. based their work on empirical datasets and propose a parametric model for the influence of ballot stuffing on election results. They include countries whenever election data for small enough territorial units is available. The data used is the number of votes of the winning party in every unit, the number of eligible voters in that unit and the number of valid votes.
The data was rescaled as it could be shown that an appropriate way of rescaling the data ensures that the distributions of votes and turnout follow a Gaussian distribution. [Borghesi and Bouchaud, 2010]

Denote the number of votes of the winning party as W_i where i is the electoral unit and the number of voters in this unit as N_i. According to Borghesi and Bouchaud the rescaling function is given by $\nu_i = \log \frac{N_i - W_i}{W_i}$ which is a function not defined in case $W_i = N_i$ which eliminates units where the percentage of the winning party is at 100% and the turnout is as well. The results were presented as so-called "fingerprints" of the elections which is the term used for the two-dimensional vote-turnout distributions.

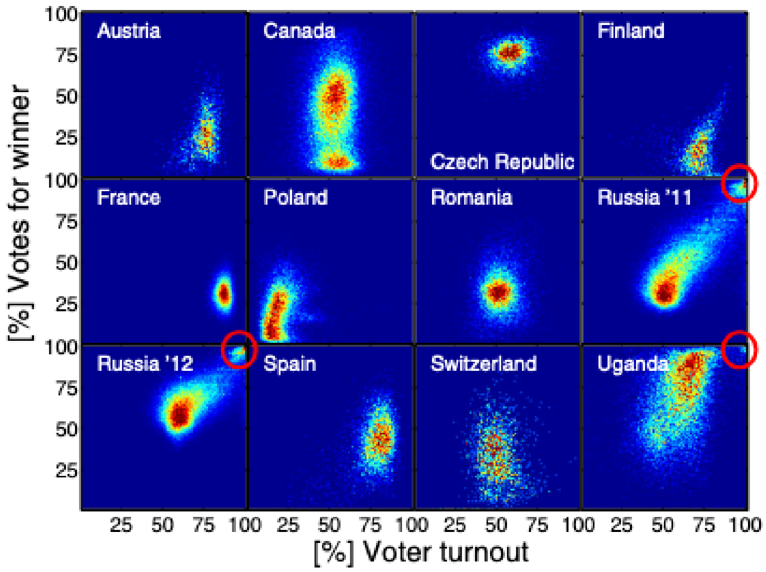

Figure 5: Election Fingerprints [Klimek et al., 2012]

In this plot the colours represent the number of units with corresponding vote and turnout percentages. Note that in the elections in Uganda and Russia the distributions are bimodal and while one cluster is at the average level of turnout and votes, the second peak is positioned in the upper right corner close to the point where vote percentage and turnout are at 100%. Klimek et al. suggest two different types of fraud being present which they name as incremental and extreme fraud. Incremental fraud is the activity of adding votes for the winning party while taking away votes from the other party. Extreme fraud is the practice of reporting complete turnout and almost all votes in favour of the winning party. Regarding notation, we say that incremental fraud occurs in fraction f_s of the units and extreme fraud occurs in fraction f_c of the units as incremental fraud is associated with the entire fingerprint being "smeared" to the upper right corner, while extreme fraud is associated with the second cluster close to the upper right corner. [Klimek et al., 2012]

Model

Assume for simplicity that voter preferences can be represented by a Gaussian distribution with mean and standard deviation taken from the election sample. Klimek et al. calculate the empirical turnout distribution which is $\frac{V_i}{N_i}$ where V_i is the number of valid votes in unit i and the empirical vote distribution which is $\frac{W_i}{N_i}$. The goal of the model is to estimate the two fraud parameters f_s and f_c.

The goal is reached by assuming a model for the election under which no fraud is present. Model turnout and vote rates for unit i are drawn from normal distributions where the mean of the model turnout is estimated from the election data as the value which maximizes the empirical turnout distribution and the same is done for the vote rate by maximizing the empirical vote distribution. This procedure is similar to the Maximum-Likelihood method. The model variances are estimated by looking at the empirical distributions. [Klimek et al., 2012]

f_s represents the probability that votes are taken away from the opposition and added to the winning party's votes. It is estimated from the data itself. The same is true for f_c which is the probability that almost all votes from the opposition party and the non-voting group were added to the winning party.

The interesting aspect of this model is how the two parameters f_s and f_c are estimated. Klimek et al. calculate the values by reverse engineering the data, by which they mean that model voter-turnout distributions are calculated according to the model described above, while testing for each combination of (f_s, f_c) where f_s and $f_c \in \{0, 0.01, 0.02 \ldots 1\}$.

So they test all different values of the two fraud parameters for determining which combination fits the model generated above best to the observed data. For all combinations the point-wise sum of the squared differences between the model and the observed vote distributions is calculated and the pair with the minimal difference between model and data is chosen to be (f_s, f_c). This procedure is iterated 100 times (as randomness is involved in the model construction), therefore 100 pairs of parameters are calculated and Klimek et al. report the average values of these pairs and their standard deviations. Under the null hypothesis that no fraud is present, both parameters are assumed to be $f_s = f_c = 0$ and deviations of that are seen as evidence for ballot stuffing. Out of the 100 simulations one also obtains the mean and standard error of the fraud parameters and can therefore construct confidence intervals for the parameters. These are used to decide on whether the deviations from 0 are significant.
The test is therefore not a classical inference test in the statistical meaning, but a collection of simulation results which in their joint appearance allow for statements about the magnitude of the fraud parameters.

Another parameter estimated via a goodness-of-fit procedure is the deliberate wrong counting parameter α which is important in the analysis of concrete elections in the next section. The idea behind $\alpha \in \{0, 0.1, \ldots 5\}$ is to measure "to which extent the ballot-stuffing process in the parametric model is combined with a deliberate wrong-counting or recasting of ballots." [Klimek et al., 2017]

Results

The estimated parameters in the analysed elections are zero or close to zero in all elections except for Russia and Uganda. That shows a property that Benford's Law lacked in that it doesn't indicate fraud in elections where one would assume no fraud to be present. The correct interpretation for the values of f_s and f_c being significantly different from 0 remains to be discussed.
To explain the smearing to the upper right corner in the elections in Russia, the incremental fraud parameter f_s has to be set to around 64% in 2011 and to 39% in 2012 which the authors interpret as fraud being present in the respective percentages of electoral units in 2011 and 2012. Klimek et al. best explain the second peak close to 100% turnout and votes for United Russia by setting $f_c = 0.033$ in 2011 and $f_c = 0.021$ in 2012 which they interpret as extreme fraud being present in 2 to 3% of all electoral units.
The parameters for the Uganda 2011 presidential election are at $f_s = 0.49$ and $f_c = 0.011$. These results showed the property of being independent of level of aggregation and no other country than Russia or Uganda had comparably high for f_s and f_c on any aggregation level. [Klimek et al., 2012]

The model indicates ballot stuffing occurring in Russia and Uganda. The election fingerprints also take on non-standard forms for other countries e.g. Finland and Canada, but for different reasons. In Canada the authors explained the second cluster as being the elective result among the population of the French speaking voters whose voting behaviour differs from the remaining voters. The result in Finland is explained by the effect of successful voter mobilization. In the analysed election the far-right party "True Finns" succeeded in mobilizing across the country except for the capital Helsinki where other parties performed better. Successful mobilization also causes correlation between turnout and the result of the winning party which results in a non-standard looking election fingerprint. [Klimek et al., 2012]

Although there exist plausible reasons for the fingerprint to attain strange forms, it is hard to find potential explanations for the territorial units in which the turnout is close to 100% and all votes being in favour of the winning party. The claim is supported by looking at a different analysis. Without constructing and analysing a complex model, the simple look at the distributions of the logarithmic vote rate shows that Uganda and Russia have non-standard forms.

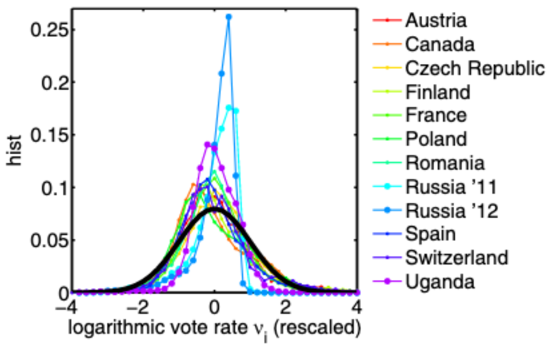

Figure 6: Logarithmic vote rate distribution [Klimek et al., 2012]

The reference point in this plot is the standard normal distribution which is expected for fair elections. For elections outside Russia and Uganda the distribution of the logarithmic vote rate is close to the standard normal distribution.

Another analysis to underline the claims by Klimek et al. is the cumulative number of votes of the winning party as a function of the turnout. This visualization demonstrates the presence of voting irregularities. The cumulative distribution functions is expected to reach a plateau at some value in fair elections and remain at this level for high turnout rates. An indicator for the presence of ballot stuffing is an increase of the winning party's vote share for high turnout rates which is not explainable by common electoral theory. The plot below indicates the same conclusions as stated above. The only elections classified as outliers are two Russian elections and one election in Uganda where a late increase of the curve instead of the curve forming a plateau is observed.

The late increase can only be explained by an unusual accumulation of districts in which turnout is high and the winning party is performing surprisingly well. Such accumulations are rather unlikely in fair elections and Klimek et al. see it as an indication for potential ballot stuffing. [Klimek et al., 2012]

The shape of the cumulative vote distribution as a function of turnout itself can only give a harsh indication on whether ballot stuffing is present, but in combination with the model results and the unusual forms that the logarithmic vote distribution takes for the elections in Uganda and Russia, the

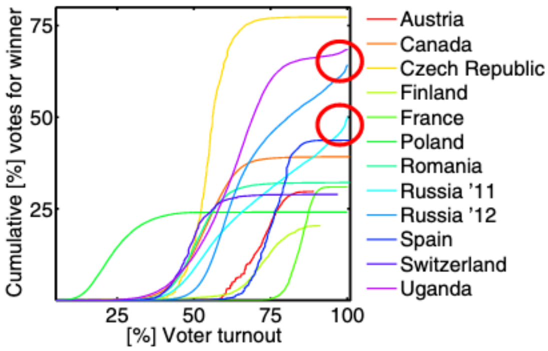

Figure 7: Cumulative votes as function of turnout [Klimek et al., 2012]

combined results are reasonably strong and indicate ballot stuffing. The authors conclude that if the fraction f_c is substantially different from zero and/or the curves discussed above take on a different shape than the usual sigmoid form expected to be observed, it is likely that "an election does not represent the will of the people". [Klimek et al., 2012]

Applying Ballot-Stuffing-Test

To illustrate how the proposed test can be applied to election data, the 2017 Turkey constitutional referendum which was subject of an election forensic analysis published in 2017 is discussed in this section. [Klimek et al., 2017]

In this referendum the population of Turkey was asked to vote on a proposed constitutional reform which would substitute Turkey's parliamentary system with a presidential system, transferring power from parliament towards the president. The result of the referendum was a narrow win for the reform which caused the opposition to immediately question whether the referendum was subject to irregularities. 51.4% of voters voted "Yes" at the referendum which corresponds to a win margin of 1.38 million votes. Until that referendum Turkey was never assumed to have major concerns with election irregularities, but shortly after the referendum took place, videos of improper behaviour were published.

Klimek et al. only include results from polling stations in Turkey in their analysis, as the number of eligible voters in other countries is hard to determine. Note that a referendum has different properties than usual general elections. The analysis does not focus on the percentage of the winning party, but the percentage of "Yes" votes. A referendum therefore shows properties of an election in a two-party system. The first visual analysis that the authors conduct, is a look at the cumulative vote percentage as a function of the turnout with the result being that only electoral units with turnout close to 100% caused the percentage of "Yes" votes to surpass the 50% mark.

The next step in the analysis is a look at the election fingerprint. A circular or elliptical symmetric fingerprint is expected in case no ballot stuffing was present. Indeed, the fingerprint of the referendum takes on a different form than the fingerprints of the elections analysed in the paper by Klimek et al.. Observe that the fingerprint is smeared to the upper right for high votes and turnouts which is expected under the presence of ballot stuffing. In addition to that the fingerprint is not symmetric. The results appear narrow if one only analyses the turnout, but the two-dimensional histogram is spread out along the vote dimension. The specific form of the fingerprint is shown below. [Klimek et al., 2017]

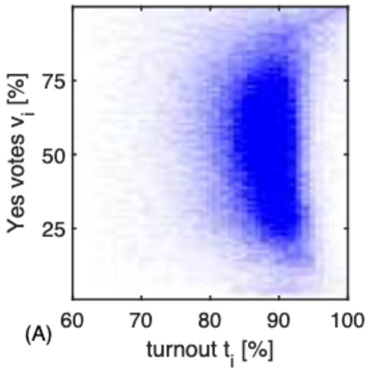

Figure 8: Election Fingerprint Turkey referendum 2017 [Klimek et al., 2017]

The model is applied to determine the values of the fraud parameters. Note that there is no evidence that some polling station's results were inflated up to 100% in vote of the winning party and turnout and therefore the parameter $f_c = 0$ in this election. The parameter of interest is f_s which is the probability that votes of the opposition party (here: "No"-votes) were added towards the vote count of the winning party. Under fair elections $f_s = 0$ is expected which is the null hypothesis of the test. For the election data Klimek et al. calculate a non-zero fraud parameter $f_s = 0.058 \pm 0.019$ which is about three standard deviations away from zero. The parameter can be interpreted as the presence of incremental fraud in f_s% of the polling stations.

The deliberate wrong counting parameter α which is a shape parameter is calculated as well. If this parameter is greater than 1, it indicates that ballot stuffing dominates over the process of deliberate wrong counting. In the Turkish constitutional referendum of 2017 the shape parameter is $\alpha = 1.3 \pm 0.2$. [Klimek et al., 2017]
The results indicate the presence of ballot stuffing in the referendum. According to Klimek et al. the result is weak compared to recent Russian elections, but it is still systematic and statistically significant. [Klimek et al., 2012]

In the analysis of the Turkish referendum the authors also looked at a potential result that Benford's Test would have yielded. However, they also mention that it is not clear how deviations from Benford's Law can be related to election fraud. The second significant digit is analysed which is exactly the test discussed in the last section. The result was that the second digits are significantly different from Benford's Law on all aggregation levels. According to Klimek et al. the magnitude of the deviation constitutes for a highly irregular observation. [Klimek et al., 2012]

3.6 Modelling Voter Rigging

The next type of election fraud is so-called voter rigging. This practice is classified as part of the class of turnout manipulations. To be more concise, voter rigging is the practice of systematic coercion and/or intimidation of voters. [Klimek et al., 2017] [Jimenez et al., 2017b]
Voter rigging is easiest in small polling stations where voters are known to the officials supervising the electoral process. Voter rigging includes practices such as forcing voters to vote for a specific party, forcing them to participate in an election or preventing them from participating in an election which can be decisive.

In this section the framework introduced by Jimenez et al. for testing for irregularities associated with voter rigging is analysed. The analysis is centred around the claim that voter rigging leads to different election results in small polling stations. Note that small polling stations do not necessarily coincide with small districts which eliminates some potential geographic factors disturbing the analysis. Jimenez et al. developed a test for determining whether the voting behaviour in small polling stations is significantly different from the results in large voting stations. This raises the problem of data availability as data is needed for fine aggregation levels and voter rigging is hard to be detected and easy to be masked in high aggregation levels.

A look at the electoral fingerprint is also interesting in context of voter rigging, but the analysis is based on the construction of the "standardized election fingerprints". The idea is to look at the deviation of the standardized fingerprints in small polling station from the standardized fingerprints in large polling stations. A test determines whether the deviations are significant and can be used as an indicator for the presence of voter rigging.

For calculating the standardized election fingerprints the notation is changed slightly. t represents the turnout percentage in a given electoral unit, and vw is the voting percentage of the winning party in the same unit. The electoral fingerprint is a useful tool when detecting ballot stuffing, but it also appears to have disadvantages as irregularities do not necessarily feature evidence for electoral fraud as successful voter mobilization could be one example causing correlation between vw and t. Such obvious reasons potentially cause the fingerprint to look different than expected under fair elections.

The detection of voter rigging is even more complicated as the distinction between actual voters and voters who are forced to vote for a certain party is hard. Jimenez et al. introduce a method which is more robust against non-fraudulent mechanisms. The idea for the standardized election fingerprints is to compare the turnout and winning party vote percentage to the electoral neighbourhood of the unit which accounts for geographic anomalies. A neighbourhood is the "smallest available administrative division to which the electoral unit belongs." [Jimenez et al., 2017b]
The developed method works with Z-scores where the Z-score of unit i is given by the formula

$$Z_t(i) = \frac{t(i) - \mu_t(i)}{\sigma_t(i)} \text{ and correspondingly } Z_{vw}(i) = \frac{vw(i) - \mu_{vw}(i)}{\sigma_{vw}(i)}$$

Note that $\mu_t(i)$ and $\sigma_t(i)$ correspond to the mean and standard deviation observed in the neighbourhood of electoral unit i. The definition of the neighbourhood varies depending on the available level of electoral units in the analysed country (counties, municipalities, districts etc.). The quantities μ_{vw} and σ_{vw} refer to the observed values of the winning party vote in the neighbourhood of unit i. Instead of drawing the two-dimensional histogram of the winning party vote and the turnout, the histogram of the standardized values Z_t and Z_{vw} is plotted which is referred to as "standardized election fingerprints". [Jimenez et al., 2017b]

Of particular interest is the development of a test determining whether voter rigging is present in an election. The irregularities the test shall be able to detect is the presence of voter rigging in small polling stations and it is based on the comparison of the standardized election fingerprints of small and large electoral units. To measure that difference Jimenez et al. introduce a discrepancy measure $D(p)$. The parameter p determines if an electoral unit is considered to be "small" such that electoral units with fewer potential voters than the p-th percentile of the usual number of electoral voters in an electoral unit of this size, is classified as a small unit while the complementary set is classified as large units.

Denote the set of electoral units in election k which fulfil the criterion of being small described above as $S(k,p)$ and the complementary set as $L(k,p)$. Cases in which the two sets contained less than ten elements are excluded. The test exploits that in elections with present voter rigging, the Z-scores of the set $S(k,p)$ differ significantly from the Z-scores in the set $L(k,p)$. The standardized election fingerprints of small and large units, are calculated and their proximity is assessed considering the distance between their centres. For simplicity the authors chose to use the Euclidean distance to compare the centres of the two curves and the median to estimate the centre points. [Jimenez et al., 2017a]

$$m_t^S(k,p) = median\left[Z_t(i), i \in S(k,p)\right]$$

yields the turnout coordinate of the curve for small electoral units in election k. The other coordinates are calculated analogously and given as $m_{vw}^S(k,p)$, $m_t^L(k,p)$ and $m_{vw}^L(k,p)$. The estimator for the distance of the two centres in election k is

$$D_k(p) = \sqrt{\left[m_t^S(k,p) - m_t^L(k,p)\right]^2 + \left[m_{vw}^S(k,p) - m_{vw}^L(k,p)\right]^2}$$

The values for $D_k(p)$ are calculated for different elections and different values of $p \in \{0.5, 1, 1.5, 2 \ldots 90\}$. The modified Thompson τ test is used to detect outliers within these elections and determines all elections not classified as outliers as "trustworthy elections".[Jimenez et al., 2017a]

The modified Thompson τ test is a test for detecting outliers within a dataset. Suppose that x is a vector containing n elements of which the mean and the standard deviation are denoted as $m(x)$ and $sd(x)$. The test uses the t distribution with $n-2$ degrees of freedom to determine the critical threshold of the test. $t_{\alpha/2}$ denotes the 1-$\frac{\alpha}{2}$ - quantile of the respective t-distribution. The rejection threshold is given by

$$r = \frac{t_{\alpha/2} \cdot (n-1)}{\sqrt{n(n-2+t_{\alpha/2}^2)}}$$

An indicator for whether this observation is an outlier is calculated by using the vector $\Delta = (\Delta_1, \ldots \Delta_n)^t$ such that $\Delta_i = |x_i - m(x)|/sd(x)$. The observation with the highest value for Δ_i is the first candidate for being an outlier and is considered as such if $\Delta_i > r$. If the observation is classified as an outlier, it is removed from x and the procedure is iteratively applied again on the remaining observations. It is stopped once all remaining observations fulfil $\Delta_i \leq r$. [Jimenez et al., 2017a]

$D_k(p)$ contains the estimated distances between the centres of the standardized election fingerprints of all elections considered in the analysis. The application of the modified Thompson τ test yields a reference set of trustworthy elections for a fixed value of p which is denoted R. The final distance value of an election is calculated by using Euclidean distance.

$$\delta_k(p) = \frac{D_k(p) - mean(D_i(p), i \in R)}{std(D_i(p), i \in R)}$$

The mean and standard deviation in the formula are calculated over the set of regular elections. $\delta(p)$ is the standardized discrepancy measure of an election. As the mean is subtracted in the numerator and divided by the standard deviation similar to how the Z-scores are calculated, values far from zero are considered as indicators for voter-rigging in small polling stations. Jimenez et al. construct the rejection region by applying the modified Thompson τ test on $\delta(p)$. They state that "if $\delta(p)$ lies outside of this region for a wide range of small values of p and, additionally, the centres of the SEFs of small units are inside the upper right region of the plot, the outcome of the corresponding election is compatible with the hypothesis of large-scale distortion of small units with respect to their electoral neighbours, which we only explain by some type of voter rigging." [Jimenez et al., 2017b]

This method has advantages such as the simplicity and the intuitiveness. On the other hand the method relies on the efficiency and power of the Modified Thompson's τ test and this test should be used with caution when detecting outliers in a dataset. It might be of interest to classify outliers based on a different procedure which for example relies on the Interquartile Range such as Tukey's fences [Tukey, 1977] and compare the results of the two procedures.

Results

Elections of different countries and whether the centres of the standardized election fingerprints coincided for small and large electoral units were analysed. The results can be seen below. [Jimenez et al., 2017b]

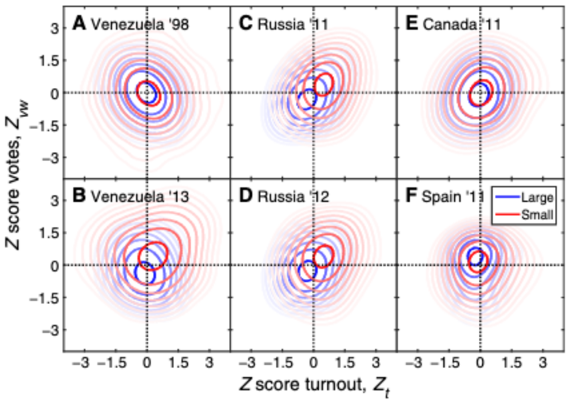

Figure 9: Comparing standardized election fingerprints [Jimenez et al., 2017b]

The centres coincide for elections of Venezuela 1998, Spain and Canada. Canada is of particular interest as the French and English speaking population usually differ in their vote, but these deviations do not affect the comparison of the standardized fingerprints for small and large polling stations. However, deviations to the upper right corner are observable for the Venezuelan election in 2013 and the Russian elections in 2011 and 2012 (presidential and parliamentary elections). [Jimenez et al., 2017b]

The fact that the fingerprint for smaller polling stations is moved to the upper right corner shows that in these electoral units the turnout and the percentage of the winning party are higher than in other electoral units in the neighbourhood. This is expected to happen in case voter rigging was present. To further assess the results, the distance measures of a larger set of elections are discussed.
The most significant results were obtained for Venezuelan elections between the years 2006 and 2013 in which the values of $\delta(p)$ reached 10 for $p = 0.05$ (i.e. 10 standard deviations away from 0). But even if the choice of p is altered, the results stay significant for a wide range of values of p. The values for the mentioned Russian elections lie between 4 and 7 (depending on p). It is noteworthy about Venezuelan elections that earlier elections (such as 1998) are not noticeable at all. Results were considered to be concerning in case deviations were higher than three standard deviations which shows the dimension of the high values for Russia and Venezuela. Other significant results were obtained for Uganda and earlier Russian elections, yet their values of $\delta(p)$ were lower than for the other elections mentioned.[Jimenez et al., 2017b]

To check the robustness of the results, Jimenez et al. performed a leave-one-out cross validation which affects the construction of the region for which the results for $\delta(p)$ are considered trustworthy. The only result that changes in the cross validation procedure is for the Venezuelan 2006 election which was significant for some choices of p and becomes insignificant in 9 of 21 cases. Other removals in the analysis conducted by Jimenez et al. did not result in significant changes. The mentioned elections in Venezuela and Russia remained significant across all 21 leave-one-out cross validations.

In addition to the cross validation the authors conducted tests for robustness. They selected elections for which the result was significant and randomly permuted the electorate sizes while the percentages of the winning party and the turnout remained fixed. The consequence is that the sizes of the electoral units are randomly assigned. The effect of the deviation is expected to disappear, because if it does not, voter rigging in small polling stations cannot be the reason for the deviations in the standardized election fingerprints. Indeed, the result obtained for the Russian and Venezuelan elections falls into the category of elections with insignificant deviations between small and large polling stations. [Jimenez et al., 2017b]

Two additional robustness tests were performed where they permuted the turnout t of the electoral units in the first test, and the percentage of the winning party vw in the second test. The three robustness tests can be summarized as permutation tests, as in each of them one of the three variables size, turnout and winning percentage is permuted while the other two are held constant. As a consequence each of the tests preserves the correlation structure between two of the variables while eliminating the correlation structure in relation to the third variable.
The result Jimenez et al. obtained was that the permutation of the turnout did not affect the significance of the elections. The permutation of the winning percentage vw resulted in two of the elections becoming insignificant (Uganda, Russia 2003), but the elections significant before remained significant. [Jimenez et al., 2017b]

These tests are methods of confirming the results obtained above. None of the tests provided any reason not to believe in the results obtained using the Modified Thompson's τ test. The central claim around this analysis is that the size of electoral units has an impact on the percentage of the winning party. Jimenez et al. conducted one last analysis where they focused on the three most significant results from above, which were the elections in the countries Venezuela, Russia and Uganda.

The principle is to sort the electoral units according to their size and calculate the cumulative percentage of the winning party up to rank i. The electoral units are sorted in descending order and the largest electoral units are included from the start while smaller units are added at the end. Assume the percentage of the winning party to be independent from the size of the electoral unit. Then the slope of the cumulative curve is expected to be 0 for small units as the addition of these small electoral units does not have an impact on the overall election result. The result of this analysis was that interesting patterns for elections in Russia and Venezuela were found. Both of these curves show an increase for small polling stations which shows that the smallest units also had an effect on the election outcome. What is particularly concerning is that in the Venezuela 2013 presidential election, the smallest electoral units caused the overall percentage of the winning party (Nicolas Maduro) to rise above 50% i.e. the small polling units for which our the presence of voter rigging is indicated were decisive for the outcome of the election. [Jimenez et al., 2017b]

The authors added that this method could also be applied in case other forms of voter rigging were present. Other forms include the intimidation of ethnic groups not to vote, but that would lead to corre-

lation between turnout and the winning percentage (potentially being negative). In that case turnout is expected to drop, while the margin of victory of the government party in districts where a certain ethnic group is intimidated, is expected to increase. That causes the standardized election fingerprint of small polling stations to shift to the upper left corner. [Jimenez et al., 2017b]
But Jimenez et al. also mention that certain forms of voter rigging cannot be detected by applying this method. Especially cases in which the intimidated group of voters does not cluster in small or large polling centres. That would not have an effect on the standardized fingerprints and neither on the distance measure $\delta(p)$. A second limitation of the method is that the selection of the set of trustworthy elections is crucial and yet it is extremely hard to classify an election as trustworthy.

Despite these limitations, the analysis of the Turkish constitutional referendum of 2017 is continued from the perspective of voter rigging in small polling stations.

Applying Voter-Rigging-Test

Klimek et al. also addressed the issue of voter rigging. The standardized election fingerprint is calculated to account for geographic anomalies. The results were checked on whether they deviated for small polling stations. $p = 10\%$ is fixed for the purpose of the analysis as the threshold between small and large electoral units. Voter rigging is more likely to occur in small polling stations in that it is easier to identify opposite voters, there are in general fewer eyewitnesses than in large polling stations and small polling stations are usually not visited by election observers. [Klimek et al., 2017]

The plot below shows that the fingerprint for small polling stations is shifted to the upper-right which is expected if voter rigging was present. Higher turnout and higher percentage of the "Yes" votes in small polling stations are observed.

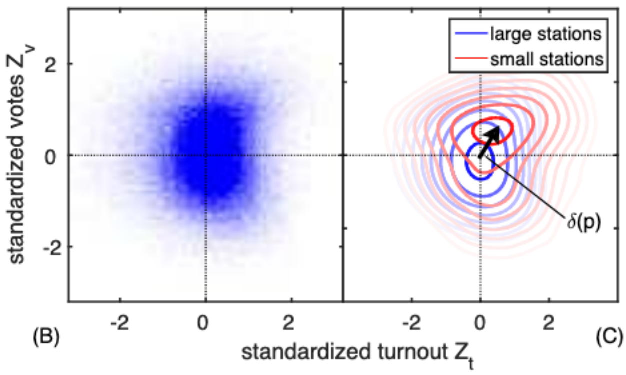

Figure 10: Standardized election fingerprint of Turkey referendum 2017 [Klimek et al., 2017]

The voter-rigging test was also conducted on the Turkish referendum. Klimek et al. were able to show that the constitutional referendum did not match the criteria needed for being considered a trustworthy election. The deviations are significant although small compared to the elections in Venezuela and Russia which is a result that is similar to the one obtained via the ballot stuffing test. But the proportion of votes added by small electoral units was crucial to push the result of the "Yes" votes above the 50% mark too. This aspect shows that the effect of these small electoral units can under no circumstances be neglected.

The authors state that voter rigging was present in the constitutional referendum in Turkey in 2017 with the magnitude of the effect being significant, but not comparable to results from Venezuela or Russia. The analysis of voter rigging also allowed to determine which regions were most affected by voter rigging. This analysis was performed by averaging the value of $\delta(p)$ for each province and $0 < p < 1$. The provinces most affected were spread equally over Turkey, but they had low population density in common. The authors state that they could find evidence for the presence of ballot stuffing and voter rigging at the constitutional referendum. [Klimek et al., 2017]

Results for Venezuela

Another example for analysing the quality of the presented methods is Venezuela. Jiménez and Hidalgo analysed various elections during the presidency of Hugo Chávez (1999-2013). Chávez was always said to have strong support among Venezuelan voters and it was therefore no surprise that he got elected. However, during his presidency questions about the integrity of the Venezuelan voting system were raised.

The advantage from an analyst's perspective is that data of similar, comparable elections is available, where the claim that some of those elections were fair while fraudulent mechanisms could have been present in others can be proved. The elections included are presidential elections in 1998, 2000, 2006 and 2012; referenda in 1999, 2004, 2007 and 2009 and parliamentary elections in 2005 and 2010. First claims of irregularities were alleged at the 2004 referendum which decided on whether Chávez should be removed from office which Chávez won with 59% of the vote. Election observers denied election fraud in the referendum, but it is still seen as a turning point as of 2004 onwards fraud allegations in Venezuelan elections increased. [Jiménez and Hidalgo, 2014]

Jiménez and Hidalgo first conducted a second digit Benford test for various aggregation levels with the result that p-values for the χ^2-test are lower for elections between 2004 and 2012 than for the period before 2004, but this result will not be interpreted as the method's weaknesses were outlined. The authors took a mixed approach between techniques discussed above in that they chose the election of 1998 and the referendum of 2004 as reference elections and determined for all other elections to which of the reference elections they were mathematically "closer". These two electoral fingerprints show big differences. [Jiménez and Hidalgo, 2014]

At that point there is no indication for elections associated with the referendum in 2004 necessarily being fraudulent elections. The analysed elections can be categorised in four different groups. The first category includes early elections in 1998 and 2000 which show similar fingerprints to the first reference election. Jiménez and Hidalgo calculate the "Mod.98" value which returns the probability of an election being associated with the 1998 presidential election which is obviously high in the first category. Elections in the second category (2004, 2006, 2009 and 2012) are characterized by low values of Mod.98 and by containing units with high turnout and high percentages for the winning party which is an indicator for potential election fraud according to Klimek et al.. The third category contains the referendum in 2007 and the parliamentary elections in 2010 with the property that they do not show any closeness towards one of the two reference elections. Jiménez and Hidalgo classify them as mixture models. The fourth and last category consists of the referendum in 1999 and the elections in 2005. They show different properties compared to the elections in the other three categories and are characterized by their low turnout and high percentage numbers for Chávez. Political observers explain that phenomenon by referring to the low turnout of opposition voters in these elections which is not necessarily an indicator for election fraud.

By having a look at the cumulative number of votes as a function of the turnout, the authors obtain a first overview. The elections in the fourth and last category showed what the function would be expected to look like. It took on a sigmoid form reaching a plateau for the votes favouring Chávez. The low turnout in these elections was visible in the plots in that the plateau is already reached for low turnout values. Elections in the first category (1998 and 2000) also took on the expected shape, but reached the plateau for moderate turnout values instead of the low ones. Elections in the second category (which are closely linked to the 2004 referendum) showed different forms than expected in fair elections as they showed a late increase close to 100% turnout which Klimek et al. consider to be an indication for ballot stuffing. The remaining elections and referenda in the third category were considered to be mixture models. The plot is shown below. [Jiménez and Hidalgo, 2014]

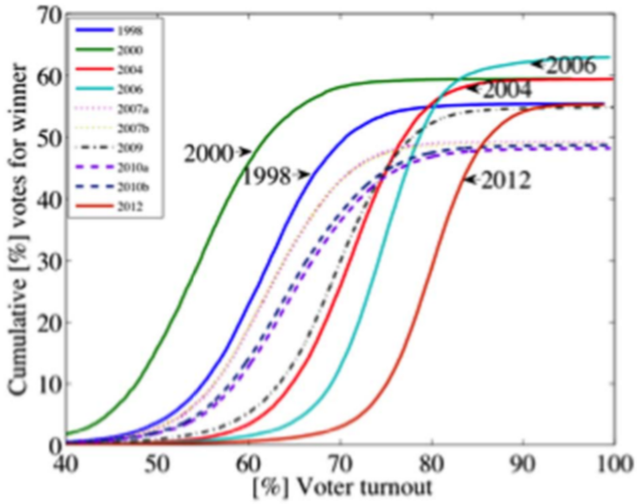

Figure 11: Cumulative votes as a function of turnout for Venezuela [Jiménez and Hidalgo, 2014]

Jiménez and Hidalgo see their initial suggestion that the referendum in 2004 represents a breakpoint confirmed. There are exceptions to that claim which are elections that occurred after the 2004 referendum, but do not take similar shapes as the referendum (e.g. 2007). However, the authors mention that there are many potential factors leading to deviations in the fingerprints, one of which being the existence of special electoral units in which the voting population votes for the winning candidate in high percentages without any special reason or the motives being of local importance only. The political situation in Venezuela is special in that there are many polarized geographical areas. Therefore it is no surprise to detect electoral units with considerably high support for Chávez or the opposition. To further examine this effect, Jiménez and Hidalgo continued to analyse the special places i.e. they analyse "atypical support (for Chávez) in electoral units, relative to the support obtained in the polling centre to which the unit belongs."[Jiménez and Hidalgo, 2014]
This effect is enabled, because voters in Venezuela can select the polling centre. In polling centres which consist out of two or more electoral units, voters are distributed to the different electoral units according to a pseudo-random criterion. Jiménez and Hidalgo therefore condition on the overall result of the polling centres and conclude that the number of votes per electoral unit is distributed according to a Hypergeometric distribution. [Jiménez and Hidalgo, 2014]

V denotes the number of votes for Chávez in an electoral unit, p the proportion of votes for Chávez divided by the number of registered voters at the entire polling centre and n and m denote the number of registered voters in the electoral unit and the entire polling centre. Modelling V given the parameters n, m and p is of interest. The Hypergeometric distribution is used in case one analyses a population with size N out of which exactly M have a certain feature. n objects are randomly sampled out of the population without replacement and the number of elements having the specified feature in the sample is modelled.

The setting is adapted here. The population size is given by the number of registered voters at the centre and the specified feature is that they voted for Chávez. The proportion is given by p. The sample size is the number of voters registered in the electoral unit (n) and the number of the voters in the sample (electoral unit) who voted for Chávez (V) is analysed.

By applying simple formulas for the Hypergeometric distribution, the expected value is $\mathbb{E}[V] = n \cdot p$ and the variance is $Var(V) = p(1-p)n(m-n)/(m-1)$. Jiménez and Hidalgo introduce a standardized measure for the regularity of the number of votes for Chávez in the electoral unit as

$$Z = \frac{V - p \cdot n}{\sqrt{p(1-p)n(m-n)/m-1}}$$

The standardization has the advantage that values of Z far away from zero imply irregular support for Chávez and rule out the case of special polling centres or special areas as these effects were already accounted for in the calculation of Z. Another useful property of the Hypergeometric distribution is that it approximates a standard normal distribution in case the sample size is high enough and the population size is still considerably higher than the sample size. There is a possibility of actions affecting the vote distribution of specific electoral units, but that does not affect the vote distribution of the entire polling centre significantly - these irregularities are considered to be non-fraudulent and can occur regularly due to the complexity of the election process which leads the authors to expect Z to have a distribution with heavier tails than the standard normal distribution. [Jiménez and Hidalgo, 2014]

Hence a student-t-distribution with three degrees of freedom ($t(3)$) is chosen to model Z. The distribution is used to simulate Z in a bootstrap, which is named the standardized differences. In a fair election Z scores far away from zero are due to chance and fall into the category of non-fraudulent irregularities. A set M_k which consists of k electoral units which all have Z scores far away from zero is introduced. The null hypothesis for the upcoming analysis is that all electoral units have the same probability of being part of M_k which is the case in fair elections.

r_k is the proportion of votes for Chávez over all valid votes in the set M_k and R is the same proportion, but calculated over all electoral units. At electoral unit i, T_i is the total number of votes and V_i the number of votes for Chávez. K is the total number of electoral units. Consequently the estimated variance within the k units included in M_k is

$$s_k^2 = \frac{1}{k-1} \sum_{i \in M_k} (V_i - r_k T_i)^2$$

.

μ is defined as the average of valid votes per electoral unit. The estimated variance of r_k is

$$S_k^2 = \left(1 - \frac{k}{K}\right) \frac{1}{\mu^2} \frac{s_k^2}{k}$$

which is due to the definition of r_k. [Jiménez and Hidalgo, 2014]

Under the null hypothesis that all electoral units have the same probability of being contained in M_k, the quantity of the standardized difference $\xi_k = \frac{r_k - R}{S_k}$ is approximately distributed according to a standard normal distribution if k, K and the difference between them are sufficiently large. That enables to construct a test which computes the quantity ξ_k for large values of k. If the analysis shows values far away from usual standard normal confidence intervals for a broad range of k, they are considered as evidence against the null hypothesis. The model of a fair election is based on a hierarchical bootstrap which means that random samples of size K are drawn from a $t(3)$ distribution which function as simulated Z-scores. The k scores which are furthest away from zero are assigned to another random sample of units and ξ_k is calculated. That way a model under the null hypothesis can be visualized and specified. [Jiménez and Hidalgo, 2014]

The results of the simulation were compared to the results of the elections and the analysis confirmed our prior observations that the group of elections close to the election of 1998 laid within the confidence interval, but the elections of the group associated with the 2004 referendum was higher than the results returned by any of the simulations. Jiménez and Hidalgo reject the null hypothesis for all elections from 2004 onwards except for the election in 2005.

The interpretation of that rejection is not as easy as to state that these elections were fraudulent. They only indicate that the irregularities occurred on a non-random set of electoral units and on this non-random set Jiménez and Hidalgo found a significant bias in favour of Chávez (or his party in parliamentary elections/his position at a referendum). The authors state that there exist a variety of reasons why some electoral units may be more likely to be subject of non-fraudulent irregularities, but what causes suspicion is that these deviations only occurred from 2004 onwards and they voted in favour of Chávez. Therefore the analysis of the irregularity in election support also points towards the referendum in 2004 being a breakpoint in Venezuela's elective history. [Jiménez and Hidalgo, 2014]

The authors performed one last analysis of elections during the Chávez era in which they looked at irregular variations in the electoral register. As there is no reasonable argument for why the irregularities only occur in pro-Chávez electoral units, the influence of these electoral units on the overall election result is checked.

In Venezuela the electoral units assignment to polling centres can differ from one election to another which can theoretically cause a single voter to be counted towards one polling centre in one election, while being counted towards another centre in the next election. Yet many of the electoral units remain assigned to the same polling centre throughout consecutive elections and therefore Jiménez and Hidalgo focus on polling centres instead of remaining on the lowest aggregation level.

One major concern regarding elections in Venezuela are the ongoing allegations of the opposition regarding the potential manipulation of the voter register by the government. The Venezuelan voting population grew about 60% between the first and the last election within the analysis, whereas the actual population only grew by about 16%. [Jiménez and Hidalgo, 2014]
This thesis is not meant to explain the deviations, but by analysing methods developed for that purpose, the correlation structure between the growth of the voting register and the election results for Chávez can be checked.

For that purpose Jiménez and Hidalgo focused on the inter-annual growth of votes in polling centres. Notation-wise, $m(t)$ is used to denote the number of registered voters in a polling centre in year t. The authors use t^- to denote the last election year before t which is included in the available data. The inter-annual growth in that voting centre in year t is given by

$$G(t) = 100 \cdot \frac{log(m(t)) - log(m(t^-))}{t - t^-}$$

.

This measure compares the number of registered voters of two selected years and divides through the length of the time span in-between. It is compared to the annual population growth of Venezuela which was at about 1.5% at the time of the paper being published and at 1.3% in 2017. These values only correspond in theory as the values of $G(t)$ may take values that are 20 times larger than the population growth. [Jiménez and Hidalgo, 2014]
Another fraud allegation brought up by the opposition during the second half of the Chávez presidency is the relocation of already registered voters to different polling centres which affects the turnout among these groups as the next polling centre may be easier/more difficult to reach. In terms of $G(t)$ that may generate negative values of the inter-annual growth rate which is why the absolute value of $G(t)$ is considered for the inter-annual variation at centre t.

To connect this measurement with the votes favouring Chávez, one computes the proportion of votes in favour of Chávez as a function of $|G(t)|$. The authors compute the percentage of votes favouring Chávez for each value of the inter-annual variation from polling centres with this value for $|G(t)|$ or lower. To make the analysis comparable the curves were centred by subtracting the overall percentage of valid votes that Chávez received in that election (R). [Jiménez and Hidalgo, 2014]

The result was that the elections showed no common curve, but differed in their shape. While elections prior to the year 2004 and one election in 2005 show a positive fluctuation, other elections such as 2006, 2009, 2010 show a negative fluctuation. The elections of another category (2004, 2007, 2012) even show non-linear negative relationships. Jiménez and Hidalgo conclude that there is an increase of the winning margin close to extreme values of $|G(t)|$ and they outline the special role that the 2004 recall referendum played in the electoral history of Venezuela. The growth rate for this referendum was two times larger than the usual growth rate in the years before.

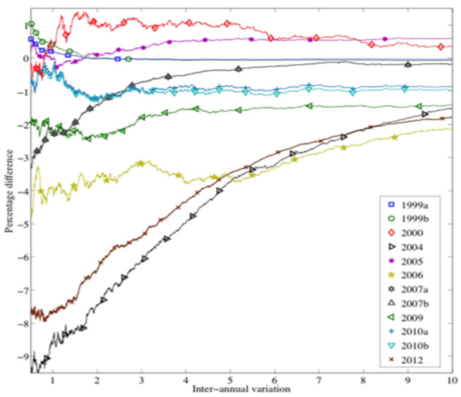

Figure 12: Difference in percent as a function of inter-annual variation [Jiménez and Hidalgo, 2014]

One potential explanation for the effect is a government program initiated by Chávez which helped people in poor and rural areas to register for voting. Before the year 2004 the poorer population usually voted for Chávez in large numbers, but political scientists agree that this was not the reason for his win in

2004. In 2004 Chávez performed well across almost all demographic groups and socio-economic statuses and the shape of the 2004 curve cannot be explained by the government program. The authors state that "both, 2004 and 2012, show an irregular pattern that suggests a strategic inter-annual variation in the electoral roll. Furthermore, this variation was decisive for winning the 50% majority." [Jiménez and Hidalgo, 2014]

By analysing different types of potentially outcome determining phenomena, the initial suggestion that the 2004 referendum marked a turning point in the electoral history of Venezuela is confirmed. The reason for extensively discussing the paper by Jiménez and Hidalgo was that they showed numerous methods for detecting phenomena not analysed before. An important point is their detailed method of analysing the irregularities in voter support for Chávez which underlines the previously only supposed effect by looking closer at the neighbourhood of electoral units.

The paper also added an additional component as it analysed potential manipulation of the registering process. The closer look at the growth rate of the electoral population proofed to be interesting in the example of Venezuela and can be used in other countries as well.

3.7 Other fraudulent mechanisms

In the introduction of this section different forms of election fraud were mentioned. In this section additional methods to detect and analyse other forms of fraud than ballot stuffing or voter rigging are discussed.

One possibility to do so is by comparing election results to the results of an exit poll. That way no concrete type of fraud can be discovered, but it yields an indication on potentially fraudulent electoral units if the election outcome and the result of the exit poll differ widely. One disadvantage is that this method relies on the quality of the exit poll and a potential bias in the sample of the exit poll leads to wrong conclusions.

Another example is not to base the analysis on the question of fraud being present, but to analyse the impact fraudulent votes potentially could have on the overall election result. The perfect example for illustrating that approach is the Austrian presidential election of 2016. It is known that Austria is not one of the countries having problems with electoral manipulations, but the occurrences in 2016 gathered the media attention to potentially fraudulent election results in Austria.

The potentially fraudulent election is the run-off between Alexander Van der Bellen, former leader of the Austrian Green Party, and Nobert Hofer of the Austrian Freedom Party (FPÖ). These candidates had the best results in the first round of the presidential election, but both of them failed to gain a majority of 50% needed for a win. The second election was known to be a close contest, resulting in a win for Van der Bellen by a margin of about 30,000 votes (50.35% of the overall vote). After the election allegations of irregularities in the voting process were raised which climaxed in the decision of the Austrian constitutional court to order a repetition of the election. [Neuwirth and Schachermayer, 2016]

The discussion was centred around the claim that about 77,000 votes were "contaminated" for different reasons and could therefore not be seen as regular votes although they were part of the final result. As the quantity of the contaminated votes was higher than the actual winning margin, the court

considered the occurrences around these votes in their entirety as potentially decisive for the election outcome.
An important aspect of the discussion was the question of how likely it would have been that these contaminated votes changed the outcome of the election.

For that purpose Neuwirth and Schachermayer developed a statistical model to determine this probability. This is not a method to detect election fraud, but it is a method for quantifying the degree of influence of a known number of "non-regular" votes. Among these non-regular votes were votes which were opened too early, some were opened at moments when not all members of the election committee were present and others stem from problems with the system of mail voting. [Neuwirth and Schachermayer, 2016]

It is of course theoretically possible that the potential miscounting of the non-regular votes was decisive for the outcome of the election.
Voters in Austria have the possibility to vote in two different ways. They can either vote personally at polling stations or via mail before election day. The claim of Hofer and his party was that polling centres had opened the mail votes too early and therefore these votes should have been invalid. The constitutional court investigated these claims and concluded that violations of the law had occurred in 11 of 117 voting districts which affected around 77,000 votes and therefore ordered a repetition of the run-off. Note that there was no concrete proof of fraud, it was only proven that the legal procedure had been violated. [Neuwirth and Schachermayer, 2016]

The first analysis Neuwirth and Schachermayer conducted was a comparison of the 106 regular districts and the 11 non-regular districts. For that purpose they plotted Hofer's ballot percentages for all voting districts against his percentages among the mail votes. The non-regular districts were highlighted in the plot. It showed no spectacular behaviour of the non-regular districts, but it demonstrated general facts about this election. The proportion of Hofer voters using the possibility of the mail vote is smaller than the proportion of Van der Bellen's voters doing so. Hence Hofer performed better among the ballot voters than among the mails voters. [Neuwirth and Schachermayer, 2016]
But what the plot could show was that there is a linear relationship between the percentages among the ballot votes and the percentages among the mail votes.

In a second analysis the non-regular votes were excluded and only the ballot votes for the 11 districts marked as non-regular were included i.e. the mail votes of these districts were neglected. The degree by which the regular votes differed from the non-regular votes could be determined and the impact of these non-regular votes on the result as well. The missing votes of the non-regular districts were then assigned proportionally to the missing districts and equally divided among the candidates. That scenario improves Hofer's performance in those districts, but these districts form a group of outliers in the modified plot. That being said Neuwirth and Schachermayer classify this scenario as unlikely. [Neuwirth and Schachermayer, 2016]

For the further analysis the assignment of the missing votes to districts is irrelevant. The only variable of interest is the total sum of the missing, non-regular votes and the probability that these results influenced the election outcome in a way such that Hofer would have won the election. To model that quantity Neuwirth and Schachermayer use a weighted linear regression which allows for heteroscedasticity as the variance depends on the number of overall votes in the district. For the non-regular votes the authors assume that they follow a similar behaviour to the regular votes. [Neuwirth and Schachermayer, 2016]

In the model the total number of voting districts is $N = 106$ with the number of valid votes being t_n in district n. In a district n there were v_n votes for Hofer and $\tilde{v}_n$ votes for Van der Bellen such that they sum up to t_n. The total votes can be divided on whether the vote was cast via mail or in person. Notation-wise, Neuwirth and Schachermayer denote the ballot votes in district n as b_n and the mail votes as m_n. These votes can be split according to the candidate the vote was for: $v_{b,n}, v_{m,n}, \tilde{v}_{b,n}$ and $\tilde{v}_{m,n}$. The parameters of interest are $v_{b,n}$ and $v_{m,n}$ which are the votes for Hofer. While the ballot votes are in the data, the mail votes are treated as realizations of the random variable

$$V_{m,n} = k \cdot v_{b,n} + \epsilon_n, n = 1 \ldots N$$

[Neuwirth and Schachermayer, 2016]

The linear relationship was already discussed above. In the formula, k is an unknown fixed number, while ϵ_n are independent, centred Gaussian random variables. The variance of this random variable is $\sigma^2 m_n$ and is therefore dependent on the total number of mail votes and positive. [Neuwirth and Schachermayer, 2016]

Hofer's mail votes depend on the amount of mail votes cast in this voting district. The consequences are that it has to allow for heteroscedasticity. The authors conduct the regression and obtain the estimators $\hat{\sigma}^2$ and $\hat{k}$. The estimator $\hat{k}$ is treated as a random variable which follows a t-distribution which can be specified via the election data.

This procedure was used to model votes in the regular districts, but can be adjusted to deal with the $M = 11$ non-regular districts. The key assumption is that theses votes follow the same model as the votes in the regular districts if a fair election is assumed. Hence they can be written as

$$V_{m,j} = k \cdot v_{b,j} + \epsilon_j, j = 1 \ldots M$$

The true value of k is unknown in these districts and is replaced by the estimator obtained in the regular districts $\hat{k}$. The new noise terms ϵ_j are assumed to be independent of the noise terms in the regular districts and their variance depends on the total number of mail votes in the non-regular districts i.e. $Var(\epsilon_j) = \sigma^2 m_j$. The parameter of interest is the total number of mail votes for Hofer in the non-regular districts which is given by [Neuwirth and Schachermayer, 2016]

$$\hat{V} = \sum_{j=1}^{M} \hat{V}_{m,j}.$$

Denote $v_b = \sum_{j=1}^{M} v_{b,j}$ as the total number of ballot votes for Hofer in the 11 districts and the total number of mail votes in the non-regular districts as m. According to the model

$$\hat{V} = \hat{k} v_b + \epsilon$$

From common regression knowledge it follows that if σ^2 was known, $\hat{V}$ would be normally distributed with known mean and variance. As σ is unknown, Neuwirth and Schachermayer replace it with the estimator $\hat{\sigma}$. Inserting the parameters the standardised form

$$\frac{\hat{V} - \hat{k}v_b}{\hat{\sigma}\sqrt{\frac{v_b^2}{\sum_{n=1}^{N}\frac{v_{b,n}^2}{m_n}} + m}}$$

is obtained which follows a t-distribution with 105 degrees of freedom. [Neuwirth and Schachermayer, 2016]
A test is constructed by comparing the value of V to the critical value $\tilde{V}$ which is the number of votes Hofer needs to win the election. The calculation of $\mathbb{P}[V \geq \tilde{V}]$ yields the probability that the non-regular votes would have changed the outcome of the election.

The results are used to compute confidence intervals. The model is given in standard form by $y = X\beta + \epsilon$ where the variance-covariance matrix of the errors is $VC(\epsilon) = \sigma^2 W$, where W is positive definite. The best linear unbiased estimator for the parameter β is the generalized least squares estimator (GLS) which is given by $\hat{\beta} = \left(X'W^{-1}X\right)^{-1} X'W^{-1}y$. This estimator is unbiased with variance-covariance matrix $VC(\hat{\beta}) = \sigma^2(X'W^{-1}X)^{-1}$.

In the concrete example of Neuwirth and Schachermayer only one regressor is included and therefore X simplifies to a $N \times 1$ matrix, W is a diagonal matrix with the quantities m_n on the diagonal with the off-diagonals being zero. The parameter β simplifies to the scalar k which was estimated by using $\hat{k}$. Therefore the variance-covariance matrix of $\hat{\beta}$ simplifies to the variance of $\hat{k}$ which was already mentioned to be [Neuwirth and Schachermayer, 2016]

$$Var(\hat{k}) = \sigma^2\left(\frac{1}{\sum_{n=1}^{N}\frac{v_{b,n}^2}{m_n}}\right)$$

The parameter of interest, $\hat{V}$, is a simple linear transformation of $\hat{k}$ and ϵ such that

$$\mathbb{E}\left[\hat{V}\right] = \mathbb{E}\left[\hat{k}v_b + \epsilon\right] = \mathbb{E}\left[\hat{k}v_b\right] + \mathbb{E}\left[\epsilon\right] = kv_b$$

as ϵ has an expected value of zero. The variance is then given by

$$Var\left[\hat{V}\right] = Var\left[\hat{k}v_b + \epsilon\right] = Var\left[\hat{k}v_b\right] + Var\left[\epsilon\right] = \sigma^2\left(\frac{v_b^2}{\sum_{n=1}^{N}\frac{v_{b,n}^2}{m_n}} + m\right)$$

If σ^2 is replaced with its estimator $\hat{\sigma}^2$ the t-distribution specified above is obtained for the standardized term. This term is calculated by subtracting its expected value of $\hat{V}$ and dividing by its standard deviation. [Neuwirth and Schachermayer, 2016]

The model is specified and a confidence interval is calculable. To determine the threshold number of votes that Hofer needs in the non-regular districts to change the election outcome, one more calculation is performed. Hofer had 34,479 votes in these districts and would have needed an additional 15,432 which results in a threshold of $\tilde{V} = 49,911$. According to the model the probability of him reaching that value is at very low $p = 1.322 \cdot 10^{-10}$. [Neuwirth and Schachermayer, 2016]

It is not the concrete result which indicates that the non-regular votes with high certainty did not change the outcome of the election that is of particular interest, but the strength of the method. Its simplicity and interpretability cause the range of possible applications to be wide. Suppose special districts are subject to extensive voter rigging. One could treat the votes in these districts as non-regular and check the probability of these votes changing the overall election results. This statement does not imply that minor cases of voter rigging should be neglected, but usually systematic fraud occurs on a level such that it is affecting the overall result as that is the goal of performing election fraud.

The same is true for the analysis of ballot stuffing as there may be examples where fraud is known to have occurred in certain areas. These areas could be classified as non-regular votes in case the election observer has some prior knowledge about the plausibility of an occurrence of ballot stuffing.

However, the use of the model is justified by the linear relationship between the ballot votes and the mail votes which may be an Austria specific phenomenon. Therefore the application of the model may lead to difficulties in other problem settings. The assumption about the linear relationship is still not a strong one and may be fulfilled in other countries as well.

3.8 Discussion

In this section different approaches of tackling the problem of election fraud in a data-driven way were discussed. The assumption was to be in the position of an election observer who had access to election results on low aggregation levels. The introduced methods turned out to be capable of detecting various forms of anomalies in the data.

The different types of election fraud were discussed and divided in three categories with no claim of completeness. Each type of election fraud has different origins and therefore also the ways in which they can be analysed or detected differ. There is no single measurement to determine whether any type of fraud is present in an election and therefore various methods needed to be introduced, each of which was tackling a specific type of election fraud.

The first method discussed was the second-digit Benford test which assumes a distribution of the second significant digit in election results. The intention with Benford's Law is to detect numbers which were made up by an election official. A test for checking whether the data is distributed according to Benford's Law was presented which is seen as an indication of election fraud in case the numbers do not follow the Benford distribution. It turned out that this procedure can not deliver anything more but an initial indication. Significant results of the second digit Benford test do not proof the presence of election fraud and if the test indicates that no fraud is present, no conclusion about the election being fair can be made. The second digit Benford test is therefore not applicable as simple examples of when the test fails were shown. The example of Deckert et al. where they simulate one fair election and one election where they insert a fraud factor in the model demonstrated this problem. The result was that the second-digit Benford test's results were of no use and problematic in case they were used by election observers.

The second method analyzed was the election fingerprint. The intention of this method is to discover the practice of ballot stuffing where votes for the winning party are added towards the ballot boxes until the turnout in these districts reaches numbers close to 100%. It turned out that the electoral fingerprint is a useful tool for discovering the presence of ballot stuffing. It exploits that ballot stuffing causes correlation between turnout and the percentage of the winning party. The fingerprint was introduced by Klimek et al. who applied it to a wide range of different elections. They found that elections in Russia in 2011 and 2012 and an election in Uganda were subject to ballot stuffing. Their model estimated two different parameters via a goodness of fit procedure. One parameter described the presence of extreme fraud which is present if the winning party and the turnout reach results close to 100%, the second parameter described the presence of incremental fraud which is the practice of taking away votes from the opposition party and adding them to the total vote count of the winning party. Extreme fraud is characterized by a small second cluster in the upper right corner of the election fingerprint. Incremental fraud causes the fingerprint to be smeared to the right.

Another model type discussed was a model to detect the practice of voter rigging. Thereby the intimidation and coercion of voters to force a certain elective behaviour was meant. The method published by Jimenez et al. was discussed which compares the standardized election fingerprints of small electoral units with those of larger electoral units. This method is based on the observation that voter rigging is more common in small polling stations than in large ones. Standardized election fingerprints account for the neighbourhood of the electoral unit and therefore remove geographic anomalies which cause strange forms of the usual election fingerprint which thereby makes them more robust. The authors also found evidence for this practice in Russia and Venezuela. They designed a test which is based on the modified Thompson τ test for detecting outliers. Thereby a pool of elections was considered to be trustworthy and if elections were not to be found in a certain range of the trustworthy elections, they were determined to have been subject to voter rigging.

To illustrate the procedures, different real-world applications were described. The first one was the referendum deciding on the transformation of the constitution in Turkey which resulted in a narrow win for the president and fraud accusations being raised soon after. Klimek et al. conducted an analysis and checked the referendum for ballot stuffing and voter rigging. The result was that the referendum was significantly influenced by incremental ballot stuffing, but not in the same magnitude as the results from Russia. The voter rigging test also showed a significant result. Combining these two results indicates that the election was problematic as the result was close and the two types of fraud were potentially decisive for the outcome.

A closer look at the paper of Jiménez and Hidalgo on elections in Venezuela during the Chávez presidency was conducted. Different analyses were performed which came to the conclusion that there was a breakpoint in the election history of Venezuela in the year 2004. Before that year elections appear to be fair, but afterwards significant anomalies appear. The combination of these approaches showing similar results also strengthens the claim of the referendum being a breakpoint election in Venezuela.

The last model took a different approach. Neuwirth and Schachermayer looked at the Austrian presidential election of 2016 in which the result was close as well and the constitutional court ruled that because of irregularities in eleven voting districts the overall result could have been influenced decisively. The model analyses the probability of the non-regular votes having caused a different election outcome. They showed that the probability of that scenario is extremely small. The strength of the analysis is that it could potentially be used to determine the influence of districts in which a higher probability of fraud occurring is assumed. It could be a valuable tool in determining the magnitude of potential election fraud.

It is expected that in the next years the presented procedures will be tested on different elections for justifying the models in case they show credible results on a wide range of elections. Still the practical use for election observers seems to be decades away. Essential for the right adaptation of the presented models is the publication of election data for low aggregation levels which is often problematic in countries with tendencies towards election fraud.

4 Conclusion

In this final part of the thesis the practices of gerrymandering and election fraud will be compared. By outlining similarities and differences, connections between the two parts shall be shown.

Gerrymandering and election fraud both undermine the fairness of an election, but the two methods intervene at different points. While election fraud is directly influencing the voting results, gerrymandering does not do so and is therefore considered legal. Gerrymandering biases the transition between the voting results and the seat distribution in the elected legislative body which can also be decisive in determining which party wins.

One of the biggest differences is the way they are conducted. Gerrymandering theoretically lies in the hands of a single person who tries to design a district plan benefiting his party. In practice the gerrymanderer does not design the district plan on his own, but the group of people doing so is comparably small. This gerrymanderer tries to give his party the maximum benefit, but possibly without raising suspicion which can be seen as a constraint in the optimization problem. The approach when performing ballot stuffing is similar. The goal of the party performing fraud is to win the election. In order to not raise suspicion, the party tries to perform the minimum amount of fraud necessary to win the election such that the fraud is hard to detect for election observers and the opposition. It is seen as one of the most important criteria of a stable democracy to hold free and fair elections and therefore autocratic regimes try to convince the international community that elections justifying their reigning are held in a fair manner.
Both practices have to find balance between the best possible result for the party without raising doubts about the integrity of the election. As the practice shall be hidden, it causes the detection to be even more difficult.

Methodically, modelling gerrymandering and election fraud also shows differences. In both cases the procedures only rely on few variables. While in case of ballot stuffing or voter rigging the focus is on the size of electoral units, turnout and the percentage of the winning party, the gerrymandering models assume to have some prior knowledge about the party preference of individual voters. Gerrymandering is stated as an optimization problem to find the best strategy which is not possible in the analysis of election fraud. An "optimal strategy for performing election fraud" does not exist as magnitude of fraud can always be augmented, but the result's credibility is suffering. Both procedures are strongly influenced by geographical factors. Gerrymandering has the constraint that districts have to be connected and in case of ballot stuffing and voter rigging different regions are known to have different voting behaviour. For a fraud practice such as voter rigging the consequences are that the governing party persuades people in party-friendly regions to take part in an election while making it as hard as possible for regions in which the opposition is expected to perform well.

The geographical constraints do not only influence the way the practices are enforced, but also the way models try to explain gerrymandering and election fraud. In case of gerrymandering these constraints prevent the gerrymanderer from creating the ideal districting for his party as house blocs cannot be split in different voting districts although the gerrymanderer may have information about different voting behaviours. In case of ballot stuffing the procedure of the election fingerprint takes on non-standard forms for many different elections. These forms can often be explained by accounting for geographical factors such as the Canadian election in which the French population voted entirely different than the English speaking population or the election in Finland where one party was successful in mobilizing people across the country, but failed to do so in the capital. These observations could all be explained by comparing the electoral units to their "neighbourhood" which leads to another major concern in the analysis.

For analysing gerrymandering and election fraud, the availability of election data for small aggregation units is essential. For gerrymandering voter preference should be known as detailed as possible to realize the optimal strategy. When analysing a district plan, party-leaning areas have to be known to the observer to realize whether these blocs were shifted to a different voting district as a result of a packing gerrymandering strategy. The illegal aspect of gerrymandering is to use it for denying minorities their proper representation in parliament. Therefore information about minorities, the areas where their votes are cast and their voting behaviour are of high relevance for the analysis.
In case ballot stuffing or voter rigging are present these practices are known to be hidden easily in large aggregation levels. Ballot stuffing can only be discovered if turnout numbers and results of the winning party are known for fine aggregation levels. These fine aggregation levels have the advantage that the fingerprint can be calculated in more detail, but most importantly electoral units with "special" voting behaviour can be detected.
The detection of voter rigging is based on the differences between small and large polling stations. The methods are not applicable if no data about the size of the electoral units is published.

Not only do the methods for modelling these practices show similarities, but also the ways of detecting them. One intuitive procedure to detect partisan gerrymandering is to simulate the election outcome for different randomly assigned districtings, forming a null model with those and comparing it to the actual election result. Large deviations in favour of the party in charge of the redistricting process could be seen as an indication for gerrymandering.
A similar procedure was applied on election fraud. As correlation between the percentage of the winning party and turnout indicates fraudulent mechanisms, robustness tests were conducted by randomly assigning the turnout towards the voting units to check whether significant results would disappear then, corresponding to the null model in the gerrymandering example.

That leads to the interesting discussion about the complexity of the detection mechanisms. Many of the presented approaches to detect gerrymandering are intuitive and comparably simple. The main focus is to look at the symmetry of the results under the districting plan. If one party is able to win 80% of the seats with 60% of the popular vote, this also has to be true for the other party. The most intuitive approach was to analyse the wasted votes which was ideal for checking which party benefited from the districting. The method showed weaknesses when focussing on the effect of homogeneous swings in the popular vote which changed the outcome of the detection method while the same districting plan was in place.
The methods for deciding on the presence of ballot stuffing or voter rigging were more complex. Via a maximum likelihood method a model for a fair election in a country is constructed depending on the data and the realization of a normal distribution. A goodness of fit test was used to determine concrete values of fraud parameters for which many different combinations of parameter values had to be inserted. That procedure is computationally more expensive and the results are not as easy to interpret. One advantage of this procedure is that it is not possible to construct simple counterexamples of when the detecting method fails. In all elections discussed, the method returned plausible results.

Yet one has to admit that the results obtained are volatile in both cases. In gerrymandering, different assumptions resulted in different optimal strategies and simple counterexamples showed that there are cases in which the detection methods did not work. In election fraud the entire procedure is based upon simulation and there are many insecurities involved in the construction of the null model. One approach which might be helpful and was presented in the section on election fraud, is to focus on the question which possible influence fraudulent votes had on the overall election result. In the context of gerrymandering one could declare areas where potentially gerrymandering for political purposes was present, as fraudulent votes and check their influence. By doing so, the robustness of the results would be raised.

It remains to be discussed how practices of election bias can be prevented from happening in future which is a question with two different answers for gerrymandering and election fraud. Gerrymandering could be prevented by not allowing political parties to draw district lines any more. The problem seems to be that political parties themselves have to take that decision and consciously give up privileges in favour of election fairness.
Different types of election fraud are harder to be prevented. One way is to make results of data-driven approaches known to a more general public to raise pressure on autocratic regimes for organizing fair elections. Nevertheless it is still likely that in the following years different additional types of election fraud will appear which require new methods of analysis.

This thesis concludes that the parallels between gerrymandering and actual election fraud are a strong argument for analysing these phenomena together in future works. The difficulties in detecting the practices are similar and therefore accomplishments in detecting either gerrymandering or election fraud may yield benefits when trying to model the other form of election bias. This thesis shall give a broad overview towards the topic of election bias and will hopefully be built upon by publications in this area. Election data for fine aggregation levels and many different elections will be essential to assess and further calibrate the presented methods in this thesis.

5 References

Bernd Beber and Alexandra Scacco. What the numbers say: A digit-based test for election fraud using new data from nigeria. In *Annual Meeting of the American Political Science Association, Boston, MA*, 2008.

Frank Benford. The law of anomalous numbers. *Proceedings of the American philosophical society*, pages 551–572, 1938.

Arno Berger and Theodore P Hill. Benford's law strikes back: no simple explanation in sight for mathematical gem. *The Mathematical Intelligencer*, 33(1):85, 2011.

Arno Berger, Theodore P Hill, et al. A basic theory of benford's law. *Probability Surveys*, 8:1–126, 2011.

Robin E Best, Shawn J Donahue, Jonathan Krasno, Daniel B Magleby, and Michael D McDonald. Considering the prospects for establishing a packing gerrymandering standard. *Election Law Journal*, 17(1):1–20, 2018.

Nigel Blundell. *The world's greatest mistakes*. Conran Octopus, 1980.

Christian Borghesi and J-P Bouchaud. Spatial correlations in vote statistics: a diffusive field model for decision-making. *The European Physical Journal B*, 75(3):395–404, 2010.

Henry E Brady. Comments on benford s law and the venezuelan election. *MS dated January*, 19:2005, 2005.

Jonathan R Cervas and Bernard Grofman. Tools for identifying partisan gerrymandering (pa). *Political Geography*, 2019.

Jowei Chen, Jonathan Rodden, et al. Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science*, 8(3):239–269, 2013.

Stephen Coate and Brian Knight. Socially optimal districting. Technical report, National Bureau of Economic Research, 2005.

Stephen Coate and Brian Knight. Socially optimal districting: An empirical investigation. Technical report, National Bureau of Economic Research, 2006.

Joseph Deckert, Mikhail Myagkov, and Peter C Ordeshook. Benford's law and the detection of election fraud. *Political Analysis*, 19(3):245–268, 2011.

John N Friedman and Richard Holden. Towards a theory of optimal partisan gerrymandering. *Mimeograph, Harvard University*, 2005.

John N Friedman and Richard T Holden. The rising incumbent reelection rate: What's gerrymandering got to do with it? *The Journal of Politics*, 71(2):593–611, 2009.

Delia Grigg and Jonathan N Katz. The impact of majority-minority districts on congressional elections. In *annual meeting of the The Midwest Political Science Association, Chicago, IL*, 2005.

Bernard Grofman and Gary King. The future of partisan symmetry as a judicial test for partisan gerrymandering after lulac v. perry. *Election Law Journal*, 6(1):2–35, 2007.

Faruk Gul and Wolfgang Pesendorfer. Strategic redistricting. *American Economic Review*, 100(4):1616–41, 2010.

Leroy C Hardy. Considering the gerrymander. *Pepperdine Law Review*, 4(2):3, 1977.

Raúl Jiménez and Manuel Hidalgo. Forensic analysis of venezuelan elections during the chávez presidency. *PloS one*, 9(6):e100884, 2014.

Raúl Jimenez, Manuel Hidalgo, and Peter Klimek. Supplementary materials for "testing for voter rigging in small polling stations", 2017a. URL `https://advances.sciencemag.org/content/advances/suppl/2017/06/26/3.6.e1602363.DC1/1602363_SM.pdf`.

Raúl Jimenez, Manuel Hidalgo, and Peter Klimek. Testing for voter rigging in small polling stations. *Science advances*, 3(6):e1602363, 2017b.

Peter Klimek, Yuri Yegorov, Rudolf Hanel, and Stefan Thurner. Statistical detection of systematic election irregularities. *Proceedings of the National Academy of Sciences*, 109(41):16469–16473, 2012.

Peter Klimek, Raul Jimenez, Manuel Hidalgo, Abraham Hinteregger, and Stefan Thurner. Election forensic analysis of the turkish constitutional referendum 2017. *arXiv preprint arXiv:1706.09839*, 2017.

Fabrice Lehoucq. Electoral fraud: Causes, types, and consequences. *Annual review of political science*, 6(1):233–256, 2003.

Michael D McDonald and Robin E Best. Unfair partisan gerrymanders in politics and law: A diagnostic applied to six cases. *Election Law Journal*, 14(4):312–330, 2015.

Walter R Mebane. Comment on "benford's law and the detection of election fraud". *Political Analysis*, 19(3):269–272, 2011.

Walter R Mebane Jr. Election forensics: Vote counts and benford's law. In *Summer Meeting of the Political Methodology Society, UC-Davis, July*, pages 20–22, 2006.

John Morrow. Benford's law, families of distributions and a test basis. 2014.

Ulrich Mücke. Elections and political participation in nineteenth-century peru: The 1871–72 presidential campaign. *Journal of Latin American Studies*, 33(2):311–346, 2001.

Erich Neuwirth and Walter Schachermayer. Some statistics concerning the austrian presidential election 2016. *arXiv preprint arXiv:1609.00506*, 2016.

Guillermo Owen and Bernard Grofman. Optimal partisan gerrymandering. *Political Geography Quarterly*, 7(1):5–22, 1988.

Lauren Payne-Riley. A deeper look at gerrymandering, 2017. URL `https://www.policymap.com/2017/08/a-deeper-look-at-gerrymandering/`.

Daniel D Polsby and Robert D Popper. The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale L. & Pol'y Rev.*, 9:301, 1991.

Ralph A Raimi. The first digit problem. *The American Mathematical Monthly*, 83(7):521–538, 1976.

Ernest C Reock. A note: Measuring compactness as a requirement of legislative apportionment. *Midwest Journal of Political Science*, 5(1):70–74, 1961.

John A Rice. *Mathematical statistics and data analysis*. Cengage Learning, 2006.

Nicholas R Seabrook. The limits of partisan gerrymandering: Looking ahead to the 2010 congressional redistricting cycle. In *The Forum*, volume 8. De Gruyter, 2010.

Kenneth W Shotts. Gerrymandering, legislative composition, and national policy outcomes. *American Journal of Political Science*, pages 398–414, 2002.

Nicholas O Stephanopoulos and Eric M McGhee. Partisan gerrymandering and the efficiency gap. *U. chi. l. Rev.*, 82:831, 2015.

John W Tukey. *Exploratory data analysis*, volume 2. Reading, Mass., 1977.

Peter Wagner. Beginning of the end for 'prison-based gerrymandering', 2012. URL `https://www.washingtonpost.com/opinions/beginning-of-the-end-for-prison-based-gerrymandering/2012/07/13/gJQAJP7fiW_story.html`.

Samuel S-H Wang. Three tests for practical evaluation of partisan gerrymandering. *Stan. L. Rev.*, 68: 1263, 2016.

Gregory S Warrington. Quantifying gerrymandering using the vote distribution. *Election Law Journal*, 17(1):39–57, 2018.

LC Washington. Benford law for fibonacci and lucas-numbers. *Fibonacci Quarterly*, 19(2):175–177, 1981.

List of Figures